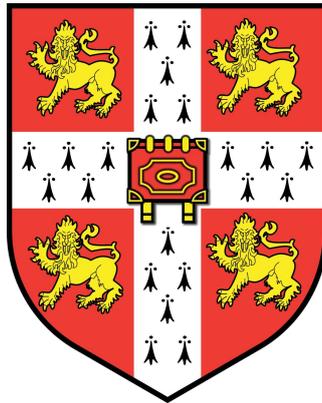


RAYLEIGH-TAYLOR MIXING: CONFINEMENT BY STRATIFICATION AND GEOMETRY

ANDREW GEORGE WEIR LAWRIE

SELWYN COLLEGE



A DISSERTATION SUBMITTED FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF APPLIED MATHEMATICS AND THEORETICAL PHYSICS

THE UNIVERSITY OF CAMBRIDGE

JUNE 2009

Preface

This thesis, which is submitted for the degree of Doctor of Philosophy at Selwyn College, University of Cambridge, describes work carried out from October 2005 to April 2009 in the Department of Applied Mathematics and Theoretical Physics, University of Cambridge. This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration, except where specifically indicated in the text. No part of this work has been, or is being submitted for any other qualification at this or any other university.

Acknowledgements

I would like to thank my supervisor, Dr. Stuart Bruce Dalziel, for his advice, foresight, support and kindness throughout, and for the many exciting opportunities he has placed in my path. I would also like to thank him for his tolerance of my verbosity, especially when it reaches its zenith on Tuesday evenings around 6pm. Others whom I'd like to acknowledge for their help in various ways are Prof. Malcolm Andrews, Dr. Nikos Nikiforakis, Dr. Robin Williams, Prof. David Youngs, Mr. David Page-Croft, Dr. Jean-Luc Weitor, and the funding agencies NERC and AWE.

Abstract

Rayleigh-Taylor instability has been an area of active research in fluid dynamics for the last twenty years, but relatively little attention has been paid to the dynamics of problems where Rayleigh-Taylor instability plays a role, but is only one component of a more complex system. Here, Rayleigh-Taylor instability between miscible fluids is examined in situations where it is confined by various means: by geometric restriction, by penetration into a stable linear stratification, and by impingement on a stable density interface. Water-based experiments are modelled using a variety of techniques, ranging from simple hand calculation of energy exchange to full three-dimensional numerical simulation. Since there are well known difficulties in modelling unconfined Rayleigh-Taylor instability, the confined test cases have been sequenced to begin with dynamically simple benchmark systems on which existing modelling approaches perform well, then they progress to more complex systems and explore the limitations of the various models. Some work on the phenomenology of turbulent mixing is also presented, including a new experimental technique that allows mixed fluid to be visualised directly, and an analysis of energy transport and mixing efficiency in variable density flows dominated by mixing.

Contents

1	Overview	1
1.1	Introduction	1
1.2	Historical perspective	2
1.2.1	Analytical studies	2
1.2.2	Experimental studies	4
1.2.3	Numerical studies	6
1.3	Ongoing research	8
1.4	Thesis outline	9
2	Experimental methods	10
2.1	Introduction	10
2.2	Low-aspect-ratio Rayleigh-Taylor instability	11
2.2.1	Equipment history and development	11
2.2.2	Operational details	16
2.3	High-aspect-ratio Rayleigh-Taylor instability	21
2.3.1	Equipment history and development	21
2.3.2	Operational details	23
2.4	Producing linear stratifications	24
2.4.1	Theoretical background	24
2.4.2	Practical implementation	25
3	Computational Methods	29
3.1	Introduction	29
3.2	Numerical implementation	31

3.2.1	High order advection	31
3.2.2	Source terms	35
3.2.3	Velocity projection and pressure correction	36
3.2.4	Boundary conditions	38
3.3	Algorithmic details	41
3.3.1	Multigrid convergence acceleration	41
3.3.2	Parallel implementation	44
4	Unconfined Rayleigh-Taylor instability	46
4.1	Introduction	46
4.2	Models for Rayleigh-Taylor instability	46
4.2.1	Early-time growth	46
4.2.2	Potential flow	48
4.2.3	Buoyancy-drag balance	49
4.2.4	Energy budget	51
4.2.5	Gradient diffusion	53
4.3	Summary	55
5	Direct Visualisation of Mixing	56
5.1	Introduction	56
5.2	Reactive Light Induced Fluorescence (RLIF)	56
5.2.1	Previous attempts	56
5.2.2	Chemical behaviour	58
5.2.3	Optical Decomposition	61
5.3	Experimental and numerical comparison	66
5.3.1	Qualitative observations	66
5.3.2	Growth profile	68
5.3.3	Molecular mixing	71
5.3.4	Turbulent structure	80
5.4	Summary	83
6	Mixing in confined geometries - simple stratifications	84
6.1	Introduction	84

6.2	Modelling approaches	86
6.2.1	Similarity model	87
6.2.2	Non-linear gradient-diffusion model	89
6.2.3	Full three-dimensional simulation	91
6.3	Validation against experiment	93
6.3.1	Overtuned configuration	93
6.3.2	Static reservoir configuration	95
6.4	Summary	98
7	Mixing in confined geometries - linearly stratified upper layer	99
7.1	Introduction	99
7.2	Model developments	101
7.2.1	Similarity model	101
7.2.2	Mass diffusion model	102
7.2.3	Energy transport model	103
7.2.4	Three-dimensional simulation	104
7.3	Comparison with experiment	105
7.3.1	Model cross-validation	105
7.3.2	Variation with Atwood number	111
7.4	Summary	115
8	Mixing confined by stable linear stratifications	117
8.1	Introduction	117
8.2	Growth law	118
8.3	Mixing efficiency	121
8.3.1	Concepts	121
8.3.2	Everywhere-unstable systems	123
8.3.3	Partially unstable systems	124
8.4	Comparison with simulation and experimental ensemble	127
8.4.1	Scalar transport	127
8.4.2	Mixing efficiency	130
8.5	Summary	137

9	Mixing confined by a stable density interface	138
9.1	Introduction	138
9.2	Previous theoretical work	139
9.3	Comparison of scalar transport and molecular mixing	142
9.3.1	Initial observations	142
9.3.2	Profile self-similarity	145
9.3.3	Time evolution of concentration field	146
9.3.4	Interfacial growth	147
9.4	Summary	153
10	Conclusions	155
10.1	Review	155
10.2	Themes and ideas	156
10.3	Remarks on future directions	160
10.4	Final thoughts	164
A	Validation exercises on MOBILE	165
A.1	Introduction	165
A.2	Grid convergence verification	166
A.2.1	Single-mode Rayleigh-Taylor instability	166
A.2.2	High-aspect-ratio Rayleigh-Taylor instability	170
A.3	Validation on other mixing problems	173
A.3.1	Stratified Kelvin-Helmholtz instability	173
A.3.2	Lock-exchange gravity currents	184
A.4	Summary	188

List of Figures

2.1	Low aspect-ratio Rayleigh-Taylor box, showing its barrier fully withdrawn.	11
2.2	Image sequence showing cross-section of the composite barrier while being withdrawn.	12
2.3	Image sequence showing composite barrier withdrawal from the tank. The chequered pattern is fixed in the reference frame of the cloth, and this remains stationary with respect to the tank. . . .	13
2.4	Iso-surface of density visualised with a semi-opaque cocktail of dyes. The nylon cloth covering the barrier can be seen in (a). The Atwood number is $A = 7 \times 10^{-4}$	15
2.5	Comparison of interface perturbations with two barriers of the same nominal design: (a) shows an incorrectly manufactured barrier which releases a Bickley jet into the Rayleigh-Taylor instability as it is withdrawn, and (b) shows a correctly manufactured barrier achieving a near-random perturbation.	16
2.6	Detail of the barrier seal, an O-ring (coloured black) which seals two PVC blocks (coloured grey) held together with clips (coloured red) against hydrostatic pressure when the barrier is closed. . . .	17
2.7	Sketch showing the user's technique for withdrawing the barrier.	18
2.8	Dissolved air is removed from solution by depressurisation in these stainless steel cylinders.	20
2.9	The high-aspect-ratio Rayleigh-Taylor instability experimental apparatus.	22

2.10	The density profile used to investigate stratification-confined Rayleigh-Taylor instability.	24
2.11	Comparison of experimentally measured stratification and numerical model at a range of double bucket recirculation rates.	27
3.1	Schematic showing how high order gradient is chosen and the quantity of fluid (shown in black) fluxed across an upwind face in time Δt	34
3.2	Stencil pattern for upwinded u momentum. The hollow circle is the point being updated, the black dots and vestigial arrows depict data needed when $u > 0$ and $v > 0$	35
3.3	Data in the black square in the domain of block C is needed by block A as ghost data for its corner. A two stage process shown by the arrows enables this data to be retrieved without any explicit connectivity between block A and C.	41
3.4	Schematic of the simple V-cycle multigrid algorithm, where each dot represents a smoothing operation. The optimal algorithm uses three iterations on each grid level, except on the coarsest grid, which is converged to tolerance.	43
3.5	Schematic showing in one space dimension the linear interpolation scheme used when projecting a coarse residual solution $\widehat{e}(\mathbf{x})$ onto a finer non co-located grid. Compared with a co-located scheme this yields a 12% reduction in overall cost.	44
4.1	Blue curves plot the buoyancy-drag growth model for several perturbation wavelengths, while the red curve is the parabolic envelope. Length- and time-scales are arbitrary.	50
4.2	Diagram illustrating the simple box model for the evolving density profile. There is an implied assumption in this model that fluid in the mixing zone quickly becomes well mixed.	52

5.1	Standard PLIF imaging yields no sub-pixel information about the local mixing state. The same intensity value can be obtained from fully mixed, and fully unmixed fluid.	57
5.2	Chemical structure of Acridine.	58
5.3	Recirculating pump, batch tank, output filter and settling bucket for production of saturated Acridine solution.	59
5.4	Recipe for two-layer RLIF experiments.	60
5.5	Bayer-filter spectra for CCD sensor used in the Uniq-Vision UC1830CL colour camera.	61
5.6	Spectral response of Perkin-Elmer arc-lamp fitted with a UV filtered parabolic reflector.	62
5.7	Normalised colour spectra for both blue and cyan-green emission states of Acridine.	63
5.8	Adsorption spectra for blue and cyan-green emission states of Acridine.	64
5.9	Map of Acridine spectral response as a function of volume fraction.	65
5.10	CCD sensor response to Acridine as a function of volume fraction.	66
5.11	Comparison between (a) traditional PLIF, and (b) new RLIF experimental diagnostics for visualising molecular mixing. Both images are taken at non-dimensional time $\tau = 0.44$ and the Atwood numbers are (a) $A = 1.5 \times 10^{-3}$, and (b) $A = 1 \times 10^{-3}$	67
5.12	Direct visualisation of molecular mixing using Acridine. The Atwood number is $A = 1 \times 10^{-3}$	69
5.13	Initial numerical scalar concentration fields for horizontal mid-planes of three simulations, with a length-wise transect to illustrate the perturbation in each case. The plot geometry is to scale.	71
5.14	Growth profiles from experimental ensemble, numerical simulations with various initial conditions, compared against an envelope of parabolae $0.03 < \alpha < 0.08$	72

5.15	Local molecular mixing fraction $\Theta_p(z, t)$ as calculated from planar slices taken from a MOBILE simulation ($z > 0.25$), and experiment ($z < 0.25$). Ideal quadratic growth is also shown for reference.	73
5.16	Global molecular mixing fraction $\hat{\Theta}_g(t)$ calculated using equation 5.9 on a recent simulation, a recent experiment, archived experimental data from Dalziel <i>et al.</i> (1999), and compared with published curves.	75
5.17	Typical histogram of volume fraction at the early stages of a mixing process. Basis functions for equation 5.9 are also shown.	76
5.18	Global molecular mixing fractions $\hat{\Theta}_a(t)$ using calculated from Fluorescein (Dalziel <i>et al.</i> (1999)) and Acridine experiments and a common simulation with synthetic Fluorescein and Acridine diagnostics.	78
5.19	Scatter plot of cross-sectional surfaces and cross-sectional volumes for an ensemble of Acridine experiments and a MOBILE simulation with a synthetic Acridine diagnostic.	80
6.1	Schematic depicting the complete evolution of $h(t)$, from the early linear stage through the non-linear self-similar regime to the region in which geometric confinement restricts Rayleigh-Taylor development.	85
6.2	Scalar concentration taken on the vertical mid-plane of a $64 \times 64 \times 2560$ MOBILE simulation. Sections are taken at (a) $t = 0$, (b) $t = 20s$, (c) $t = 40s$, (d) $t = 60s$, (e) $t = 80s$, (f) $t = 100s$, (g) $t = 120s$. The Atwood number is $A = 1.5 \times 10^{-3}$	90
6.3	The dashed line indicates the simulated reservoir, which is 5.7% by volume of the reservoir used in the experiment. Reservoir and tall tube are depicted approximately in perspective.	91
6.4	Scalar concentration taken on the vertical mid-plane of the reservoir. The Atwood number is $A = 7.3 \times 10^{-3}$. The greyscale colour scheme illustrates well how efficiently dye disperses in the computational reservoir.	92

6.5	Time-series image of the horizontally averaged scalar concentration as a function of height in the computational reservoir. The blue superimposed curve shows the time-evolution of mean density at the interface between reservoir and tall tube. $A = 7.3 \times 10^{-3}$.	94
6.6	Time-series images of horizontally averaged scalar concentration from (a) MOBILE , and (b) overturned tank experiment of Dalziel <i>et al.</i> (2008). Theoretical prediction of $h(t) \sim t^{\frac{2}{5}}$ is shown in black, scaled according to the Atwood number. $A = 1.5 \times 10^{-3}$.	95
6.7	Time-series images of horizontally averaged scalar concentration from (a) one-dimensional model, and (b) MOBILE simulation in the static reservoir configuration. The white curves show $h(t)$ calculated using the zero-dimensional model. The boundary condition for the one-dimensional model is taken from figure 6.5. $A = 7.3 \times 10^{-3}$.	96
6.8	Time-series images showing horizontally averaged scalar concentration from (a) one-dimensional model, and (b) experiment in the static reservoir configuration. The white curve shows the zero-dimensional prediction of $h(t)$. $A = 1.02 \times 10^{-2}$.	97
7.1	Initial stratification for stratified high-aspect-ratio experiment.	100
7.2	Initial stratification for run with initial Atwood number $A_0 = 10^{-2}$. The blue curve indicated the measured stratification, the red dashed lines indicate the neutral buoyancy height for reservoir fluid.	105
7.3	Time-series of horizontally averaged scalar concentrations as a function of height from (a) one-dimensional model and (b) three dimensional simulation. The white curve is the penetration height $h(t)$ obtained from the one-dimensional model, shown in both plots for comparison. A boundary condition correction for the small simulated reservoir of (b) has been made in (a).	106
7.4	Time-series image of scalar concentration in the reservoir predicted by MOBILE . The blue curve denotes the horizontally averaged density at the interface between the tall tube and the reservoir.	107

7.5	The red curves are one-dimensional predictions of the vertical scalar concentration profile, scaled for height and end-point concentrations. The blue curves come from the three-dimensional simulation.	108
7.6	Time-series of horizontally averaged scalar concentration as a function of height from (a) one-dimensional model, and (b) experiment. The white curve is the penetration height $h(t)$ obtained from (a), and plotted over (b) for comparison. A correction for dye non-linearity has been made in the experimental time-series.	109
7.7	This plot shows compares all the model predictions with the experimental timeseries, and a $t^{\frac{2}{5}}$ curve is plotted for reference. . .	110
7.8	One-dimensional predictions of envelope profiles $h(t)$ for experiments at various Atwood numbers.	111
7.9	Time-series extract showing double-diffusive steps when primarily salt-driven mixing occurs in the presence of a temperature mismatch between the water temperatures and laboratory ambient. $A_0 = 0.5 \times 10^{-2}$	113
7.10	One-dimensional model predictions of (a) idealised initial conditions across a range $2.5 \times 10^{-4} < A_0 < 0.1$, and (b) self-similar collapse by scaling with non-dimensional time.	114
8.1	Diagram illustrating the initial density profile. Note that in this case the neutral buoyancy height z_n coincides with z_{\max}	118
8.2	Experimental image sequence showing the Rayleigh-Taylor instability confined by linear stratification. The Atwood number is $A = 2 \times 10^{-3}$	120
8.3	This schematic shows the various permitted pathways for energy to be exchanged from one type to another.	122
8.4	Rayleigh-Taylor growth $h(t)$ evolving from an unstable interface sitting between two stable linear stratifications, theoretical and numerical predictions compared with an experimental ensemble.	129

8.5	Initial and final stratifications for an experiment with a stratification-confined Rayleigh-Taylor instability, compared with equivalent simulation and theoretical end-state gradients.	131
8.6	Time evolution of energy distribution in a stratification-confined simulation.	133
8.7	Time evolution of energy fluxes during (a) a stratification-confined simulation, and (b) a homogenous-layer simulation.	134
8.8	Time evolution of instantaneous and cumulative mixing efficiency for both homogenous-layer and stratification-confined Rayleigh-Taylor instability. The Atwood number of the interface is common to both simulations.	136
9.1	Diagram illustrating the initial density profile.	139
9.2	Middle layer scalar transport in the stable interface Rayleigh-Taylor problem. The field of view does not include the very bottom of the lower layer. The unstable interface Atwood number is $A = 2 \times 10^{-3}$, and stable interface Richardson number sits in the range $1.66 < Ri < 2.44$	144
9.3	Self-similar collapse of density profiles at various times, experimental and numerical comparison.	146
9.4	Decay of maximum concentration \hat{C} , plotted as an inverse quantity to compare analytical prediction with experiment and simulation.	147
9.5	Visualisation of molecular mixing across the unstable interface in Rayleigh-Taylor instability confined asymmetrically by a stable interface. The initial stratification and the times at which images are shown are identical to those in figures 9.2 and 9.6.	149
9.6	Visualisation of molecular mixing across the stable interface in Rayleigh-Taylor instability confined asymmetrically by this stable interface. Note that the tank-filling process induced mixing prior to release of the unstable interface and is thus identified by the diagnostic. The initial stratification and the times at which images are shown are identical to those in figures 9.2 and 9.5.	151

9.7	Integral measure of interfacial growth, comparison of experiment and simulation.	153
A.1	Single-mode bubble growth with Atwood number $A = 0.25$. Scalar concentration is represented in greyscale, and spatial resolution is $32 \times 32 \times 256$	167
A.2	Horizontal section of scalar concentration at mid-height plane, showing patterns of increasing complexity with time. The Atwood number is $A = 0.25$	169
A.3	Convergence of single-mode instability growth evolution with increasing resolution.	170
A.4	Convergence of single-mode instability growth velocity with increasing resolution.	171
A.5	Plot of instability growth against instability velocity.	172
A.6	Convergence of instability growth evolution with increasing resolution, and compared with experiment.	173
A.7	Demonstration of self-similarity of scalar concentration profile from MOBILE simulation at three resolutions, and compared with experiment.	174
A.8	Convergence of mean normalised scalar concentration profile with increasing resolution. MOBILE simulations are compared with experiment.	175
A.9	MOBILE simulations of Kelvin-Helmholtz instability comparing no interfacial perturbation (upper image of each pair) with one (lower image of each pair) perturbed to closely match an experiment. For ease of presentation the tank inclination is not shown.	177
A.10	Comparison of experiment and MOBILE simulation of Kelvin-Helmholtz instability. The tank tilt angle is 7.58° and the upper and lower layer densities in both cases are $\rho_u = 998.2 \text{kgm}^{-3}$, $\rho_l = 1012.5 \text{kgm}^{-3}$. The experimental sequence runs down the left column, and simulation down the right.	180

A.11 Time-evolution of total potential, background potential, kinetic and available energy in the tilting tank experiment, measured in the reference frame of gravity once the tank is stationary in its tilted state.	182
A.12 Time-evolution of total potential, available potential and background potential energy in Kelvin-Helmholtz billows, measured in the reference frame of the tank.	183
A.13 Boussinesq half-depth gravity current evolution predicted by MOBILE simulation, $\rho_{\text{left}} = 1000\text{kgm}^{-3}$, $\rho_{\text{right}} = 1003\text{kgm}^{-3}$	185
A.14 Time-series image of an asymmetric lock release with 4% depth ratio, showing progress of the front and internal characteristics reflected off the endwall.	187
A.15 Front velocity as a function of time for three depth ratios, and compared with appropriate theoretical predictions for early and late time.	188

Chapter 1

Overview

1.1 Introduction

The classic view of Rayleigh-Taylor instability is the unbounded growth of any perturbations initially present on an interface between initially quiescent fluids of different density when the interface is subjected to an acceleration normal to the interface in the direction from the light to the heavy fluid. The instability leads to the distortion of the interface and in miscible fluids, their mixing, driven by the baroclinic generation of vorticity which arises from the misalignment of pressure and density gradients, as the inviscid variable density vorticity equation,

$$\frac{D\boldsymbol{\omega}}{Dt} - (\boldsymbol{\omega} \cdot \nabla) \mathbf{u} = -\frac{1}{\rho^2} \nabla \rho \times \nabla p, \quad (1.1)$$

makes clear. The vorticity generated by the instability is initially concentrated at the interface between the two fluids, but in miscible fluids undergoing molecular mixing, the region over which there are large density gradients enlarges with time, and so too does the region generating vorticity.

Rayleigh-Taylor instability is one of the purest fluid systems in which molecular mixing can be initiated, since the instability exists independently of boundary conditions. Despite its apparent geometric simplicity, a horizontal density interface in free space gives rise to some of the most complex, important and least well-understood phenomena in classical mechanics. Once vorticity has been generated baroclinically, it non-linearly self-advects according to equation 1.1, progressively increasing its spatial complexity and leading to turbulence.

Lord Rayleigh identified an instability between fluids with different densities in a now famous publication of experimental work, [Rayleigh \(1883\)](#). Theoretical studies by G.I.Taylor were published some 67 years later ([Taylor \(1950\)](#)), and unaware of Rayleigh's previous work, Taylor's name became synonymous with the instability of two fluids of different density in an acceleration field; only more recently has the historical connection with Rayleigh's original work been recognised and in his honour Rayleigh-Taylor instability has been so named.

1.2 Historical perspective

1.2.1 Analytical studies

G.I. Taylor's work ([Taylor \(1950\)](#)) considered the following simplified thought experiment - two inviscid incompressible two-dimensional fluids in an unbounded domain, separated by a perfectly sharp, nominally horizontal interface, with small sinusoidal perturbations. If the amplitude were chosen to be much smaller than the wavelength of the instability, the equations of motion could be linearised. He showed that if the acceleration field is directed from the light to the heavy fluid, the perturbations grow exponentially in time. Following experimental work designed to verify Taylor's theoretical model, a diffuse interface correction was developed by [Duff *et al.* \(1962\)](#), which accounted for the reduced observed growth rate.

It was observed that the instability did not conform to linear theory beyond its very earliest stages, and a model was proposed by [Davies & Taylor \(1950\)](#) and [Layzer \(1955\)](#) to explain the development in the non-linear regime at later times. Using a potential flow approximation, this model studied the motion of a rising air bubble in a tall tube, and generalised this for a sinusoidal initial interface perturbation. This was the first body of work to show that rising bubbles reach an asymptotic velocity, and by extension Rayleigh-Taylor instability might tend towards linear growth.

Interest in Rayleigh-Taylor instability was revitalised when [Youngs \(1984b\)](#) predicted numerically that the instability follows quadratic growth in time (*i.e.* a linear increase in velocity), and [Read \(1984\)](#) confirmed this experimentally. On dimensional grounds this seems rather obvious: there is no alternative combination of reduced

gravity g' and time t which could yield a length scale h that does not grow as t^2 . By convention we write

$$h = \alpha A g t^2, \quad (1.2)$$

where α is a scaling factor, A is non-dimensional density,

$$A = \frac{\rho_u - \rho_l}{\rho_u + \rho_l}, \quad (1.3)$$

and is related to a reduced gravity in the Boussinesq limit by

$$g' = \frac{\Delta\rho}{\rho} g = 2Ag. \quad (1.4)$$

Despite the simplicity of this scaling (see §4.2.4 for a more rigorous derivation), superficially at least it contradicts Layzer's earlier prediction that $h \sim t$.

More recently Layzer's potential flow ideas have been developed to gain further insight into the phenomenology of Rayleigh-Taylor instability. [Linden *et al.* \(1994\)](#) simplified the potential flow approach to a simple force-balance relationship, accounting for the action of mutual buoyancy and drag forces on fluid bodies. Such models, known as 'buoyancy-drag' models, make assumptions about the shape of the interpenetrating structures and compute the force balance accordingly. These models predict exponential growth at early times when the drag force is negligible relative to the buoyancy, and thus matches Taylor's original theory. At later times, however, the drag becomes more significant and eventually exactly balances the drag, leading to an asymptotic velocity and associated linear growth. Layzer's and Taylor's ideas are thus reconciled, but neither appears to predict quadratic growth.

The most important distinction between the conditions that give rise to late-time quadratic growth and those that give linear growth is the modal character of the initial perturbation. In general, real interfacial perturbations are multi-modal and the (non-linear) interaction of these modes when they develop changes the scale of the most-unstable mode, and hence changes the growth rate. A number of multi-mode models have been constructed which attempt to capture some of the experimentally observed modal dynamics. So-called 'bubble competition' and 'bubble merger' models (*e.g.* [Zufiria \(1988\)](#); [Ofer *et al.* \(1996\)](#); [Rikanati *et al.* \(2000\)](#)) assume the development of each mode can be described independently by a buoyancy-drag model, and a

chosen interaction mechanism couples the modes together and provides a mechanism for the dominant scale to change. Some models use a ‘takeover’ analogy whereby smaller bubbles are engulfed, others demand that the bubble structures maintain geometrical self-similarity throughout their growth and in some way ‘merge’. By choosing carefully how the modes interact, the correct (quadratic) overall growth behaviour can be recovered, and some insight into the Rayleigh-Taylor phenomenology can be deduced.

1.2.2 Experimental studies

Lord Rayleigh’s initial attempts at observing the instability experimentally (Rayleigh (1883)) used warm salty dyed water supported by a porous membrane above cold fresh water. Thermal diffusion of heat to the atmosphere then allowed the upper layer to become cool, and reach parity of buoyancy with the lower layer. Under further cooling, the upper layer started to migrate into the lower layer in thin vertical finger-like strands. Rayleigh believed he was observing the baroclinically driven process we know today as Rayleigh-Taylor instability. However, he had inadvertently (see Schmitt (1995) for a historical review) discovered a diffusion driven process which we know today as salt fingering. The finger-like structures occur because heat diffuses more efficiently than salt between upper and lower layer fluid, making lower layer fluid warm and relatively buoyant in between the fingers, and upper layer fluid in the fingers colder and heavier and thus perpetuating the instability.

The earliest experimental work, Lewis (1950); Emmons *et al.* (1959), that faithfully demonstrated the baroclinically driven instability shortly followed Taylor’s analytical study (see §1.2.1). Interest was not revived until there became a technological imperative to develop the basic science. The linear regime, which predicts exponential growth of the instability, is valid where perturbations have large wavelength relative to their amplitude, and this describes the Rayleigh-Taylor flow only at very early times in its evolution. The non-linear behaviour that develops at later times presents a much richer scientific problem and recent attention has focussed on this latter regime.

One of the many complications in the non-linear regime is that self-advection

of vorticity leads to the interface (interfacial region in miscible fluids) becoming increasingly convoluted in space. Small initial perturbations become accentuated due to the baroclinic vorticity they supply to themselves, and they organise into self-advected structures known as ‘bubbles’ if they contain relatively less dense fluid and are rising upwards, and ‘spikes’ if they are more dense. Bubble and spike structures eventually overturn due to their own vorticity, and acquire a mushroom-like appearance, before individual structures break down into much less well-defined shapes at later times. While the bubbles and spikes remain coherent structures, the growth rate of the instability decreases - as predicted by the theoretical models balancing buoyancy and drag. Only under very carefully controlled conditions, eg. [Waddell *et al.* \(2001\)](#); [Wilkinson & Jacobs \(2007\)](#), can this process be observed in isolation; a much more common occurrence is to have a wide spectrum of modes in the initial perturbation, and this growth rate decrease is not observed.

[Read \(1984\)](#) made the first experimental study of Rayleigh-Taylor instability that demonstrated a growth law in the non-linear regime of the form $h = \alpha A g t^2$, and work since then has been dominated by attempts to quantify the rate coefficient α . Read’s work used rocket propulsion to reverse the acceleration field on a quiescent gravitationally stable density stratification. Similar work has been done more recently with a linear electric motor (*e.g.* [Dimonte & Schneider \(1996\)](#)), compressed gas driving a piston (*e.g.* [Nevmerzhitsky *et al.* \(1994\)](#)) and using high-energy lasers (*e.g.* [Robey *et al.* \(2001\)](#)). Simply using gravity itself to act on an initially unstable stratification has also been thoroughly investigated in the G. K. Batchelor Laboratory (*e.g.* [Linden *et al.* \(1994\)](#); [Dalziel *et al.* \(1999\)](#)).

Despite - or perhaps because of - the range of experimental approaches taken, there has as yet been no convincing settlement on a unique value of α . As numerical simulations have developed over the same period, they have been predicting ever lower values of α , and concern has grown that our scientific understanding of the fluid mechanics is in some way deficient. The simple arguments used to construct the $h = \alpha A g t^2$ growth law embed assumptions about the nature of the turbulent mixing, particularly that on a horizontally averaged basis, kinetic energy and density fields have a self-similar (invariant with h) form. [Dalziel *et al.* \(1999\)](#) establishes

experimentally that these assumptions are indeed valid, though the experimental configuration used was not appropriate for making measurements of α with a high degree of statistical confidence. Snider & Andrews (1994) built a water tunnel with two inlet streams (one for each density), from which a spatially evolving steady state Rayleigh-Taylor instability was created and much more precise growth rate measurements have been made with this apparatus. Even with such accurate measurements, the discrepancy between numerical and experimental estimates of α has continued to widen; only very recently has it been appreciated that the modal structure of the initial conditions plays a significant role in determining the subsequent growth of the instability. Inevitably in experiments this is poorly controlled, whereas in simulations as mesh resolutions have increased, prescribed perturbation wavenumbers have been increasing in consort.

1.2.3 Numerical studies

Numerical simulation has become an increasingly useful tool in studying Rayleigh-Taylor instability, and its use on the problem dates back to Youngs (1984b), which predicted quadratic growth and matched the parallel experimental work of Read (1984). The growth of computing capacity in the intervening years has allowed simulation to explore the instability in ever-greater detail, particularly in regimes that cannot be easily reached by experiment. However there has been a persistent, and increasing concern that simulated growth rates were generally lower than those measured in experiments. The growth rates in the published literature lie in the range $0.01 < \alpha < 0.07$ - variation almost of an order of magnitude, it should be noted - with experiments congregating at the upper end. In an attempt to resolve these concerns, an international collaboration was established, known as the ‘Alpha Group’ Dimonte *et al.* (2004), which simulated a series of simple well-defined test problems aimed at definitively establishing a universal rate coefficient. On these test cases, with few exceptions, the computer codes predict similar growth rates, and this confirmed that in general the spread of α values was not caused by algorithmic variation.

Indeed as computing capacity expanded, the disparity between numerics and

experiment appeared to widen, and a view has only recently emerged [Ramaprabhu *et al.* \(2006\)](#); [Mikaelian \(2008\)](#) that the spectral distribution of the initial condition might have a role in determining the rate coefficient. As computational grid sizes have increased, the resolvable wavenumber range has widened, so this view is gaining credibility.

What makes such assertions hard to verify, are the respective limitations of experiment and simulation. Interfacial perturbations cannot generally be prescribed *a priori* in experiments, so any direct comparison with simulation is marred by uncertainty in matching initial conditions. Experimentally measurable features of the initial condition have in [Dalziel *et al.* \(1999\)](#) been used to improve the match with simulation, but this approach necessarily has its limitations. A more subtle problem arises when simulating Rayleigh-Taylor instability. As the instability grows, its Reynolds number increases with the size of the largest structures, and the length scale of the smallest eddies decreases in consort. Quite quickly, these become inadequately resolved, and this modifies the turbulent energy flux that drives energy from large scales to dissipation at small scales. Direct Numerical Simulation (DNS) aims to completely resolve all these scales, but the overwhelming computational cost of doing so in any reasonable scenario restricts this method exclusively to scientific rather than practical utility.

A whole field of research has grown to try and mitigate the net (large scale) effects of under-resolving dynamical processes at small scales, and thus maximise the utility of computation. However the algorithmic details are handled, the general aim is to get the right rate of energy flux leaving the resolved scales. Approaches range in complexity from inventing small scale velocity fields using a self-affine projection of the resolved fields (see [Meneveau & Katz \(2000\)](#)), through traditional local eddy-viscosity methods ([Smagorinsky \(1963\)](#)) to doing no explicit modelling of the small scales ([Margolin *et al.* \(2006\)](#)). The ‘do nothing’ approach is known formally as Implicit Large Eddy Simulation (ILES), and is increasingly acquiring acceptance in the Rayleigh-Taylor community since it gives surprisingly accurate predictions despite its simplicity.

1.3 Ongoing research

Current work on the Rayleigh-Taylor instability seeks to build on the vast literature and datasets that now exist, and refine our understanding of the growth behaviour so that predictive models can be applied with confidence. Long-held assumptions are being challenged - notably that α may not be a universal constant, and may depend (as discussed in §1.2.3) on the initial perturbation, the density difference (Wilkinson & Jacobs (2007)), the miscibility of the fluids (Mikaelian (2008)), and the Reynolds number (Cabot & Cook (2006)). Indeed $h \sim t^2$ may not be the correct functional form at all according to Cook *et al.* (2004); one of the highest resolution simulations (Cabot & Cook (2006)) has identified a four-stage growth process: independent modal growth, weak turbulence, mixing transition, strong turbulence. Current numerical simulations are only sufficiently well-resolved to glimpse the third stage, so it seems likely that there remain answers to be uncovered to questions we have not yet even posed.

Given that the growth behaviour - the most elementary statistic of the instability - has not been quantified to our satisfaction, it is less surprising that other critical questions also remain unresolved. Predicting the degree of molecular mixing taking place in Rayleigh-Taylor instability has long been one of the motives for its study, and while in a limited parameter regime this is straightforward to measure experimentally (*e.g.* Linden *et al.* (1994); Lawrie & Dalziel (2006a); Andrews *et al.* (2007)), in some applications we do not have the capability to do such measurements. How mixing behaviour is changed by density gradients surrounding a Rayleigh-Taylor unstable interface is another important question that has hitherto received little attention (excepting Jacobs & Dalziel (2005)), and one which this thesis works towards addressing.

That such a geometrically simple and elegant problem has given rise to decades of intensive research, and yet even now the most basic laws are a matter of serious debate, is a testament to the phenomenal complexity of Rayleigh-Taylor instability.

1.4 Thesis outline

The remainder of this thesis is organised as follows: chapter 2 introduces the experimental apparatus and basic laboratory techniques used to examine the Rayleigh-Taylor instability and charts the history of their development. Chapter 3 explains the numerical algorithms used in MOBILE, the Implicit Large Eddy Simulation tool developed for and used in the study of Rayleigh-Taylor instability throughout this thesis. Chapter 4 provides an introduction to existing modelling approaches for Rayleigh-Taylor instability that will be used and developed in subsequent chapters.

Chapter 5 discusses a new experimental diagnostic technique called RLIF, which has been developed to directly visualise molecular mixing in Rayleigh-Taylor instability. Comparison with existing experimental methods and with simulation is also considered here. Chapter 6 investigates the Rayleigh-Taylor instability confined by geometry in a high-aspect-ratio domain, a benchmark test case selected because it reduces dynamical complexity and a variety of modelling approaches perform well. Again in a high-aspect-ratio domain, chapter 7 investigates Rayleigh-Taylor instability penetrating into a stable linear stratification, a scenario widely believed to pose a challenge for numerical models. Chapter 8 considers the Rayleigh-Taylor instability sandwiched between two stable linear stratifications in a low-aspect ratio domain, and examines energy transport and mixing efficiency in the system. Chapter 9 studies asymmetric Rayleigh-Taylor instability, where it is confined on one side by a stable density interface.

Chapter 10 reviews the thesis, draws together the main ideas and outlines relevant further work. The appendix contains four numerical studies designed to demonstrate grid convergence and validate MOBILE on a range of fluid problems, non-Boussinesq single-mode Rayleigh-Taylor instability, high-aspect-ratio Rayleigh-Taylor instability, stratified Kelvin-Helmholtz instability, and lock-exchange gravity currents.

Chapter 2

Experimental methods

2.1 Introduction

The rationale behind all experiments contributing this thesis is this: the techniques used to create the flow, and the techniques used to observe the flow should avoid modifying the physics of the problem being studied.

The regime in which the Rayleigh-Taylor instability has been studied here was chosen with care to be tractable for mathematical modelling. In this regard, density differences were selected to study the Rayleigh-Taylor instability under conditions of ‘high’ Reynolds number turbulence while still satisfying the Boussinesq criterion. Further, geometric simplicity has been preserved, both by apparatus design and avoidance of intrusive measurement techniques.

The main experimental techniques used here were Planar Light-Induced Fluorescence (PLIF) and so-called Dye-attenuation. Both are optical diagnostics that exploit passive scalar transport of a dye and are appropriate for revealing aspects of the mixing process - the theme of this thesis.

Subsequent sections are organised as follows: §2.2.1 discusses the evolution of the apparatus that has historically been used to study Rayleigh-Taylor in the G. K. Batchelor Laboratory, focussing on scientific and phenomenological aspects; §2.2.2 covers the details pertaining to modern experimental techniques in this apparatus. §2.3 introduces apparatus used to study Rayleigh-Taylor instability under conditions of geometric confinement, and §2.4 analyses the techniques used in this thesis to



Figure 2.1: Low aspect-ratio Rayleigh-Taylor box, showing its barrier fully withdrawn.

create complex vertical profiles of density for experiments in both sets of apparatus.

2.2 Low-aspect-ratio Rayleigh-Taylor instability

2.2.1 Equipment history and development

The PLIF experiments in this thesis were conducted using developments of apparatus which have been applied previously by Dalziel (1993); Dalziel *et al.* (1999); Jacobs & Dalziel (2005). Some methods (e.g. Waddell *et al.* (2001)) can induce a finite, systematic and repeatable perturbation of the interface, but historically the research aim of experiments in this apparatus has been to study Rayleigh-Taylor instability in its purest form - where the interface is minimally and randomly perturbed over a wide spectrum of wavelengths.

In the simplest configuration to give rise to a Rayleigh-Taylor instability driven flow, relatively dense fluid sits above a horizontal barrier located at the mid-plane position of the tank. Relatively less dense fluid sits beneath the barrier in a chamber of equal volume. When fully closed, the barrier sits in a tight-fitting horizontal groove on three of the four sides, and is pulled out of the tank through a letterbox-shaped seal on the fourth side to initiate the experiment. The acrylic container used

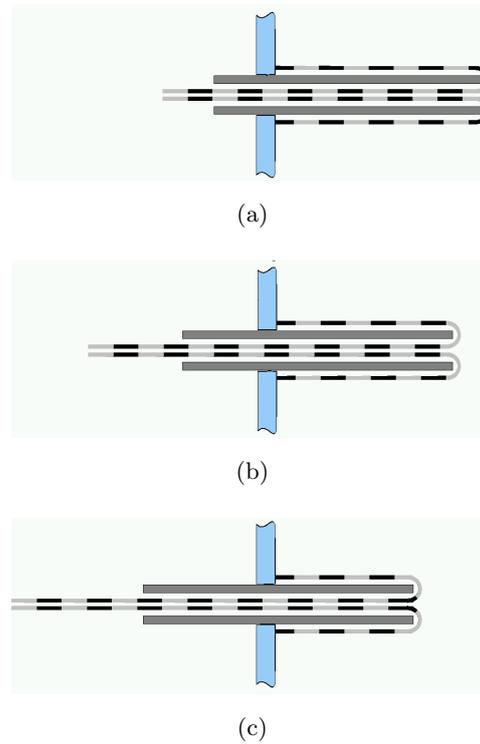


Figure 2.2: Image sequence showing cross-section of the composite barrier while being withdrawn.

for the experiment is shown in figure 2.1.

Unfortunately the withdrawal of a horizontal barrier has been shown to have a significant and unwelcome influence on the evolution of the flow, and much development work since this method of initiating Rayleigh-Taylor instability was adopted has been focussed on mitigating the unwanted behaviour.

A thin metal barrier has been used in the past, notably by [Linden *et al.* \(1994\)](#). Unfortunately, while the barrier is being withdrawn from the tank, momentum of fluid adjacent to it becomes non-zero, and under the action of viscosity, diffuses outwards to form a boundary layer above and below the barrier. This contaminates the Rayleigh-Taylor interface with momentum. Additionally however, this boundary layer (of moving fluid) is stripped off the barrier surface by the seal at the point where the barrier leaves the tank, and is deflected by the vertical endwall. By considering continuity, it is obvious that the action of withdrawing the barrier induces a circulation cell in both upper and lower chambers of the tank.

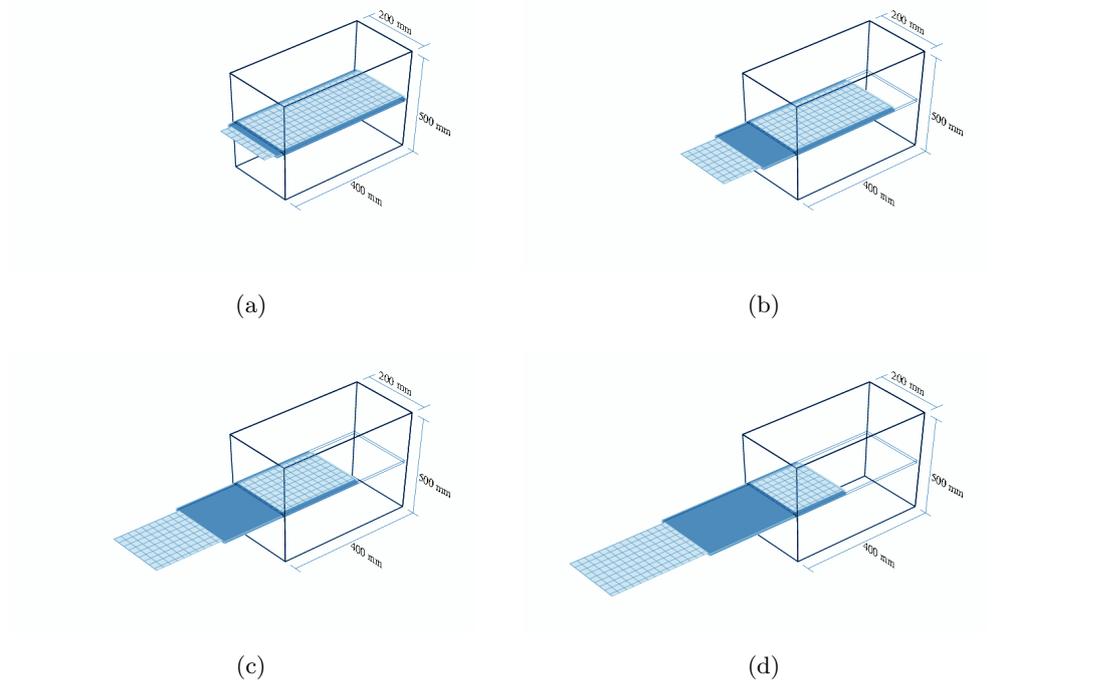


Figure 2.3: Image sequence showing composite barrier withdrawal from the tank. The chequered pattern is fixed in the reference frame of the cloth, and this remains stationary with respect to the tank.

One desires a method of withdrawing the barrier that avoids inducing boundary layer development, and a composite barrier solution was devised by Lane-Serff (1989). The metal part of the barrier comprises two horizontal sheets separated vertically to form a wide flat tube. A sheet of polyester fabric is fixed to the end face of the tank structure through which the barrier is withdrawn, lies on the upper surface of the tube, and folds inside it, passing out of the tank through the barrier at the withdrawing end. A second piece of fabric treats the lower surface in the same way (see diagram 2.2). Shear due to withdrawing the barrier is confined between fabric and metal elements, rather than at the surface exposed to the fluid, which, apart from at the barrier's very tip, is fixed in the reference frame of the tank throughout the withdrawal. A sketch of the process is indicated in figure 2.3.

An unavoidable complication is the effect that a barrier of non-negligible thickness (2.4mm) has on the overall dynamics of the flow. The inviscid response of the flow to the removal of a finite volume from the mid-height of the box is to establish a

pressure gradient, and this induces upper layer fluid to descend - the upper surface being left free to air to allow the adjustment - and it acquires momentum as it does so. The fluid beneath, being incompressible, cannot respond symmetrically, and this gives rise to a net circulation. Figure 2.4(a) illustrates the asymmetry at early time. That this circulation persists at later times too is simply a statement of angular momentum conservation. The strength of the circulation is controlled by the barrier velocity profile. If the initial acceleration is too sudden, the angular momentum generated can come to dominate the flow evolution and lead at later times to a bulk overturning of the flow, so it may exhaust its potential energy when the Rayleigh-Taylor instability has still only partially evolved. Incidentally, deceleration at the end of travel needs also to be smooth; bumping the barrier on the endstops of its guide rails sends a highly disturbing pressure wave into the fluid. Optimally, the barrier removal takes 2-4 seconds and by controlling the tension on the polyester cloth and resisting this with pressure on the metal part of the barrier, a smooth velocity profile is achieved.

Manufacturing the composite, hollow barrier to achieve an approximately random initial perturbation proved to be a further complication. In figure 2.5 a high amplitude single mode can be observed, which is compared and contrasted with a close to random perturbation generated from a correctly manufactured barrier. As the barrier is withdrawn it is squeezed through the endwall seal, and if its internal volume is significant, then any flexibility in the metal leads to small quantities of fluid inside being forced out through its open end, into the tank as a small flow-rate planar jet. The jet has a characteristic modal oscillation (see [Bickley \(1937\)](#)), and the Rayleigh-Taylor interface is unstable to this mode.

The design of the sealing for the barrier has recently (in 2007) undergone considerable modification to greatly reduce its leakage. Traditionally, neoprene rubber sliding seals were pressed against the moving surfaces where the barrier exits the tank, but it proved impossible to seal the sharp corners, resulting in a steady leak. With more complex stratifications being studied, the barrier leakage gives rise to a very important component of error in the vertical density profile. A secondary seal was developed so that the neoprene rubber was not being asked to withstand

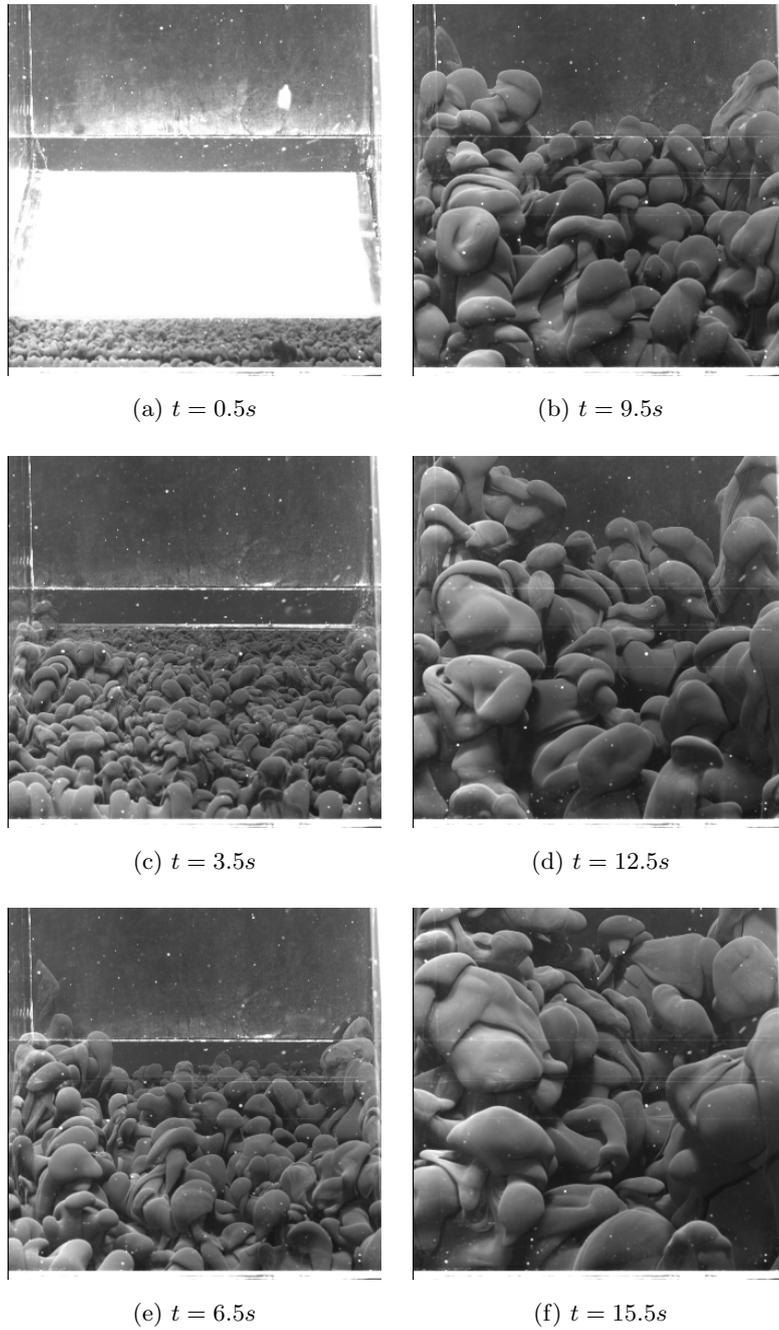
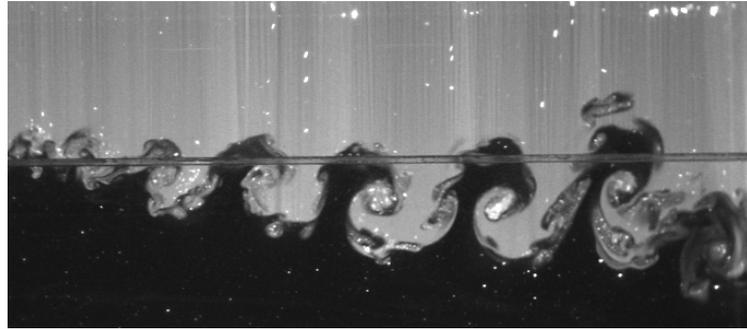


Figure 2.4: Iso-surface of density visualised with a semi-opaque cocktail of dyes. The nylon cloth covering the barrier can be seen in (a). The Atwood number is $A = 7 \times 10^{-4}$.

the full hydrostatic pressure head. This seal is a rubber O-ring encircling the barrier (see figure 2.6), held in a groove on a PVC block attached to the barrier, and



(a)



(b)

Figure 2.5: Comparison of interface perturbations with two barriers of the same nominal design: (a) shows an incorrectly manufactured barrier which releases a Bickley jet into the Rayleigh-Taylor instability as it is withdrawn, and (b) shows a correctly manufactured barrier achieving a near-random perturbation.

sealing against another PVC block fixed to the tank. When the barrier is clipped shut for long periods (as required when creating complex stratifications) only the baroclinic head can drive a leakage flow - through either the tank sidewall grooves or the endwall sliding seal, and since this obviously scales with the density difference, the Boussinesq problems studied here tend not to suffer from significant leakage.

2.2.2 Operational details

The photograph in figure 2.7 shows how the low-aspect-ratio Rayleigh-Taylor apparatus is operated. The user's right hand pulls the nylon cloth from the tank, and the force this exerts on the metal part of the barrier causes it to slide out in a manner akin to a first-order pulley system. The user's left hand provides a resistive (into



Figure 2.6: Detail of the barrier seal, an O-ring (coloured black) which seals two PVC blocks (coloured grey) held together with clips (coloured red) against hydrostatic pressure when the barrier is closed.

tank) force to maintain tension in the nylon cloth during the removal.

The sides of the tank were covered with a matt black book-binding material on both internal and external surfaces, since this limited the optical contamination due to total internal reflection of incident beams. The viewing sidewall and endwall were not permanently covered, though the viewing wall not in operation was covered in matt black insulating tape on the inside surface wherever the light sheet impinged, and a removable cover fitted to the outside face. Experiments contributing to this thesis used overhead illumination, and the base of the tank is partially covered in a Maltese cross arrangement. Thus both longitudinal and transverse oriented vertical light sheets could pass through the base of the tank without reflecting, meanwhile stray reflections back into the tank were minimised by having black covering elsewhere on the base. A mirror directed the exiting light sheet away at an oblique angle.

Waves on the free water surface, initiated by the removal of the barrier, were found to have a negligible effect on the dynamics of the flow, but a non-negligible effect on the illumination. To achieve a consistent refraction of the incident light from above as it entered the tank, the surface needed to be covered, yet allowed to adjust to the change in volume caused by removing the barrier. A transparent



Figure 2.7: Sketch showing the user's technique for withdrawing the barrier.

floating lid was found to be adequate although for practical reasons a tank-fixed transparent plate, only partially covering the surface and submerged by 3mm was used instead. Comparison of these devices with a free surface showed no significant change in the internal flow. Light shielding was used to prevent peripheral light, both ambient, and directly from the beam, from entering the top of the tank.

The light source was a 700W xenon arc lamp mounted vertically downwards producing light across a spectrum ranging from 300nm to 1200nm and passed through a 350nm long-pass dichroic reflector for safety reasons. An optical path of 2000mm, was the maximum obtainable given constraints of ceiling height. Beam divergence, in this case, was such that the light sheet thickened from 0.25mm at the optical guide slit resting on the submerged transparent plate, to 4mm over a ray path dis-

tance of 500mm. Conveniently, the change in refractive index between air and water tended to reduce the divergence angle. The light sheet thickness at the position of the barrier was 2.1mm.

Data acquisition was via a UNIQVision 1830C-12B-CL series digital CCD Bayer mosaic colour camera, which used the CameraLink interface protocol and a BitFlow R64 framegrabber PCI card to acquire images. DigiFlow software passed the data to a RAID hard disk array in real time. Typically f0.95 25mm focal length C-mount Vortex lenses were used; fast lenses permitted the use of thin light sheets, and/or less efficient fluorophores. The 25mm focal length was appropriate when viewing the whole domain; for close-up detail photography, a 50mm lens was used. Depth-of-field and parallax problems were not encountered since detailed imaging work concentrated on imaging a single, thin light sheet.

The removal of the finite-thickness barrier induces a very low pressure at its tip, and this can cause dissolved air to come out of solution, form rising bubbles, and interfere with optical diagnostics. To prevent this, the water was depressurised in two 54-litre stainless-steel cylinders (see figure 2.8). Each has a acrylic lid and was sealed with a rubber gasket and a smearing of silicon grease. A venturi pump was connected to a small air space left between the water surfaces and the acrylic lids.

Optical differences between layers can pose a major problem when conducting PLIF experiments, since fluoresced light from the illuminated plane may have to pass a large distance through fluid that does not necessarily have a constant refractive index. Dissolved salt - the customary agent of stratification in water - increases both the density and refractive index. Propan-2-ol has the interesting property that it reduces the density while *increasing* the refractive index. Thus a small quantity of Propan-2-ol added to the less dense layer successfully matches refractive indices of the two layers, and so the fluoresced light is not appreciably refracted as it traverses through the tank. Where PLIF was performed on flows with density profiles based on multiple linear stratifications, each stratification was stably stratified in salt and unstably stratified in alcohol. Approximately 3ml Propan-2-ol matches the refractive index change caused by 1g salt in each litre of water.

Stratification profiles were measured using a conductivity probe calibrated to read



Figure 2.8: Dissolved air is removed from solution by depressurisation in these stainless steel cylinders.

density (where measured, salt was the sole stratifying agent). The probe had two coaxial electrodes, and water was drawn by capillary siphon up through a 0.3mm hole in an electrically insulating acetyl tip to contact the internal electrode. The resistivity of the water was measured on the path between internal and external electrode, but since the field lines focus at the small hole, the measurement is dominated by the conductivity of the water being drawn through the hole. The probe traversed downwards through the fluid driven by software through a PCI DAQ card using a constant speed motor. Microswitches on an end buffer cut the motor drive and the probe output signal to the PCI card.

The stratification was achieved using a double-bucket arrangement (Fortuin (1960)), filling slowly through a tube suspended just above the tank base, and with the least dense water entering first. An analysis of the issues is detailed in §2.4.

2.3 High-aspect-ratio Rayleigh-Taylor instability

2.3.1 Equipment history and development

The dynamics of Rayleigh-Taylor instability are profoundly affected by geometric confinement. When unconfined, one feature of Rayleigh-Taylor instability is that the characteristic bubble and spike (aspect-ratio $O(1)$) structures increase in length scale with time. Once these structures reach parity with the domain horizontal length scale, they cannot continue to develop, and the growth rate of the instability is thus constrained.

This effect was observed when conducting experiments in a tall thin tube ($5\text{cm} \times 5\text{cm} \times 2\text{m}$). Previous work in this apparatus (Dalziel *et al.* (2008)) has focussed on overturning a two-layer stable density stratification in the tube, to initiate Rayleigh-Taylor instability. The work for this thesis focussed on the erosion of a linear stable density stratification, so the thin tube, containing the stratification, was held static and the instability was initiated by releasing fluid from a large (165 litre) reservoir beneath (see figure 2.9) which was less dense than the tube base density. The reservoir was dyed red with food colouring (3ml Red Fiesta), and the attenuation of a backlight illumination due to the dye as it progressed up the tube was measured.

The release mechanism was a simple swivelling gate, initially clipped shut against the full hydrostatic head. The disturbance introduced by opening the gate inevitably affected the flow at early time, but only until the dominant length scale of the instability reached parity with the domain horizontal length scale. Thereafter, the dynamics become increasingly insensitive to the initial conditions. At very late time - the main focus of this work - the separation of time scales is so vast that the initialisation method has a negligible effect on the flow.

Some unexpected dynamics were observed at late time, particularly in experiments with a relatively low initial Atwood number. One would expect the interface marking the maximum progress of dyed reservoir fluid up the tube to decelerate towards zero when approaching the vertical height in the stratification whose density corresponds to that of the reservoir. In fact, the interface actually accelerated again, driven by exchange flows between thin layers present in the stratification. The for-



Figure 2.9: The high-aspect-ratio Rayleigh-Taylor instability experimental apparatus.

mation of these layers was caused by the differential diffusion of heat and salt (see [Turner \(1974\)](#)). Unfortunately there were a number of possible causes inherent in the experimental environment:

1. Heat was added to the fluid in the tube by the pumping apparatus used to set up the stratification.
2. Heat may have been added by red dye absorbing IR radiation from the back-light.
3. Water used in the experiment, particularly in the reservoir, may not have reached thermal equilibrium with ambient before use.

4. Ambient temperature variation in the laboratory was approximately 1°C over a period of an experiment, and a thermal stratification of 2.5°C existed over the 2m height of the tube.

The tube inevitably has a high surface to volume ratio, so fluid could respond to the thermal gradients across its surfaces within the time-scale of an experiment. Any combination of these factors would give rise to a double-diffusive staircase. To mitigate the ambient thermal stratification, desk fans were used to circulate air around. To reduce incident radiation, the backlight was moved a good distance from the tank and an optical diffuser and heat shield placed in between. An exhaustive study of the time water was left at various points in the experimental procedure to equilibrate with ambient temperature - with and without evaporation permitted - showed that double-diffusive effects were minimised when the time water was exposed to laboratory conditions was minimised. However, the effects of double-diffusion were never completely eliminated.

2.3.2 Operational details

Images were acquired with an Atmel-Grenoble Cameila 8M CCD monochrome camera with an LCD timing shutter between the lens and the CCD sensor. The CameraLink protocol was used to transfer data from the camera to an R2 BitFlow framegrabber PCI card, and DigiFlow software was used to write the data to a RAID array of hard disks in real time. The camera was positioned 10m away from the tube to minimise image-processing errors due to parallax.

The stratification was achieved using a double-bucket arrangement (Fortuin (1960)), filling slowly from the base with the least dense water entering first. Air escaped from the tube via a small release valve in the lid. Obtaining consistent stratifications was problematic, and an analysis of the issues is detailed in §2.4.

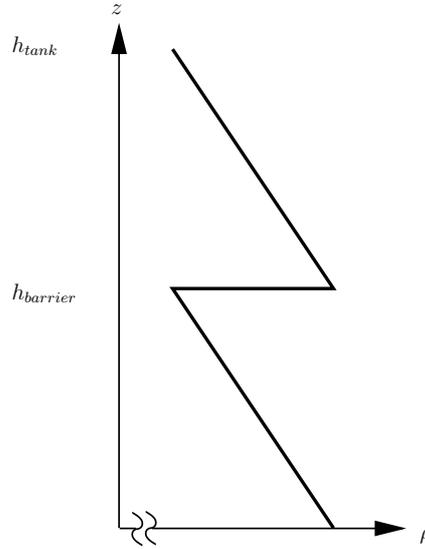


Figure 2.10: The density profile used to investigate stratification-confined Rayleigh-Taylor instability.

2.4 Producing linear stratifications

2.4.1 Theoretical background

Two layer Rayleigh-Taylor instability has been studied extensively in various contexts, and some previous work has been done on multi-layer configurations (Jacobs & Dalziel (2005)). Rayleigh-Taylor instability had not been investigated in configurations where it is confined by stable linear stratifications. As figure 2.10 shows, the configuration studied involves two identical linear stable stratifications arranged one above the other, with a Rayleigh-Taylor unstable interface between them.

The linear stratifications were created using a technique first developed by Fortuin (1960). The principle works as follows: the hydrostatic pressures in two equally filled geometrically identical buckets supported at the same height are equal (neglecting any baroclinic head), so provided the two buckets are connected, removing volume at a constant rate Q from one bucket - here called the mixing bucket - to fill a tank causes a reduction in water surface height, and therefore a loss in hydrostatic pressure at the base of the bucket. A pressure gradient is therefore maintained to drive a flow at a rate $\frac{Q}{2}$ from the other bucket - here called the supply bucket. This new volume of fluid may contain a tracer, possibly salinity, at concentration

c_s , and provided there is sufficient agitation of the mixing bucket the average tracer concentration in the mixing bucket c_m has incremented.

The above process leads to a pair of differential equations describing the volume of the mixing bucket and the average quantity of the tracer in the mixing bucket as functions of time:

$$\frac{\partial V c_m}{\partial t} = \frac{1}{2} Q c_s - Q c_m \quad (2.1)$$

$$\frac{\partial V}{\partial t} = -\frac{1}{2} Q. \quad (2.2)$$

Expanding equation 2.1 with the product rule, integrating the volume equation, with integration constant V_0 , and rearranging, we have

$$\left(V_0 - \frac{1}{2} Q t \right) \frac{d c_m}{d t} = \frac{1}{2} Q (c_s - c_m). \quad (2.3)$$

Written as

$$\frac{d c_m}{d t} + \left(\frac{-\frac{1}{2} Q}{V_0 - \frac{1}{2} Q t} \right) c_m = \left(\frac{\frac{1}{2} Q c_s}{V_0 - \frac{1}{2} Q t} \right), \quad (2.4)$$

this is in the standard form,

$$\frac{d y}{d x} + p(x) y = q(x) \quad (2.5)$$

that has a well-known solution

$$y = \frac{\int e^{\int p(x) dx} q(x) dx}{e^{\int p(x) dx}}, \quad (2.6)$$

Thus we arrive at a linear function,

$$c_m = -c_s + k \left(V_0 - \frac{1}{2} Q t \right), \quad (2.7)$$

for $c_m(t)$, where k is constant of integration. The next section explores the practical implementation of this principle.

2.4.2 Practical implementation

Previous stratified flow experiments in the G. K. Batchelor Laboratory (*e.g.* Higginson *et al.* (2003); Scase & Dalziel (2004, 2006)) used double buckets designed to operate as described in §2.4.1. A constant flow rate pump provided both kinetic energy for mixing, and a head for an outlet flow. When the bucket is shallow, the

maximum flow rate at which the pump can operate is limited by the entrainment of air by a vortex at the pump intake. When the bucket is full, the volume flux through the pump is a very small proportion of the volume, and this sets a minimum bound on the rate at which mixing can take place. For stratifications to be approximately C_2 continuous, the pump flow rate and control valve settings must be constant while stratifying.

When filling a tank from the bottom, the pump has to overcome an increasing hydrostatic head (2m in the case of the high-aspect-ratio Rayleigh-Taylor tank), and this decreases the flow rate into the tank and changes the rate of recirculation in the mixing bucket. It proved impossible even with a range of control valves to select a pump flow rate that both delivered fluid against a 2m hydrostatic head and satisfied the maximum bound on inlet flow rate. To avoid this issue, a separate peristaltic pump fed directly from the mixing bucket was used to deliver a well-controlled volume flow rate to the tank irrespective of the head, and the main pump was dedicated to recirculation only.

Unfortunately the stratifications produced using this method did not always behave as expected. In particular, although the mean gradient was consistent between nominally identical experiments, the absolute values of the densities were not. This was especially noticeable in the high-aspect-ratio Rayleigh-Taylor experiments where density differences were relatively high, and further investigation of these inconsistencies was warranted.

The analysis in §2.4.1 assumes instantaneous mixing of fluid in the mixing bucket, and clearly this is not an accurate model. A starting point for a modification could be to use a previous value of c_m as an outflux tracer concentration:

$$\frac{\partial V c_m(t)}{\partial t} = \frac{1}{2} Q c_s - Q c_m(t - \tau), \quad (2.8)$$

where τ is representative of the mixing time. To track this intermediate concentration analytically, a third ODE is required, and to circumvent the need to solve an increasingly complex system, a simple numerical representation of the turbulent mixing process in the bucket - impractical to model directly - was conceived. The experimental apparatus used a constant-volume-flux recirculating pump to provide kinetic energy for the required mixing. It was modelled as a single flow path from

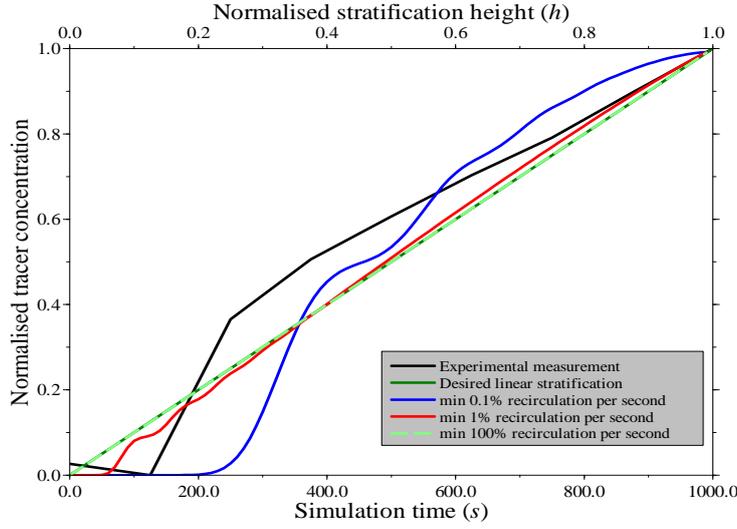


Figure 2.11: Comparison of experimentally measured stratification and numerical model at a range of double bucket recirculation rates.

the outlet of the pump, through the bucket to the pump intake, thus the problem was reduced to one dimension. Inflow from the supply bucket is assumed to mix with fluid directly leaving the pump outlet and extraction of fluid to fill the tank was taken off at the pump intake.

Numerically, a time-explicit finite-volume representation of the bucket tracks both bucket volume and tracer concentration. The bucket volume adjusts uniformly over each cell during each timestep, and the tracer is advected passively, hence the update equation has the form

$$c_i^{n+1} = \frac{1}{V_{new}} \left(V_{old} c_i^n + \Delta t \left(F_{i-\frac{1}{2}} - F_{i+\frac{1}{2}} \right) \right) \quad (2.9)$$

where V_{new} and V_{old} are new and old cell volumes respectively. The fluxes are upwinded according to the velocity, but because dV leaves the bucket per unit time, and only $\frac{1}{2}dV$ enters, to satisfy continuity the velocity must increase linearly along the bucket flow-path. Hence the fluxes are defined as

$$F_{i+\frac{1}{2}} = c_i^n \left(\frac{\Delta R}{\Delta t} + \frac{1}{2} \frac{\Delta V}{\Delta t} \left(1 + \frac{l_i}{l_{max}} \right) \right) \quad (2.10)$$

where ΔR is the volume recirculated per unit time, and l_{max} is the total flow-path

length. The upstream and downstream boundary fluxes are

$$F_{upstream} = \frac{1}{2}c_s \frac{\Delta V}{\Delta t} + c_{i=max} \frac{\Delta R}{\Delta t} \quad (2.11)$$

$$F_{downstream} = c_{i=max} \left(\frac{\Delta V}{\Delta t} + \frac{\Delta R}{\Delta t} \right) \quad (2.12)$$

Evolving the model in time reveals how a curved stratification profile can be expected, since the proportion of fluid volume being driven through the pump increases as the total volume in the mixing bucket falls. At low - practical - rates of recirculation $O(1\%$ of initial volume per unit time), damped oscillations of the profile are observed in simulation. These correspond in frequency to those observed in experiment. See figure 2.11 for a comparison of the model and a sample stratification from an experiment.

From this it is clear that the double bucket technique does not deliver perfect linear stratifications, and since the oscillations identified above arise from a highly turbulent mixing process in the bucket, it is unsurprising to find that there exists some variability in the initial density being fed to the tank.

Chapter 3

Computational Methods

3.1 Introduction

A computer program called MOBILE was developed to model the experimental configurations studied in this thesis. The following observations help identify a suitable modelling strategy. Fluids used in the experiment were liquid-based, and hence, at least to an acceptable approximation, incompressible. Rayleigh-Taylor instability is driven by spatial variations in density, so in simulation the density field must be tracked. Since the fluids used in the experiment have low surface energy when in contact with one another, they easily inter-diffuse and are termed ‘miscible’. This greatly simplifies the computational problem, since any interfacial surface has a negligible influence on the dynamics of the flow, and therefore need not be tracked explicitly. A suitable governing equation set is therefore

$$\begin{aligned}\frac{\partial c_k}{\partial t} + \frac{\partial (c_k u_j)}{\partial x_j} &= 0, \\ \frac{\partial (\rho u_i)}{\partial t} + \frac{\partial (\rho u_i u_j)}{\partial x_j} &= -\frac{\partial p}{\partial x_i} + \rho g_i, \\ \rho \frac{\partial u_i}{\partial x_j} \frac{\partial u_j}{\partial x_i} &= -\frac{\partial^2 p}{\partial x_i^2}, \\ \Sigma c_k &= 1, \\ \rho &= \Sigma (c_k \rho_k),\end{aligned}\tag{3.1}$$

where c_k is the volume fraction of the k 'th advected tracer which has a prescribed density ρ_k . Note that the Boussinesq approximation has been made and that the

constraint of incompressibility is enforced by adjustment of the pressure field to conserve volume implicitly.

For many years it has been argued (*e.g.* Margolin *et al.* (2006)) that the quite considerable existing numerical technology developed for accurately capturing shock systems in the compressible Euler equations can be applied when studying variable density incompressible turbulence, since there is a coincidence of requirements to provide sharp resolution of steep gradients. Typically numerical schemes whose truncation error scales with a high power of the grid cell size capture such steep gradients, but without special treatment they induce entropy violating oscillations in the surrounding region.

A class of methods known as ‘Total Variation Bounded’, prevent spurious oscillations by smoothly interpolating between a high order upwind-biased scheme in regions of the domain which are deemed sufficiently smooth, and a first order upwind scheme where flow variable curvatures are high. Inevitably where the scheme behaves as first order, there is some compromise in accuracy locally, and formally the global accuracy cannot be higher than first order. However, the additional numerical dissipation arising from using the lower order scheme is localised around regions of high curvature, and in nature viscous forcing $\mu \nabla^2 \mathbf{u}$ is significant only in these regions, so it is becoming accepted in some fields to employ numerical error as a valid proxy for physical viscosity. This is the semi-empirical basis for a class of numerical approaches known as Implicit Large Eddy Simulation (ILES), and this method has been adopted in the current work.

Despite having been used on this empirical basis for fifteen years, only very recently has there been a determined effort to formalise the above rationale for using ILES. Its attraction is algorithmic simplicity and correspondingly low computational cost, since explicit viscous terms that scale with $\frac{1}{(\Delta x)^2}$ are not computed, and unlike conventional LES models (*e.g.* Smagorinsky (1963); Pullin (2000); Meneveau & Katz (2000)), there is no semi-empirical sub-grid-scale model. Particularly in applications with a wide separation of scales and where it is therefore unfeasible to directly compute the smallest of them, such as meteorology (Smolarkiewicz *et al.* (2007)) and Rayleigh-Taylor driven flows (*e.g.* Dalziel *et al.* (1999); Hahn & Drikakis (2005);

Thornber *et al.* (2008); Rider (2007)), ILES has a distinct advantage. Inevitably it is less suitable for viscous boundary layers and other diffusion dominated flows, and performs less well than Direct Numerical Simulation (DNS) on the rare problems where DNS is affordable.

3.2 Numerical implementation

3.2.1 High order advection

A fractional step approach, closely following Andrews (1995) has been used to decouple the governing equation set 3.1 into hyperbolic and elliptic components; for both, well-developed numerical techniques exist. This section discusses the discretisation of the hyperbolic element.

The weak form of the conservation laws,

$$\iiint \frac{\partial \phi}{\partial t} dV + \iint \phi \mathbf{u} \cdot \mathbf{n} dS = 0 \quad (3.2)$$

is enforced locally over each grid cell for each conserved variable ϕV . By Stokes theorem, a polyhedral finite-volume discretisation over the domain enforces the weak form globally. The upwind flux function is evaluated at each cell boundary face, and subsequently these values are used to advect conserved variables from cell to cell over the course of one timestep. The accumulated contributions from neighbouring cells are included in a volume-weighted average of the conserved variable over the cell. This is an unsplit approach and remains numerically stable provided the total volume fluxed over a timestep from neighbouring cells does not exceed the volume of a grid cell. This method has a truncation error that scales with the timestep size - known as a first order scheme.

For hexahedral discretisation, as implemented in MOBILE, the multi-dimensional problem can be decomposed into a sequence of one-dimensional update instructions. This is known as a fractional step method. The most recently updated values of the conserved variable are used to compute fluxes in the next direction in the sequence. Strang (1968) noticed that the particular sequence X-Y-Z-Z-Y-X has a truncation error that scales with the square of the timestep, provided the complete sequence takes place within one timestep (*i.e.* before application of boundary conditions,

source terms and pressure corrections). If the Strang splitting were implemented literally, global numerical stability could only be achieved by running each directional splitting at or below half of its optimal timestep. Although not formally second order, a popular adaptation involves running at the optimal timestep, and spreading the Strang splitting over two timesteps. Empirical evidence suggests that in this case truncation error accumulates less quickly. MOBILE uses this adaptation.

The one-dimensional update instructions each individually must satisfy the weak form of the linear advection equation in each cell:

$$\phi_i^* = \left(\phi_i^n V_i^n + F_{i-\frac{1}{2}}^n - F_{i+\frac{1}{2}}^n \right) / V_i^{n+1}, \quad (3.3)$$

where in this case superscript n represents time level, and subscript i represents spatial location. The updated value of the conserved variable, ϕ_i^* , is used as input for the next one-dimensional step in the sequence. The flux function F is evaluated at the cell faces (located at $i + \frac{1}{2}$ and $i - \frac{1}{2}$ in space) by treating the piecewise constant cell average values as initial conditions for a local Riemann problem on the cell face. For the linear advection equation, here shown in its weak form,

$$\iiint \left(\frac{\partial \phi}{\partial t} + u \frac{\partial \phi}{\partial x} \right) dV = 0 \quad (3.4)$$

the wave system reduces to a single wave (the contact discontinuity) so Godunov's first order solution (Godunov (1959)) to the Riemann problem is evaluated directly using a local estimate of the advection velocity u and the values of ϕ available on either side of a cell face.

Once again, in strong parallel with shock-capturing compressible schemes, extension to higher order is achieved here by solving the Riemann problem on a modified set of left and right ϕ states. Such schemes employ a piecewise polynomial reconstruction of the spatial field $\phi(x)$ from the cell mean values. The simplest higher order scheme uses piecewise linear reconstruction, and the overall scheme accuracy in this case scales with the square of the grid cell size. The method is known as MUSCL extrapolation (Collela (1985)).

The most obvious gradient to choose for a linear reconstruction in the cell at x_{i-1} has the form

$$m_{i-1} = \frac{\phi_i - \phi_{i-2}}{2\Delta x}. \quad (3.5)$$

To maintain conservation, the mean value of ϕ_{i-1} must remain unchanged. The reconstructed profile is advected with fluid velocity $u_{i-\frac{1}{2}} > 0$ in time Δt , so the actual value of ϕ which is fluxed across the cell face is the mean value over a small distance $u_{i-\frac{1}{2}}\Delta t$ on the upwind side of the $x_{i-\frac{1}{2}}$ cell face.

A yet higher order estimate of the gradient, due to Youngs (1984a), can be constructed by using the fluxed volume $u_{i-\frac{1}{2}}\Delta t$ as a weighting to bias the m_{i-1} gradient as far towards a central difference over the cell face as stability permits. This more sophisticated gradient has the form

$$m_{i-1} = \left(1 + \frac{u_{i-\frac{1}{2}}\Delta t}{\Delta x}\right) \frac{\phi_{i-1} - \phi_{i-2}}{3} + \left(2 - \frac{u_{i-\frac{1}{2}}\Delta t}{\Delta x}\right) \frac{\phi_i - \phi_{i-1}}{3}. \quad (3.6)$$

Unfortunately, as with all higher order schemes employing a *linear* combination of local values of ϕ (see Godunov (1959) for background), oscillations arise around steep gradients because the truncation error *modifies* the equation the discretisation was intended to solve. To illustrate this, consider the following conditionally stable discretisation of the linear advection equation, which converges at a rate $O(\Delta x^2, \Delta t^2)$, and is nominally equivalent to the scheme developed from equation 3.5 :

$$\frac{3\phi_i^n - 4\phi_i^{n-1} + \phi_i^{n-2}}{\Delta t} + \frac{3\phi_i^n - 4\phi_{i-1}^n + \phi_{i-2}^n}{2\Delta x} = 0. \quad (3.7)$$

By approximating the terms ϕ_{i-1}^n , ϕ_{i-2}^n , ϕ_i^{n-1} and ϕ_i^{n-2} with Taylor's expansions in space and time as appropriate, then substituting those into the update scheme 3.7, it is clear that early terms in the series cancel out, but later terms do not. Noting that for sufficiently smooth ϕ , higher time derivatives can be recast as spatial derivatives, we see that the modified equation has the form

$$\frac{\partial \phi}{\partial t} + u \frac{\partial \phi}{\partial x} = \left(\frac{2u\Delta x^2}{3} + \frac{2u^3\Delta t^2}{3}\right) \frac{\partial^3 \phi}{\partial x^3} + \dots, \quad (3.8)$$

which has a dispersive source term at leading order. The approximate solution to the linear advection equation with this numerical scheme will be contaminated with oscillations around sharp gradients, and this has become known as Gibbs behaviour. Physically, its occurrence implies that such systems can become increasingly ordered, and clearly this violates the Second Law of Thermodynamics. To circumvent this we locally select a numerical method that satisfies the entropy inequality (*i.e.* dissipative leading order error) wherever we expect oscillatory behaviour to emerge.

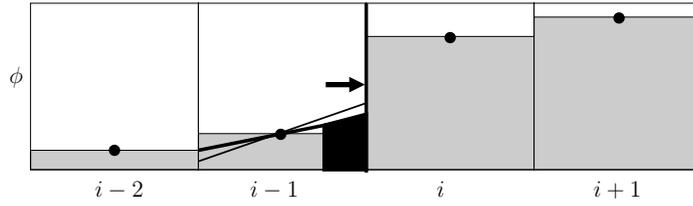


Figure 3.1: Schematic showing how high order gradient is chosen and the quantity of fluid (shown in black) fluxed across an upwind face in time Δt .

Interpolation between low and high order schemes is achieved in the current work by preventing the MUSCL reconstructed field $\hat{\phi}(x)$ from having new extrema relative to the piecewise constant function $\bar{\phi}(x)$ in neighbouring cells. Figure 3.1 illustrates a geometrical argument for selecting a function that interpolates between numerical schemes whose accuracy scales with high and low order powers of the grid cell size. In the figure, the piecewise constant function $\bar{\phi}(x)$ is shown in grey fill, and a higher order, MUSCL reconstructed field is shown for the upwind cell $i - 1$. The medium thickness line is a second order slope that preserves the total quantity of some scalar ϕ in the cell, but this produces an unphysical breach of monotonicity at $i - \frac{3}{2}$. A modified slope, which conserves the cell volume integral and also satisfies monotonicity, is shown by the heavy line. The higher order estimate of the quantity of ϕ fluxed across the cell face in time Δt is given by the volume shown in black, and is a function of velocity, as indicated by the length of the arrow.

The above scheme for multi-dimensional advection is employed to update both scalar quantities and components of momentum, though to avoid exciting the $2\Delta x$ grid mode while maintaining high order accuracy, the discrete velocity components are defined on the cell faces, while the scalars are defined at cell centres. Thus, $\rho V u$ is conserved in cells of volume V indexed by $(x_{i-\frac{1}{2}}, y_j, z_k)$, $\rho V v$ is conserved in cells $(x_i, y_{j-\frac{1}{2}}, z_k)$, and $\rho V w$ in $(x_i, y_j, z_{k-\frac{1}{2}})$. Thus the full computation uses four mutually staggered, distinctly indexed grids to solve all velocity and scalar advection. The transverse derivative terms, *e.g.* $\frac{\partial(\rho w)}{\partial x}$ incorporate velocity and density values computed on all other grids into the current grid, hence fully coupling adjacent values and reducing the gain of the $2\Delta x$ mode. A schematic showing the grid configuration and the stencil for upwind horizontal momentum advection is shown in figure 3.2. The grid of solid lines depicts the scalar cell boundary, and the grid of dashed lines

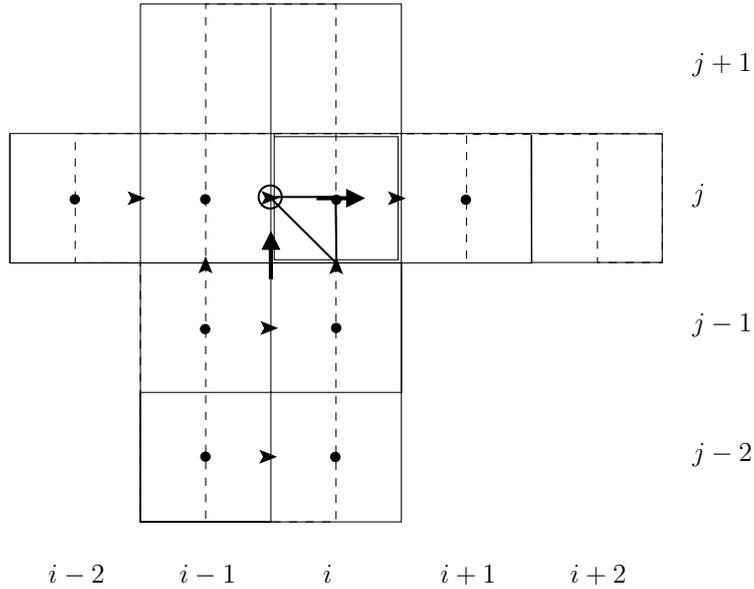


Figure 3.2: Stencil pattern for upwinded u momentum. The hollow circle is the point being updated, the black dots and vestigial arrows depict data needed when $u > 0$ and $v > 0$.

depicts the offset grid of u momentum cells. The black coloured dots and vestigial arrows show where density and velocity information is retrieved from the data array, in the case where the flow direction is from the bottom left corner. The cell outlines show the extent of the stencil should the upwind direction be reversed. The large arrows show the fluxes of u momentum that are computed using the current stencil information for a given (i, j) ordinate, and the large hollow circle shows the position that will receive an updated u momentum value from the newly computed fluxes. The small triangle indicates how the data is mapped from cell faces and cell centres to an (i, j) index location in the data array.

3.2.2 Source terms

The fluid is forced by the application of source terms to the (face centred) velocity field. Obviously the internal forcing is due to the pressure field applied across the faces S of the cell,

$$u_{new} = (\rho V u_{old} + \Delta t S (p_{i-1} - p_i)) / (\rho V), \quad (3.9)$$

and for variable density problems, the volumetric forcing due to buoyancy

$$u_{new} = (\rho V u_{old} + \Delta t V g_x (\rho - \rho_b)) / (\rho V), \quad (3.10)$$

over the cell volume V . No Boussinesq approximation is made, but to reduce the numerical error associated with performing arithmetic operations on numbers with large differences in magnitude, the mean hydrostatic pressure is debited from the computed pressure field, and only the remaining baroclinic head forces the flow.

For problems in which the reference frame rotates (see appendix A) with an angular velocity $\boldsymbol{\Omega}$, volumetric forcing due to coriolis and centripetal terms is included,

$$\mathbf{u}_{new} = (\rho V \mathbf{u}_{old} + 2\Delta t \rho V (\boldsymbol{\Omega} \times \mathbf{u}_{old})) / (\rho V) \quad (3.11)$$

$$\mathbf{u}_{new} = (\rho V \mathbf{u}_{old} + \Delta t \rho V (\boldsymbol{\Omega} \times \boldsymbol{\Omega} \times \mathbf{x})) / (\rho V). \quad (3.12)$$

Where $\boldsymbol{\Omega}$ is non-constant, an additional update,

$$\mathbf{u}_{new} = \left(\rho V \mathbf{u}_{old} + \Delta t \rho V \left(\frac{\partial \boldsymbol{\Omega}}{\partial t} \times \mathbf{x} \right) \right) / (\rho V), \quad (3.13)$$

is required. The updated velocity field now incorporates changes due to momentum advection and the additional accelerations caused by pressure gradients, buoyancy, and rotation of the reference frame. However, this new field is very unlikely to be exactly divergence free, and discussion of how this is restored follows.

3.2.3 Velocity projection and pressure correction

When solving the equations 3.1 numerically, the hyperbolic fractional step is constrained to conserve mass and momentum, but does not necessarily conserve volume. The velocity field \mathbf{U}^* obtained from the hyperbolic fractional step and the source term updates must be constrained to satisfy $\nabla \cdot \mathbf{U} = 0$, which is an elliptic problem. Therefore a projection is required to map the intermediate vector field \mathbf{U}^* onto the space of divergence free fields \mathbf{U}^{div} . There is no unique projection, and by definition none can conserve linear momentum, but one that conserves angular momentum is highly desirable, since the projection is transparent to the flow vorticity. For this reason the Helmholtz Decomposition is used:

$$\mathbf{U}^* = \nabla \Phi + \nabla \times \Psi \quad (3.14)$$

$$\mathbf{U}^* = \mathbf{U}^{no\ curl} + \mathbf{U}^{no\ div} \quad (3.15)$$

where the decomposed fields have the following properties,

$$\nabla \cdot \mathbf{U}^{no\ div} = 0 \quad (3.16)$$

$$\nabla \times \mathbf{U}^{no\ curl} = 0, \quad (3.17)$$

and we can thus define a velocity potential,

$$\nabla \Phi = \mathbf{U}^{no\ curl}. \quad (3.18)$$

Taking the divergence of equation 3.15 we have

$$\nabla \cdot \mathbf{U}^* = \nabla^2 \Phi + \nabla \cdot \mathbf{U}^{no\ div}. \quad (3.19)$$

By property 3.16 we have a Poisson equation for the velocity potential Φ . In the continuum limit, Φ can be characterised as a Lagrange multiplier on the divergence-free constraint. In an incompressible fluid there is no constitutive relation from kinetic physics to fix the value of pressure, but given an arbitrary value at some point in the domain, a field can be defined, and its gradient provides a force on the fluid. This force accelerates the fluid in a pressure-like manner such that volume is conserved, according to,

$$\frac{\partial \mathbf{U}}{\partial t} = -\nabla \Phi. \quad (3.20)$$

Algorithmically, we wish to satisfy volume conservation discretely,

$$\begin{aligned} & \frac{\Delta t \Delta y \Delta z}{\Delta x \Delta y \Delta z} \left(u_{i+\frac{1}{2}}^{n+1} - u_{i-\frac{1}{2}}^{n+1} \right) + \\ & \frac{\Delta t \Delta x \Delta z}{\Delta x \Delta y \Delta z} \left(v_{j+\frac{1}{2}}^{n+1} - v_{j-\frac{1}{2}}^{n+1} \right) + \\ & \frac{\Delta t \Delta x \Delta y}{\Delta x \Delta y \Delta z} \left(w_{k+\frac{1}{2}}^{n+1} - w_{k-\frac{1}{2}}^{n+1} \right) = 0, \end{aligned} \quad (3.21)$$

and obtain the required velocities by applying pressure as a surface force to each (staggered) momentum cell. The acceleration update we require is of the form

$$u_{i-\frac{1}{2}}^{n+1} = u_{i-\frac{1}{2}}^* + \frac{\Delta t \Delta y \Delta z}{\frac{1}{2}(\rho_{i-1} + \rho_i) \Delta x \Delta y \Delta z} (p_i^{n+1} - p_{i-1}^{n+1}), \quad (3.22)$$

where $u_{i-\frac{1}{2}}^*$ is the intermediate velocity field obtained after the hyperbolic step, and the p^{n+1} field is implicitly defined. Assuming we have a reasonable initial estimate of the pressure field p^n we can express the correction to the pressure as

$$p^{n+1} = p^n + \Delta p. \quad (3.23)$$

Thus we can substitute our expression for u^{n+1} into our equation for volume conservation:

$$\begin{aligned}
& \frac{\Delta t \Delta y \Delta z}{\Delta x \Delta y \Delta z} \left(u_{i+\frac{1}{2}}^* + \frac{\Delta t \Delta y \Delta z}{\frac{1}{2}(\rho_{i+1} + \rho_i) \Delta x \Delta y \Delta z} (p_{i+1}^n + \Delta p_{i+1} - p_i^n - \Delta p_i^n) \right. \\
& \quad \left. - u_{i-\frac{1}{2}}^* - \frac{\Delta t \Delta y \Delta z}{\frac{1}{2}(\rho_{i-1} + \rho_i) \Delta x \Delta y \Delta z} (p_i^n + \Delta p_i - p_{i-1}^n - \Delta p_{i-1}^n) \right) + \\
& \frac{\Delta t \Delta x \Delta z}{\Delta x \Delta y \Delta z} \left(v_{j+\frac{1}{2}}^* + \frac{\Delta t \Delta x \Delta z}{\frac{1}{2}(\rho_{j+1} + \rho_j) \Delta x \Delta y \Delta z} (p_{j+1}^n + \Delta p_{j+1} - p_j^n - \Delta p_j^n) \right. \\
& \quad \left. - v_{j-\frac{1}{2}}^* - \frac{\Delta t \Delta x \Delta z}{\frac{1}{2}(\rho_{j-1} + \rho_j) \Delta x \Delta y \Delta z} (p_j^n + \Delta p_j - p_{j-1}^n - \Delta p_{j-1}^n) \right) + \\
& \frac{\Delta t \Delta x \Delta y}{\Delta x \Delta y \Delta z} \left(w_{k+\frac{1}{2}}^* + \frac{\Delta t \Delta x \Delta y}{\frac{1}{2}(\rho_{k+1} + \rho_k) \Delta x \Delta y \Delta z} (p_{k+1}^n + \Delta p_{k+1} - p_k^n - \Delta p_k^n) \right. \\
& \quad \left. - w_{k-\frac{1}{2}}^* - \frac{\Delta t \Delta x \Delta y}{\frac{1}{2}(\rho_{k-1} + \rho_k) \Delta x \Delta y \Delta z} (p_k^n + \Delta p_k - p_{k-1}^n - \Delta p_{k-1}^n) \right) = 0.
\end{aligned} \tag{3.24}$$

The pressure correction field Δp is the only unknown, and so the form of the equation becomes evident:

$$\begin{aligned}
& W_{i+\frac{1}{2}} (\Delta p_i - \Delta p_{i+1}) + W_{i+\frac{1}{2}} (\Delta p_i - \Delta p_{i-1}) + \\
& W_{j+\frac{1}{2}} (\Delta p_j - \Delta p_{j+1}) + W_{j-\frac{1}{2}} (\Delta p_j - \Delta p_{j-1}) + \\
& W_{k+\frac{1}{2}} (\Delta p_k - \Delta p_{k+1}) + W_{k-\frac{1}{2}} (\Delta p_k - \Delta p_{k-1}) = \nabla \cdot \mathbf{U}^{**}
\end{aligned} \tag{3.25}$$

where W is a weighting function based on local density, cell face area and cell volume, and \mathbf{U}^{**} is a notational contraction for the \mathbf{U}^* velocity field updated by accelerating with the existing pressure field p^n (as implied by equation 3.24). Noting that i, j, k and n indices have been suppressed where possible, it is clear that equation 3.25 is elliptic, with a naturally arising density-weighted 7 point stencil in three dimensions. In a uniform density flow, the weights are equal, and equation 3.25 reduces to a standard discrete Poisson equation. Solving by Successive Over Relaxation (SOR), or similar, is an obvious solution approach.

3.2.4 Boundary conditions

Several simple types of boundary condition have been written, namely

- Slip wall / symmetry

- No-slip wall
- Inflow
- Transmissive
- Periodic / internal ,

to maximise the flexibility of the code. These conditions are implemented as suitably chosen values in three rows of ‘ghost cells’ sitting outside the domain. Each block, or grid, is initialised with updated ghost cell data at the beginning of each timestep. The slip wall condition imposes zero velocity through the outermost face of the last cell, a Dirichlet condition, and supplies a pressure in the cell immediately outside, equal to the pressure in the cell immediately inside, a von Neumann condition. This is self-consistent, and when sources terms come to be updated, the velocity remains unchanged. The pressure correction weight is also set to 0, which is equivalent to specifying that solid walls are infinitely massive. This effectively re-weights the stencil from a central difference approximation to the laplacian with seven points to a one-sided difference with six. By construction, corners are correctly handled.

The no-slip condition operates identically to the slip condition, except that an equal and opposite wall-parallel velocity is imposed immediately outside the domain. Velocities are defined on cell faces, so the interpolated value at the wall satisfies no-slip exactly. This is not especially useful in an ILES code, since transverse gradient does not activate the limiters and thus make no contribution to dissipation. There is no material difference in results between slip and no-slip condition, except in a rotating reference frame where it is less trivial to calculate an appropriate boundary condition. With a physical viscosity implemented, the no-slip condition would perform satisfactorily.

The inflow boundary condition has been implemented primarily for testing the advection routines, and specifies a pressure distribution, velocity vector and incoming scalar concentrations from which a density is computed. The pressure correction weight across the inflow face is specified from the density field. The transmissive boundary condition maps the pressure, scalar and velocity field from the edge of the domain to the ghost cells outside. While outflow conditions have well-known

problems and remain an active research area (see [Giles \(1990\)](#) for background), a naive implementation works well on simple problems, provided that the mean exit velocity is sufficiently large that flow never re-enters the domain.

Periodic and internal boundaries do not require assumptions to be made when populating the ghost cells. These boundaries are packaged up and delivered to the matching block that unwraps the data (re-ordered according to orientation) and places in the ghost cells. The upwinding algorithm for wall-normal advection uses information from no more than two cells centres and three cell faces beyond the domain, in both wall-normal and wall-parallel directions. As can be deduced from [figure 3.2](#), in calculating u momentum at the corner cell, scalar information is required from the cell diagonally outside the domain from the corner. This cell, and others in the vicinity, needs to be initialised. In a square arrangement of blocks, named from A to D as shown in [figure 3.3](#), information needed for the ghost cell corner of block A , say, can be found inside the domain of the diagonally opposite block C . Manually specifying the grid corner connectivity is tedious and error-prone, but by careful sequencing of boundary copying operations, the information can be transferred automatically. Firstly the ghost faces of all grids are packed, delivered and unpacked, and then secondly the ghost edges (the corners) of all grids are packed, delivered and unpacked. Hence domain-internal information in block C (that is needed by block A for its ghost corner) is passed through the $B - C$ face connectivity to the face ghost cells of block B , then in the second step the ghost corner of grid A is filled directly through the $A - B$ connectivity. Of course, all the other overlapping corners are filled in the same way by the single sequence of operations.

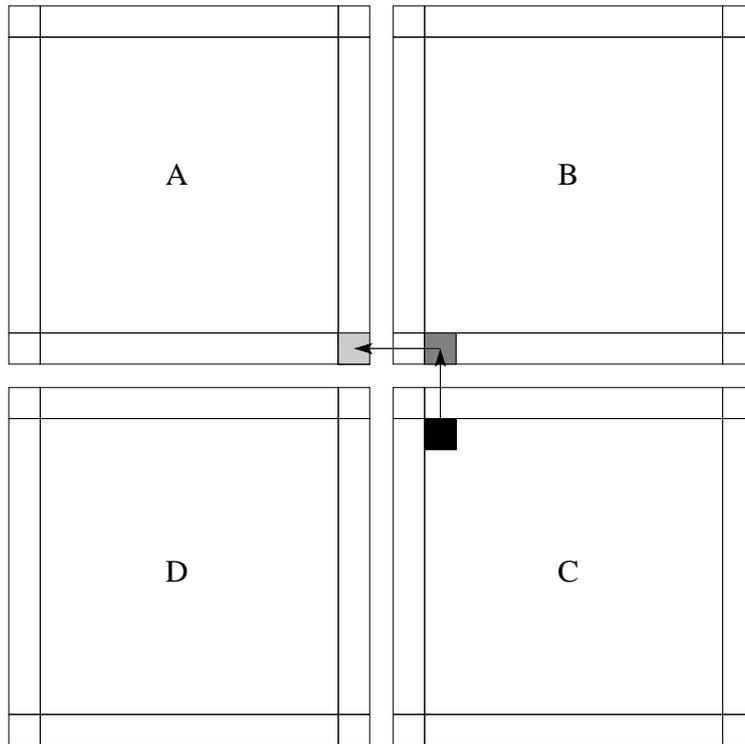


Figure 3.3: Data in the black square in the domain of block C is needed by block A as ghost data for its corner. A two stage process shown by the arrows enables this data to be retrieved without any explicit connectivity between block A and C.

3.3 Algorithmic details

3.3.1 Multigrid convergence acceleration

Many standard tools have been developed to solve the sparse linear systems that emerge from discrete elliptic problems. Direct methods such as Gaussian elimination are accurate to machine precision and some efficient methods have been developed. The operation count of naive direct methods are of $O(N^6)$, while the best methods, using Fast Fourier Transforms and block decomposition, scale as $N^3 \log N$, where N is a representative index dimension of one side of the computational domain. However to remain efficient these methods tend to be restricted to domain sizes $N = 2^\chi, \chi \in \mathcal{N}$, which makes them somewhat application-specific. Relaxation methods, on the other hand, are easy to implement, and very general. A naive relaxation method, such as Jacobi or SOR iteration scales as N^6 ; this is uncompetitive, but applying a multigrid methodology this can be reduced to $O(N^3 \log N)$, as fast as the fastest

direct methods. Additionally we only require to converge equation 3.25 to the level of truncation error accrued in the hyperbolic step, so direct methods are wasteful.

Equation 3.24 corresponds to a sparse linear system of the form

$$Ah = b \quad (3.26)$$

where b is the divergence denoted $\nabla \cdot \mathbf{u}^{**}$, h is the vector of pressure corrections Δp and A is a septa-diagonal matrix of density weighting coefficients of the form $W_{i-\frac{1}{2}} = \frac{\Delta t \Delta y^2 \Delta z^2}{\frac{1}{2}(\rho_{i-1} + \rho_i) \Delta x \Delta y \Delta z}$. A relaxation iteration using the Jacobi method has the form

$$g_{i+1} = (L + U)g_i + D^{-1}b \quad (3.27)$$

where $L + U = A - D$ with D a diagonal matrix, and g is the current approximation to h . Multigrid exploits the fact that relaxation operations with a Poisson-like stencil are inherently parabolic with respect to the iteration index, and ‘dissipate’ (smooth) most efficiently on a scale $O(\Delta x)$. By defining a residual vector r and an error vector e according to

$$r = b - Ag, \quad (3.28)$$

$$e = h - g, \quad (3.29)$$

it is clear that

$$Ae = r = b - Ag, \quad (3.30)$$

which implies that by mapping the residual vector r , rather than our current estimate g onto more coarse grids, we solve directly for the error e , which is more likely than the vector g to have high amplitude modes of $O(\Delta x)$. The reduction of e over length-scales $O(q\Delta x)$ is handled most efficiently on a grid with a mesh spacing $O(q\Delta x)$, so it makes sense to spread the calculation over several resolutions. The multigrid approach is applied recursively, halving the grid index dimension until the coarsest possible grid, then mapping the solution estimate successively from coarse grids to fine. This form of the multigrid method is called the V-cycle. One period of the simple V-cycle multigrid algorithm is shown in figure 3.4, where each dot represents an iteration sweep.

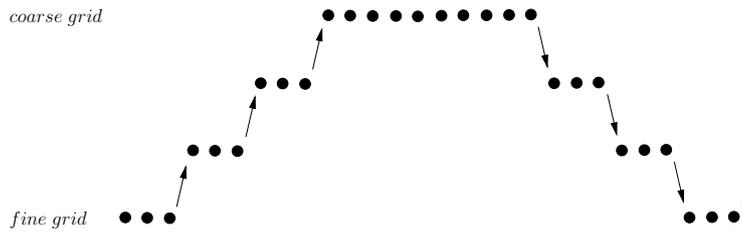


Figure 3.4: Schematic of the simple V-cycle multigrid algorithm, where each dot represents a smoothing operation. The optimal algorithm uses three iterations on each grid level, except on the coarsest grid, which is converged to tolerance.

The fine-coarse mapping used in MOBILE is a finite-volume formulation, and is achieved by integrating the residual over all the fine-grid cells that contribute to the corresponding coarse-grid cell. The problem for the error at the coarsest grid is solved iteratively to a very tight tolerance using the Jacobi scheme. This eliminates the longest wavelength mode, since this was found to considerably reduce the number of V-cycles required. The updated error vector e is projected from coarse grid to fine grid using a compact stencil linear interpolation that is transparent to modes with wavelength $O(2\Delta x^{fine})$ and greater. Figure 3.5 shows in one space dimension that in a cell- rather than node-centred algorithm, there are no points which are co-located at all levels of refinement, thus modes of twice the frequency are transferred from coarse to fine grid than can be achieved with a co-located scheme. The fine grid smoothing operates more efficiently, and tests indicated a 12% reduction in computational cost overall. Once the error e has been projected onto the fine grid, it is then added to the existing fine-grid estimate g , as implied by equation 3.29. In performance testing, the optimal algorithm was found to use only three relaxation iterations before moving to a finer or coarser grid, since the efficiency of the smoothing degraded approximately exponentially with the number of iterations.

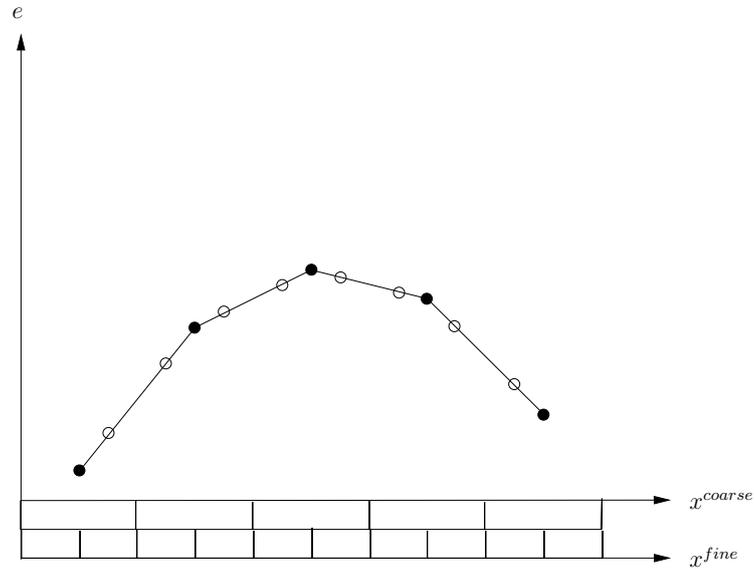


Figure 3.5: Schematic showing in one space dimension the linear interpolation scheme used when projecting a coarse residual solution $\widehat{e}(\mathbf{x})$ onto a finer non co-located grid. Compared with a co-located scheme this yields a 12% reduction in overall cost.

3.3.2 Parallel implementation

The Message Passing Interface (MPI) protocol is used to enable a single problem to be spread over multiple processors with distributed memory. The computational domain is decomposed into cuboidal blocks, and boundary information is passed between blocks in advance of the algorithm requiring updated information at the boundary of a block. The code is structured such that all processes are aware of the existence, location and size of every block in the domain, and messaging operations are organised in mutual `MPI_Isend()`-`MPI_Irecv()` pairs. Separate, contiguous send and receive buffers are allocated for each internal boundary on a grid and these are filled with the (in general) non-contiguous data from the grid before the communication is posted. Boundary information scales with $O(N^2B)$, where again N is the index dimension, and B is the number of blocks.

The quantity of data that must be passed is proportional to the half-width of the numerical stencil. For the hyperbolic step, this includes the two upwind cells needed to reconstruct the MUSCL gradients. The algorithm requires a staggered grid arrangement to advect the cell-face velocities, and therefore the velocity field

is defined on all block boundary surfaces. While this is advantageous for efficiently defining external boundary conditions, two upwind cells are required to update these values on internal domain decomposition boundaries, and ghost values for these cells are indexed by $x_{i-\frac{3}{2}}$ and $x_{i-\frac{5}{2}}$. The arrays for the primitive variables and working storage arrays are interleaved, so that memory access is localised and the cache hit rate is maximised, an increasingly important consideration for modern processor architectures. Thus a convention has been adopted whereby values on the mutually staggered grids are mapped to an index location in memory as follows: $x_{i,j,k}, x_{i-\frac{1}{2},j,k}, x_{i,j-\frac{1}{2},k}, x_{i,j,k-\frac{1}{2}} \rightarrow i, j, k$. Consequently, three index locations are required outside the domain in each direction to provide enough ghost cell space to completely capture the numerical stencil at an internal boundary.

The elliptic step uses a much smaller stencil, and this needs only one row of ghost cells at a block boundary, which improves the parallel efficiency. However, if parity is maintained between parallel and serial algorithms, communication is required in advance of every relaxation step, which is inefficient. The parallel efficiency is even poorer on coarser multi-grids, since the execution time of a relaxation step scales as $O(N^3)$ while the communication overhead scales as $O(N^2)$. Relatively little can be done to ameliorate this, since the apparently plausible approach of restricting blocking to finer grids and handling coarse grids on one processor actually raises the communication overhead scaling to $O(N^3)$. On small test problems of $(32 \times 32 \times 96)$ cells, with two internal block boundaries and sixteen periodic domain boundaries, the communication overhead was 40% of the total wall clock time on a dual-core CPU, but of course this scales favourably as problem size increases. The principal aim of parallelising this code was to permit the solution of large problem sets ($>4\text{GB}$) on available 32-bit operating systems, and on large problems the development costs of algorithmic refinement outweigh the execution time saved.

Chapter 4

Unconfined Rayleigh-Taylor instability

4.1 Introduction

Theoretical work on gravitationally unstable density interfaces began with [Rayleigh \(1883\)](#), and continues to this day. Such difficulty has been encountered in trying to quantify the most elementary statistic about evolution of a Rayleigh-Taylor unstable flow - the height of the energised layer - that almost all the modelling work in the literature to date focuses on this issue. This chapter introduces some of these modelling approaches, applied here in the context of the idealised Rayleigh-Taylor problem, and where appropriate these will be developed in later chapters to help understand Rayleigh-Taylor instability evolution in situations where it is confined by density stratification or geometric restriction.

4.2 Models for Rayleigh-Taylor instability

4.2.1 Early-time growth

It was noted in [Taylor \(1950\)](#) that two superposed fluids with an interface normal to the acceleration field could exist in a state of unstable equilibrium. However if small corrugations were to exist on the interface, a growing standing wave pattern could be expected. By considering an unbounded two-dimensional domain, incompressible

potential flow, and a sinusoidal interface, the velocity potentials ϕ in upper and lower layers are given by

$$\phi_u = ae^{-kz+nt} \cos(kx), \quad (4.1)$$

$$\phi_l = -ae^{kz+nt} \cos(kx), \quad (4.2)$$

with an interfacial surface displacement

$$\zeta = a \frac{k}{n} e^{nt} \cos(kx), \quad (4.3)$$

where k is the wavenumber of the surface, a is an arbitrary initial constant and n is a time-evolution growth parameter, which is as yet unknown. If \bar{p} is the mean pressure at the interface, then the upper and lower pressure fields are

$$p_u = \bar{p} - g\rho_u z + \rho_u \dot{\phi}_u \quad (4.4)$$

$$p_l = \bar{p} - g\rho_l z + \rho_l \dot{\phi}_l \quad (4.5)$$

and since pressure must be continuous across the interface, the following relation must hold:

$$-g(\rho_l - \rho_u)\zeta = (\rho_l + \rho_u)nae^{nt} \cos(kx). \quad (4.6)$$

Solving for the rate parameter n , we have

$$n^2 = kg \frac{\rho_u - \rho_l}{\rho_u + \rho_l}. \quad (4.7)$$

Choosing the positive solution for n and substituting into equation 4.3 gives an exponential growth of the surface displacement. As one would expect, the growth is a function of gravity, of a non-dimensional density $\frac{\rho_u - \rho_l}{\rho_u + \rho_l}$ (known as the Atwood number and usually denoted A), and more interestingly, the interface wavenumber. It is clear from this analysis that the Rayleigh-Taylor instability grows most quickly at higher wavenumbers, and that the rate is unbounded as $k \rightarrow \infty$. However, if one accounts for the viscosity of a real fluid, a stability analysis will yield a wavelength of maximum instability, which scales as

$$\lambda \sim \left(\frac{\nu^2}{Ag} \right)^{\frac{1}{3}}, \quad (4.8)$$

simply on dimensional grounds. Surface tension also acts to stabilise the Rayleigh-Taylor instability - to the extent that interface perturbations actually decay above

a critical wavelength - but for the miscible fluids studied in this thesis, this effect is negligible.

4.2.2 Potential flow

Development of unstable interfaces beyond the linear phase was first studied by [Davies & Taylor \(1950\)](#), and generalised somewhat in [Layzer \(1955\)](#). Both papers begin by considering a two-dimensional incompressible potential flow in a vertical cylinder with a initially flat free under-surface. This is Rayleigh-Taylor unstable with an Atwood number $A \rightarrow 1$. The velocity potential must satisfy Bernoulli's equation in cylindrical polar ordinates:

$$\frac{\partial\phi}{\partial t} - \frac{1}{2} \left(\frac{\partial\phi}{\partial z} \right)^2 - \frac{1}{2} \left(\frac{\partial\phi}{\partial r} \right)^2 - z = \alpha(t) \quad (4.9)$$

where $\alpha(t)$ is arbitrary. Although a velocity potential that everywhere satisfies equation [4.9](#) was not found,

$$\phi = F(t)e^{-z}J_0(r) \quad (4.10)$$

is an appropriate linearised functional form in the region surrounding the point of maximum displacement of the interface. $F(t)$ must be chosen to satisfy the free-surface condition near this point. In the late-time limit, the Bernoulli equation leads to

$$\frac{\partial F(t)}{\partial t} - \int F(t)dt = 0. \quad (4.11)$$

Thus

$$F(t) = e^{t+c} \quad (4.12)$$

where c is a constant of integration. When substituted into equation [4.10](#), the corresponding late-time velocity potential clearly represents a steady state in a reference frame moving upwards with a constant velocity. Generalising for an array of such bubbles is trivial, by the natural symmetry conditions for potential flow at a wall, and thus Layzer's analysis describes single-mode Rayleigh-Taylor instability and predicts a terminal velocity.

4.2.3 Buoyancy-drag balance

The complexity of directly evaluating a potential flow to calculate growth behaviour can be avoided. From first principles, the following equation describes how buoyancy and drag compete to accelerate a body, which in this case is a bubble of lower layer fluid penetrating the upper layer:

$$(\rho_l V + C_a \rho_u V) \frac{du}{dt} = (\rho_u - \rho_l) V g - \rho_u S u^2. \quad (4.13)$$

The term $C_a \rho_u V$ accounts for the inertia of the fluid displaced around the body under the action of a potential flow, a so-called ‘added mass’. S and V are body surface and volumes respectively. If the surface to volume ratio were constant (set by the wavelength λ of the instability) we would have

$$(\rho_l + C_a \rho_u) \frac{du}{dt} = (\rho_u - \rho_l) g - \frac{C_d}{\lambda} \rho_u u^2 \quad (4.14)$$

with both coefficients C_a and C_d being geometrically determined. Clearly in the early stages while drag is much less significant than buoyancy, we have exponential growth, whereas at late time, the terminal velocity is given by

$$u_\infty = \sqrt{\frac{1}{C_d} \left(\frac{2A}{1+A} \right) g \lambda} \quad (4.15)$$

where A is the Atwood number. Interestingly, this suggests that the terminal velocity increases with λ and hence decreases with k , which superficially contradicts Taylor’s earlier work.

[Linden *et al.* \(1994\)](#) proposed a more refined buoyancy-drag model of the form,

$$(2 + E) \frac{d^2 h}{dt^2} = Ag(1 - E) - \frac{C_d}{\lambda} \left(\frac{dh}{dt} \right)^2 \quad (4.16)$$

$$E = e^{-\frac{6\pi h}{\lambda}},$$

where the parameter E modifies the effective mass and the buoyancy to account for the change in aspect-ratio of the bubble structure. As the height becomes large in comparison to its wavelength, the baroclinic torque becomes bigger since the bubble has long near-vertical sides (hence a greater buoyancy force), and the mass of displaced fluid becomes a smaller proportion of the overall mass (hence a reduced added mass component). [Figure 4.1](#) illustrates the behaviour of this model for a

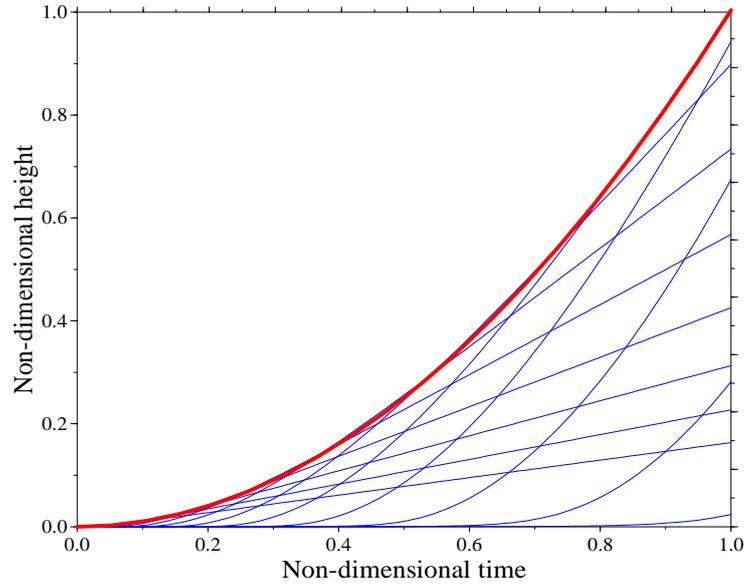


Figure 4.1: Blue curves plot the buoyancy-drag growth model for several perturbation wavelengths, while the red curve is the parabolic envelope. Length- and time-scales are arbitrary.

range of wavelengths λ . It is clear that the functional form reconciles a) Taylor’s prediction of initial exponential growth; b) small wavenumbers growing most quickly; c) Layzer’s prediction of a terminal velocity; and d) terminal velocity increasing with λ .

In real-world situations, well-defined initial perturbations are difficult to achieve, and single-mode perturbations can only be created in carefully chosen circumstances (Waddell *et al.* (2001); Wilkinson & Jacobs (2007)). Interface perturbations with a wide spectrum of wavenumbers are more commonly encountered. Some insight into the evolution of such a system can be obtained by regarding the evolution of a broad-band perturbation spectrum as an ensemble of independent single modes. The envelope of this ensemble appears parabolic in shape in figure 4.1, and this, as identified by Dalziel (2001), correlates with the experiments of Read (1984) and simulations of Youngs (1984b). The next section illustrates a simple way to recover this behaviour by analysing the energetics of the system.

4.2.4 Energy budget

By considering energy exchange between potential and kinetic as the instability develops, we can derive a scaling law for its evolution. We assume here that the initial state is at rest, and all the potential energy released after time t_0 is transformed into kinetic energy, and is not dissipated on a timescale over which the model is valid. Following the notation of [Jacobs & Dalziel \(2005\)](#), we have

$$\int \overline{\overline{\rho_0}} g z dz = \int \overline{\overline{\rho}} g z dz + \int \overline{\overline{\frac{1}{2} \rho \mathbf{u}^2}} dz, \quad (4.17)$$

where $\rho = \rho(z, t)$, $\rho_0 = \rho(z, 0)$ and the symbol $\overline{\overline{\phi}} = \frac{\iint \phi dx dy}{\iint dx dy}$ indicates a fluid property averaged over a horizontal plane. Under Boussinesq conditions, we can ignore changes in kinetic energy due to density variation, and assign a representative density ρ_b . Thus

$$\int (\overline{\overline{\rho_0}} - \overline{\overline{\rho}}) z dz \sim \frac{1}{2} \frac{\rho_b}{g} \int \overline{\overline{\mathbf{u}^2}} dz. \quad (4.18)$$

Assuming the density and velocity profiles remain self-similar across the energised zone throughout the time evolution of the flow, we construct a similarity variable

$$\zeta = \frac{z}{h(t)}, \quad (4.19)$$

where h is the current height of the energised zone, and functions of space and time are separated:

$$(\overline{\overline{\rho_0}} - \overline{\overline{\rho}}) = \hat{\rho}(t) r(\zeta) \quad (4.20)$$

$$\overline{\overline{\mathbf{u}^2}} = \hat{\mathbf{u}}^2(t) s(\zeta). \quad (4.21)$$

The integral equation then becomes

$$\hat{\rho} g h \int r(\zeta) \zeta d(\zeta) \sim \frac{1}{2} \rho_b \hat{\mathbf{u}}^2 \int s(\zeta) d(\zeta). \quad (4.22)$$

The integrals on both sides are constants, so

$$\hat{\rho} g h \sim \rho_b \hat{\mathbf{u}}^2. \quad (4.23)$$

It is reasonable to assume that $\frac{\partial h}{\partial t}$ varies directly with $\hat{\mathbf{u}}$, and to define the representative inertial density as $\rho_b = \frac{1}{2} (\rho_u + \rho_l)$ under a Boussinesq approximation. The

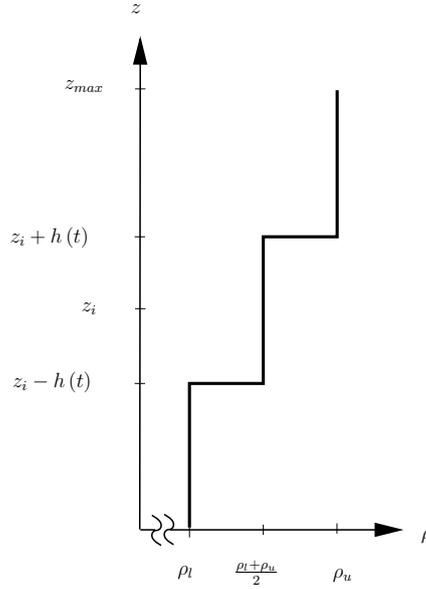


Figure 4.2: Diagram illustrating the simple box model for the evolving density profile. There is an implied assumption in this model that fluid in the mixing zone quickly becomes well mixed.

amplitude function $\hat{\rho}(t)$ does not vary with time in this problem, and an obvious choice is $\hat{\rho}(t) = \rho_u - \rho_l$ so that $-\frac{1}{2} < r(\zeta) < \frac{1}{2}$. Hence

$$(\rho_u - \rho_l)gh \sim \frac{1}{2}(\rho_u + \rho_l) \left(\frac{\partial h}{\partial t} \right)^2 \quad (4.24)$$

Noticing that the Atwood number $A = \frac{\rho_u - \rho_l}{\rho_u + \rho_l}$ is contained in this equation, we can integrate both sides in time to give

$$h \sim \frac{1}{2}Ag(t - t_0)^2, \quad (4.25)$$

where t_0 is a constant of integration and behaves like a virtual time-origin. This scaling is constructed on the assumption that *all* the potential energy released by the instability goes into advancing the front h of the energised zone. In fact h grows much more slowly, which implies that only a small proportion of the released potential energy is spent on moving the front. We introduce an empirically derived constant α to account for this, and arrive at the well-known relationship

$$h = \alpha Agt^2. \quad (4.26)$$

4.2.5 Gradient diffusion

A variety of methods can be used to deduce an $h \sim t^2$ relationship in Rayleigh-Taylor instability. We have seen how this was achieved using both buoyancy-drag and energy models; here a method is discussed which aims to parameterise the influence of small-scale turbulent mixing on the large-scale overall dynamics. Prandtl (1925) first proposed that scalar transport in a zero-mean flow, while complex and non-linear locally, could be modelled in aggregate by linear diffusion,

$$\frac{\partial \phi}{\partial t} = \frac{\partial}{\partial z} \left(\kappa \frac{\partial \phi}{\partial z} \right) \quad (4.27)$$

where ϕ is a scalar being diffused. Prandtl noticed that the factor κ has dimension $\frac{L^2}{T}$ and proposed that this could be expressed as

$$\kappa = \gamma l_{turb} u_{turb}, \quad (4.28)$$

with l_{turb} and u_{turb} representative length and velocity scales respectively, and γ an arbitrary constant. Prandtl argued that $\gamma = \frac{1}{3}$, but in general this value is determined by calibration. There are no universal choices for l_{turb} and u_{turb} , since these parameters are strongly problem-dependent and may change throughout the evolution of the system.

Here, u_{turb} and l_{turb} must be carefully chosen to reflect the specific nature of turbulence in a Rayleigh-Taylor context. The variable density incompressible vorticity equation,

$$\frac{\partial \omega}{\partial t} + (\mathbf{u} \cdot \nabla) \omega = -\frac{1}{\rho^2} (\nabla p \times \nabla \rho) + (\omega \cdot \nabla) \mathbf{u} + \nu \nabla^2 \omega, \quad (4.29)$$

is possibly the clearest starting point. We aim to estimate the magnitude of the important terms and consider the balance of forces in the non-linear regime of Rayleigh-Taylor instability.

Both the advective term $(\mathbf{u} \cdot \nabla) \omega$ and the vortex stretching term $(\omega \cdot \nabla) \mathbf{u}$ have dimensions of $\frac{u^2}{l_{turb}^2}$, and the diffusion term $\nu \nabla^2 \omega \sim \nu \frac{u}{l_{turb}^3}$. Under Boussinesq conditions, the hydrostatic pressure dominates the pressure gradient, so $\nabla p \approx \rho g$. Hence the baroclinic torque scales as

$$-\frac{1}{\rho^2} (\nabla p \times \nabla \rho) \sim -\frac{1}{\rho^2} \rho g \frac{\Delta \rho}{\Delta x}. \quad (4.30)$$

In the more natural dimensions of force per unit mass, and noting that if eddies have aspect-ratio $O(1)$, then horizontal density gradients will have the same magnitude as vertical density gradients, we have the following groupings:

$$\begin{aligned} \text{Inertia} &\sim \frac{u^2}{l_{turb}}, \\ \text{Viscosity} &\sim \nu \frac{u}{l_{turb}^2}, \\ \text{Buoyancy} &\sim \frac{1}{\rho} g l_{turb} \frac{\Delta \rho}{\Delta z}. \end{aligned} \quad (4.31)$$

Arguably viscosity has only a negligible influence in the non-linear regime of Rayleigh-Taylor instability, since inertia scales with u^2 , viscosity only with u , and u is linearly increasing with time. Therefore the balance between inertia and buoyancy is likely to determine flow conditions. Rearranging, and using appropriate notation, we can deduce a scaling for the turbulent velocity,

$$u_{turb} = \sqrt{\frac{l_{turb}^2 g \frac{\partial \rho}{\partial z}}{\rho}}, \quad (4.32)$$

provided we know a length-scale l_{turb} . In Rayleigh-Taylor instability, the largest turbulent eddies are likely to scale with the height of the energised zone, h . Following [Inogamov *et al.* \(2001\)](#), the diffusion equation 4.27, becomes

$$\frac{\partial \rho}{\partial t} = \frac{1}{3} l_{turb}^2 \sqrt{g} \frac{\partial}{\partial z} \left(\frac{1}{\sqrt{\rho}} \left(\frac{\partial \rho}{\partial z} \right)^{\frac{3}{2}} \right), \quad (4.33)$$

where the diffused scalar is the density ρ . In Rayleigh-Taylor evolution there are no variable parameters in z or t so according to this model, the diffusion profile must be self-similar. Inogamov's similarity variable was

$$\epsilon = \gamma \frac{z}{t^\delta} \quad (4.34)$$

and transforming the diffusion equation to this similarity variable, we have

$$-\epsilon \frac{\partial \rho}{\partial \epsilon} = \frac{1}{3} \sqrt{g} l_{turb}^2 \frac{\gamma}{\delta t^{\frac{5\delta}{2}-1}} \frac{\partial}{\partial \epsilon} \left(\frac{1}{\rho^{\frac{1}{2}}} \left(\frac{\partial \rho}{\partial z} \right)^{\frac{3}{2}} \right). \quad (4.35)$$

This is a second order ODE provided

$$\frac{l_{turb}^2}{t^{\frac{5\delta}{2}-1}} = \text{const.} \quad (4.36)$$

Inogamov argues that one would not expect the physical behaviour of the system to exceed second order since we are modelling scalar transport as a diffusive process, so it is reasonable to enforce condition 4.36. There is only one value of δ which both satisfies condition 4.36 and permits the scaling $l_{turb} \sim h$ which we know occurs in the physical system. This value is $\delta = 2$, and for a constant value of the similarity variable ϵ , this implies $z \sim t^2$ as we have come to expect from previously discussed methods.

Inogamov *et al.* (2001) also explored the possibility of having a constant l_{turb} , and this yields $\delta = \frac{2}{5}$, giving a power law growth profile which scales with $h \sim t^{\frac{2}{5}}$. This is a result that shall be developed in detail in chapters 6 and 7 where geometry constrains the turbulent length-scale.

4.3 Summary

A review of existing approaches for modelling the growth of Rayleigh-Taylor instability is presented, beginning with Taylor's potential flow prediction of exponential growth in the small amplitude regime, then outlining Layzer's prediction of a terminal velocity in a potential flow by generalising the case of a bubble in a vertical tube. Similar results were reached by considering the force-balance on a penetrating body, and this reconciled the apparent paradox that terminal velocity decreases with wavenumber but the most-unstable mode is at high wavenumber. The observation that the ensemble of several wavelengths leads to a parabolic envelope growth profile explains the connection between the single-mode analyses and the classic $h = \alpha A g t^2$, which is shown to be obtainable in several ways. In preparation for developments in later chapters, both an energy-budget and gradient-diffusion approach is outlined here.

Chapter 5

Direct Visualisation of Mixing

5.1 Introduction

The molecular mixing behaviour of density driven flows has historically proven difficult to quantify, in part due to the phenomenal complexity of the turbulent processes that lead to mixing at small scales, but also since molecular mixing itself is very hard to accurately measure - surely a first step in acquiring a more complete understanding. This chapter presents a new technique for measuring mixing induced by Rayleigh-Taylor instability, and uses a chemical indicator to highlight regions of flow that have become molecularly mixed. The indicator fluoresces with one of two colours depending on the surrounding hydrogen ion concentration, and the boundary between the two colours is an unambiguous definition of a volume fraction contour in the fluid. This experimental technique is applied herein to mixing induced by the classical two-homogenous-layer Rayleigh-Taylor instability, and compared with numerical simulation.

5.2 Reactive Light Induced Fluorescence (RLIF)

5.2.1 Previous attempts

Planar Light Induced Fluorescence (PLIF) is a common technique for observing scalar transport, and has successfully been used in many fields to obtain some insight into mixing processes (*e.g.* [Dahm & Dimotakis \(1987\)](#); [Lester & Clemens](#)

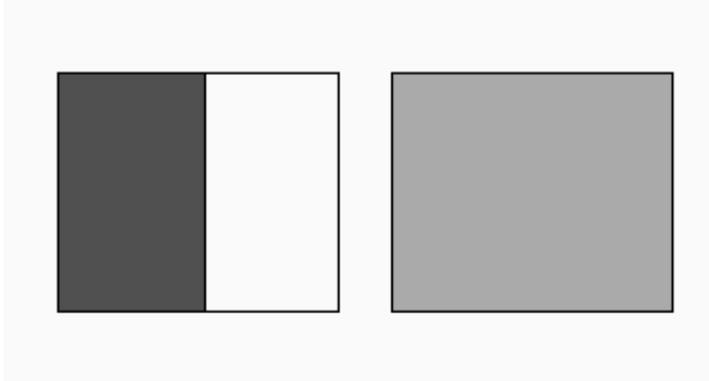


Figure 5.1: Standard PLIF imaging yields no sub-pixel information about the local mixing state. The same intensity value can be obtained from fully mixed, and fully unmixed fluid.

(2003); v Vliet *et al.* (2004)). Used conventionally, however, it does not yield any absolute quantification of mixing at a point in the observed plane. It merely places a maximum bound on the molecular mixing at such a point (see figure 5.1). The light intensity incident on one pixel of a CCD sensor is formed from the volume integral of all light sources and absorbers on the incident ray paths to that pixel. Thus in a PLIF context, and neglecting parallax and adsorption, this integral is composed of light from a small voxel bounded by the projection of the pixel in the PLIF plane and the thickness of the light sheet itself. There is no sub-pixel information about mixing, since an arbitrary normalised intensity value created from homogeneously mixed fluid is indistinguishable from that created from appropriate proportions of wholly unmixed fluid. Until imaging hardware is able to reach well below the Batchelor scale, spatially precise molecular mixing information cannot be obtained in high Reynolds number fluid mechanics with standard PLIF techniques.

Only by exploiting the molecular mixing itself as a diagnostic tool can progress be made. Light/Laser Induced Fluorescence has been used in the past to study mixing in miscible shear flows, *e.g.* Koochesfahani & Dimotakis (1985), where disodium fluorescein, known to lose fluorescence intensity when mixed with acid, was used to observe Kelvin-Helmholtz instability. Also, more recently, a hybrid fluorescence/phosphorescence technique in gaseous shear flows (Hu & Koochesfahani (2002)) was used to observe mixing in jets. However, in Rayleigh-Taylor based flows there seem not to be comparable experiments. Experiments using non-fluorescent

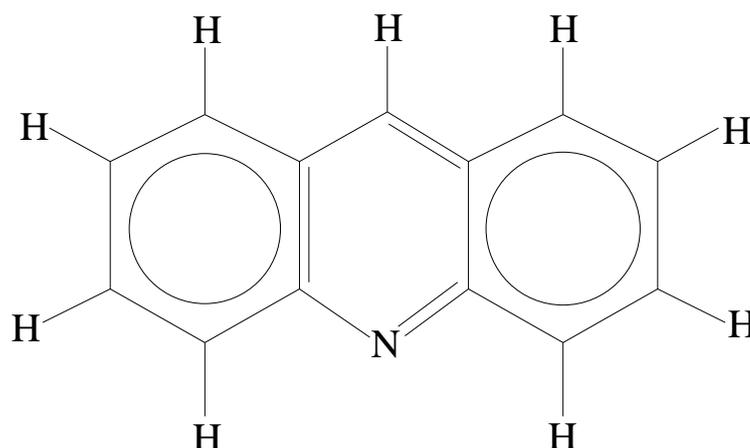


Figure 5.2: Chemical structure of Acridine.

pH indicators to quantify aggregate mixing have been performed (*e.g.* Linden *et al.* (1994); Andrews *et al.* (2007)), but direct visualisation of the mixing process, and the detailed fluid structure giving rise to it, has not previously been achieved. For the work presented in this thesis, a pH sensitive fluorescent dye was found with the special property that it maintains emission intensity on mixing, but sharply changes colour, thus yielding spatially accurate passive tracer and molecular mixing information from a single experimental realisation.

5.2.2 Chemical behaviour

The chemical, $C_{13}H_9N$, commonly known as Acridine, is an organic fluorophore with two benzene rings, as shown in figure 5.2. Incident light excites the de-localised electrons associated with the rings, and raises their energy state. When they collapse back to the ground state they emit a proportion of that absorbed energy as light. Since Planck's constant directly relates energy and emission wavelength, the emission must be at a longer wavelength. Conveniently for the present purposes, Acridine is sensitive to pH , and when H^+ ions congregate around the Acridine molecule, their influence reduces the proportion of incident energy that can be released upon collapse to ground state. Thus the emission spectrum is a function of pH , and to a good approximation behaves like a Heaviside function since the electron excitation rapidly switches between energy states. At high pH , the emission is blue, and at low pH the emission is cyan-green.

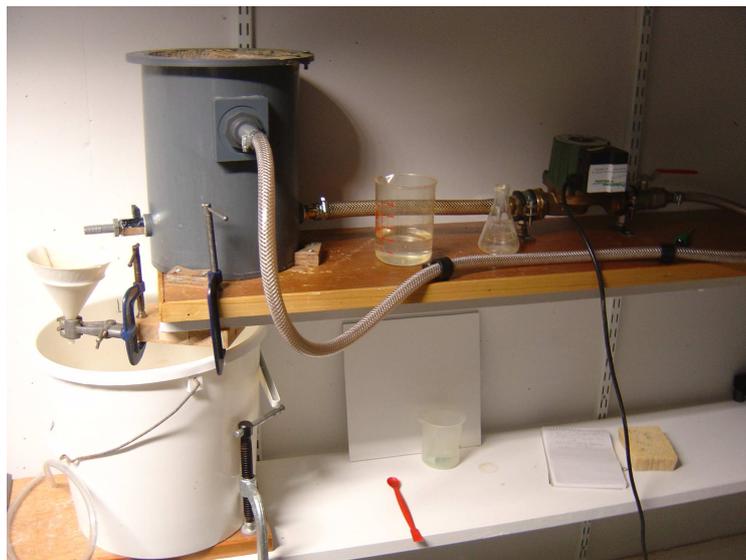


Figure 5.3: Recirculating pump, batch tank, output filter and settling bucket for production of saturated Acridine solution.

Acridine is only sparingly soluble in water, so producing a usable fluorescent solution was a non-trivial task. Agitating water and raw granular Acridine with a magnetic stirrer, until an acceptable level of dissolution had been obtained, then filtering the solute from the suspension achieved test-tube scale production. However, this process is hugely inefficient and labour intensive for the large volumes required in experiment. It was found that, with agitation, Acridine dissolves well in Propan-2-ol, but precipitates when the solution is diluted with water. The particle size is extremely small though, and therefore the surface area for dissolution to occur is many orders of magnitude greater than directly dissolving raw granules in water. A batch production system comprising a 9-litre tank and circulating pump, output filter and collection bucket (see figure 5.3) was assembled to agitate large volumes of the suspension and encourage dissolution over a period of time (usually 24 hours). The secondary effect of the pump was to heat the water and this was thought to aid the process of dissolution, as well as to accelerate the rate at which dissolved air contained in the water left solution. Fine-scale filtering of the remaining particulates in each 9-litre batch was found to be impractical, and a sequence of storage tanks was relied upon to let the suspension settle over several days at laboratory temperature. Finally, a saturated solution of Acridine was obtained.

	Two-layer Rayleigh-Taylor instability
Top layer	0.5 litres 2M hydrochloric acid 19.5 litres de-aerated fresh softened water
Bottom layer	2 litres Acridine solution 0.15 litres propan-2-ol 17.85 litres de-aerated fresh softened water

Figure 5.4: Recipe for two-layer RLIF experiments.

Some unwanted chemical interaction was noted in refining the experimental technique. In particular, salt (*Na Cl*), customarily used to control density differences, was found to significantly reduce the fluorescent signal from Acridine. Supporting evidence was found in the chemistry literature ([Geddes \(2001\)](#)) where it has been noted that halide ions indeed quench fluorescence by reducing the quantum yield when excited electrons return to ground state. To circumvent the need for salt as a stratifying agent, the experimental configuration was inverted, with relatively light Acridine solution placed in the lower layer, and relatively dense hydrochloric acid (*H Cl*) placed in the upper layer. Other acids were considered, but for laboratory safety in the event of reaction products, and the longevity of the tank inside surfaces, hydrochloric acid was chosen. Refractive indices were matched by adding (relatively less dense) alcohol as required to the lower layer. In turn, to limit the effect on data quality of dye attenuating the light sheet intensity, the illumination was mounted downwards from the ceiling, maximising optical contrast at the Rayleigh-Taylor interface. It is noted in passing that the fluorescent emission from di-sodium fluorescein solution must presumably also suffer a similar loss of efficiency in the presence of halide ions, although investigators exploiting it in their studies of mixing have not appeared to identify this as a source of error.

A further effect attributed to chemical interaction was observed when one or other layer contained incompletely de-aerated water. Tiny bubbles were sometimes observed in clouds at the tip of the barrier and it is thought that dilution of the acid by mixing could lead to residual dissociation of the H^+ and Cl^- ions, providing a small heat source just sufficient to induce air to leave solution and form these

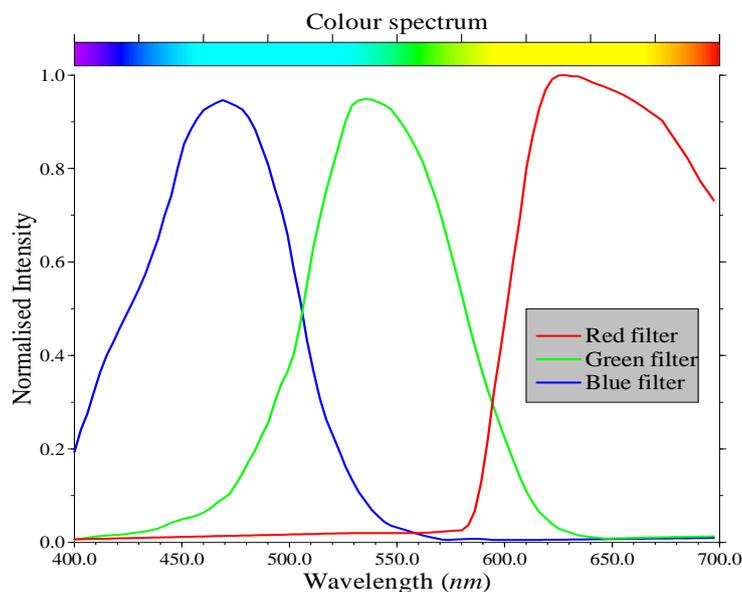


Figure 5.5: Bayer-filter spectra for CCD sensor used in the Uniq-Vision UC1830CL colour camera.

bubbles. At the acid concentrations used, this could only be a very small heat source, and certainly of negligible dynamic significance and has only a minor and short-lived impact on the optical measurements. The recipe for an exemplar two-layer experiment is given in the table 5.4.

5.2.3 Optical Decomposition

By eye, Acridine emits fluoresced light in two very distinct colours, depending on the pH of the solution. Unfortunately a camera with colour space discretised into Red, Green and Blue CCD sensors cannot so easily discriminate between the colours, since the spectra of the fluoresced colours overlap. Further complication arises since Bayer filter CCD sensors use filters with overlapping spectra. Accumulated intensity values in the camera can thus be modelled as a linear superposition of all visible wavelengths, after various filtering operations, at each pixel location. A suitable

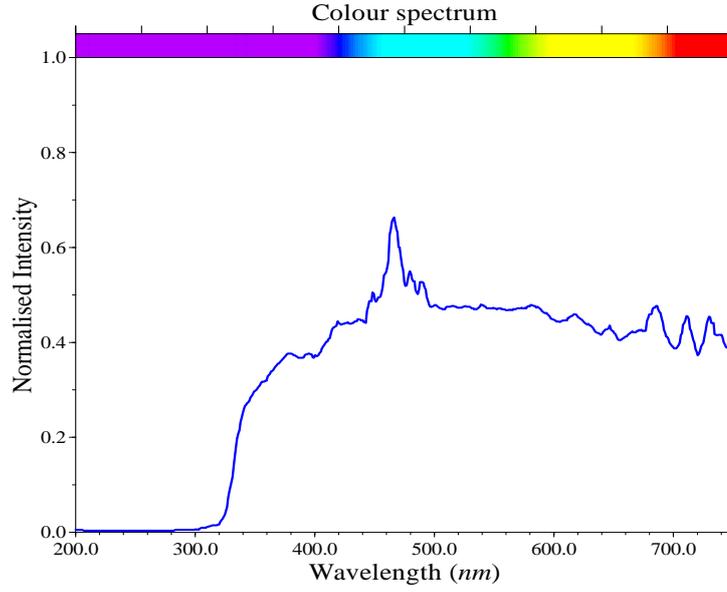


Figure 5.6: Spectral response of Perkin-Elmer arc-lamp fitted with a UV filtered parabolic reflector.

model for sensor values might be

$$\begin{aligned}
 R &= \int_{400nm}^{700nm} \phi E(\lambda, pH(\phi), I(\lambda)) F_R(\lambda) d\lambda \\
 G &= \int_{400nm}^{700nm} \phi E(\lambda, pH(\phi), I(\lambda)) F_G(\lambda) d\lambda \\
 B &= \int_{400nm}^{700nm} \phi E(\lambda, pH(\phi), I(\lambda)) F_B(\lambda) d\lambda,
 \end{aligned} \tag{5.1}$$

where ϕ is the volume fraction (of acid, say), I is the incident spectrum, F_R , F_G and F_B are the Bayer filter spectra, and E is the Acridine emission profile as function of λ , pH and I . Note also that this model neglects any losses from fluoresced light being attenuated when travelling through from the RLIF imaging plane to the side of the tank.

Since our objective is to interpret camera RGB intensity values as a volume fraction contour, we need to deconvolve the filter functions to recover ϕ for each pixel location. Clearly, as presented above, this problem is infinitely under-determined, but we can replace the integral definition of equation 5.1 by a discrete sum and the functions by vectors, and construct a closely equivalent linear system. Recognising that the aim is to classify fluid by its mixedness, we can define a fluid state M representing mixed fluid, and a corresponding state U for unmixed fluid, according to

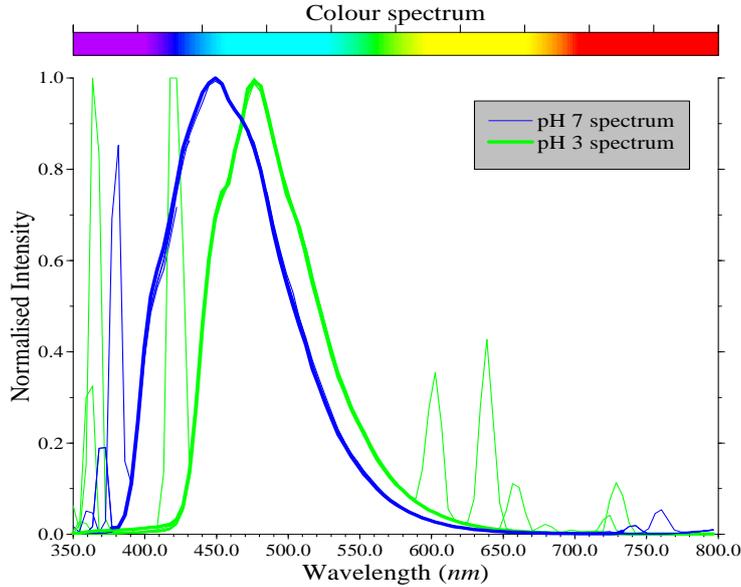


Figure 5.7: Normalised colour spectra for both blue and cyan-green emission states of Acridine.

some volume fraction threshold. From calibration data we also know that Acridine colours have a low spectral density above $600nm$, so there is little useful information in the Red-filtered CCD sensor, and we can discard this, yielding a correctly determined system for each pixel location. If we consider a discretisation with n wavelengths and p volume fractions, then

$$(2 \times n) (n \times p) (p \times 2) (2 \times 1) \quad (2 \times 1)$$

$$[\mathbf{F}(\lambda)] [\mathbf{E}(\lambda, \phi)] [\mathbf{T}(\phi)] \begin{bmatrix} U \\ M \end{bmatrix} = \begin{bmatrix} B \\ G \end{bmatrix}, \quad (5.2)$$

where $\mathbf{T}(\phi)$ is a transition matrix which classifies fluid according to volume fraction into mixed or unmixed. To be well-conditioned, the volume fraction threshold at which \mathbf{T} is biased towards U or M must match the transition threshold in the experimental Acridine-Acid system, and similarly B CCD data must correspond to classification U (and G to M).

To populate \mathbf{F} , \mathbf{E} and \mathbf{T} , a calibration study was performed with the assistance of chemist Dr. Jean-Luc Weitor, using a fluorimeter, spectrometer and a glass prism. The matrix \mathbf{F} was obtained by taking a prism and separating approximately white arc-lamp light into its spectrum, projecting through a semi-opaque diffusive screen,

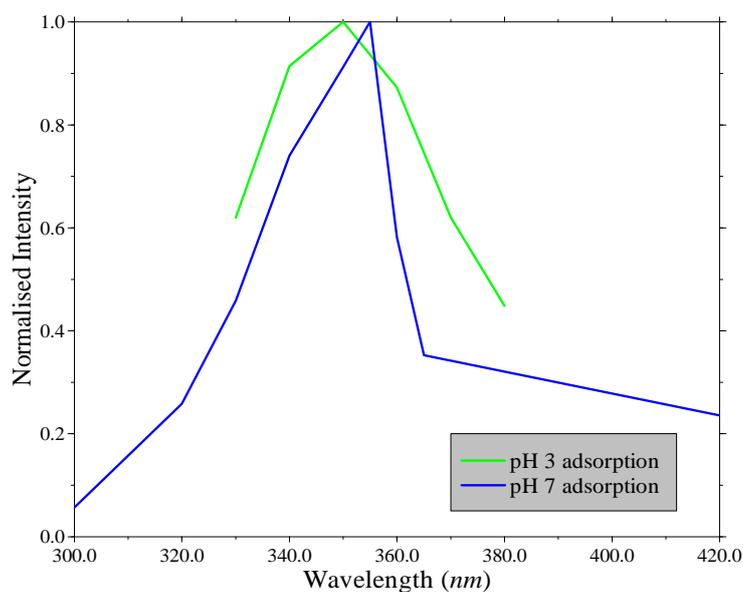


Figure 5.8: Adsorption spectra for blue and cyan-green emission states of Acridine.

and photographing the projection. The $R(\lambda)$, $G(\lambda)$, $B(\lambda)$ curves obtained for the Uniq-Vision UC1830CL are shown in figure 5.5, and correction for the non-uniformity of the arc-lamp light intensity across the visible spectrum has been made according to Perkin-Elmer supplied data, shown in figure 5.6.

The matrix \mathbf{E} was obtained by performing fluorimetry at various pH values and concentrations. These tests established not only that the normalised fluorescent emission spectra (*i.e.* the perceived colours) were independent of excitation wavelength and concentration, but also that the adsorption spectrum varies with pH . Figure 5.7 cross-plots the colour spectra, with multiple curves from various excitation wavelengths collapsing under normalisation to one of the two colours, depending on pH . The sharp peaks are artefacts of the fluorimetry, and are caused by the sensor picking up scattered incident light (which is concentrated in a narrow band of wavelengths) and harmonics of the incident light. Some peaks are larger than others, with high peaks relative to the fluorescent emission indicating low fluorescent efficiency at that incident wavelength. The true fluoresced signal of Acridine on both sides of the colour transition is indicated in figure 5.7 by the more heavily weighted blue and green lines. Figure 5.8 shows the normalised incident energy adsorption

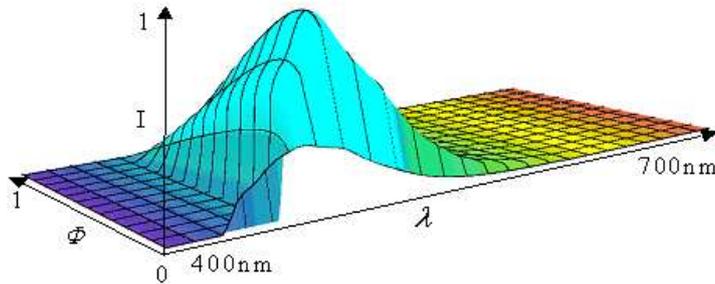


Figure 5.9: Map of Acridine spectral response as a function of volume fraction.

spectra for each colour, as calculated from several measurements around the peak adsorption wavelength.

Compiling this information to express the spectral response of Acridine as a function of volume fraction produces the data-set to populate $\mathbf{E}(\lambda, \phi)$, depicted in figure 5.9. In a two-fluid system, pH is directly related to volume fraction ϕ (where $\phi = 0$ corresponds to unmixed Acridine; $\phi = 1$ corresponds to unmixed Acid solution), and it is convenient that pH varies logarithmically with ϕ , so around the threshold at which Acridine changes emission colour, large variations in pH correspond to small variations in ϕ . Thus the Acridine colour transition occurs at low values of ‘mixedness’, making this an especially sensitive diagnostic for the occurrence of molecular mixing. The current experimental implementation has the advantage that both outer boundaries of the mixing region are well defined (the non-reacting boundary of the mixing region is also well defined in an Rayleigh-Taylor flow since it behaves like a conventional passive fluorescent dye penetrating into an un-dyed region). For other conceivable applications, particularly where chemical reaction rates are related to specific ratios of species, other volume fraction contours may be sought, and with this technique these can in principle be observed simply by varying acid concentration. The visualisation and illumination technology used herein is not sufficiently refined to have much flexibility in selecting the volume fraction contour and still return a high signal-to-noise ratio video image, but this is merely a practical issue, and not a limitation of the method.

Convolving \mathbf{E} with the CCD sensor filters \mathbf{F} predicts the camera response as a function of volume fraction, shown in figure 5.10. For completeness, $R(\phi)$ is included,

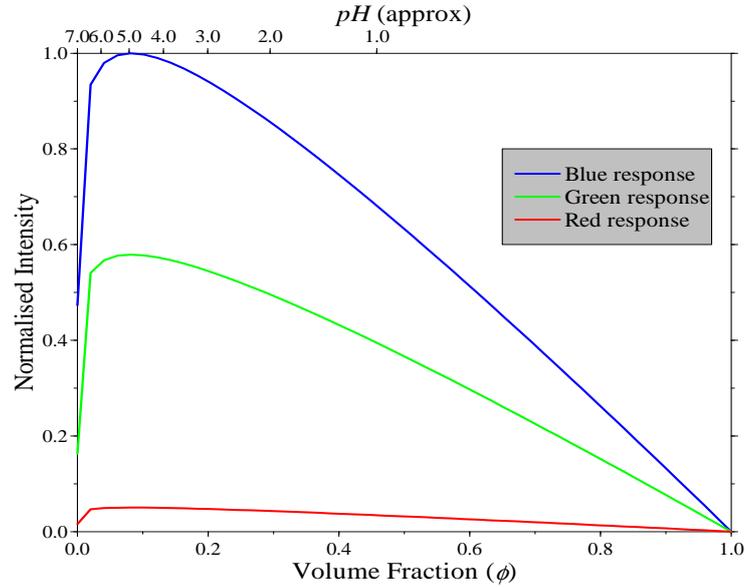


Figure 5.10: CCD sensor response to Acridine as a function of volume fraction.

but clearly $R(\phi) \approx 0$, hence the earlier statements about ensuring the linear system is well conditioned. By extracting numerical values from the figure, it can be seen that there is a sudden change in the ratio $\frac{B(\phi)}{G(\phi)}$ in the range $0 < \phi < 0.1$: $\frac{B(0)}{G(0)} = \frac{0.48}{0.17} = 2.82$; $\frac{B(0.1)}{G(0.1)} = \frac{1.0}{0.57} = 1.75$. This is the Acridine colour transition.

The spectra of figure 5.8 show how much broader the peak adsorption band becomes when Acridine is in its low- pH state. This means that more incident energy is absorbed when the emitted colour is of longer wavelength (and lower energy), and this is reflected in the step rise of fluoresced intensity values during colour transition from blue state to green state. The subsequent reduction in fluoresced intensity tends towards a linear decrease, as a consequence of dilution until at $\phi = 1$ there is no Acridine remaining, so no fluorescent emission at all.

5.3 Experimental and numerical comparison

5.3.1 Qualitative observations

Having refined a technique for generating large quantities of Acridine, and measured its spectral properties, an ensemble of experiments was conducted, replicating the classic Rayleigh-Taylor problem and, for consistency and convenience, using the same

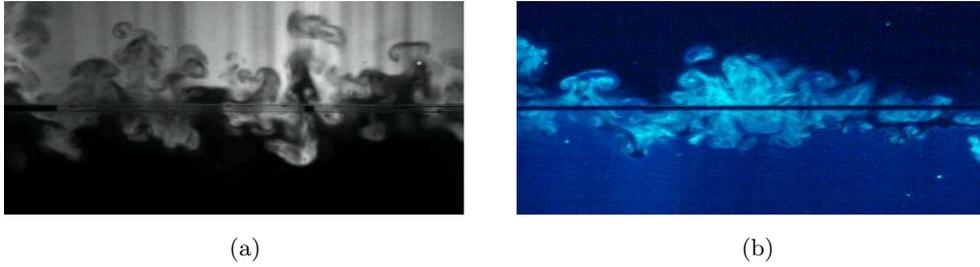


Figure 5.11: Comparison between (a) traditional PLIF, and (b) new RLIF experimental diagnostics for visualising molecular mixing. Both images are taken at non-dimensional time $\tau = 0.44$ and the Atwood numbers are (a) $A = 1.5 \times 10^{-3}$, and (b) $A = 1 \times 10^{-3}$.

apparatus as Dalziel *et al.* (1999). Initial observations suggest the new diagnostic does not interfere with the Rayleigh-Taylor instability, and figure 5.11 shows snapshots taken at the same non-dimensional time $\tau = \sqrt{Ag/Ht} = 0.44$ (where H is the tank height) in a Rayleigh-Taylor flow with the same Atwood number, using both traditional PLIF with a Fluorescein tracer and the new RLIF using Acridine.

A sequence of colour images in figure 5.12 shows the development of the flow at various times, where the time origin is determined by the time at which the end of the barrier passes through the imaging plane (note that the $t = 0$ image is brighter above the barrier when it is closed, because the nylon cloth reflects and scatters incident light from above). The advantage of the RLIF diagnostic technique is the ability to directly visualise the molecular mixing and examine the evolution of the advancing surface which bounds the mixed region. At $t = 4s$ the mixed region has few very well-defined shapes. This is consistent with small length-scale, short time-scale processes early in the instability development, and with the imparting by the barrier of some small but finite initial kinetic energy on the flow during withdrawal. As the instability develops, the eddy turnover time-scales grow and can be more easily captured at the camera frame rate. The imaging plane sits in the geometric centre of the tank and is oriented across the barrier. As discussed in §2.2.1, the removal of the barrier induces a net overturning circulation, and there appears to be a systematic upward velocity component at the tank mid-plane that persists some time after withdrawal. This is not uniform across the width of the tank, but is

largest in amplitude at the middle, and therefore must arise from a low wavenumber component of the initial perturbation. As one would expect from the analysis of §4.2.3, this does not manifest itself immediately, but from observation it appears to become significant when the instability has grown to a height comparable with the domain width (around $t = 12s$). Ultimately it dominates the flow's evolution. It is clear from images at $t = 16s$ and $t = 20s$ that unmixed lower layer fluid remains unmixed for a considerable distance inside the large rising bubble, and the front of mixing propagates from the sides of the bubble inwards over time, until eventually all of the bubble-internal fluid crosses the mixing threshold. Such is the unpredictable nature of Rayleigh-Taylor turbulence that even with similar initial perturbation this general behaviour is not consistent. Analysis of other video sequences shows that in some instances even at moderate stages of development ($t = 10s$) very little mixing occurs on the sides of the bubbles, and the majority occurs on the bubble and spike heads.

5.3.2 Growth profile

The simplest and most fundamental statistic related to the mixing process is the envelope growth profile that, according to the analysis of §4.2.4, should be quadratic if the mixing region has self-similar density and kinetic energy profiles. It is acquiring greater acceptance in the research community that the spectral profile of the initial condition strongly influences the subsequent instability development, and here three MOBILE simulations with a range of initial conditions are presented in comparison with the experimental ensemble. Figure 5.13 shows a colour image indicating the structure of the density/scalar field at the interfacial mid-plane. The image is split into three patches with different spectral profiles for each initial condition, from left to right corresponding to (1) ($x < 0.133m$) a random amplitude, random phase (idealised) density perturbation, (2) ($0.133m < x < 0.267m$) a constant amplitude, random phase (broadband) density perturbation, and (3) ($x > 0.267m$) a high wavenumber, random phase (narrowband) density perturbation. A transect through the centreline is shown superimposed on the image to show the wave-form of each perturbation.

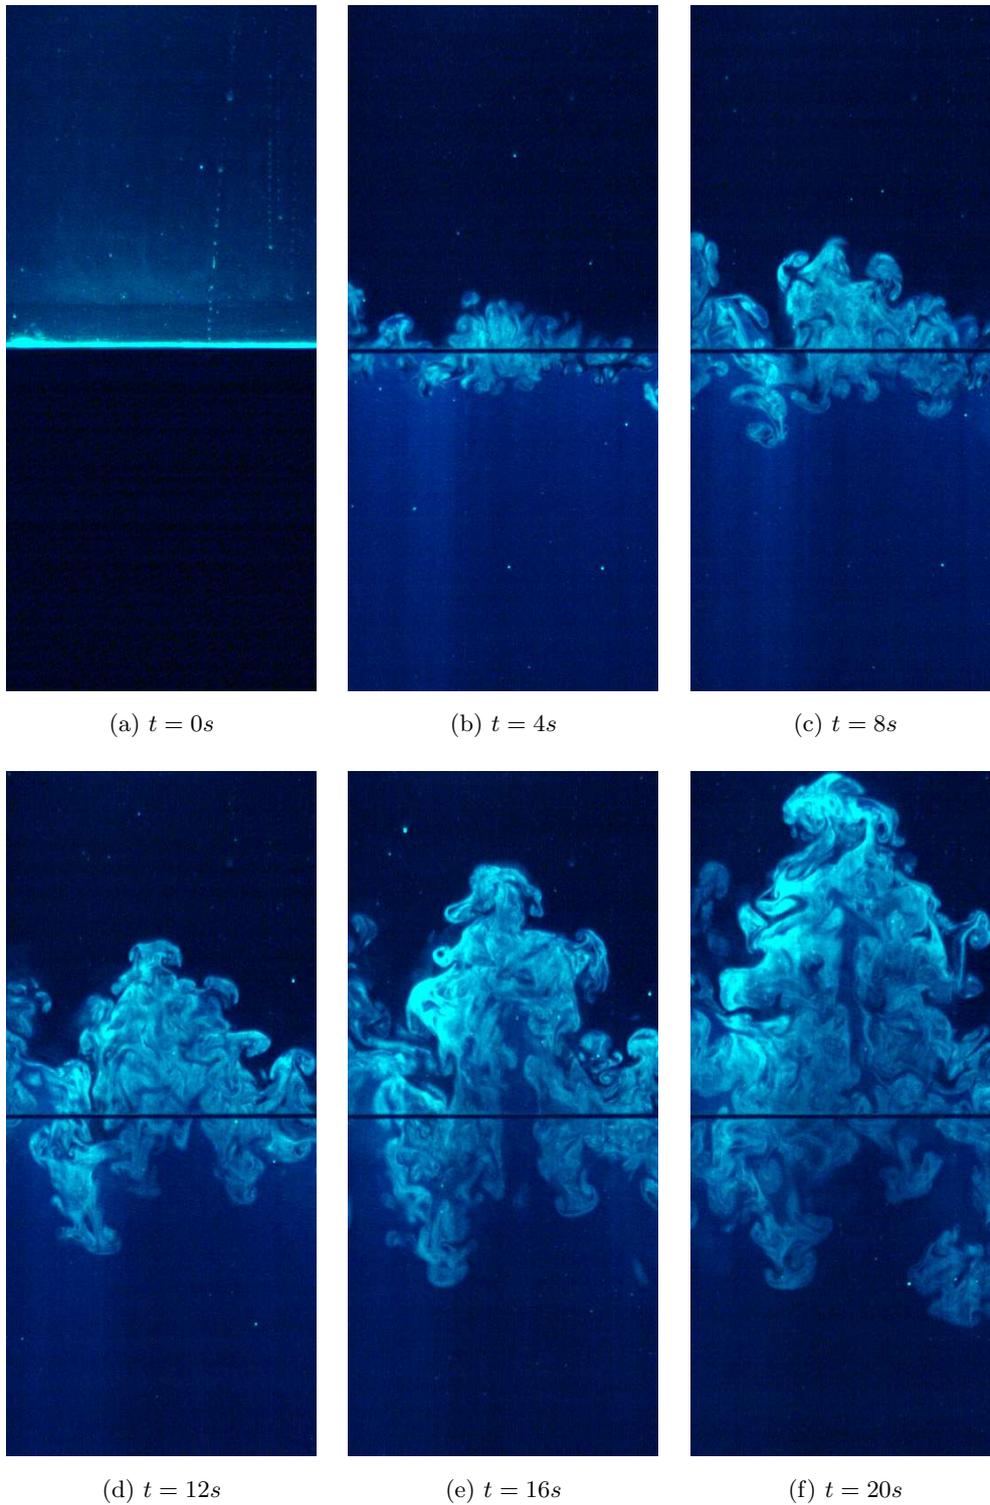


Figure 5.12: Direct visualisation of molecular mixing using Acridine. The Atwood number is $A = 1 \times 10^{-3}$.

The simulated growth profiles with these differing initial conditions are compared with an experimental ensemble in figure 5.14. The simulations are remarkably consistent, but although the upward trend is generally accelerating, there is an apparently systematic deviation from quadratic growth. This behaviour is also in evidence in simulations by Youngs published in Dalziel *et al.* (1999) when they are processed with a similar measure. Here, the vertical trajectory with time of the 2% / 98% horizontally averaged volume fraction contour is plotted, for consistency with experiments. Historically more sophisticated measures have been used, *e.g.* $\int \bar{\phi} dz$ and $\int \bar{\phi}(1 - \bar{\phi}) dz$ (where $\bar{\bullet}$ is a horizontally averaged quantity), since these enhance the visual appearance of the numerical results in resembling quadratic growth. The well-known discrepancy between numerical and experimental growth rates discussed in §1.2.3 is also very clear in the figure. The experimental ensemble does display quadratic growth, but with a displaced origin. It is believed that this arises from the barrier withdrawal, which imparts significant kinetic energy on the flow before Rayleigh-Taylor instability gets underway. A non-zero initial kinetic energy modifies equation 4.24 to

$$(\rho_u - \rho_l)gh \sim \frac{1}{2}(\rho_u + \rho_l) \left(\frac{\partial h}{\partial t} \right)^2 + E_0, \quad (5.3)$$

where E_0 is an initial kinetic energy associated with the barrier withdrawal, and integrating this through yields $h(t)$ with the normal quadratic relationship, but with an additional linear term related to the initial kinetic energy, in the form

$$h = \alpha Agt^2 + \widehat{E}_0 t, \quad (5.4)$$

where \widehat{E}_0 incorporates E_0 and the extra integration constants. In figure 5.14, the upper bound on the growth envelope includes a linear term. The convention for non-dimensional time-scale adopted in this figure and henceforth has been to define $\tau = \sqrt{Ag/H}t$ with, for later convenience in chapter 8 and in contrast to previous work, H defined as the half-height of the tank ($H = 0.25m$), rather than the full height.

The value of α which very closely matches the experimental mean is $\alpha = 0.05$, which is slightly lower than other previous experimental measurements (*e.g.* Snider & Andrews (1994) which measured $\alpha = 0.070 \pm 0.011$). In obtaining the recent figure,

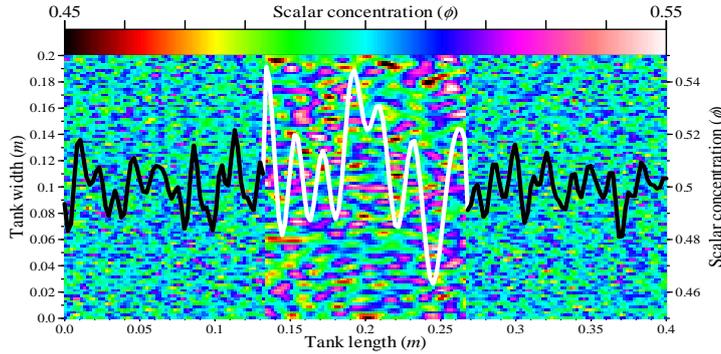


Figure 5.13: Initial numerical scalar concentration fields for horizontal mid-planes of three simulations, with a length-wise transect to illustrate the perturbation in each case. The plot geometry is to scale.

the linear correction for contamination by the initial kinetic energy caused by the barrier withdrawal has been included in the calculation. In the previous work there is a similar source of kinetic energy (from boundary layer formation on the splitter plate dividing the two flows), and no similar linear correction was used. This may explain the small discrepancy.

5.3.3 Molecular mixing

More revealing than the envelope growth for examining the internal structure of a scalar mixing process is the so-called molecular mixing fraction. This quantity is defined over a horizontal plane as

$$\Theta_p(z, t) = \frac{\overline{\overline{\phi(1-\phi)}}}{\overline{\overline{\phi(1-\overline{\phi})}}}, \quad (5.5)$$

where $\overline{\bullet}$ is a horizontally averaged quantity and ϕ is the volume fraction. The function $\phi(1-\phi)$ has a maximum at $\phi = 0.5$, which represents well mixed fluid. Taking the ratio of the two distinct evaluations of the horizontal planar average helps to distinguish, for instance, a sinusoidally perturbed interface that remains unmixed, from a well-mixed layer where both cases have the same arithmetic mean $\overline{\phi}$ (at $z = z_i$). In the first case $\overline{\overline{\phi(1-\phi)}} = 0$, in the second, $\overline{\overline{\phi(1-\phi)}} = 0.25$, hence when normalised by $\overline{\overline{\phi(1-\overline{\phi})}}$ the mixing parameter lies in the range $0 < \Theta_p < 1$. However, as figure 5.1 is intended to illustrate, the domain over which the averaging operation $\overline{\bullet}$ is performed cannot be arbitrarily reduced when experimental data has

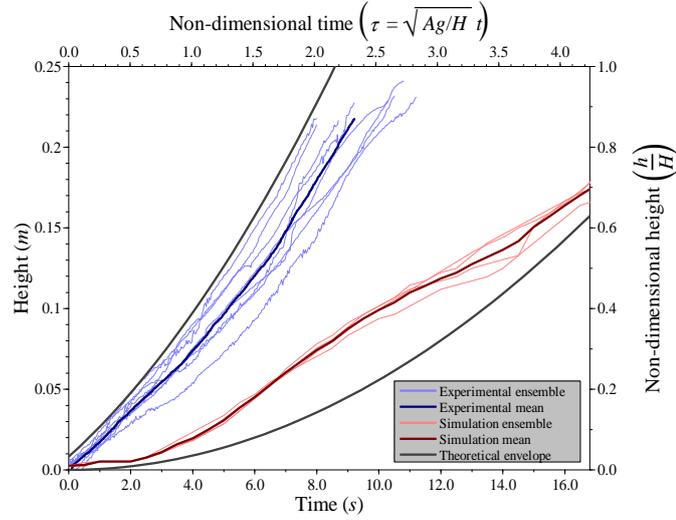


Figure 5.14: Growth profiles from experimental ensemble, numerical simulations with various initial conditions, compared against an envelope of parabolae $0.03 < \alpha < 0.08$.

finite resolution. Therefore in practice $\phi(1 - \phi)$ is not unambiguously discriminatory between mixed and unmixed fluid, rather it provides an upper bound for the mixing fraction.

Figure 5.15 shows how $\Theta_p(z, t)$ varies as the instability evolves (where the denominator of equation 5.5 is very close to zero, the image is coloured black). The bottom half of the plot ($z < 0.25$) shows experimental results from the bottom half of the tank, and the top half of the plot ($z > 0.25$) shows a numerical prediction generated by MOBILE using the broadband initial condition, and processed in a consistent manner by taking the vertical mid-plane at $z > 0.25$. To make a fair comparison, the time-scales have been adjusted to account for the well-known discrepancy in estimates of α between simulation and experiment. The idealised quadratic profile required to reach the tank extremities at the same time is shown in white, with $\alpha = 0.13$ for the experiment (with no linear correction term) and $\alpha = 0.03$ for the simulation.

The simulation has a statistically homogenous initial condition, so out-of-plane scalar transport is insignificant. This is not the case in the experiment because of the initial barrier perturbation, and growing horizontal streaks appear in the time-series image as spikes distorted from a downward path penetrate the viewing plane.

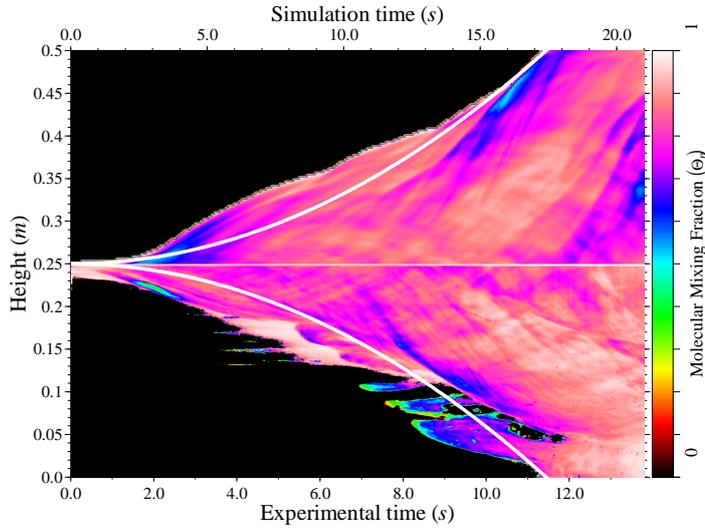


Figure 5.15: Local molecular mixing fraction $\Theta_p(z, t)$ as calculated from planar slices taken from a MOBILE simulation ($z > 0.25$), and experiment ($z < 0.25$). Ideal quadratic growth is also shown for reference.

This is most noticeable at later time ($t_{exp} = 8s$), where less well-mixed fluid from the middle of a large spike has appeared in the image plane. The distortion of bubble and spike paths is systematic so an ensemble of such experiments (Dalziel *et al.* (1999)) behaves in the same manner. The data presented here is from a single experiment, since it reveals rather more of the internal structure of the mixing zone than could be observed in an ensemble average. This diagnostic reveals the vertical transport of parcels of fluid both upwards and downwards, which are distinguishable as a diagonal criss-cross because they undergo a relatively small amount of additional mixing as they move. Numerical resolution imposes a limit on how many of these individual diagonal transport paths can be captured in the simulation, but their spatial structure is well modelled. The detail of the envelope profile is less well modelled, as one would expect without directly simulating the barrier removal, but one notable common feature is the reduction in molecular mixing close the boundary as the instability reaches the vertical extremities of the tank. The molecular mixing is a little lower at late time ($t_{sim} > 20s$) in the simulation than in the experiment, though the simulation values are consistent with the equivalent measure,

$$\Theta_s(z, x) = \frac{\overline{\phi(1-\phi)}}{\overline{\phi}(1-\overline{\phi})}, \quad (5.6)$$

(where $\bar{\bullet}$ is a temporal mean) computed by [Wilson & Andrews \(2002\)](#) in their steady state (spatially evolving in x) water-tunnel experiments. They reported $\Theta_s \approx 0.7$ with very little spatial variation either in z or x within the Rayleigh-Taylor mixing region.

A convenient parameter analogous to Θ_p that captures the aggregate mixing state of the Rayleigh-Taylor system can be defined. The global mixing fraction,

$$\Theta_g(t) = \frac{\int_{-\infty}^{\infty} \overline{\overline{\phi(1-\phi)}} dz}{\int_{-\infty}^{\infty} \overline{\overline{\phi}}(1-\overline{\overline{\phi}}) dz}, \quad (5.7)$$

indicates of the proportion of fluid that has become mixed, how well mixed it has become, and again by construction lies in the range $0 < \Theta_g < 1$. [Linden *et al.* \(1994\)](#) introduced this measure, and [Dalziel *et al.* \(1999\)](#) presents a modified formula,

$$\hat{\Theta}_g(t) = \left\langle \frac{\int_{-\infty}^{\infty} \overline{\phi(1-\phi)} dz}{\int_{-\infty}^{\infty} \overline{\phi}(1-\overline{\phi}) dz} \right\rangle, \quad (5.8)$$

where $\langle \bullet \rangle$ is an ensemble average. Having investigated archive data from experiments contributing to the paper and reverse engineered the associated processing code, it appears that the $\bar{\bullet}$ operator is not simply a horizontal average (which is sensitive to any spatial inhomogeneity of the Rayleigh-Taylor instability), but is a more robust measure based on a probability density function $P(\phi)$ of the volume fraction evaluated over the whole image plane. The algorithm used is as follows,

$$Q(\phi) = F_s(\langle P(\phi) \rangle) \quad (5.9)$$

$$\hat{\Theta}_g(t) = \left(\int_0^1 \hat{\phi}(1-\hat{\phi})Qd\hat{\phi} \right) / \left(\left(\int_0^1 \hat{\phi}Qd\hat{\phi} \right) \left(1 - \int_0^1 \hat{\phi}Qd\hat{\phi} \right) \right), \quad (5.10)$$

where F_s is a low-pass smoothing filter, and $\hat{\phi}$ is a dummy variable looping over the range $0 < \hat{\phi} < 1$ of the probability density function Q . To confirm that this is indeed the algorithm used to generate figure 24 of [Dalziel *et al.* \(1999\)](#), the archived experimental ensemble has been re-processed, and this is shown in figure 5.16.

The re-processed mean and the originally published curve display the same trend, although there is a small offset, possibly caused by alternative choices made in defining boundary conditions for the smoothing filter F_s . The published curve is not shown before $t = 2s$ and this is the time taken for the barrier to be removed from the tank. There is a very high degree of scatter in the individual measurements, despite

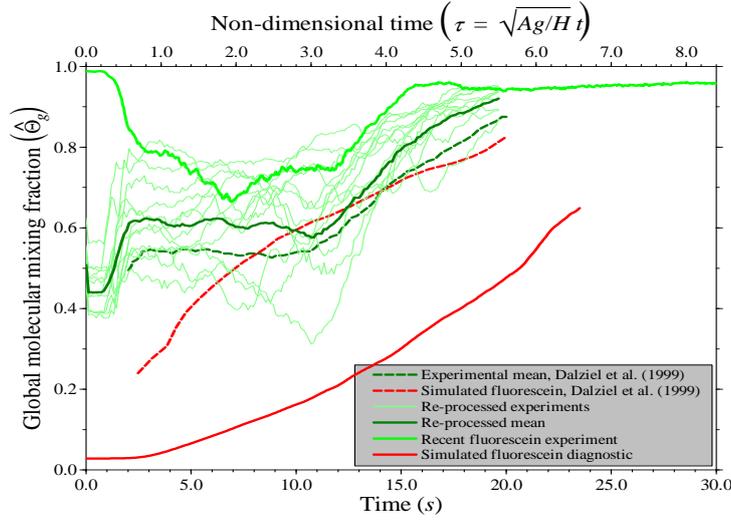


Figure 5.16: Global molecular mixing fraction $\hat{\Theta}_g(t)$ calculated using equation 5.9 on a recent simulation, a recent experiment, archived experimental data from Dalziel *et al.* (1999), and compared with published curves.

nominally identical barrier withdrawal and quiescent initial conditions. Visually too, the archived video images show that dye transport is remarkably consistent between elements of the ensemble, and only the central (well-illuminated) portion of the tank was used to gather data, thereby maximising the signal-to-noise ratio. It is unclear why the scatter is so wide, and it is also unclear why the early time data ($t \leq 2s$) does not grow from 0. The recent experiment was conducted in the ‘transverse’ orientation where the end of the barrier passes through the imaging plane at a single instant, rather than the ‘along-tank’ orientation used in Dalziel *et al.* (1999), where the barrier can be observed being withdrawn in the viewing plane. Apart from the early time ($t \leq 2s$) where for these reasons $\hat{\Theta}_g$ are not directly comparable, the data from the recent ‘transverse’ experiment predominantly sits within the range of scatter of the ‘along-tank’ experiments. The marked transition around $t = 10s$ in the ensemble mean from $\hat{\Theta}_g \approx 0.6$ to $\hat{\Theta}_g \rightarrow 1.0$ coincides with the time at which the Rayleigh-Taylor instability reaches the top of the tank and can no longer continue to develop freely. The transition point is noted slightly earlier in the recent experiment than in the majority of the archive ensemble, but there is an ambiguity over where a virtual time-origin should be placed to account for the change of orientation. The time-origin chosen for the ‘transverse’ experiment is the time at which the barrier passes

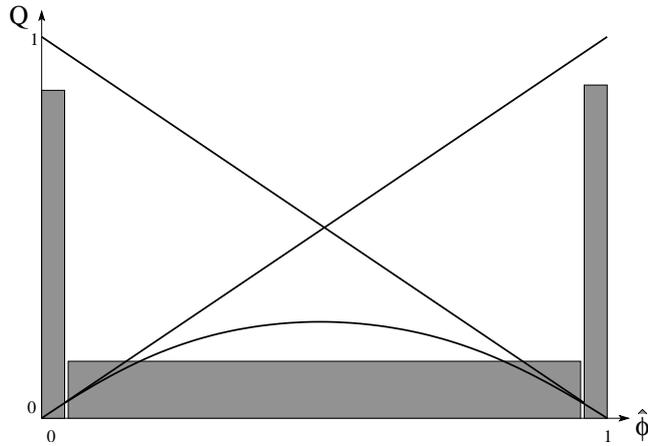


Figure 5.17: Typical histogram of volume fraction at the early stages of a mixing process. Basis functions for equation 5.9 are also shown.

through the imaging plane; in the archive experiments, the time-origin is taken as the instant of barrier release.

One of the simulations from Dalziel *et al.* (1999) is also plotted in figure 5.16 and, broadly speaking, this predicts a steady increase in molecular mixing fraction rather than an early-time ($t < 10s$) plateau transitioning to a late-time plateau ($t > 10s$). However, firm conclusions cannot be drawn because the methodology used to plot simulation results in Dalziel *et al.* (1999) is not consistent with that used for experiments. Instead of ensemble averaging across several simulations, averaging is performed across several vertical slices in the same simulation. It is not clear how this would affect the results, but there is evidence (discussed in more detail later in this section) that the tank dimensions influence the structure of the Rayleigh-Taylor instability once $h(t)$ has grown to be comparable with the tank width and this influence could be masked by averaging vertical slices at a variety of spatial locations. A simulation using MOBILE is also plotted for comparison, using a methodology consistent with the re-processed experimental data, and this shows significantly less mixing than either previous simulations or experiments, though in common with the earlier simulation predicts a steady increase in molecular mixing.

The parameters $\hat{\Theta}_g(t)$, $\Theta_p(z, t)$ and $\Theta_s(z, x)$ are appropriate measures of molecular mixing when the experimental diagnostic is correlated with volume fraction. Acridine, on the other hand, provides a resolution-independent diagnostic for iden-

tifying where fluid has mixed beyond a pre-determined volume fraction, and this region is bounded on one side by an iso-surface $\phi_{crit}(x, y, z, t)$. The other bounding surface can be defined in the same manner as a threshold for a passive scalar, and for symmetry this is here selected to be $(1 - \phi_{crit})$. To establish some degree of consistency between the two diagnostic techniques, the Fluorescein experiments were processed to classify fluid as being in one of three states: (1) unmixed undyed, (2) mixed, (3) unmixed dyed, intended to match the equivalent volume fraction states that are obtained naturally from the Acridine experiments. In this new definition, $\hat{\Theta}_a$, say, a histogram with three unequally sized bins replaces the probability density function Q , as illustrated in figure 5.17. The normalised basis functions $\hat{\phi}$, $(1 - \hat{\phi})$ and $\hat{\phi}(1 - \hat{\phi})$ used in equation 5.9 are also shown, to illustrate why $\hat{\Theta}_a$ increases as the histogram bin classifying mixed fluid becomes populated with more pixels from the image plane. As discussed earlier, all variants of the measure Θ have a maximum value of unity since the quadratic basis function in the numerator equals the product of the linear basis functions in the denominator when $\hat{\phi} = 0.5$. When Q is replaced with a three-state histogram, the value in each bin is calculated as the mean of the occurrence frequency of volume fractions covered by that bin. When the ‘mixed’ bin dominates the range of volume fractions (the Acridine thresholds are approximately $\phi = 0.02$ and $\phi = 0.98$), Θ no longer has a maximum value of unity because fluid with a homogenous volume fraction of $\phi = 0.5$ is processed as having a mean occurrence frequency spread evenly across the bin and this is interpreted as being less than fully mixed. By making appropriate changes to the basis functions, this deficiency could be alleviated, but to avoid further complication and uncertainty in interpreting the data, the original basis functions have been used.

Figure 5.18 shows the closest representative comparison between the Acridine and Fluorescein experiments on this modified measure $\hat{\Theta}_a$. The experimental ensembles (and simulations) have separate time-scales because the Acridine experiments were conducted at a slightly lower Atwood number than the Fluorescein experiments, and the results have been scaled to have a consistent non-dimensional time-scale. Compared with $\hat{\Theta}_g$, the new measure is more robust, and the scatter in the Fluorescein experiments is substantially reduced. It appears that there is a very distinct tran-

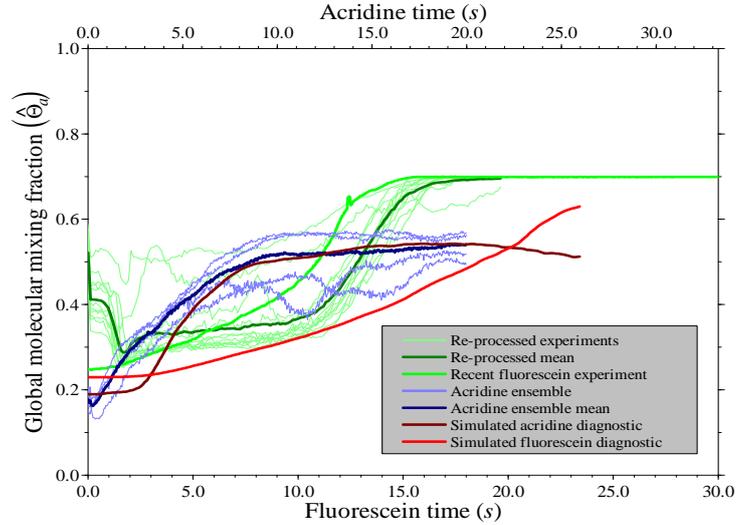


Figure 5.18: Global molecular mixing fractions $\hat{\Theta}_a(t)$ using calculated from Fluorescein (Dalziel *et al.* (1999)) and Acridine experiments and a common simulation with synthetic Fluorescein and Acridine diagnostics.

sition between predominantly unmixed and predominantly mixed (fully mixed fluid would return a mixing fraction of $\hat{\Theta}_a = 0.7$) occurring at $t_{Fluorescein} = 10s$, and this corresponds to the same transition point in figure 5.16. The recent Fluorescein experiment with a transverse orientation of the image plane sits more closely within the spread of archive ensemble data, but the ambiguity over the location of the virtual time-origin mentioned above is very clear here, since the transition time is shifted by $\sim 2s$.

The Acridine ensemble has a distinctly different trajectory to the Fluorescein ensemble despite considerable care being taken to match volume fraction thresholds. The argument illustrated in figure 5.1 would suggest that the Fluorescein measure of molecular mixing should be an upper bound on the actual value, which in turn should be more closely matched by the Acridine measure. Clearly the results in figure 5.18 do not support this hypothesis, but unfortunately there are too many variables between the experimental ensembles to make conclusive statements of this sort. The major, and at present unavoidable, inconsistency between the measurements is the correction made for light attenuation due to adsorption by the dye. Conventional PLIF data can be corrected by assuming that incident light sheet intensity decays exponentially in a homogeneously dyed medium according to the empirical Lambert-

Beer rule,

$$\frac{\partial I}{\partial s} = -\eta(\phi)I, \quad (5.11)$$

where I is the light intensity field, s is a light ray path, and η is an adsorption coefficient which varies in some way with dye concentration ϕ (usually taken as a linear variation). By inverting the decay function along each light ray, a correction for the attenuation can be made so that the corrected image more accurately represents ϕ . As with many inverse problems, the method is poorly conditioned, but it works reasonably well for simple problems, and the archive data from Dalziel *et al.* (1999) was processed in this way. Since Acridine is a less efficient fluorophore than Fluorescein, Acridine dye concentrations had to be maximised to achieve an acceptable signal-to-noise ratio with the (sub-optimal) arc-lamp light source, and therefore attenuation is a more significant problem. Unfortunately the attenuation correction is much more complex to implement when adsorption of incident light by the dye may be dependent both on the wavelengths of incident light and the excitation state of the dye. While the signal-to-noise ratio was generally satisfactory for making threshold-based measurements, the technological hardware is not yet sufficiently refined for more detailed data extraction. The measurements of $\hat{\Theta}_a$ from RLIF experiments have been taken without any attenuation correction, and this possibly explains the difference in trend shape compared with Fluorescein experiments, and its lower terminal value.

To shed some light on the influence the experimental diagnostic has on the results, a MOBILE simulation post-processed in two ways, to resemble both Acridine and Fluorescein experiments. An empirically calibrated estimate of light attenuation was used to modify incident light, and the calibration shown in figure 5.10 was used to map scalar concentration/volume fraction ϕ to an *RGB* image. The resulting simulated Acridine curve is shown as a dark red line in figure 5.18, and the comparison with the Acridine ensemble mean is good except for the early time ($t_{Acridine} < 3s$) while the barrier is being withdrawn. The synthetic Fluorescein diagnostic (red line) is not so successful at replicating the relevant experimental behaviour, but this is unsurprising given the large discrepancy in the corresponding (red) curve in figure 5.16 where the whole probability density function is used to calculate $\hat{\Theta}_g$. In many ways $\hat{\Theta}_a$ is a much less severe test of a numerical simulation than $\hat{\Theta}_g$, but it is the

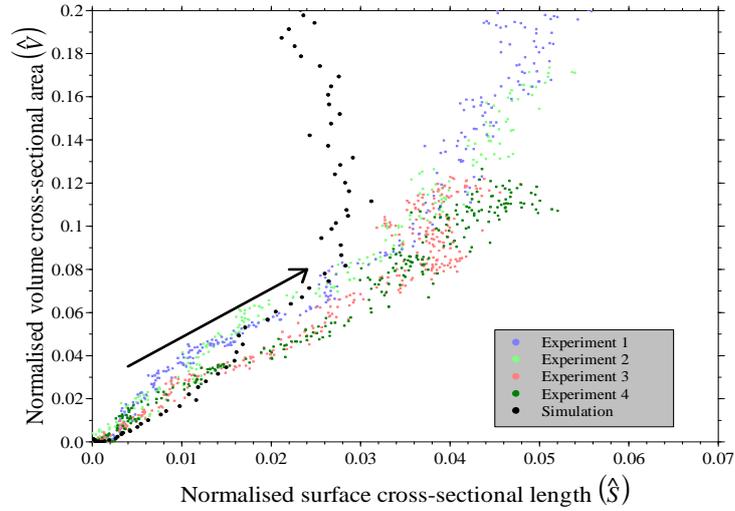


Figure 5.19: Scatter plot of cross-sectional surfaces and cross-sectional volumes for an ensemble of Acridine experiments and a MOBILE simulation with a synthetic Acridine diagnostic.

only reasonable way at present to compare Fluorescein and Acridine experiments.

5.3.4 Turbulent structure

The Acridine experimental ensemble is best suited, given the technological limitations, to providing information about pre-determined iso-surfaces of volume fraction and the volume of mixed fluid that these surfaces bound, and some insight is given by this into the structure of Rayleigh-Taylor induced turbulence. For dimensionally consistent nomenclature in the forthcoming discussion, the imaging plane cuts through a volume of mixed fluid yielding a volume cross-sectional area, and the cuts through the bounding surfaces yield cross-sectional lengths.

We expect from our understanding of Rayleigh-Taylor instability discussed in §4.2.3 that bubble and spike structures grow in length-scale as the instability develops. We therefore expect certain geometric relationships between the surface area of these structures and their volumes. Consider a single-mode sinusoidal perturbation that does not grow isotropically, but instead maintains a constant wavelength and increases in amplitude. Modelling each bubble or spike as a cylinder, their surface areas would scale as $S = \pi\lambda h(t)$ once $h(t) \gg \lambda$, and the inter-penetrating volume would scale as $V = \frac{1}{4}\pi\lambda^2 h(t)$. The surface to volume ratio would in this case be con-

stant and independent of $h(t)$. Consider now the case where the bubbles and spikes maintain geometric self-similarity (compare with bubble merger models discussed in §1.2.1). Irrespective of whether the most appropriate model for an individual structure is a cylinder or a sphere, $\lambda \sim h$, so $S \sim h^2$, and $V \sim h^3$. However, as bubbles merge and/or engulf one another and become larger, fewer can be accommodated across any horizontal plane, so their numbers reduce by $\frac{1}{h^2}$. From this analysis, we would expect the ratio $\frac{S}{V} \sim \frac{1}{h}$, and since $h \sim t^2$ and $V \sim h \iint dxdy \sim t^2$, that would suggest that the iso-surface area remains constant in time. This runs counter to our understanding of turbulence and its role in enhancing the rate at which inter-diffusion of species can occur.

It is generally believed that the stirring of a fluid interface at high Reynolds number induces a net transfer of kinetic energy from large length-scales to successively smaller ones. Any scalar interface in the turbulent field will be advected around by vorticity associated with this energy, and the vorticity tends to organise in structures that preferentially stretch scalar interfaces, steepening the gradients by reducing the length-scale over which scalars change concentration. Diffusion occurs at very low rates per unit surface area, but in a turbulent field the available surface tends towards being volume-filling, and therefore interleaving at a molecular level can occur rapidly.

The degree to which iso-surfaces become space-filling can be quantified by their fractal dimension, a parameter which has been used extensively in the past to describe Rayleigh-Taylor instability and other interfacial fluid flows (*e.g.* Linden *et al.* (1994); Dalziel *et al.* (1999); Maxworthy (1987)). One convenient property of fractal objects is that the volume enclosed by a fractal surface of dimension $D \in \mathcal{R}$ has dimension $D + 1$, so if the experimental cross-sections of surface and volume are fractal objects then we would expect the surface cross-sectional length (\hat{S}) to volume cross-sectional area (\hat{V}) ratio to behave like

$$\frac{\hat{S}}{\hat{V}} = \frac{r^{1+\delta}}{r^{2+\delta}} \quad (5.12)$$

where $D = \text{floor}(D) + \delta$, $0 < \delta < 1$, and r is simply a length-scale parameter. For $\delta = 0$, $\hat{S}(\hat{V})$ is a quadratic function, and for values of δ in the valid range the deviation from quadratic is small. Previous work of Dalziel *et al.* (1999) and Lawrie

& Dalziel (2006b) have both obtained $\delta \approx 0.5$.

We have now developed two estimates for the relationship between cross-sectional surfaces and volumes, one argued on the basis of geometric self-similarity that $S(t) = \text{const.}$ and one from scale-invariant self-affine self-similarity given by equation 5.12. Figure 5.19 shows a scatter plot of an ensemble of Acridine experiments and a matching (synthetic diagnostic) MOBILE simulation. The arrow indicates the direction of time evolution. The normalisation of \hat{V} is by the total image-plane area, and the \hat{S} contour is computed at the imaging resolution and is represented by pixels tagged where the volume fraction thresholds are traversed. The total length of the contour is calculated from the number of tagged pixels and normalised again by the total image-plane area. Testing at reduced resolution shows that the relationship between \hat{S} and \hat{V} is independent of the resolution at which the contour is computed. In the early stages of the instability growth (when \hat{V} is small), \hat{S} and \hat{V} are quite clearly linearly related, and this represents an intermediate state between geometric self-similarity and fractal self-similarity. At later time (though well before the instability reaches the top of the tank) \hat{S} grows more slowly per unit increase in \hat{V} . This trend - although it occurs somewhat earlier - is particularly marked in the numerical simulation, and that it occurs in both experiments and simulation suggests that the trend is independent of initial conditions. Examination of the video images of both experiment and simulation shows that the change in behaviour occurs when the instability height $h(t)$ is comparable with the width of the domain, and this coincides with the emergence of an instability mode with a wavelength the size of the domain width. It would appear that this disrupts the self-similar development that appears at early stages when the instability is unconfined by the tank dimensions. It seems reasonable to conclude from the experimental evidence that where the Rayleigh-Taylor instability is unconfined, a clean, well defined linear surface to volume relationship is observed, and while there is no obvious model for this that fits our pre-conceptions of self-similarity, these results are consistent with some intermediate form of self-similarity.

5.4 Summary

In this chapter a new diagnostic technique developed to be able to visualise directly the molecular mixing taking place in miscible Rayleigh-Taylor instability. Until we can resolve down to the Batchelor scale, conventional PLIF measurements can only provide an upper bound on the mixing taking place; Reactive Light Induced Fluorescence (RLIF) uses a chemical indicator of mixing and is therefore a resolution-independent diagnostic. A detailed calibration of the chemical and the optical filtering process is presented. By inverting the filtering, a method for recovering volume fraction from optical data is suggested. Experimental results using RLIF are compared with MOBILE simulations. As in previous studies, simulations under-predict the growth rate of Rayleigh-Taylor instability. Notwithstanding this discrepancy, numerical predictions of the local molecular mixing fraction compare well with conventional PLIF experiments, and with data from steady-state water-tunnel experiments. The global molecular mixing fraction as defined and applied historically was found not to be robust, and an alternative measure which could be used to directly compare molecular mixing information yielded by the RLIF diagnostic with conventional PLIF has been developed. The technology used in RLIF measurements is not yet sufficiently refined to demonstrate satisfactory comparison, though by carefully matching the post-processing of simulations with the experimental diagnostic, it is clearer where the discrepancies lie. RLIF excels at identifying iso-surfaces of volume fraction, and this information has been used to identify a form of self-similarity in Rayleigh-Taylor instability that is intermediate between a fractal and geometric self-similarity.

Chapter 6

Mixing in confined geometries - simple stratifications

6.1 Introduction

Rayleigh-Taylor instability has been extensively studied in configurations where geometry has played little role in modifying its evolution. When unconfined, one feature of Rayleigh-Taylor instability is that the characteristic bubble and spike structures (both $O(1)$ aspect-ratio at low Atwood number) increase in length scale with time. Once these structures reach parity with the horizontal length scale of the domain, they cannot continue to develop in a structurally self-similar fashion. The flow features then no longer have the appearance of bubbles and spikes, but more closely resemble turbulent eddies that commonly arise in constant density flows.

Changing the nature of the turbulent motion affects the growth rate of the instability. We know from analyses in chapter 4 that the ideal Rayleigh-Taylor process begins with exponential growth of initial interface perturbations, but remains in this state for only a short period. Thereafter the instability moves into the classic, non-linear, structurally self-similar regime with quadratic growth. When the turbulent length-scales are restrained by geometry, the instability decelerates and grows as a rational power of time. A schematic description of the Rayleigh-Taylor lifecycle is indicated in figure 6.1. Depending on the precise geometric configuration and initial conditions, the relative time-scales of the development stages vary. In the experi-

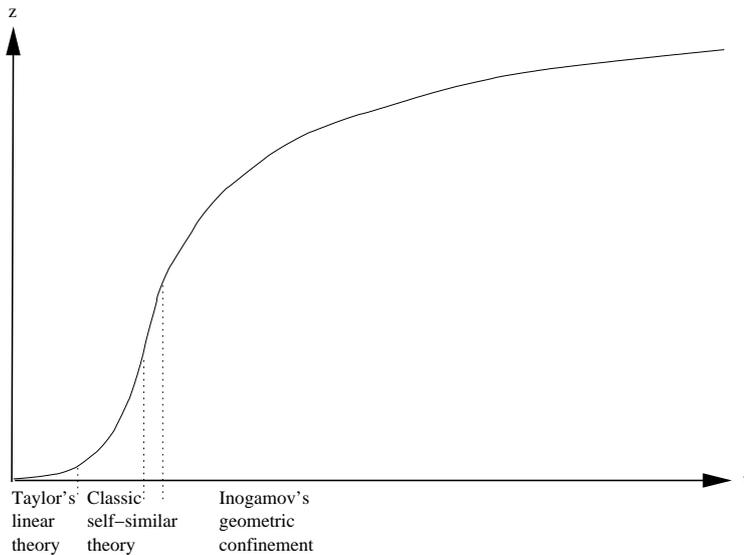


Figure 6.1: Schematic depicting the complete evolution of $h(t)$, from the early linear stage through the non-linear self-similar regime to the region in which geometric confinement restricts Rayleigh-Taylor development.

mental configuration studied in this chapter, the maximum eddy size is restricted by the walls of a tall thin tube. Instead of motion in the tank dissipating after 2-3 minutes, as one observes in a small low aspect-ratio domain, motion in the high-aspect-ratio case continues for 2-3 hours, and except for a neglectably small time at the beginning, all of this motion is in the final rational power regime.

In this chapter, an analytical model, making the assumption of profile self-similarity and thus termed a ‘zero-dimensional model’, is proposed in §6.2.1 to describe the flow. A one-dimensional numerical model is developed in §6.2.2 that makes more detailed predictions of the system, and then numerical simulations using MOBILE, considered in §6.2.3, provide three-dimensional modelling. Two experimental configurations are studied, the classic case studied experimentally in Dalziel *et al.* (2008), in which a two-layer statically stable density profile in the tall tube is overturned, and a static case where less dense fluid rises into the tall tube from a large reservoir. This second configuration enables more complex initial stratifications to be studied in chapter 7. Comparison with experiments in both cases shows that these models predict with surprising accuracy the functional form of the fluid behaviour. The high-aspect-ratio Rayleigh-Taylor instability is therefore a sufficiently simple benchmark problem from which modelling of more complex problems can be

explored.

6.2 Modelling approaches

The modelling approaches used in this chapter develop from the gradient-diffusion method described in §4.2.5. To reiterate briefly, Prandtl (1925) first proposed that scalar transport in a zero-mean flow, while complex and non-linear locally, could be modelled in aggregate by a simple scalar diffusion,

$$\frac{\partial \phi}{\partial t} = \frac{\partial}{\partial z} \left(\kappa \frac{\partial \phi}{\partial z} \right), \quad (6.1)$$

where ϕ is a scalar being diffused. Prandtl noticed that the factor κ has dimension $\frac{L^2}{T}$ and proposed that this could be expressed as

$$\kappa = \gamma l_{turb} u_{turb}, \quad (6.2)$$

with l_{turb} and u_{turb} representative length and velocity scales, respectively, and γ an arbitrary constant to be determined by calibration. Definitions for l_{turb} and u_{turb} are problem-specific, but in the Rayleigh-Taylor case there are natural choices.

Here, we assume that the overall dynamics of the system can be adequately represented by the interaction of buoyancy, inertial and viscous forces, as described by equations 4.31. In the non-linear regime of Rayleigh-Taylor instability, molecular viscosity is relatively unimportant, so the force balance determining the evolution of the system is between inertia and buoyancy. A velocity scale,

$$u_{turb} = \sqrt{\frac{l_{turb}^2 g \frac{\partial \rho}{\partial z}}{\rho}}, \quad (6.3)$$

can be deduced from this balance, provided a representative length-scale l_{turb} is known. When unconfined, the length l_{turb} , (a measure of some representative eddy size) would normally grow as a function of the mixing zone height, whereas in a narrow tube, this cannot happen. It has been observed (Dalziel *et al.* (2008)) that the velocity field contains approximately circular structures, so the length-scale in all directions is therefore limited by the lateral constraint of the tube. Hence in the growth phase before the geometry is a serious constraint (a negligibly short time-scale), $l_{turb} \sim t^2$ until it is limited by the tube and thereafter remains constant.

With these choices for l_{turb} and u_{turb} , we can proceed to examine the Rayleigh-Taylor instability when it is laterally constrained.

6.2.1 Similarity model

Experimental observation of Rayleigh-Taylor flows suggests that the shape of the horizontally averaged density profile is approximately invariant while the instability is growing. To make progress analytically, we invoke this observation and develop a model in which the vertical profile of a conserved scalar ϕ can be expressed in the form

$$\phi(z, t) = \hat{\phi}(\zeta) h(t), \quad (6.4)$$

where ζ is a non-dimensional height. Following [Dalziel *et al.* \(2008\)](#), in which the Rayleigh-Taylor instability was initiated by overturning a stable density interface at $t = 0$, a model for the transport of ϕ can be developed. If the initial interface is chosen to coincide with $z = \zeta = 0$, and ϕ is a proxy for relative density and sits in the upper half of the tube at $t = 0 + \epsilon$, then the total quantity of ϕ that has migrated across $z = 0$ after some time t is proportional to the height $h(t)$ of the instability. The scalar flux across $z = 0$ is driven by turbulent diffusion caused by the instability, and it might seem reasonable that this could be modelled as a linear function of the scalar gradient at $z = 0$ as follows,

$$\iiint_{z < 0} \frac{\partial \phi}{\partial t} dV = \kappa b^2 \left. \frac{\partial \phi}{\partial z} \right|_{z=0}. \quad (6.5)$$

This step invokes an assumption of global self-similarity, and reduces the problem to determining the time-evolution of flux at a single point, $\zeta = 0$. For this reason this similarity model is considered to be ‘zero-dimensional’. Substituting the separated variables form for ϕ and rearranging, we have

$$h \frac{dh}{dt} = \kappa b^2 \frac{\left. \frac{\partial \hat{\phi}(\zeta)}{\partial \zeta} \right|_{\zeta=0}}{\int_{-\infty}^0 \hat{\phi}(\zeta) d\zeta}, \quad (6.6)$$

for a square tube of width b . By self-similarity, $\left. \frac{\partial \hat{\phi}(\zeta)}{\partial \zeta} \right|_{\zeta=0}$ and $\int_{-\infty}^0 \hat{\phi}(\zeta) d\zeta$ terms are constant and can thus be ignored, giving the compact description

$$h \dot{h} \sim \kappa, \quad (6.7)$$

provided the forces of buoyancy and inertia remain in balance. The diffusivity κ can be replaced with definitions for the representative turbulent length and velocity scales,

$$h\dot{h} \sim l_{turb}u_{turb}, \quad (6.8)$$

and in this simple scaling argument these scales can be evaluated at the $z = 0$ plane. Assuming a self-similar scalar profile, an approximation to the density gradient $\frac{\partial \rho}{\partial z}$ at $z = 0$ is given by $\frac{\Delta \rho}{\Delta z}$, with

$$\begin{aligned} \Delta z &= h \\ \Delta \rho &= \rho_u - \rho_l. \end{aligned} \quad (6.9)$$

Thus we can now evaluate a velocity scale of the form of equation 6.3 and hence

$$h\dot{h} \sim b\sqrt{\frac{gb^2}{\rho} \frac{\Delta \rho}{\Delta z}}. \quad (6.10)$$

At no point is there a need to solve for the functional form $\hat{\phi}(\zeta)$. Where ρ appears in the denominator of equation 6.10, a Boussinesq reference density $\rho_b = \frac{1}{2}(\rho_l + \rho_u)$ is chosen. This leads to the relationship

$$h\dot{h} = \sqrt{\frac{2gb^4 A_0}{h}} \quad (6.11)$$

where A_0 is the initial Atwood number. We solve an integral equation of the form

$$\int \sqrt{\frac{h^3}{A_0} \frac{dh}{dt}} dt \sim \int C dt \quad (6.12)$$

where $C = \sqrt{2gb^4}$, and the functional dependence $h(t)$ is easily obtained:

$$\begin{aligned} \int h^{\frac{3}{2}} \frac{dh}{dt} dt &\sim t \\ h &\sim t^{\frac{2}{5}}. \end{aligned} \quad (6.13)$$

This result provides a useful rule-of-thumb estimate of the growth rate of the instability in confined geometries, and agrees with the somewhat less succinct derivation of Inogamov *et al.* (2001). Inogamov also showed, as reiterated in §4.2.5, that $h \sim t^2$ is recovered if $l_{turb} \sim h$, and the same result can be obtained from the above exposition. It should also be noted that previous work on these high-aspect-ratio Rayleigh-Taylor problems (Debacq *et al.* (2001, 2003)) has modelled κ as a constant, which gives a

growth scaling with a different rational power $h \sim t^{\frac{1}{2}}$ and this appears consistent with their measurements in a very narrow tube. A more detailed discussion of this is presented in Dalziel *et al.* (2008).

6.2.2 Non-linear gradient-diffusion model

It is difficult to increase the sophistication and predictive capacity of models while still pursuing analytical solution. Recourse to numerical evaluation permits a more complete model to be constructed, which can inform us easily about behaviour in situations where analytical models are difficult to construct.

Motivated by the similarity solution presented in §6.2.1, a numerical model is developed which defines, as before, the diffusion parameter κ using Prandtl's mixing length hypothesis with l_{turb} fixed by geometry, but in this case the velocity scale u_{turb} is determined according to the *local* density gradient. The turbulent flux function is hence evaluated locally and no assumption of global self-similarity is invoked. Density diffusion is modelled according to equation 6.1 and is approximated using a time-explicit finite-volume model,

$$\rho_i^{n+1} = \frac{1}{\Delta z_i b^2} \left(\Delta z_i b^2 \rho_i^n + \Delta t \left(F_{i-\frac{1}{2}}^n - F_{i+\frac{1}{2}}^n \right) \right). \quad (6.14)$$

This model has been used with the second experimental configuration mentioned in §6.1, namely the static tall tube overlying a large reservoir. To provide boundary conditions appropriate for this case, ghost cell values selected to provide a zero-flux boundary condition at $z = z_{max+\frac{1}{2}}$ and maintain an imposed density boundary condition at $z = z_{min-1}$. The reservoir was assumed to be large enough not to be affected either in density or dye concentration by discharge of dense water from the tube (acceptable since the tube is less than 4% of the reservoir volume) and as a first approximation the initial reservoir density was imposed as a boundary condition at $z = z_{min-1}$.

The flux definition for mass diffusion was taken as

$$F_{i+\frac{1}{2}} = -u_{turb} l_{turb} \frac{\partial \rho}{\partial z} \quad (6.15)$$

with the definition for u_{turb} taken as equation 6.3 where this is consistent with the potential energy available. As discussed in §6.2.1, l_{turb} is taken as the tube width b .

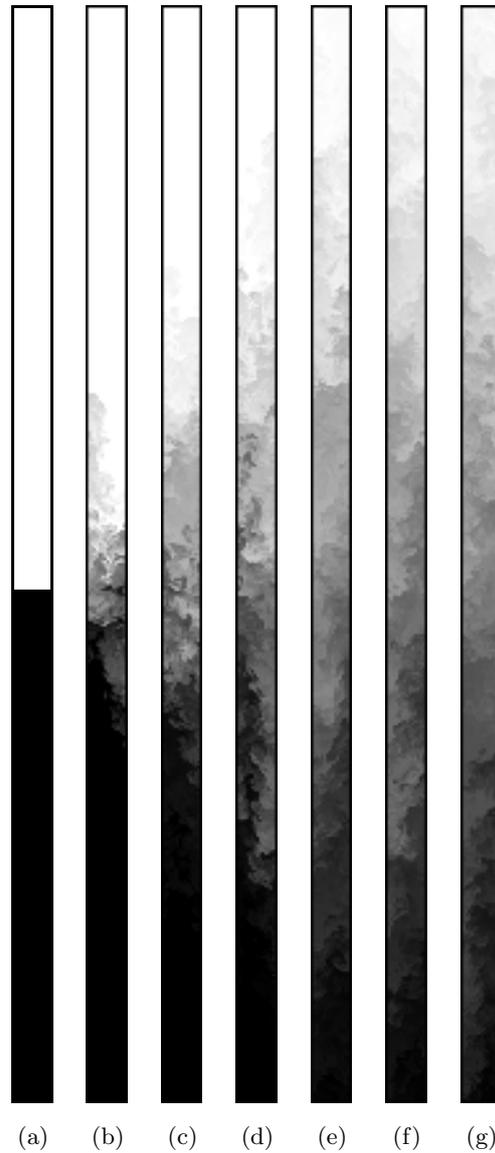


Figure 6.2: Scalar concentration taken on the vertical mid-plane of a $64 \times 64 \times 2560$ MOBILE simulation. Sections are taken at (a) $t = 0$, (b) $t = 20\text{s}$, (c) $t = 40\text{s}$, (d) $t = 60\text{s}$, (e) $t = 80\text{s}$, (f) $t = 100\text{s}$, (g) $t = 120\text{s}$. The Atwood number is $A = 1.5 \times 10^{-3}$.

The only adjustable constant in this model is γ , and in the absence of any *a priori* estimate, this was set to unity. While dye and density fields are linearly related in the current, very simple density stratification, this is not true in general, so a second scalar equation with the same diffusion coefficient κ and appropriate boundary/initial conditions was evolved in tandem with the density field to represent the behaviour

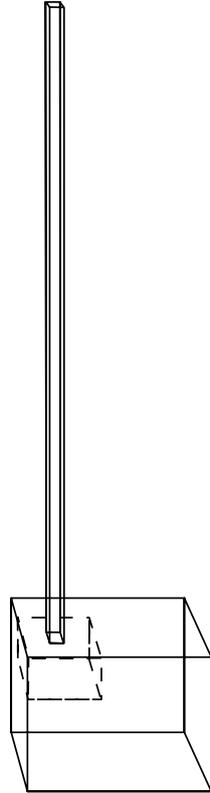


Figure 6.3: The dashed line indicates the simulated reservoir, which is 5.7% by volume of the reservoir used in the experiment. Reservoir and tall tube are depicted approximately in perspective.

of dye in the experiment.

6.2.3 Full three-dimensional simulation

The MOBILE numerical model was used to simulate both the overturned tank and static/reservoir cases. Many available research codes are unable to model such configurations, however MOBILE was developed with this case in mind, and uses a multi-block structured approach to handle the geometry of the reservoir. The algorithm uses the ILES methodology, so while the small-scale details and wall boundary layers are unlikely to be accurately captured, the flow's bulk properties could be expected to be correctly estimated. Figure 6.2 shows a sequence of vertical slices through a simulation modelling the classic two-layer overturned configuration.

Unfortunately the memory requirements for simulating a large reservoir comparable with the experiment were beyond that available, if the flow inside the tall tube

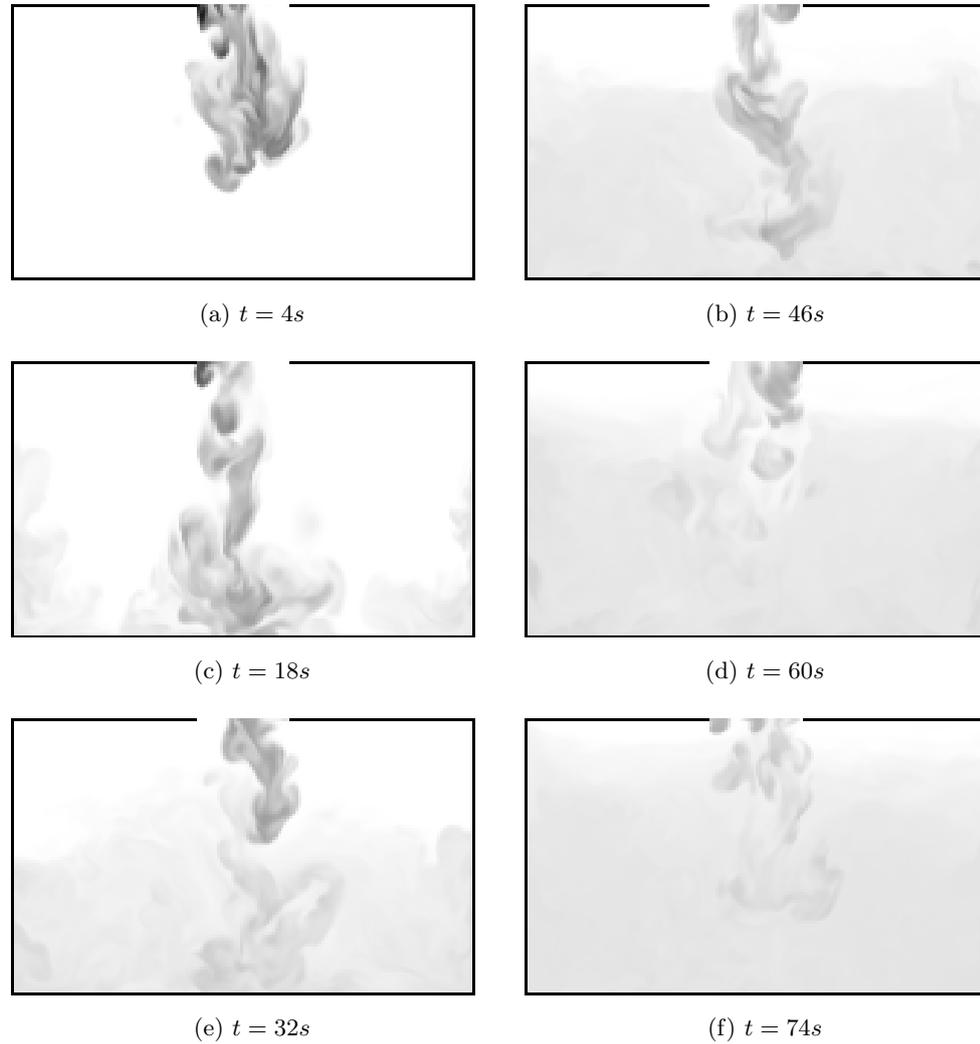


Figure 6.4: Scalar concentration taken on the vertical mid-plane of the reservoir. The Atwood number is $A = 7.3 \times 10^{-3}$. The greyscale colour scheme illustrates well how efficiently dye disperses in the computational reservoir.

were to be modelled at a meaningful resolution. To preserve mesh density at affordable cost, the size of the computational domain was reduced, as shown in figure 6.3.

Since an excessively small reservoir might significantly alter the dynamical behaviour inside the tube and introduce a modelling error in the region of interest, the flowfield inside the reservoir was checked to quantify the error. As can be seen from figure 6.4, the box fills with dense fluid from the bottom, thus for an extended period

the supply of relatively unmixed light fluid is maintained at the interface with the tall tube.

The computational reservoir represents the worst-case scenario for deviation from the ‘large reservoir condition’ that is needed for the one-dimensional numerical model of §6.2.2 and is approached in the experiment, yet the boundary condition at the tube-reservoir interface is remarkably close to having a constant density even in the simulations where the volume ratio is approximately 2:1. This is illustrated in figure 6.5.

6.3 Validation against experiment

6.3.1 Overturned configuration

The classic overturning tank case was studied to verify that the model predictions reflected the experimental results obtained with a dye-attenuation technique. The tall tube was illuminated with a backlight and the camera recorded any attenuation of the backlight due to fluid dyed red (dense fluid). By design, the dye attenuation was integrated along each ray path (approximately horizontal), and by taking a horizontal average across the width of the tank, the vertical profile of mean scalar concentration was obtained. The experimental image in figure 6.6 was formed by concatenating a sequence of such one-dimensional vertical profiles taken over 600 seconds. Existing experimental data from a parallel study Dalziel *et al.* (2008) was compared against the analytical prediction and the MOBILE simulation, as illustrated in figure 6.6.

The coefficients (estimates of $\frac{\partial \hat{\phi}(\zeta)}{\partial \zeta} \Big|_{\zeta=0}$ and $\int_{-\infty}^0 \hat{\phi}(\zeta) d\zeta$) in the similarity model were set using data given in Dalziel *et al.* (2008) to retrieve the 98% scalar concentration contour. When compared against one experiment from the ensemble presented in Dalziel *et al.* (2008), the coefficient values appear not to be very accurate, though there are several possible causes for the discrepancy. The most likely explanation is that the time origin appropriate for the experiment does not coincide with the corresponding origin for the similarity model. The act of overturning the stable interface to initiate Rayleigh-Taylor instability takes time, initiates a net circulation (in the rotating reference frame), and thus distorts the initial interface by up to 81 degrees

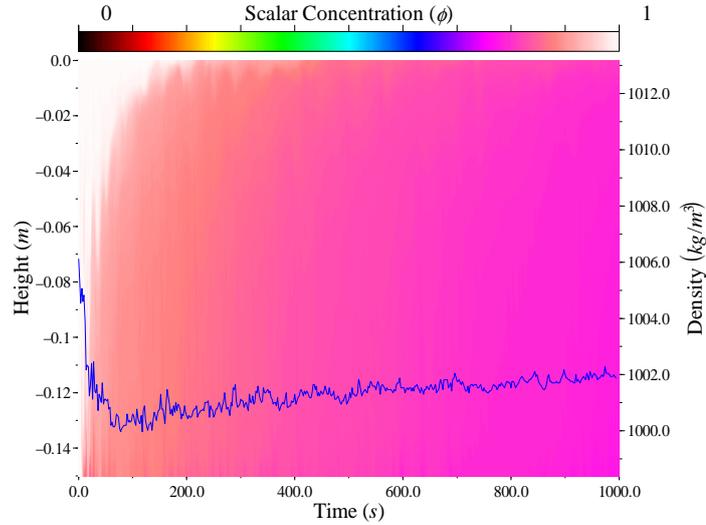


Figure 6.5: Time-series image of the horizontally averaged scalar concentration as a function of height in the computational reservoir. The blue superimposed curve shows the time-evolution of mean density at the interface between reservoir and tall tube. $A = 7.3 \times 10^{-3}$.

(see [Linden \(1977\)](#)). The point in time at which Rayleigh-Taylor instability begins is not well defined, and it happens over a finite region in space, as indicated by the initial smudge of horizontally averaged scalar concentration contours near $t = 0$.

Attempts were made to simulate the overturning by adding appropriate source terms (see §3.2.2) to the governing fluid equations 3.1. These terms result in distortion of the interface, and the standard Rayleigh-Taylor idealised initial conditions (isotropic initial random-amplitude random-phase density perturbation) becomes highly stretched in the rotation direction and inclined to the vertical. The stretching of the perturbations reduces the rate of generation of baroclinic torque, and hence the Rayleigh-Taylor growth rate, and the interfacial inclination favours shear rather than baroclinically driven instability. While both effects may to some extent occur in the initial stages of the experiment, visual observation suggests that the eddy size very quickly reaches parity with the lateral tube dimensions, and this was not observed on a comparable time-scale in the ILES simulations. It is quite possible that the velocity perturbations unavoidably introduced to the experiment by the tank overturning accelerate the breakdown of any shear-based instability into a density-driven Rayleigh-Taylor instability within the time-scale of the ‘smudge’.

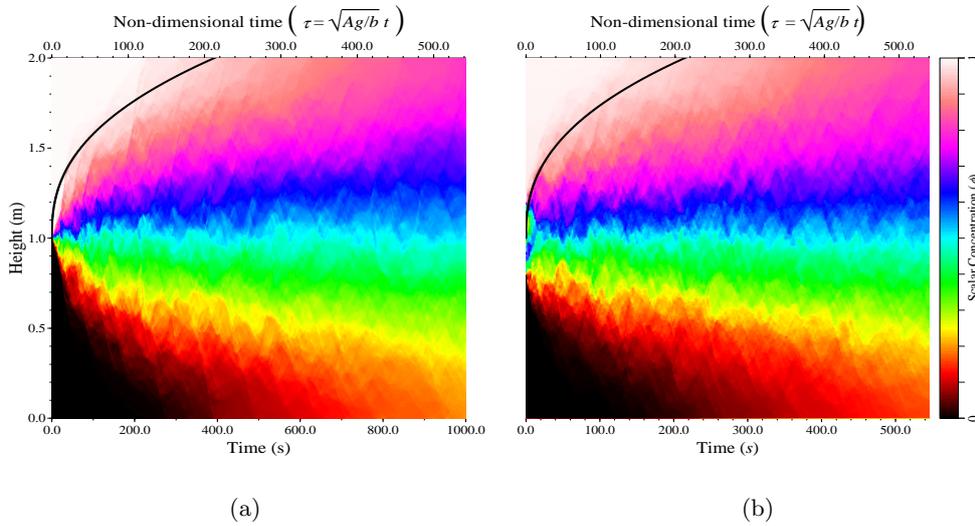


Figure 6.6: Time-series images of horizontally averaged scalar concentration from (a) MOBILE , and (b) overturned tank experiment of Dalziel *et al.* (2008). Theoretical prediction of $h(t) \sim t^{2/5}$ is shown in black, scaled according to the Atwood number. $A = 1.5 \times 10^{-3}$.

Velocity perturbations were not explored numerically, and non-overturned idealised initial condition Rayleigh-Taylor simulations were used instead to model the system, and this proved much more successful. The similarity model tracks the 98% idealised simulation contour more closely than the experiment, lending weight to the assertion that uncertainties in the experimental initial condition are responsible for the aforementioned discrepancy with the similarity model.

6.3.2 Static reservoir configuration

Previous experiments on high-aspect-ratio Rayleigh-Taylor instability have focussed on the overturned configuration, but this is unsuitable for studying the confinement of Rayleigh-Taylor instability driven mixing by stable stratification. To verify that the modelling assumptions that have been used in the overturning context remain applicable to the static reservoir configuration, two-layer Rayleigh-Taylor unstable experiments were performed in both cases, again using the dye-attenuation method and this time with light reservoir fluid being dyed red. One particular advantage of the static case is the relative insignificance of the initial condition. The experimental

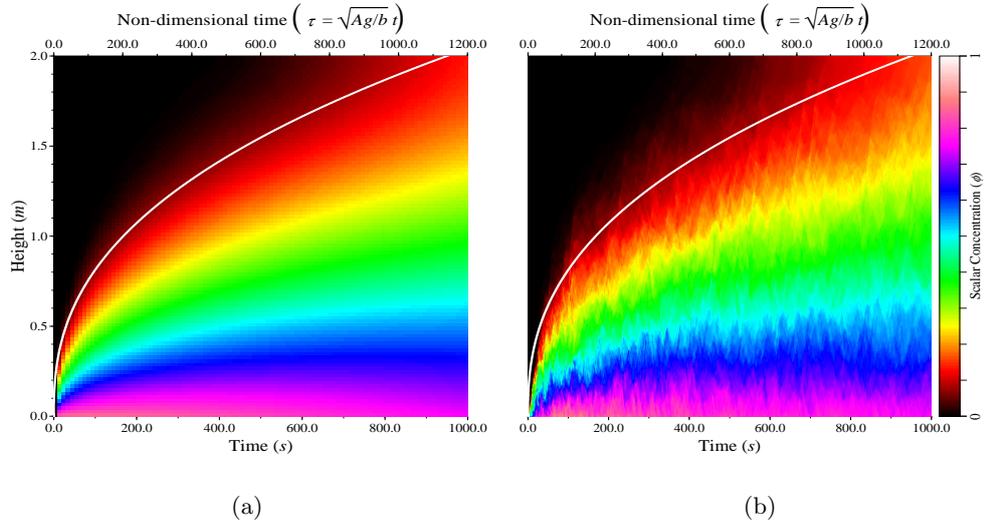


Figure 6.7: Time-series images of horizontally averaged scalar concentration from (a) one-dimensional model, and (b) MOBILE simulation in the static reservoir configuration. The white curves show $h(t)$ calculated using the zero-dimensional model. The boundary condition for the one-dimensional model is taken from figure 6.5. $A = 7.3 \times 10^{-3}$.

apparatus simply has a small plate initially supporting the fluid that is rotated out of the way to begin the experiment. Since the lateral dimensions are discontinuous at the tube-reservoir interface and the flow in the reservoir is of little intrinsic interest, the presence of the plate and any disturbance associated with its removal is insignificant on the time-scale of the experiment.

The same coefficients were used for the similarity model in figure 6.7 as for figure 6.6, but it appears that the scalar concentration contour being followed in the static/reservoir case is closer to 95% than 98%. The model parameters were determined in Dalziel *et al.* (2008) against a single overturned tank experiment and checked against experiments over a range of Atwood numbers (there is no Atwood number dependence on the coefficients). There is little reason to expect these values to change markedly between overturning and static configurations, since the non-dimensionalised density profile in one half of the overturned tube should correspond to the non-dimensionalised density profile in the static tube. However, uncertainties with the initial conditions in the overturned case may be the cause of a small calibration error in the calculated coefficient values.

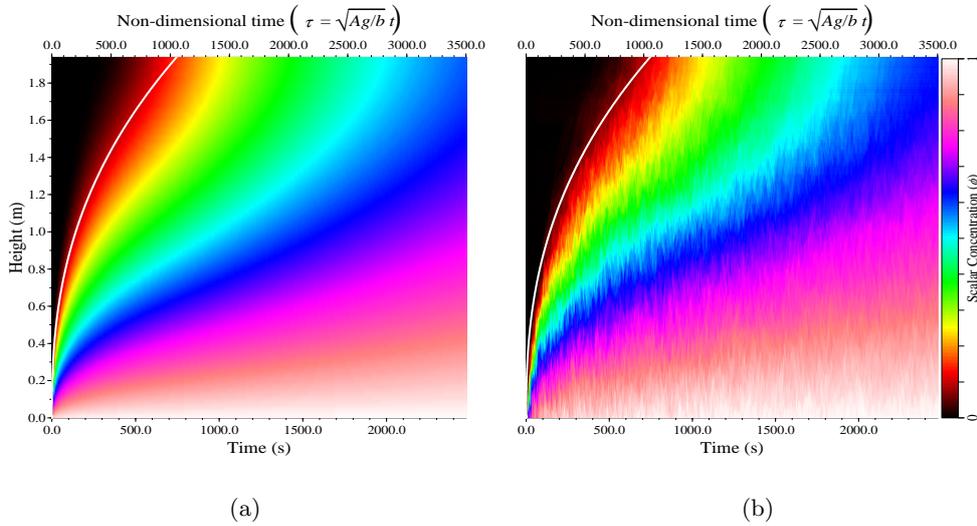


Figure 6.8: Time-series images showing horizontally averaged scalar concentration from (a) one-dimensional model, and (b) experiment in the static reservoir configuration. The white curve shows the zero-dimensional prediction of $h(t)$. $A = 1.02 \times 10^{-2}$.

To reduce potential sources for error in the comparison between the one- and three-dimensional models, the effective boundary condition at the tube-reservoir interface, shown in figure 6.5, is extracted from the ILES simulation and supplied to the diffusion model. The effective instantaneous Atwood number is reducing very slowly over time as the reservoir fills with dense fluid. Clearly this level of detail cannot easily be incorporated in the self-similarity model, so some late-time mismatch in their predictions is to be expected.

Since we have established that the zero- one- and three-dimensional models give self-consistent results in a test case where all initial and boundary conditions are known, the zero- and one-dimensional models were applied to a real experiment. Figure 6.8 illustrates this comparison. Here the zero-dimensional model tracks the scalar concentration boundary well, and the one-dimensional model predicts the experimental $\phi(z, t)$ very accurately. The value of γ used remains the *a priori* value of unity; there was no justification for modifying this. Any minor discrepancies in the concentration field can be attributed to calibration errors in mapping dye-attenuated light intensity to actual dye concentration, since this departs from a linear function

at low dye concentrations.

One particularly notable feature in figure 6.8 is the late-time behaviour. There appear to be inflexion points in scalar concentration contours, and the similarity model does not predict this. It is unclear whether or not the MOBILE simulations would, given enough time, since it was prohibitively expensive to run the calculations for so long. These inflexions are thought to be caused by the zero-flux top boundary condition, which the one- and three-dimensional models incorporate, and as seen in the figure, the location of the inflexion point propagates out from the top boundary once light fluid has contacted it. The information propagation downwards into the fluid must follow the same diffusion law as the scalars rising upwards, but at later time the assumption that buoyancy and inertia are the only forces in balance is somewhat questionable, and the evolution of the contours in this new regime has not been investigated.

6.4 Summary

This chapter has investigated the extent to which two configurations for laterally confined two-layer Rayleigh-Taylor instability can be modelled as a turbulent diffusion process. Prandtl's mixing length hypothesis was used as a closure for an analytical self-similarity model, and a more sophisticated non-linear one-dimensional model in which experimental boundary conditions can be represented. A three-dimensional ILES model was also used to approximate the full turbulent behaviour of the experiment, and the relevant comparisons are remarkably close.

Chapter 7

Mixing in confined geometries - linearly stratified upper layer

7.1 Introduction

In chapter 6 pre-existing work on high-aspect-ratio Rayleigh-Taylor instability in overturned tubes was compared with Rayleigh-Taylor instability between fluid in a static tube and a large bottom reservoir. This new configuration has two obvious advantages: the initial condition is more controllable, and the static tube can be filled with a more complex stratification profile than an overturned tube. This chapter considers the development of Rayleigh-Taylor instability into an upper layer linearly stratified in density, where the reservoir fluid has some intermediate density between the extrema of the stratification. The density profiles $\rho(z, t)$ considered herein have the following initial form:

$$\rho(z, 0) = \begin{cases} \rho_u - \beta(z - z_i), & z \geq z_i \\ \rho_l, & z \leq z_i \end{cases} \quad (7.1)$$
$$\rho_u > \rho_l$$
$$\beta \geq 0,$$

as indicated schematically in figure 7.1. On this basis representative scaling parameters can be chosen. Here an Atwood number is defined by densities ρ_u and ρ_l at

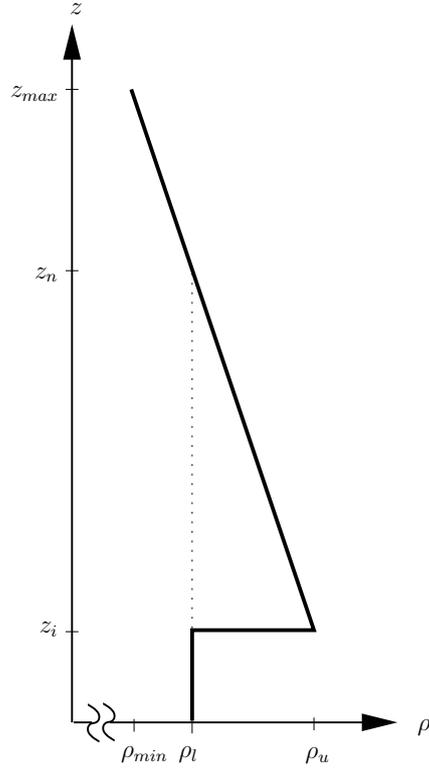


Figure 7.1: Initial stratification for stratified high-aspect-ratio experiment.

the initial interface height z_i :

$$A = \frac{\rho_u - \rho_l}{\rho_u + \rho_l}, \quad (7.2)$$

and for convenience we define a neutral buoyancy height

$$z_n = z_i + \frac{\rho_u - \rho_l}{\beta}, \quad (7.3)$$

which is the point to which reservoir fluid would reach if the initial stratification were to be adiabatically rearranged to be everywhere statically stable. The following section, §7.2, outlines how models proposed in §6.2 can be modified and extended for use in more complex density profiles. These are then compared against one another and against experiment in §7.3.

7.2 Model developments

7.2.1 Similarity model

The zero-dimensional (self-similarity) model of §6.2.1 is modified for penetration of Rayleigh-Taylor instability into stable stratifications by introducing an instantaneous Atwood number that represents the density difference over the Rayleigh-Taylor unstable region. As this region grows, by erosion of the stable stratification, the density difference gets smaller, until the upper boundary of the region reaches the neutral buoyancy height, where there is no remaining available energy to do further mixing and the instantaneous Atwood number correspondingly approaches zero. In this model the height $h(\infty)$ of the instability is necessarily given by the neutral buoyancy height: $h(\infty) = z_n - z_i$. Equation 6.10 becomes

$$\begin{aligned} h\dot{h} &= l_{turb}u_{turb} \\ &= b\sqrt{\frac{gb^2}{\rho} \frac{\Delta\rho}{h(\infty)}}, \end{aligned} \quad (7.4)$$

where Δz is the current height of the instability $h(t)$, and $\Delta\rho$ is in this case given by

$$\begin{aligned} \Delta\rho &= \rho(h(t)) - \rho_l \\ &= \rho_u + \frac{h(t)}{h(\infty)}(\rho_l - \rho_u) - \rho_l, \end{aligned} \quad (7.5)$$

for the linear stratification described by equation 7.1. Equation 7.4 simplifies to

$$h\dot{h} = \sqrt{(2gb^4A_0) \left(\frac{1}{h} - \frac{1}{h(\infty)} \right)}, \quad (7.6)$$

where A_0 is the initial Atwood number. In the forthcoming figure 7.7, the blue curve plots $h(t)$ obtained from a numerical integration of equation 7.6, using experimental values of ρ_u and ρ_l , and the neutral buoyancy point $h(\infty)$ which is known from the initial stratification. This model predicts the correct initial growth rate, since at early times the dynamics are indistinguishable from the homogenous, unstratified upper layer case, but the deceleration of the instability is underestimated, and hence the time-scale to reach the neutral buoyancy point is also underestimated. A more sophisticated model is required and the next section discusses this.

7.2.2 Mass diffusion model

The one-dimensional numerical model described in §6.2.2 uses equation 6.3 to define an appropriate local turbulent velocity scale u_{turb} . This is well defined in regions of the fluid with $\frac{\partial \rho}{\partial z} > 0$, but breaks down in a stable stratification. To resolve this anomaly, an additional energetic condition was introduced to the model. In this closed, initially quiescent system, a turbulent velocity can only be generated by baroclinic release of potential energy, so where the stratification cannot support such release, u_{turb} is identically zero. The conditional branch in the numerical model is thus

$$u_{turb} = \begin{cases} \sqrt{\frac{l_{turb}^2 g}{\rho} \frac{\partial \rho}{\partial z}}, & \frac{\partial \rho}{\partial z} > 0 \\ 0, & \frac{\partial \rho}{\partial z} < 0. \end{cases} \quad (7.7)$$

While this model is constructed to be energetically consistent, and can therefore describe mixing in an arbitrary vertical density profile with both statically stable and statically unstable regions, experimental evidence shows that a sharp interface forms between the Rayleigh-Taylor unstable region and the remaining quiescent part of the stable stratification. The detailed dynamics of the interface between these regions is complex and this simple model cannot accurately represent them.

The local physical process can be described as follows: the turbulent kinetic energy previously released from potential energy causes eddies, some of which penetrate a distance $O(b)$ into the otherwise quiescent statically stable region, engulfing this small quantity of (on average) more dense fluid into the kinetically active zone, releasing further potential energy and progressively eroding the stable stratification.

The analogous local process in the one-dimensional numerical model leads to a downward flux of mass through the unstable stratification, and away from the interface, reducing the density in the numerical cell at the top of the unstable region until it becomes less dense than the cell above. The conditional branch on the top face of this cell switches, across which diffusion can now begin, perpetuating the release of potential energy. The downward flux of mass away from the interface controls the rate of its upward progress, and this process operates on a macro-scale governed by the density gradient in the entire system, so the approximation of the interface

detail does not lead to an overall error in the predicted growth of the instability.

7.2.3 Energy transport model

An attempt to improve the one-dimensional numerical model's ability to capture the interface dynamics has been made by modelling the energetics of the system more carefully. An energy transport equation was added to permit the non-local production and dissipation of energy that would be necessary to achieve some representation of engulfment events. The energy exchange per unit volume between potential and kinetic as a result of mass diffusion can be explicitly evaluated with Prandtl's mixing model:

$$\Delta E_k \sim \rho b u_{turb}^2 = l_{turb}^2 g \frac{\partial \rho}{\partial z} \quad (7.8)$$

This additional kinetic energy ΔE_k may be allowed to diffuse around at a similar rate κ_T to the mass - subject perhaps to a turbulent Schmidt number. It must also dissipate at some rate ϵ . Unfortunately, at this point we encounter the classic problem of closure, since a transport equation for ϵ could be created, with sources and sinks that would also have to be modelled. Because we have no reliable way to model such parameters, we choose here to derive an approximate closure for ϵ .

We know that the turbulent length scale is constrained by geometry such that $l_{turb} = b$. We assume that the resulting eddies have an approximately circular structure of size b - as confirmed by observation - so the turbulent kinetic energy with a characteristic velocity u_{turb} induces an isotropic velocity gradient

$$\frac{\partial u}{\partial x} \sim 2 \frac{u_{turb}}{b}, \quad (7.9)$$

where the factor 2 indicates that in an approximately zero-mean flow by continuity the velocity difference across the tube must be twice the absolute value of u_{turb} . The shear stress generated on a vertical mid-plane in the fluid is therefore $\tau \sim 2\rho\nu\frac{u_{turb}}{b}$ and the energy dissipated as a result of this shear is

$$\iint_{\partial V} u \cdot \tau dS \sim \iiint 2\rho\nu \frac{u_{turb}^2}{b} \frac{dV}{b}, \quad (7.10)$$

where the surface integral is re-expressed as a volume integral for subsequent convenience. Noting that u_{turb} can be expressed as a function of E_k , the whole energy

process can be modelled thus:

$$\frac{\partial E_k}{\partial t} + \frac{\partial}{\partial z} \left(\kappa_T l_{turb} \sqrt{2 \frac{E_k}{\rho} \frac{\partial E_k}{\partial z}} \right) = l_{turb}^2 g \frac{\partial \rho}{\partial z} - c_\epsilon \frac{E_k \nu}{b^2} \quad (7.11)$$

where $\kappa_T = \gamma_{energy} u_{turb} l_{turb}$ is a turbulent energy diffusivity and the dissipation rate coefficient $c_\epsilon \leq 1$. Numerical investigation of the ‘second order’ system (coupled energy and mass transport) was performed with the initial aim of reducing the step-wise character the interface appears to have as it progresses from numerical cell to cell. While the local interface dynamics were not actually modelled significantly better, the exercise uncovered an interesting macro-scale feature of the system. Giving kinetic energy the freedom in the model to transport away from its point of production before it is dissipated leads to an increase in the overall kinetic energy in the system, which artificially raises the rate of diffusion and hence the instability growth rate $\dot{h}(t)$. No non-trivial combination of the free parameters in the coupled system ($\gamma_{mass}, \gamma_{energy}, c_\epsilon$) was found which could recover $h(t)$ as given by the mass transport model with $\gamma_{mass} = 1$. The trivial case where $c_\epsilon = 1$ implies instantaneous, local production and dissipation of kinetic energy, and the system reduces to the mass transport equation. Since, as illustrated later in figure 7.7, the mass transport model predicts $h(t)$ very well, this makes the important suggestion that the physical system also has an approximate *local* balance between kinetic energy production and dissipation.

7.2.4 Three-dimensional simulation

The three-dimensional simulation tool MOBILE was once again used to model the tall tube and reservoir experimental configuration, and the stratification measured from an experiment was used as an initial condition for a simulation. Development of the instability takes place over time-scales an order of magnitude larger than for the homogenous case, and since the maximum velocities that limit the numerical timestep in the simulation are $O(1)m/s$ even at relatively late times, resolving the $2m$ tall tube with $32 \times 32 \times 1280$ cells over $O(10^4)s$ would take $O(10^6)$ timesteps. This computational cost was unacceptable, so numerical investigation was limited to a single Atwood number and the simulation was truncated after $O(10^5)$ timesteps. Three

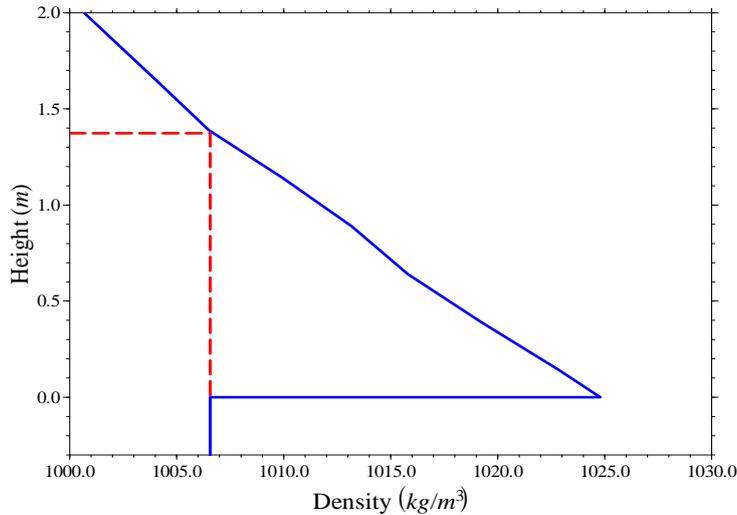


Figure 7.2: Initial stratification for run with initial Atwood number $A_0 = 10^{-2}$. The blue curve indicated the measured stratification, the red dashed lines indicate the neutral buoyancy height for reservoir fluid.

scalars were advected in this simulation, one representing the scalar concentration of reservoir fluid (dyed red in the experiment) and two others representing the dense and light extrema of the linear stratification. The baroclinic source term is evaluated from the volume-weighted scalar concentration in a computational cell. This is not numerically equivalent to advecting density directly (because of non-linearities in high order monotone fluxes), and therefore one cannot guarantee density monotonicity, though extensive tests suggest that monotonicity is achieved in practice.

7.3 Comparison with experiment

7.3.1 Model cross-validation

Having established in chapter 6 that the zero-, one- and three-dimensional models work well for unstably stratified two-layer systems, these models have been modified for stably stratified flows as detailed in §7.2, and are here used to examine the development of Rayleigh-Taylor instability into a stably stratified upper layer.

The Atwood number of the experiment selected for three-dimensional simulation was $A_0 = 10^{-2}$, and the stratification in the tall tube was very close to linear with a gradient $\beta = 12 \text{ kg/m}^4$ as shown in figure 7.2.

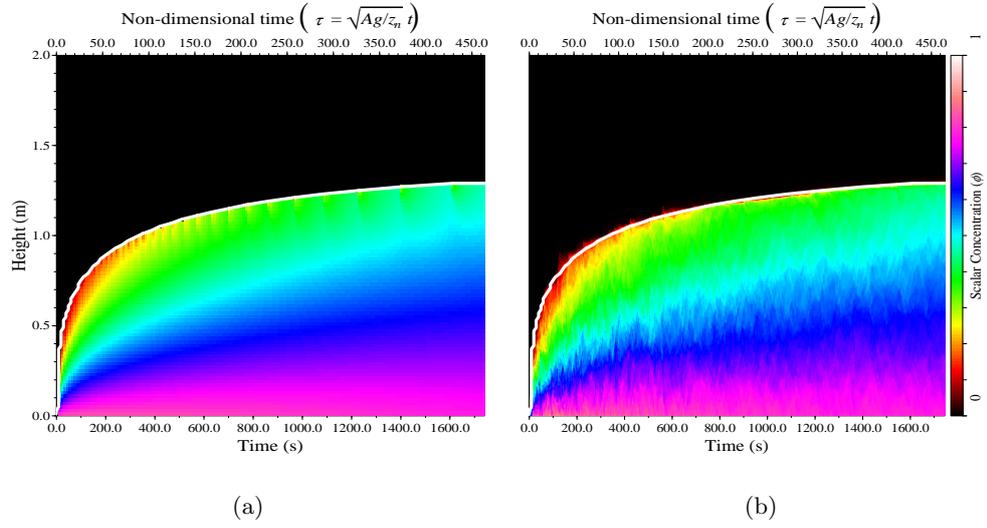


Figure 7.3: Time-series of horizontally averaged scalar concentrations as a function of height from (a) one-dimensional model and (b) three dimensional simulation. The white curve is the penetration height $h(t)$ obtained from the one-dimensional model, shown in both plots for comparison. A boundary condition correction for the small simulated reservoir of (b) has been made in (a).

The neutral buoyancy height is a rule of thumb prediction of the behaviour of the system, since there exists insufficient potential energy available to drive reservoir fluid above this point.

The prediction of the scalar concentration field made by MOBILE is shown in figure 7.3(b). As with the corresponding two-layer case, the computational reservoir that was affordable to simulate is much smaller than the reservoir used in the experiment, and although the reservoir fills with dense fluid from the bottom, the horizontally averaged density at the reservoir-tube interface,

$$\rho_{bc} = \frac{\iint \rho(z_i, t) dx dy}{\iint dx dy}, \quad (7.12)$$

sits at only a small offset from the initial reservoir density, and is a weak function of time. In figure 7.4 the boundary condition $\rho_{bc}(t)$ is plotted as a line, and the horizontally averaged density in the reservoir $\rho(z, t)|_{z < z_i}$ is shown as a time-series image.

To make a fair comparison between one- and three-dimensional models, the the interface density $\rho_{bc}(t)$ was implemented as a time-varying boundary condition for

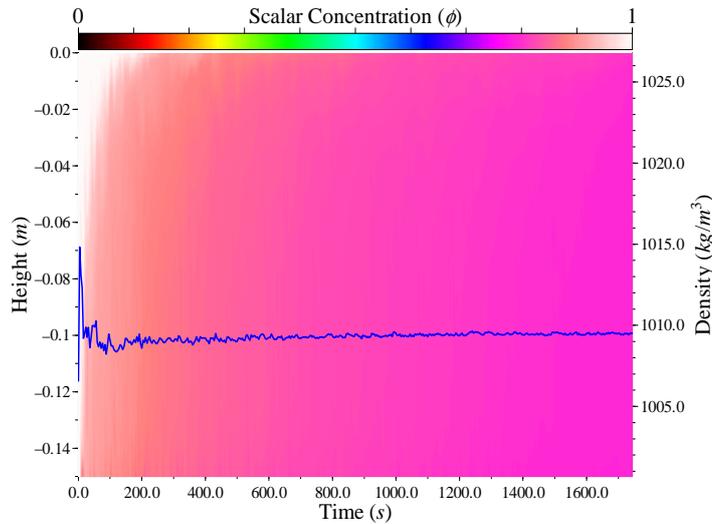


Figure 7.4: Time-series image of scalar concentration in the reservoir predicted by MOBILE . The blue curve denotes the horizontally averaged density at the interface between the tall tube and the reservoir.

the one-dimensional model. Comparison between the left and right plots of scalar concentration $\phi(z, t)$ indicates that the correction is appropriate and the envelope profiles match extremely well. There is some turbulent noise evident in the simulation, with the period and variance of the noise in any individual contour associated with motion at the time and length scales of individual eddies. There is clearly a large separation of time-scales between those governing the gross dynamics of the system and the detailed turbulent structure, and this perhaps indicates why the gross dynamics are modelled so well with such simple models, and it also explains why simulations are so expensive. There is sufficient time for potential energy acquired from more dense fluid from a point above to do mixing and/or dissipate before the structure of the density profile has changed markedly. Thus a model that enforces locality of energy production and dissipation could be expected to perform well.

As discussed in §7.2.2, the detail of the stable-unstable interfacial dynamics cannot be fully represented with the one-dimensional model. The separation of time-scales reduces the sensitivity of the system response to this error, but in figure 7.3(a), the scalar contours slightly beneath the interface are somewhat step-like. This has no great effect on the system overall, but these steps arise from the impulse supplied to the system by the conditional branch (equation 7.7) switching as the density

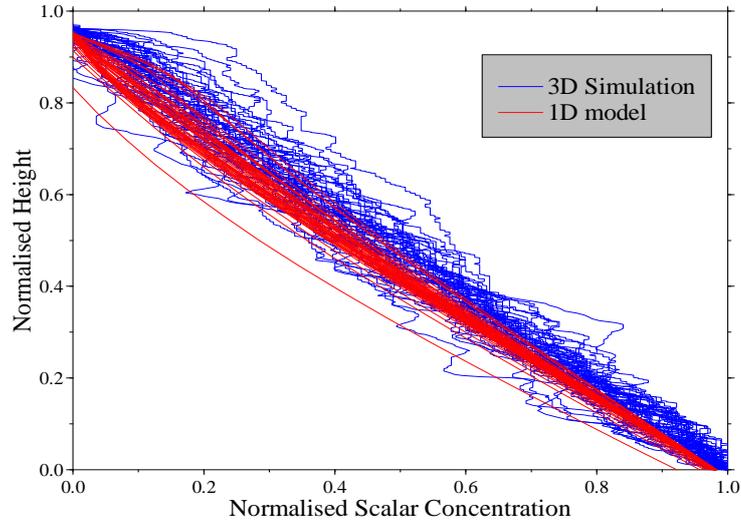


Figure 7.5: The red curves are one-dimensional predictions of the vertical scalar concentration profile, scaled for height and end-point concentrations. The blue curves come from the three-dimensional simulation.

on the interface cell reaches equality with the cell beneath. The interfacial region (unlike in the two-layer case) remains confined in vertical extent to $O(b)$ and over the range $h - b < z < h$ the vertical scalar concentration has a steep gradient and ϕ falls to zero at h . The scalar concentration at $h(t) - b$ increases over time until all the fluid more dense than the reservoir has fallen out of the tube. At this point $\phi(h(t) - b) = 1$.

Normalising the vertical scalar concentration profile in the kinetically active region with the value $\phi(h(t) - b, t)$ and normalising the height simply by $h(t)$, as in figure 7.5, it is clear that the interior profile is approximately linear. The slight curvature does change from concave to convex over time, a feature which is accentuated in the one-dimensional model by the discrete steps which mark the advance of the envelope at late times, and the high levels of kinetic energy at very early times which makes the assumption of a sharp interface confined to width $O(b)$ and clearly separated time-scales inappropriate.

Comparison of the one-dimensional model with the experiment as shown in figure 7.6 establishes that the scalar concentration field $\phi(z, t)$ and the instability growth function $h(t)$ match particularly closely with the experimental measurements. Some distortion of the experimental image occurs near the bottom of the tube, due to

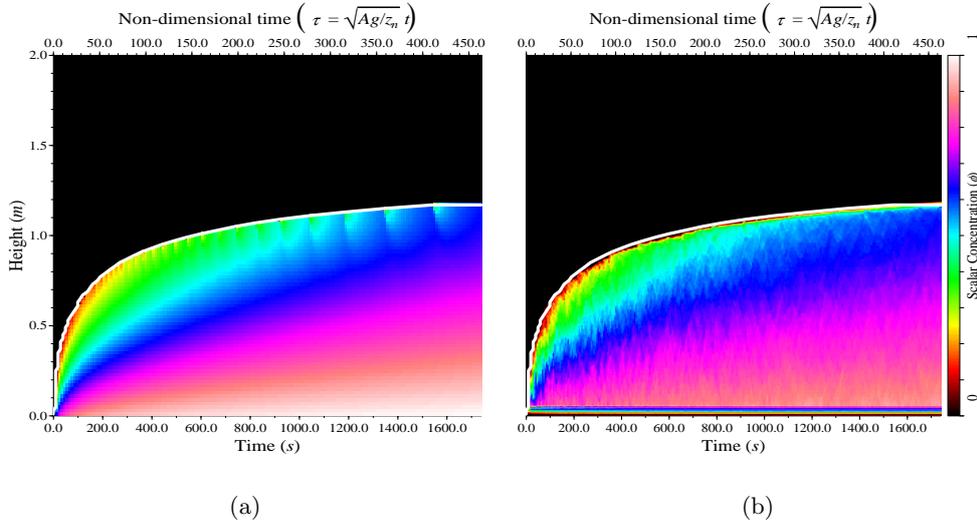


Figure 7.6: Time-series of horizontally averaged scalar concentration as a function of height from (a) one-dimensional model, and (b) experiment. The white curve is the penetration height $h(t)$ obtained from (a), and plotted over (b) for comparison. A correction for dye non-linearity has been made in the experimental time-series.

line-of-sight interference around the sliding hatch that releases the flow, and scalar concentration contours in this region cannot be regarded as accurate. Parallax correction was not used since the camera was $10m$ from the tube, and line-of-sight averages along rays of angle θ converging at the camera focal point have been mapped directly to a vertical ordinate z . The red dye used for the dye-attenuation measurements had a slightly non-linear light adsorption function when illuminated with white light, and this has been accounted for in the mapping of the attenuated light intensity field I to scalar concentration ϕ , $I(\theta, t) \rightarrow \phi(z, t)$.

An interesting qualitative comparison can be made between the period and variance of fluctuations in the scalar concentration contours in simulation and experiment. These fluctuations are associated with individual eddy structures in the flow. Although the sampling rates were severely limited by data storage capacity both in the experimental image capturing and the simulation, in each case there is a clear decay in the period of these fluctuations, indicating longer eddy turnover times and generally declining levels of kinetic energy in the system. This can be seen as encouraging confirmation that ILES modelling (in which kinetic energy leaves the system

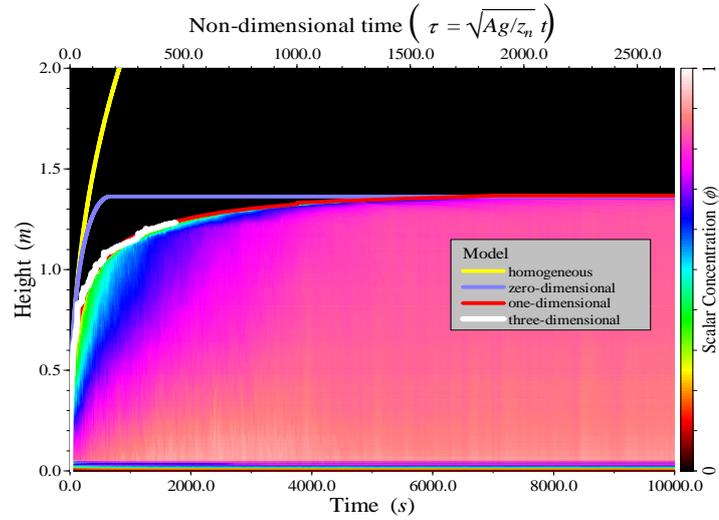


Figure 7.7: This plot shows compares all the model predictions with the experimental timeseries, and a $t^{\frac{2}{5}}$ curve is plotted for reference.

only through the enforcement of monotonicity of momentum) can predict with surprising accuracy not only the global properties of the system but some important aspects of its fine detail. It is unlikely to be coincidental that ILES performs very well when there is a clear separation of scales between detailed mixing dynamics and the instability growth $h(t)$, particularly when ILES performs less well on the low aspect-ratio Rayleigh-Taylor instability where there is no separation of mixing and instability growth scales. A separation of scales implies that kinetic energy is locally produced and dissipated, and it may well be that the ILES methodology embeds this assumption.

Figure 7.7 shows the ensemble of predictions for the experimental measurements examined in detail in this section. The $h(t) \sim t^{\frac{2}{5}}$ curve shown in yellow is the 98% contour prediction from the similarity model discussed in §6.2.1 for the unstable two-layer system, and is included to provide a comparison of time-scales. The modified model (which takes into account the time-dependence of the effective Atwood number) makes an unacceptable estimate, shown in blue, of the time-scales involved. While a scaling parameter could be included in the model to correct the predicted time-scale, the functional form of $h(t)$ does not correspond to the experiment and there is little *a priori* justification for involving arbitrary parameters

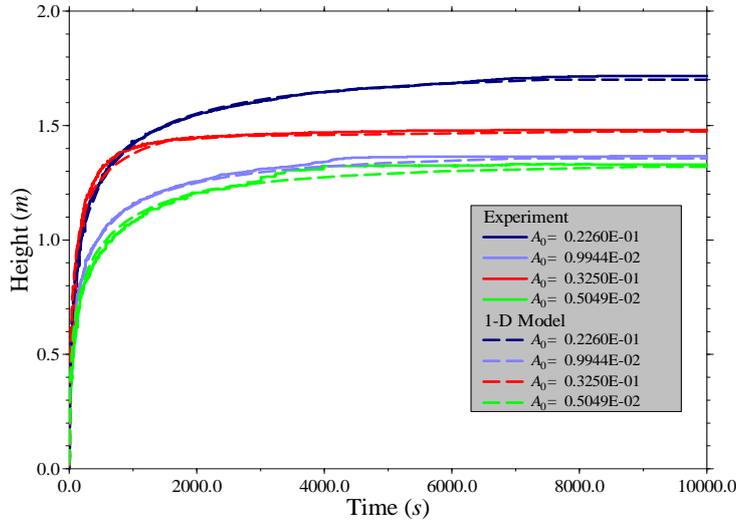


Figure 7.8: One-dimensional predictions of envelope profiles $h(t)$ for experiments at various Atwood numbers.

which would create a match. The intrinsic cause of the discrepancy is that the form of self-similarity used to develop the model is no longer valid, since the stratification gradient introduces a second length-scale to the problem. The one-dimensional model (in red) makes the best overall prediction of the envelope $h(t)$ and accurately estimates the time-scale. The MOBILE prediction of $h(t)$ is plotted in white, but the tight scalar concentration threshold used to define h is sensitive to turbulent fluctuations and is something of a disadvantage in making a fair comparison with other models. It should also be noted that no correction has been made for the small computational reservoir, and the close match of the MOBILE and experimental $h(t)$ trajectories implies that the modelling error introduced to cut the computational cost is insignificant.

7.3.2 Variation with Atwood number

Since the computational cost of modelling the full range of experimentally studied Atwood numbers using MOBILE was prohibitively expensive, the one-dimensional model, shown in §7.3.1 to be an excellent proxy for both experiment and simulation, was used to explore a wider range of parameters than those that could be tested experimentally or numerically.

Figure 7.8 shows $h(t)$ covering the experimentally feasible range of Atwood numbers, and compares experiments with corresponding one-dimensional model predictions. The model performs extremely well, except for very minor discrepancies during the transition from rapid growth rate to slow growth in two cases, and larger discrepancies at late time in the low Atwood number cases. The experiments were intended to have a common neutral buoyancy height $z_n = 1.5m$, but as the figure shows, this was not achieved. The causes of these three errors are discussed in turn below.

The errors during transition from rapid to slow growth may be due simply to inadequate sampling of the initial stratification. Measurements of the experimental stratification were taken at 10 points during filling, and these values were used to generate an initial condition for the one-dimensional model.

The errors in low Atwood number cases at late time are due to thermal effects modifying the fluid buoyancy. In these cases the potential energy released by eroding the stratification unexpectedly turned out to be of the same order as the heat transfer through the walls of the tube. There were a number of complicating factors in the experimental environment. Firstly the laboratory ambient was vertically stratified by $2^{\circ}C$, and the mean temperature had a diurnal variation, which was significant over the three-hour duration of each experiment. Secondly, part of the process of producing the stratification requires water to be pumped around a recirculating bucket (see §2.4.2 for details), thereby doing work on the water and raising its temperature above ambient. Thirdly, the reservoir was filled directly from the laboratory mains supply, around ambient laboratory temperature. The surface to volume ratio of the reservoir and the tube are very different, so even leaving the experiment set up for a long period would not have ensured uniform water temperatures, due to the diurnal variation in the ambient.

The mis-match of temperatures gave rise to an interesting phenomenon at the interface, known as double-diffusive convection, and this is illustrated in figure 7.9, a time-series excerpt taken just before $h(t) = z_n$. The mechanism for this form of convection is well known (see Huppert & Turner (1981)) and in the present example can be explained as follows.

In this system heat and salt diffuse at markedly different rates, since the large

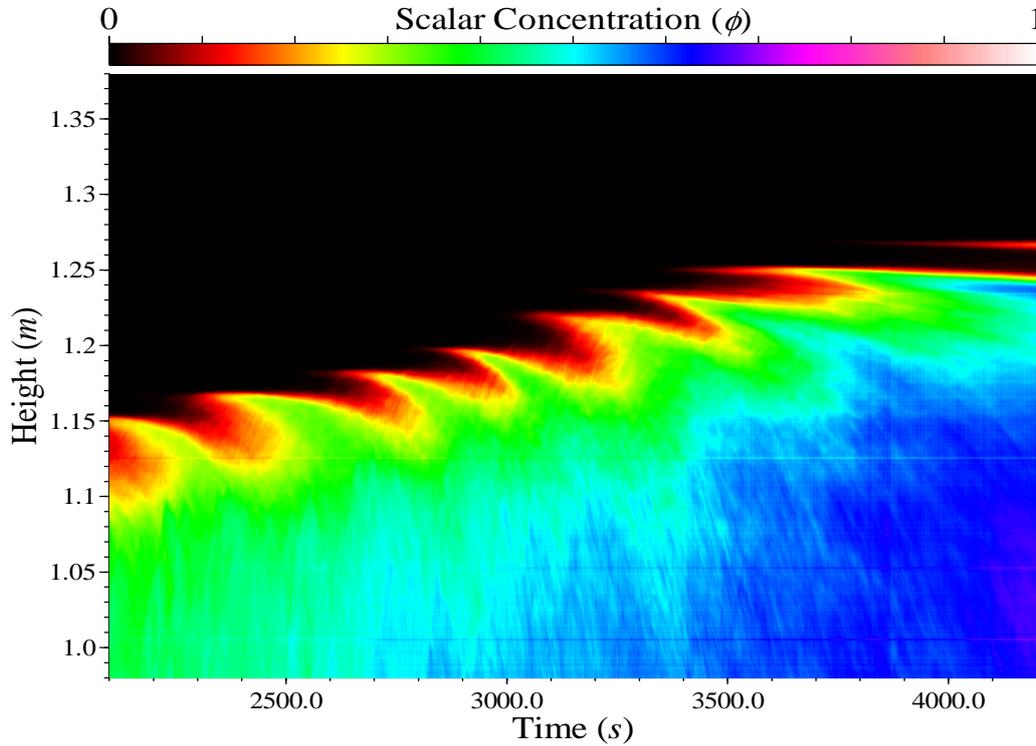


Figure 7.9: Time-series extract showing double-diffusive steps when primarily salt-driven mixing occurs in the presence of a temperature mismatch between the water temperatures and laboratory ambient. $A_0 = 0.5 \times 10^{-2}$.

relative surface area of the tube walls permits relatively quick thermal diffusion in and out of the system, while salt must remain in the system. The temperature gradient is stronger between ambient and the warmer undyed fluid from the initial stratification than the lighter, cooler, dyed fluid penetrating upwards from the reservoir. The undyed fluid, statically stable when warm, becomes relatively denser than its immediate surroundings as it cools, and needs to fall slightly to regain neutral buoyancy. By continuity, an exchange event must take place, lifting dyed fluid and giving rise to the step like progress of the dye front. The fallen fluid is undyed, and denser than anything else in the dyed region, having come from a height slightly above the dye interface (which coincides with the height of maximum density) and fallen to a height just beneath, and therefore continues to accelerate through the unstably stratified dyed region, as can be seen by small downward-pointing streaks in the scalar concentration profile. Over time molecular mixing smears out any mea-

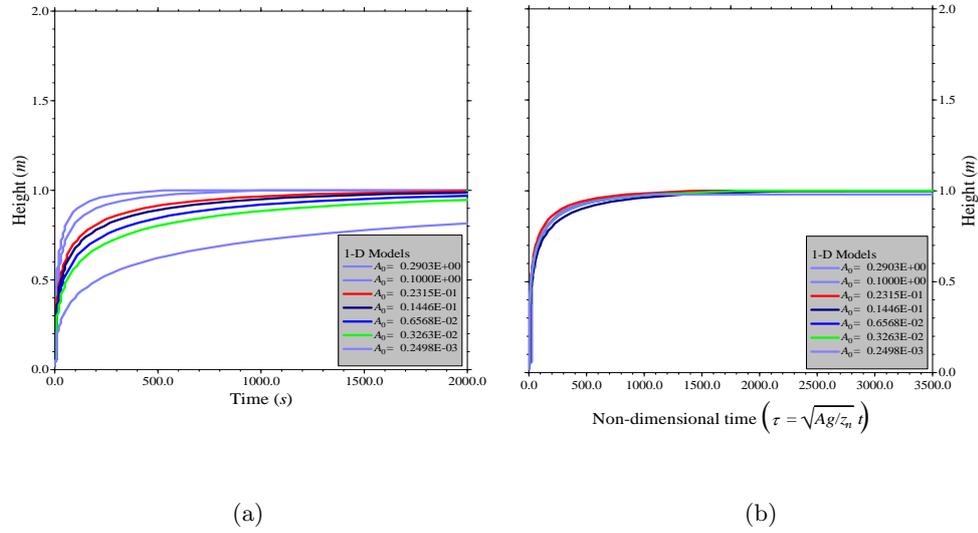


Figure 7.10: One-dimensional model predictions of (a) idealised initial conditions across a range $2.5 \times 10^{-4} < A_0 < 0.1$, and (b) self-similar collapse by scaling with non-dimensional time.

surable variation in the volume fraction of dye and so these streaks are only visible for a short time. Each discrete exchange event is a double-diffusive step, and they serve to increase the growth rate $h(t)$ of the dye front, as shown in the blue and green experimental curves of figure 7.8 but not predicted by the one-dimensional model.

The final error, the unintended variation in neutral buoyancy height, was caused by a minor experimental error when stratifying the tall tube, which was magnified by the very weak density gradient. A valve at the top of the tube permitted air (and in the event of over-filling, water) to escape benignly, and when shut would support the hydrostatic pressure of the water column. At the end of stratifying both the valve at the top and the peristaltic pump at the bottom had to be shut off, ideally simultaneously. In practice the pump was stopped just before shutting the valve to avoid over-pressurising the tube, and since the pump could not completely support the hydrostatic pressure, a safety margin of overfill was required so water could still leak through the pump at the bottom without ingesting air through the valve at the top. The tube was inconsistently over-filled with water, and this shifts the stratification and hence the neutral buoyancy point.

From the above discussion it is clear that experimental conditions were too sen-

sitive to initial stratification, thermal effects and experimental error to characterise the system's behaviour across a wide range of Atwood numbers. Additionally, these experiments were restricted in Atwood number by the solubility of salt in water, and some practical considerations due to the method of stratifying from the bottom rather than the top of the tube (see §2.4.2).

The one-dimensional model has no such restrictions. Having ascertained that it models the system well in situations for which comparable experiments exist, it was then applied over a wider range of initial Atwood numbers. The modelling assumptions remain valid even at very high Atwood numbers, since locally the density gradients are small and can still be treated as Boussinesq. Idealised initial conditions were created such that the neutral buoyancy height was $1m$ over a range of $2.5 \times 10^{-4} < A_0 < 0.1$, and the results are shown in figure 7.10, where the light-blue curves are at Atwood numbers than could not be studied experimentally, and the other colours correspond to Atwood numbers measured and shown in figure 7.8, though with a stratification chosen to achieve a consistent neutral buoyancy height. Immediate observations are that the functional form $h(t)$ seems independent of Atwood number, and the non-dimensionalisation $\tau = \sqrt{Ag/z_n}t$ is appropriate for self-similar collapse.

7.4 Summary

This chapter has investigated the extension of models, which were initially applied in chapter 6 to laterally confined two-layer Rayleigh-Taylor instability, to the case where the upper layer is stably stratified. The presence of a stable stratification has raised new issues about energetic consistency for the one-dimensional numerical model, which have been resolved and the model agrees very closely with experimental measurements. Further extension of the model to incorporate energy transport leads to the conclusion that potential energy is released into kinetic energy that is entirely dissipated locally, so there is no net transport of kinetic energy in this system. The three-dimensional ILES model performs very well on three harsh diagnostic tests, which suggests, at least in cases where turbulent eddy turnover times are well separated from the global system dynamic time-scale and energy is produced and

dissipated locally, that ILES captures many of the important features of the flow. Experimental artefacts due to the baroclinic interaction of temperature and salt were also explained, and a simple relationship between Atwood number and rise-time to neutral buoyancy was found.

Chapter 8

Mixing confined by stable linear stratifications

8.1 Introduction

The focus in this chapter is to examine the mixing behaviour of systems where Rayleigh-Taylor instability drives the flow, but additional features of the system interact with the developing instability to restrict its growth. Geometric confinement was considered in chapters 6 and 7; here we study Rayleigh-Taylor instability confined both above and below by a pair of stable linear stratifications. The initial condition is defined as

$$\begin{aligned}\rho(z, 0) &= \rho_u - \beta(z - z_i), z \geq z_i \\ \rho(z, 0) &= \rho_l - \beta(z - z_i), z \leq z_i \\ \rho_u &> \rho_l \\ \beta &> 0\end{aligned}\tag{8.1}$$

where ρ_u and ρ_l are densities just above and below the interface located at z_i , and β is a (chosen) stratification gradient. The interface density, $\rho(z_i, 0)$ is multi-valued. For convenience this is shown pictorially in figure 8.1, and an experimental illustration is shown in figure 8.2.

Subsequent sections are organised as follows: in §8.2 the energy budget analysis of §4.2.4 is extended to include stratified initial profiles, and compared with an

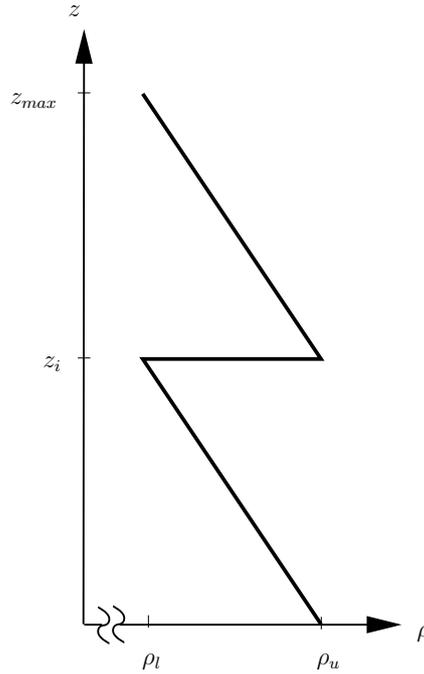


Figure 8.1: Diagram illustrating the initial density profile. Note that in this case the neutral buoyancy height z_n coincides with z_{\max} .

ensemble of experiments and a matching simulation. The concept of mixing efficiency is introduced in §8.3 and some simple relationships between the density profile and mixing efficiency are derived. These results are used to gain some understanding of the noticeable differences between the simulation and the ensemble of experiments, which are discussed in §8.4.

8.2 Growth law

Modifying the approach of §4.2.4 we can derive a growth law for this configuration. Using the same notation, we begin with the energy budget equation

$$\int \overline{\overline{\rho_0}} g z dz = \int \overline{\overline{\rho}} g z dz + \int \frac{1}{2} \overline{\overline{\rho \mathbf{u}^2}} dz. \quad (8.2)$$

However, it is difficult to progress by invoking a similarity argument, because the problem contains both the length-scale of the instability growth, and a fixed length-scale set by the stratification. We can progress by assuming a functional form for $(\overline{\overline{\rho_0}} - \overline{\overline{\rho}})(z, t)$, and evaluate directly. The simplest reasonable model of Rayleigh-Taylor instability assumes that fluid in the mixing zone of height h has mixed in-

stantaneously, and has acquired the uniform density $\frac{1}{2}(\rho_u + \rho_l)$. Thus

$$\begin{aligned} \frac{1}{S} \int_{-\infty}^{+\infty} (\overline{\rho_0} - \overline{\rho}) g z dz &= \int_{z_i}^{z_i+h} \frac{\rho_u - \rho_l}{2} - \beta(z - z_i) g z dz \\ &+ \int_{z_i-h}^{z_i} -\frac{\rho_u - \rho_l}{2} - \beta(z - z_i) g z dz \\ &= gh^2 \frac{\rho_u - \rho_l}{2} - \frac{2}{3} \beta gh^3, \end{aligned} \quad (8.3)$$

where $S = \iint dx dy$ is the area over which horizontal averaging takes place.

The kinetic energy, after making the Boussinesq approximation, becomes

$$\int_{-\infty}^{+\infty} \frac{1}{2} \overline{\rho \mathbf{u}^2} dz \sim \frac{\rho_u + \rho_l}{2} \left(\frac{\partial h}{\partial t} \right)^2 \overline{h}, \quad (8.4)$$

where we argue that kinetic energy $\overline{\mathbf{u}^2}$ scales with $\left(\frac{\partial h}{\partial t} \right)^2$. That energy is distributed over the mixing zone volume, which scales with \overline{h} , so the energy balance is therefore

$$gh^2 \frac{\rho_u - \rho_l}{2} - \frac{2}{3} \beta gh^3 \sim \frac{\rho_u + \rho_l}{2} \left(\frac{\partial h}{\partial t} \right)^2 h. \quad (8.5)$$

Noting that

$$\beta = \frac{\rho_u - \rho_l}{z_{max} - z_i}, \quad (8.6)$$

the integration to recover $h(t)$ is of the form

$$\int dt \sim \int \frac{1}{\sqrt{ghA(1-Bh)}} \frac{dh}{dt} dt, \quad (8.7)$$

where A is the Atwood number, and $B = \frac{4\beta}{3(z_{max} - z_i)}$, to which there is a solution

$$t \sim \frac{1}{\sqrt{gAB}} \tan^{-1} \left(\sqrt{\frac{B}{h - Bh^2}} \left(h - \frac{1}{2B} \right) \right) + h_0. \quad (8.8)$$

This simple model gives us some insight into the Rayleigh-Taylor instability when confined on each side by a stable stratification. Its central assumption is that, where energetically possible, fluids of different densities that are in proximity with each other reach a well-mixed state on a time-scale much faster than the rate at which the mixed layer grows. While this is a valid assumption on high-aspect-ratio flows, as seen in chapter 7, its applicability to the present case does not necessarily follow. Indeed, numerical evidence suggests that the profile in the mixed region is approximately linear with $\frac{\partial \rho}{\partial z}$ progressively reducing, and this is inconsistent with a well-mixed model.

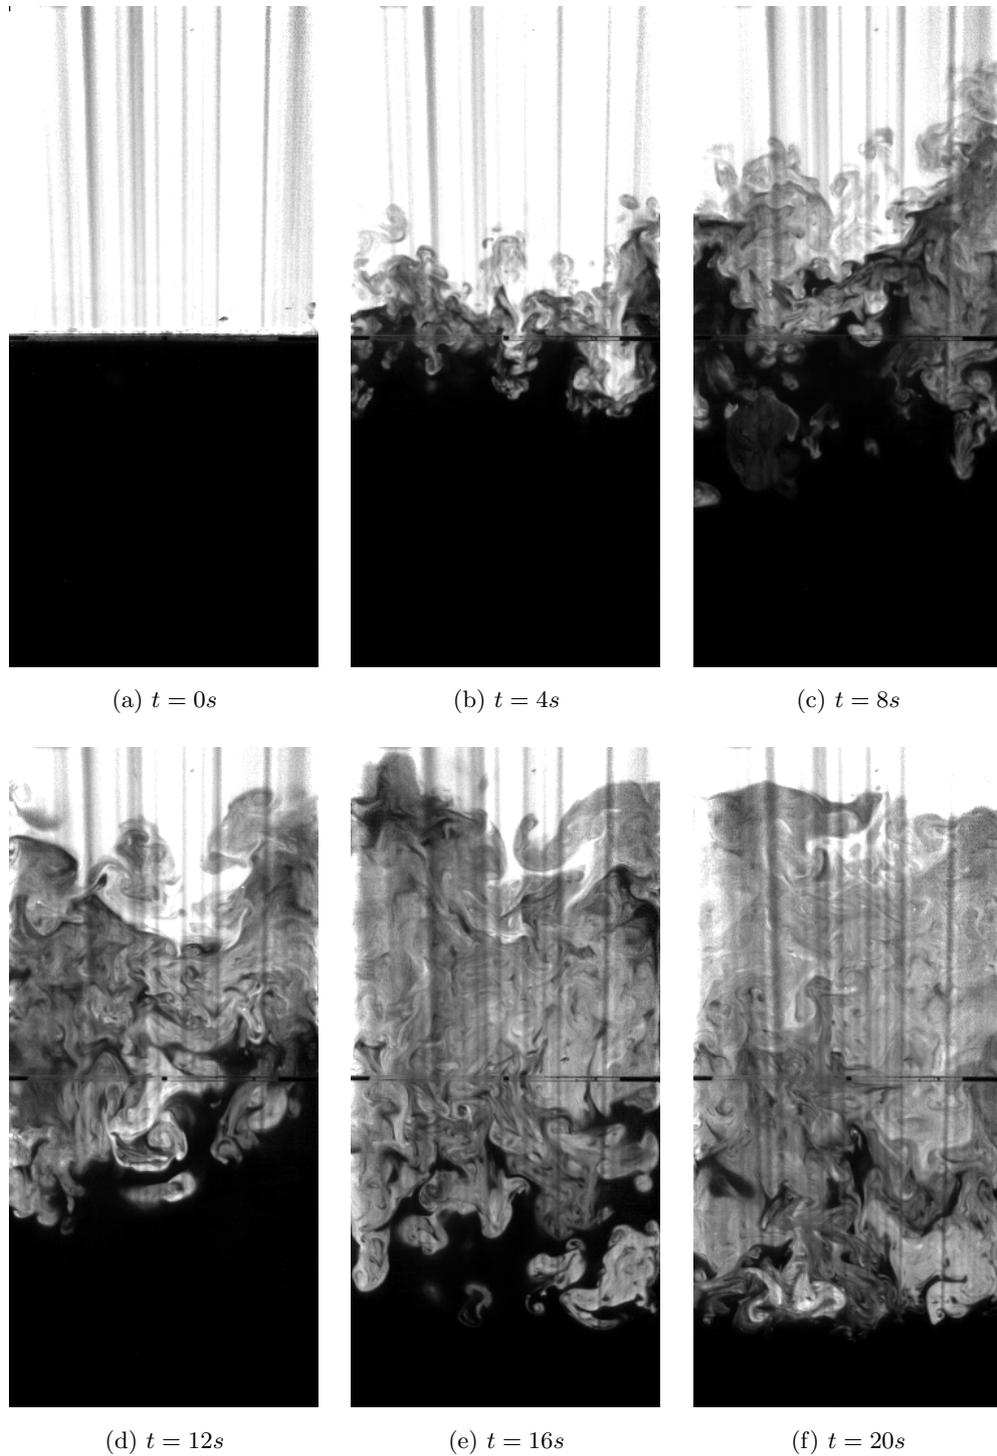


Figure 8.2: Experimental image sequence showing the Rayleigh-Taylor instability confined by linear stratification. The Atwood number is $A = 2 \times 10^{-3}$.

8.3 Mixing efficiency

8.3.1 Concepts

Mixing between miscible fluids is a two-step process by which fluid parcels in proximity are first stirred and then inter-diffuse. The stirring requires kinetic energy, and the turbulence that this creates stretches filaments of each fluid so they have a large surface area relative to their volume. The rate at which molecular diffusion takes place is very slow, but the morphological changes brought about by the turbulence greatly increase the interfacial area over which this happens, and so mixing is a rapid means of inter-diffusing two fluids.

The mixing efficiency η of a process measures the effectiveness of the kinetic energy input at increasing the interfacial area, and this helps quantify the character of the turbulence. In a variable density incompressible system initially at rest (with no external source of energy), mixing can only be achieved by converting potential energy into kinetic energy, then spending that kinetic energy in some way which gives rise to mixing. There are two ways energy can be irreversibly ‘lost’ from the variable density system: 1) by viscous dissipation to heat, and 2) by inter-diffusion to change the structure of the density field.

Following [Winters *et al.* \(1995\)](#) we decompose the total potential energy (E_{tp}) into two components, ‘Available’ and ‘Background’. Available potential energy (E_{ap}) is that component which is available to provide kinetic energy for mixing, and the Background potential energy (E_{bp}) is the remainder, which has already been used to do mixing (or was already unavailable at $t = 0$).

A density stratification at rest which has $\frac{\partial \rho}{\partial z} \leq 0$ has no potential energy available for mixing, hence a measure of the background potential energy in a system at any time is the potential energy of the *equivalent* stable stratification - that obtained if one could adiabatically rearrange the actual stratification $\rho(z)$ into the stable state $\rho(z_*)$ with minimal potential energy. Perturbing any given stratification simply by moving fluid parcels around from their original vertical position z to a new position z' does not change the background state. In an evolving flow with no mixing, $\rho(z, t)$ has a constant background state $\rho(z_*)$. Should molecular mixing of fluid parcels of

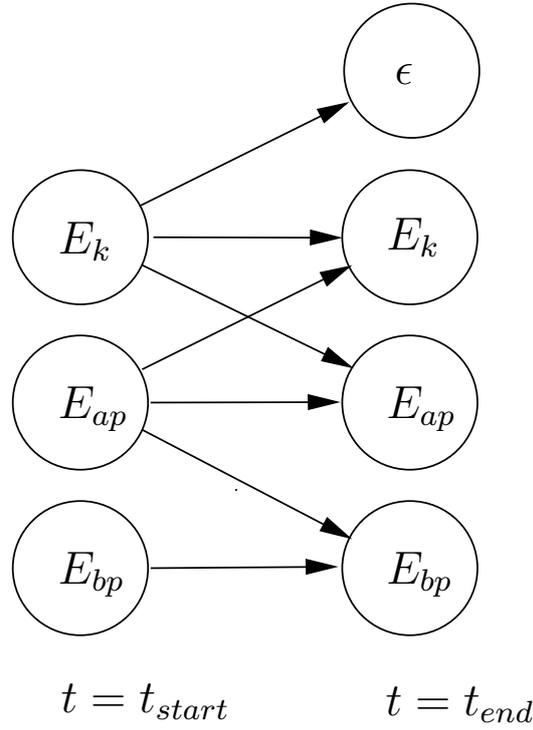


Figure 8.3: This schematic shows the various permitted pathways for energy to be exchanged from one type to another.

differing densities occur over some time Δt , then the combined fluid parcel has a new, unique density, and there is no longer a bijective mapping between $\rho(z, t + \Delta t)$ and $\rho(z_*)$. A new background state $\rho(z_*, t + \Delta t)$ must be created, and the Second Law of Thermodynamics dictates that its potential energy is greater. Hence the background state $\rho(z_*, t)$ can be viewed as a datum of zero *available* energy for the flow in its current condition $\rho(z, t)$.

Figure 8.3 illustrates the pathways by which energy can be exchanged in a flow over an arbitrary timespan Δt . A suitable differential system that mirrors this description is

$$\begin{aligned}
 \frac{\partial E_i}{\partial t} &= \varepsilon \\
 \frac{\partial E_k}{\partial t} &= -\phi - \varepsilon \\
 \frac{\partial E_{ap}}{\partial t} &= \phi - \zeta \\
 \frac{\partial E_{bp}}{\partial t} &= \zeta,
 \end{aligned} \tag{8.9}$$

where E_i is internal energy, E_k is kinetic energy, ϕ is a net adiabatic energy flux per-

mitting exchange between potential and kinetic energies, ε the power lost to internal energy by viscous dissipation, and ζ the power lost to changes in the background state. Since ζ effectively represents the energy that has gone into molecular mixing, an appropriate definition of an ‘instantaneous mixing efficiency’ would be that proportion of the total power expended by the system which was not lost as heat, *i.e.*

$$\eta_i = \frac{\zeta}{\zeta + \varepsilon} \quad (8.10)$$

Integrating numerator and denominator separately in time, we can define an ‘aggregate mixing efficiency’, which simplifies in the Rayleigh-Taylor case with zero initial velocity field, to

$$\eta_a = \frac{\Delta E_{bp}}{\Delta E_{ap}}, \quad (8.11)$$

provided one adopts the convention $\Delta E_{bp} = -\int_0^\infty \zeta dt$. Thus, solely from measurement of stratifications at $t = 0$ and $t = \infty$ (which are both quiescent conditions in the closed systems we consider here) we can determine an efficiency.

8.3.2 Everywhere-unstable systems

We define ‘everywhere-unstable systems’ to be those with an initial stratification $\frac{\partial \rho}{\partial z} \geq 0$. The simplest such case is the classic two-homogenous-layer Rayleigh-Taylor unstable problem, and previous work [Linden *et al.* \(1994\)](#) has shown that the system approaches a well-mixed homogenous end state. Defining our potential energies, for convenience, in a reference frame $\hat{z} = z - z_i$ (where z_i is the interface location), with an arbitrary upper and lower bound H , and ρ_u, ρ_l constant upper and lower layer densities, we have

$$\begin{aligned} E_{tp}^0 &= \int_{-H}^0 \rho_u g \hat{z} d\hat{z} + \int_0^H \rho_u g \hat{z} d\hat{z} = \frac{1}{2} g H^2 (\rho_u - \rho_l) \\ E_{bp}^0 &= \int_{-H}^0 \rho_l g \hat{z} d\hat{z} + \int_0^H \rho_u g \hat{z} d\hat{z} = \frac{1}{2} g H^2 (\rho_l - \rho_u) \\ E_{bp}^\infty &= \int_{-H}^H \frac{\rho_u + \rho_l}{2} g \hat{z} d\hat{z} = 0 \\ E_{tp}^\infty &= E_{bp}^\infty. \end{aligned} \quad (8.12)$$

The mixing efficiency in terms of initial and final state stratifications is

$$\begin{aligned}
 \eta_a &= \frac{E_{bp}^\infty - E_{bp}^0}{\left(E_{tp}^0 - E_{bp}^0\right) - \left(E_{tp}^\infty - E_{bp}^\infty\right)} \\
 &= \frac{0 - \frac{1}{2}gH^2(\rho_l - \rho_u)}{\left(\frac{1}{2}gH^2(\rho_u - \rho_l) - \frac{1}{2}gH^2(\rho_l - \rho_u)\right) - 0} \\
 &= \frac{1}{2}.
 \end{aligned} \tag{8.13}$$

The above result is well-known for the classic Rayleigh-Taylor problem. Further exploration shows that this generalises to all odd-function initial stratifications symmetric about an interface height z_i . All such systems approach a homogenous well-mixed end state ρ_b (as discussed in Dalziel *et al.* (2008)), so the potential energies are as follows:

$$\begin{aligned}
 E_{tp}^0 &= \int_{-H}^H (\rho_b + \Delta\rho(\hat{z})) g\hat{z}d\hat{z} \\
 E_{bp}^0 &= \int_{-H}^H (\rho_b - \Delta\rho(\hat{z})) g\hat{z}d\hat{z} \\
 E_{bp}^\infty &= \int_{-H}^H \rho_b g\hat{z}d\hat{z} \\
 E_{tp}^\infty &= E_{bp}^\infty.
 \end{aligned} \tag{8.14}$$

As before, the well-mixed state has zero energy in this reference frame, so the mixing efficiency is independent of ρ_b and is thus

$$\begin{aligned}
 \eta_a &= \frac{E_{bp}^\infty - E_{bp}^0}{\left(E_{tp}^0 - E_{bp}^0\right) - \left(E_{tp}^\infty - E_{bp}^\infty\right)} \\
 \eta_a &= \frac{0 - \int_{-H}^H -\Delta\rho(\hat{z}) g\hat{z}d\hat{z}}{\left(\int_{-H}^H \Delta\rho(\hat{z}) g\hat{z}d\hat{z} - \int_{-H}^H -\Delta\rho(\hat{z}) g\hat{z}d\hat{z}\right) - 0} \\
 &= \frac{1}{2}.
 \end{aligned} \tag{8.15}$$

8.3.3 Partially unstable systems

The statement $\eta_a = \frac{1}{2}$ implies that half of the potential energy released from a variable density system is expended as heat by viscous dissipation, and the other half is expended by doing molecular mixing. Both these pathways for energy to leave the system are micro- (*i.e.* Kolmogorov- and Batchelor-) scale processes, and this suggests that the power output (energy flux) of the system may be equipartitioned at a

micro-scale between a flux to internal energy and a flux to background potential energy. In the case of a unit Schmidt number, where Kolmogorov and Batchelor scales coincide, there seems no dynamical reason for the micro-scale events to preferentially select one energy pathway over the other, so it is a reasonable hypothesis that $\eta_a = \frac{1}{2}$ is a general property of low Schmidt number self-similar mixing. Rayleigh-Taylor experiments in water with salt as the stratifying agent have Schmidt number $Sc = O(1000)$, and there is ample evidence in the literature (*e.g.* [Linden *et al.* \(1994\)](#); [Dalziel *et al.* \(2008\)](#)) that $\eta_a = \frac{1}{2}$ even at these very high Schmidt numbers. One motivation for studying partially unstable systems is to determine whether η_a is a macro-scale property of the system, derivable from energetic considerations, or whether it is a micro-scale property, fundamental to miscible fluids. The case where Rayleigh-Taylor instability is confined between stable stratifications is investigated in the following analysis.

The relative buoyancy of fluid parcels places a bound on the extent of interpenetrating fluid above and below the initial interface. In the convenient, interface-centred reference frame, this neutral buoyancy height is

$$z_n = \frac{\rho_u - \rho_l}{\beta}, \quad (8.16)$$

and fluid motion is restricted to the region $-z_n < \hat{z} < z_n$. As with everywhere unstable systems, a well-mixed final state indicates completion of the mixing process.

Using the well-mixed model proposed in §8.2 to predict the growth profile of the

system, we can calculate an aggregate mixing efficiency. The potential energies are

$$\begin{aligned}
E_{tp}^0 &= \int_{-z_n}^0 (\rho_l - \beta \hat{z}) g \hat{z} d\hat{z} + \int_0^{z_n} (\rho_u - \beta \hat{z}) g \hat{z} d\hat{z} \\
&= \frac{\rho_u - \rho_l}{2} g z_n^2 - \frac{2}{3} \beta g z_n^3 \\
E_{bp}^0 &= \int_{-z_n}^{z_n} \left(\frac{\rho_u + \rho_l}{2} - \frac{\rho_u - \rho_l}{2} \frac{\hat{z}}{z_n} \right) g \hat{z} d\hat{z} \\
&= -\frac{\rho_u - \rho_l}{3} g z_n^2 \\
E_{tp}^\infty &= \int_{-z_n}^{-\frac{z_n}{2}} (\rho_l - \beta \hat{z}) g \hat{z} d\hat{z} + \\
&\quad \int_{-\frac{z_n}{2}}^{\frac{z_n}{2}} \frac{\rho_u + \rho_l}{2} g \hat{z} d\hat{z} + \\
&\quad \int_{\frac{z_n}{2}}^{z_n} (\rho_u - \beta \hat{z}) g \hat{z} d\hat{z} \\
&= \frac{3}{8} (\rho_u - \rho_l) g z_n^2 - \frac{2}{3} (\rho_u - \rho_l) \beta g \frac{7}{8} z_n^3 \\
E_{bp}^\infty &= E_{tp}^\infty,
\end{aligned} \tag{8.17}$$

and noting equation 8.16, the aggregate mixing efficiency is

$$\eta_a = \frac{\left(\frac{3}{8} - \frac{14}{24}\right) + \frac{1}{3}}{\left(\frac{1}{2} - \frac{2}{3}\right) + \frac{1}{3}} = \frac{3}{4}. \tag{8.18}$$

Clearly it is energetically possible to obtain a mixing efficiency greater than $\frac{1}{2}$, though it is not automatic that a real fluid will evolve to a stratification that gives rise to the maximum mixing efficiency. Given that the eddy turnover and interfacial growth time-scales are not well separated and the available energy is *finite*, it is plausible that the unstable region does not become perfectly well-mixed in the final state, and instead becomes stably stratified with some gradient γ . To examine this case we generalise the end state potential energy, introducing a parameter θ to denote the height of maximum inter-penetration as a proportion of z_n :

$$\begin{aligned}
E_{tp}^\infty &= \int_{-z_n}^{-\theta z_n} (\rho_l - \beta \hat{z}) g \hat{z} d\hat{z} + \\
&\quad \int_{-\theta z_n}^{\theta z_n} \frac{\rho_u + \rho_l}{2} g \hat{z} d\hat{z} + \\
&\quad \int_{\theta z_n}^{z_n} \left(\frac{\rho_u + \rho_l}{2} - \frac{(\rho_u - \rho_l) \left(\theta - \frac{1}{2}\right) \hat{z}}{\theta z_n} \right) g \hat{z} d\hat{z} \\
&= (\rho_u - \rho_l) g z_n^2 \left(\frac{1 - \theta}{2} - \frac{2}{3} (1 - \theta^3) - \frac{2}{3} \theta^2 \left(\theta - \frac{1}{2} \right) \right).
\end{aligned} \tag{8.19}$$

This gives a mixing efficiency

$$\eta_a = 1 - \theta^2 = 1 - \frac{1}{4 \left(1 - \frac{\gamma}{\beta}\right)^2}, \quad (8.20)$$

for non-zero initial stratification gradient β , and subject to the condition $\theta \geq \frac{1}{2}$ (for end-state static stability). With this insight, it is instructive to review the growth law in §8.2, since it is now evident that assuming a uniformly well-mixed mixing region by construction fixes the value of θ arbitrarily at $\theta = \frac{1}{2}$. This anomaly can be resolved by noting that the assumed density gradient profile in the mixed region can be set arbitrarily, and could coincide with the final stratification gradient γ . In this case the energy exchange calculation would then include a transformed gradient $\beta + \gamma$, which yields an $h(t)$ of the same form.

8.4 Comparison with simulation and experimental ensemble

8.4.1 Scalar transport

Fluorescein dye experiments were conducted to examine the Rayleigh-Taylor instability when confined between linear stratifications, with the dye present in uniform concentration throughout the lower layer, irrespective of the initial local density. The migration of this dye upwards due to Rayleigh-Taylor interpenetration is a useful diagnostic to observe the instability growth profile. The theoretical prediction of §8.2 is cross-plotted with the experimental ensemble and corresponding simulation, as shown in figure 8.4. The adjustable parameters in the theoretical model are the virtual time-origin, the time-scale constant (akin to the growth rate constant α in the classic Rayleigh-Taylor case), and the end-state penetration height θH (in the figure $\theta = \frac{1}{\sqrt{2}}$, corresponding to $\eta_a = \frac{1}{2}$). For visual guidance, $h = \alpha A g t^2$ is plotted with $\alpha = 0.05$, $A = 1.5 \times 10^{-3}$ and at early time, the predicted growth from the stratified model is indistinguishable from that for homogenous layers. As the Rayleigh-Taylor instability grows into the stratification, the effective Atwood number is reduced, so one would expect the growth to diverge from quadratic. Indeed, this is observed in

the experiments: the mean of the ensemble matches the theoretical profile in the phase where it is approximately linear.

The limitations of a well-mixed model are seen when the instability reaches its eventual penetration height θH . Each experiment exhibits a tendency to over-shoot and oscillate around θH . Much of the oscillation is likely to have been caused by a phase locking between the speed of barrier removal and the phase velocity of internal waves in the stratification, and so their strength could be expected to vary somewhat between experiments. The penetration height varies somewhat between experiments in the ensemble, even after taking into account oscillations arising from the internal waves. The cause is experimental error in the height to which the lower half of the tank was filled (a small excess above the barrier is needed to prevent air bubbles from forming on its upper surface). When stratifying from the bottom, the excess slightly shifts the density profile to have a larger ρ_l , thus adjusting the Atwood number and the penetration height.

The ILES simulation makes some questionable predictions about the internal dynamics of the flow. Regions that appear to be quiescent in the experiment (*i.e.* those above and below the mixing zone) have a small amount of kinetic energy in the simulation. Mixing takes place across the (evolving) boundary between dyed fluid and undyed fluid and so unlike the experiment, the penetration limit is not clearly defined. The reason for this apparently spurious mixing is the way in which the scalar transport is modelled. Four tracers were used, two with an associated density ρ_u and two with ρ_l , to correctly form the initial stratification as a suitably volume-weighted sum of tracer concentrations, and meanwhile distinguish between (lower layer) dyed and (upper layer) un-dyed fluid. However, once the instability develops, dyed ρ_u tracer, for instance, is in immediate proximity (inside a numerical control volume) with undyed ρ_u tracer, and there is no dynamical forcing which prevents the two tracers from swapping roles. Thus it takes an arbitrarily small kinetic energy to transport dyed tracer far into the undyed region. For this reason the simulated dye region does not have a sharp boundary, and the threshold used as demarcation for the dyed region to obtain $h(t)$, has been selected to identify the point of maximum vertical dye gradient.

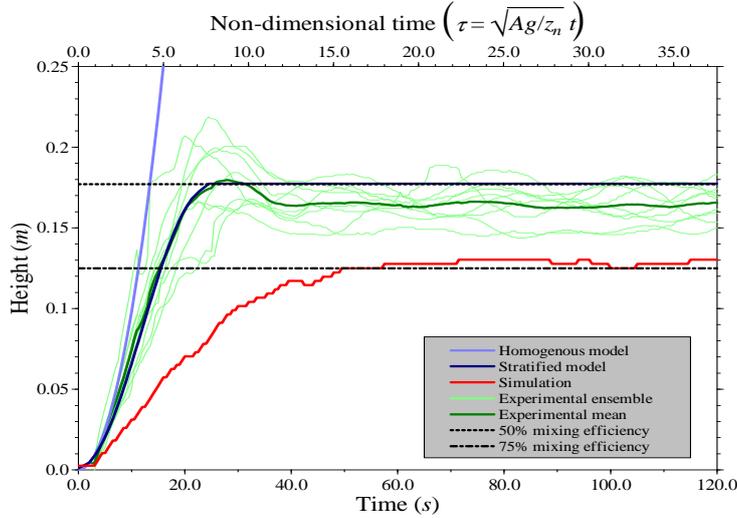


Figure 8.4: Rayleigh-Taylor growth $h(t)$ evolving from an unstable interface sitting between two stable linear stratifications, theoretical and numerical predictions compared with an experimental ensemble.

Except during the initial Rayleigh-Taylor growth phase, the simulation does not follow the experimental/theoretical trajectory for $h(t)$; once the stratification begins to impede the Rayleigh-Taylor instability, the simulated growth rate is significantly below experimental expectations. Since long wavelength modes in the initial perturbation spectrum gather increasing importance with time (see §4.2.3), it was conjectured that the ‘idealised’ random amplitude, random phase initial conditions used elsewhere in this thesis should be modified to include a bias to towards high amplitudes at low wavenumbers in order to better match the experimental initial conditions. However, when tested, neither the eddy size nor the growth rate seemed sensitive to low wavenumbers in the initial condition. This suggests that the ILES methodology itself is not adequately capturing the fluid processes in stably stratified fluids.

Among the many distinguishing factors between the current low-aspect-ratio stratified case and the high-aspect-ratio stratified case detailed in chapter 7 where ILES performs remarkably well, a poor separation of time-scales in the problem stands out as an obvious feature. The Brünt-Vaisala frequency associated with the stratification may be acting as a constraint on the eddy turnover frequency, so as length- and time-scales grow in the Rayleigh-Taylor instability, they reach a limit

imposed by the stratification, rather than by geometry as in chapter 7. A constraint on the eddy size would imply $l_{turb} \rightarrow const.$ and once again this would lead to a growth behaviour where h scales as a rational power of t . Once the stratification impedes the Rayleigh-Taylor instability, the ILES calculated $h(t)$ appears to have this functional form.

As is clear in figure 8.4, the experiment does not have the same functional form, and it remains unclear what modelling assumptions embedded in ILES reduce the quality of predictions in this flow. We know from previous work of Turner (1968) that there is a strong Schmidt number dependence on the erosion of stable interfaces, and by design ILES does not model viscous effects explicitly, and therefore cannot support Schmidt numbers except implicitly at $O(1)$. It seems plausible that this modelling omission might ultimately be responsible for predicting a lower growth rate than observed in experiment. Unfortunately, testing this assertion in the current system would require direct numerical simulation at a resolution not currently feasible.

8.4.2 Mixing efficiency

Examining the energetics of the stratification-confined Rayleigh-Taylor problem offers valuable insight into the system dynamics, and helps establish the extent to which ILES modelling is useful in such situations. The simulation presented in §8.4.1 used four transported scalars to discriminate between upper and lower layer fluid and high and low density; unfortunately, although the high order advection scheme is exactly monotonic for each individual tracer, the volume-weighted sum of more than two tracers is not necessarily monotonic, and detailed measurements based on the density field are sensitive to such errors. To study the system energetics new MOBILE simulations were performed, again initialised to match the experimental ensemble with an Atwood number $A = 1.5 \times 10^{-3}$, but using only two scalar fields to represent density and hence removing the possibility of monotonicity violation. The beginning and end state density profiles are plotted in figure 8.5. The experimental density values in this figure come from a vertical traverse of the tank by a probe. The probe has a metallic outer surface and inner core, separated by insulating material and the conducting path is from the outer surface through the water around the

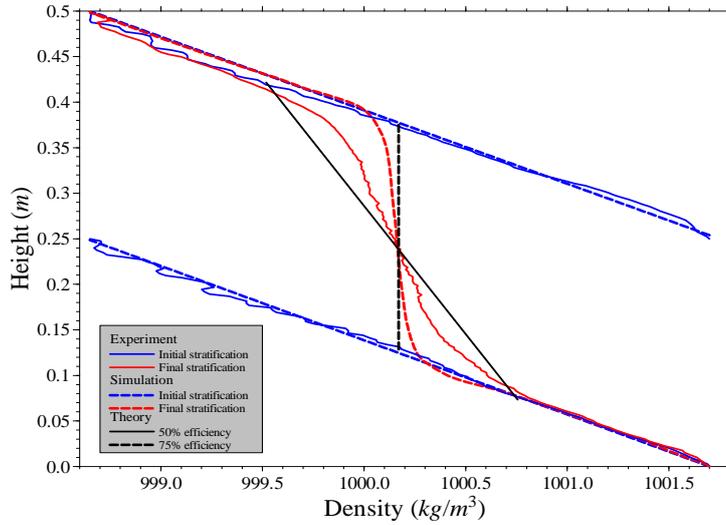


Figure 8.5: Initial and final stratifications for an experiment with a stratification-confined Rayleigh-Taylor instability, compared with equivalent simulation and theoretical end-state gradients.

probe tip to the inner core. The resistivity of the conducting path is calibrated against salt concentration and therefore measures density.

Quite clearly, although the initial conditions are very similar (even accounting for some uncertainties in the conductivity measurements and some leakage from the tank), the end state profile from the simulation does not match the experiment well. Given equation 8.20, we would expect there to be a corresponding discrepancy in the aggregate mixing efficiency. Indeed, integrating the potential energies directly, the simulation achieves $\eta_a = 0.735$ against the experiment $\eta_a = 0.490$, with a discretisation error in both cases of approximately ± 0.002 . It would appear that the simulation is very close to achieving the maximum mixing efficiency that is energetically possible in this configuration ($\eta_a = 0.75$ for a well-mixed end state). The real fluid on the other hand is very close to achieving equidistribution of energy flux into molecular mixing and viscous dissipation ($\eta_a = 0.5$), as hypothesised in §8.3.3. While the experimental results exhibit a curvature which prevents them closely matching the theoretical gradient, assuming a mixing efficiency of $\eta_a = 0.5$ leads to accurate predictions of the upper and lower boundaries of the mixing region. This is significant, since it offers a good estimate of the height across the initial interface to which a scalar is transported in this system, a parameter which remained free in the

well-mixed growth model of §8.2, and this corresponds well with the experimental evidence in figure 8.4.

While it is clear that simulation is attracted to a maximum mixing efficiency relaxation of the system and the real fluid is not, it remains instructive to examine how energy is passed around in the simulated system, particularly since the data extraction is trivial. The distribution of energy is shown in figure 8.6, where the ‘zero-energy’ datum for potential energies is taken as the initial state background, and the scale is normalised by the initial available potential energy. The kinetic energy rises to a peak by 20s and takes a much longer time to decay. Although most of the available energy ($E_{ap} + E_k$) is spent over the first 60s, internal waves are induced on the boundary between un-mixed stable stratification and the mixed zone, and these exchange some residual kinetic and available potential energy. Energy contained in internal waves is not efficiently removed by shear or creation of background potential energy, so they last for considerable time. These waves are thought to be caused by imperfect statistical distribution (in a finite-sized box) of initial random amplitude, random phase perturbations, leading to turbulence which is not perfectly self-similar. Internal waves are also noticeable in the experimental context due to initialisation effects (see discussion in §8.4.1), but in the simulations these waves are much less energetic.

The available energy clearly is spent in two ways, (1) raising the background potential, and (2) converting to fluid internal energy by viscous, or here in the case of an ILES simulation, numerical dissipation. The aggregate loss of total potential energy, assuming there is no residual kinetic energy, equates to the gain of internal energy. Thus the aggregate mixing efficiency can be read off the graph directly.

To shed some light on the processes that give rise to the aggregate quantities, we now examine the transient re-distribution of energy. The model equations 8.9 describe the system energetics, and figure 8.7 plots as functions of time the various terms in these equations. The energy fluxes ε , ϕ and ζ are labelled ‘Dissipation’, ‘Exchange’ and ‘Background’ respectively; labels for $\frac{\partial E_k}{\partial t}$ and $\frac{\partial E_{ap}}{\partial t}$ are self-explanatory. The vertical scale uses the same normalised units as figure 8.6.

The peak in available potential energy release occurs at about $t = 5s$, and corre-

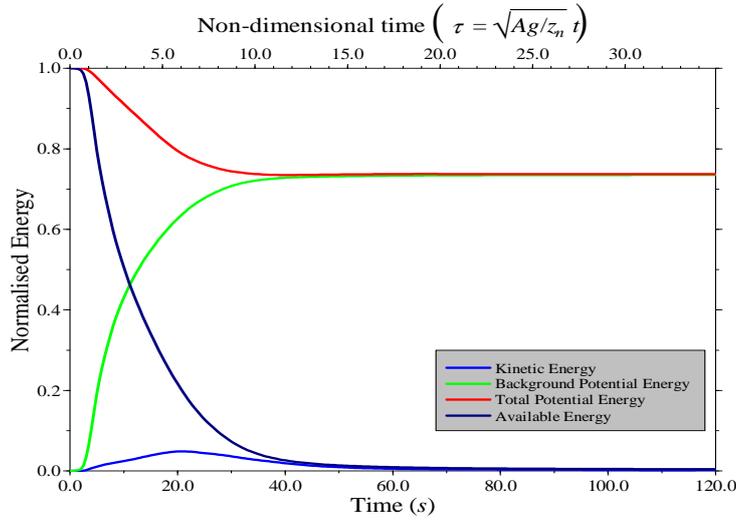


Figure 8.6: Time evolution of energy distribution in a stratification-confined simulation.

sponds, unsurprisingly, to the maximum growth rate of the instability (see figure 8.4). Equations 8.9 show that available potential energy is released into an net exchange flux ϕ and flux to background ζ . From figure 8.7(a) it is clear that the exchange flux (whose sign is negative simply by the chosen sign convention in equations 8.9) is very much smaller than the flux to background. Integrated over time, the flux to background is almost exactly three times larger than the exchange flux, giving a mixing efficiency of $\eta_a \approx 0.75$, but the graphs show that the flux to background (the proxy for molecular mixing) reaches a much higher peak when the instability growth rate is highest, whereas the exchange flux is more stable over time. This suggests that the mixing efficiency is greatest during the instability acceleration, and reduces when the stratification begins to impede the instability growth. The exchange flux to kinetic energy, ϕ , is a small proportion of the overall energy release, and surprisingly, it almost exactly balances the flux from kinetic to internal energy, ε . The kinetic energy in the system at any one time is therefore very low relative to the potential energy released, and it appears that the system is converting only just sufficient potential energy into kinetic to perpetuate the instability.

Although the flow of causality is not clear, the combination of low kinetic energy and a constrained eddy size may explain the unusually high values of mixing

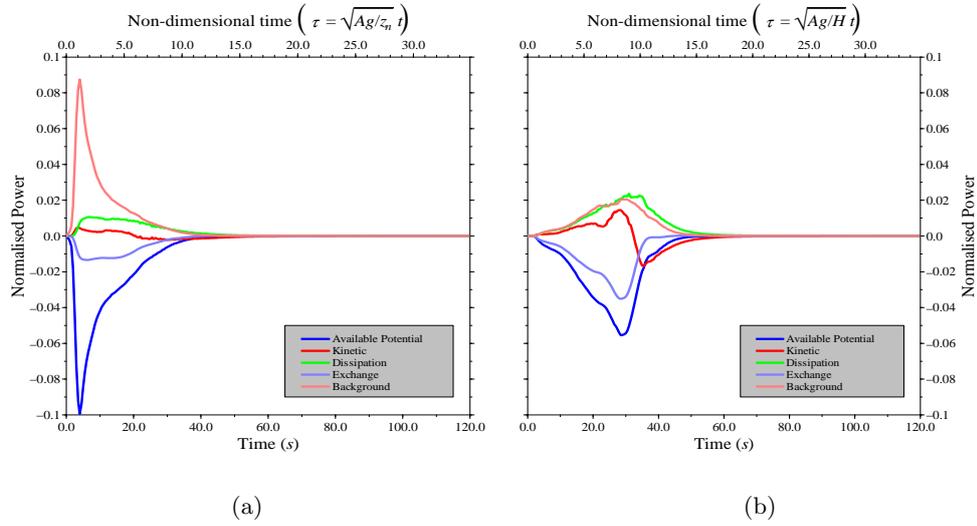


Figure 8.7: Time evolution of energy fluxes during (a) a stratification-confined simulation, and (b) a homogenous-layer simulation.

efficiency predicted by ILES. Consider the converse: if there were a higher level of kinetic energy in a Rayleigh-Taylor flow, there would be a greater likelihood of large eddies overturning fluid in bulk to reach a stable, low available energy state, without individual fluid parcels having had the chance to mix. This scenario is self-sustaining, since larger eddies would penetrate further into the quiescent region and unlock potential energy even more rapidly. Minimising the kinetic energy, whether or not this is a consequence of a constraint on eddy size, increases the mixing efficiency.

It is instructive to compare the energy fluxes calculated for the stratification-confined Rayleigh-Taylor instability shown in figure 8.7(a) with the equivalent diagnostic for the homogenous-layer Rayleigh-Taylor instability shown in figure 8.7(b), since it is generally accepted that ILES performs well on this problem. There are many important qualitative differences between the energetic characteristics of the two cases. Firstly, the flux to background and the dissipation flux almost exactly overlies one another in the homogenous-layer case, making $\eta_a \rightarrow 0.5$ inevitable. Secondly, also in this case, the exchange flux is a large proportion ($> 50\%$) of the released available potential energy, and more of this exchange remains as kinetic energy in the system rather than being quickly dissipated. Thirdly, and perhaps most strikingly, the flux of available potential energy being released during the acceler-

ating growth phase is much lower in the homogenous-layer simulation than in the stratification-confined simulation. Given that so much energy is needed to initiate the instability, this suggests that the stratification strongly inhibits the instability growth from $t = 0$, and this is certainly not observed in the experiment. It is also notable that the energy released does not end up predominantly as kinetic energy, so by equations 8.9 the only other possible pathway for this energy is into background potential energy, and one would therefore expect the instantaneous mixing efficiency to be very high.

The instantaneous mixing efficiency η_i is a useful measure to show how effective time-local processes are at doing mixing, but this does not indicate the relative contribution of any phase of the process to the overall mixing efficiency, because it does not account for fluctuations in the available energy with time. A ‘cumulative mixing efficiency’,

$$\eta_c(t) = \frac{\int_0^t \zeta(\hat{t}) d\hat{t}}{\int_0^t \zeta(\hat{t}) + \varepsilon(\hat{t}) d\hat{t}}, \quad (8.21)$$

which interpolates between the instantaneous measure at early time to the aggregate measure at late time, satisfies this requirement. The simulations show $\eta_i = \eta_c = 1$ at early time (see figure 8.8), and this is due to the initial conditions having a perturbation in density only. An infinitesimal advection of density due to buoyancy gives rise to ‘mixing’ across numerical cell boundaries with there being zero or negligible kinetic energy present to begin with. Thus no dissipation can occur, and the process is nominally 100% efficient. In the experiments a velocity perturbation is induced by the barrier, dissipation will begin immediately, and $\eta_i = 0.5$ is a more likely starting point. Indeed in previous experimental work by [Holford *et al.* \(2003\)](#) this has been shown to be the case.

To make some progress in understanding why ILES does not perform well on the stratification-confined Rayleigh-Taylor problem, the time evolution of the mixing efficiency both for this case and the classic two-homogenous-layer case are cross-plotted in figure 8.8. In both cases the simulated system relaxes towards maximum mixing efficiency, yielding an aggregate values $\eta_a = 0.735$ and $\eta_a = 0.463$ respectively. The real fluid also approaches maximum mixing efficiency in the two-layer case, as shown in [Holford *et al.* \(2001\)](#). For completeness, a single linear unstable stratification

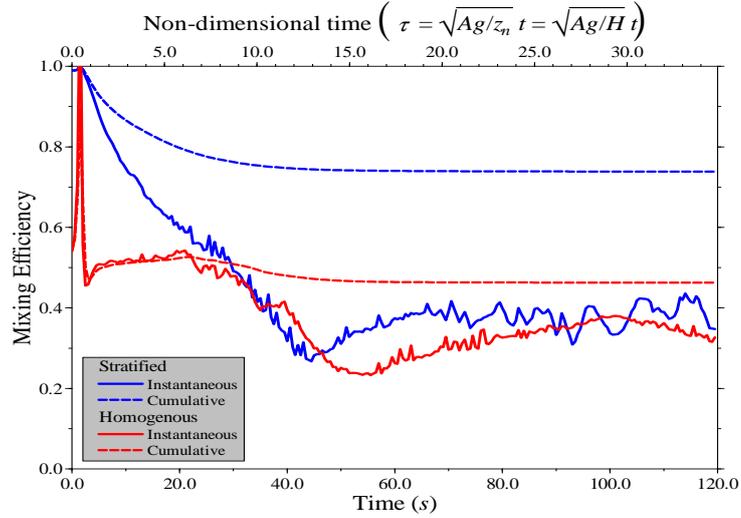


Figure 8.8: Time evolution of instantaneous and cumulative mixing efficiency for both homogenous-layer and stratification-confined Rayleigh-Taylor instability. The Atwood number of the interface is common to both simulations.

perturbed in the middle was also simulated, as a simple case of the kind considered in §8.3.2, and once again the mixing efficiency approaches the maximum possible, with $\eta_a = 0.468$. The instantaneous mixing efficiency appears to pass through two phases: a Rayleigh-Taylor growth phase and a late-time decay phase.

In both homogenous-layer and stratification-confined η_i hovers between 0.3 and 0.4 at late time ($t > 60s$), and this corresponds in time to the form of self-similar mixing we would expect for well-developed Rayleigh-Taylor instability in a box. It remains unclear why the mixing efficiency is not closer to 0.5 at this late stage, but the magnitude of the energy fluxes ζ and ε decay to almost 0 after $t = 60s$ so the mixing in this phase contributes little to the cumulative efficiency.

The mixing efficiencies η_i differ markedly between the two cases during the Rayleigh-Taylor growth phase ($t < 30s$), and this is very surprising since experiments confirm that the Rayleigh-Taylor growth is initially only weakly affected by the stratification. Given that $\eta_a \approx 0.5$ in experiments of both homogenous-layer and stratification-confined Rayleigh-Taylor instability, it seems likely that the discrepancy in ILES simulations is a modelling error. The high instantaneous mixing efficiency in the stratified case suggests that inadequate kinetic energy is present in the system and therefore dissipation is under-predicted. Abnormally low kinetic en-

ergy is consistent with a reduced overall growth rate, reduced turbulent length scales, and visual inspection of the bubble sizes compared with the two-homogenous-layer case.

8.5 Summary

This chapter has used a case study of Rayleigh-Taylor instability confined by stable stratifications above and below the unstable interface to develop our understanding of the mixing process in confined Rayleigh-Taylor instability, and in particular its efficiency. A model was derived to predict the instability growth $h(t)$, assuming that the interfacial region becomes well mixed on a much faster time-scale than its overall growth. The concept of mixing efficiency was developed from first principles by identifying the allowable energy pathways in variable density systems. Three definitions of mixing efficiency were adopted - instantaneous, cumulative, and aggregate. The aggregate mixing efficiency of the classic two-homogenous-layer Rayleigh-Taylor system was calculated from the system initial and final states as $\eta_a = 0.5$; this result was shown to generalise to an arbitrary antisymmetric initial density profile, and simulations closely approach this value. It was hypothesised that variable density mixing in general will tend towards $\eta_i = 0.5$ since this indicates that there is no preference at a molecular level for energy to be channeled into molecular mixing over viscous dissipation. In the stratification-confined case, it was shown to be energetically possible to achieve a mixing efficiency of $\eta_a = 0.75$, yet experimental results clearly show that the mixing efficiency remains close to $\eta_a = 0.5$. Using this phenomenological evidence, the height to which scalars are transported in the interfacial growth model is no longer a free parameter. Unfortunately, the ILES simulations of the system do not relax in the same way as the real fluid, and they tend towards maximum mixing efficiency. The late time behaviour of ILES is reminiscent of the tall tube simulations and experiments, and suggests that the turbulent length scale l_{turb} does not scale with the height of the mixed region, as one would expect where there is no geometric confinement.

Chapter 9

Mixing confined by a stable density interface

9.1 Introduction

Confinement of Rayleigh-Taylor instability growth has so far been achieved in two ways, firstly by limiting the domain geometry, and secondly by providing a linear stable density stratification into which the instability grows. A third technique is examined in this chapter, where a sharp, stable density interface inhibits the Rayleigh-Taylor growth in one direction. A three-layer configuration is envisaged, with densities denoted ρ_u , ρ_l on either side of the unstable interface at height z_i , and ρ_s for the fluid beyond the stable interface at some height z_s , as shown in figure 9.1. There are several possible initial arrangements for the density, with corresponding and distinct flow regimes that could be expected to evolve from them. For instance, the system is globally unstable if the volume weighted mean density of ρ_u and ρ_l is unstable against ρ_s . The arrangement

$$\begin{aligned}\rho(z, 0) &= \rho_u, & z > z_i \\ \rho(z, 0) &= \rho_l, & z_s < z < z_i \\ \rho(z, 0) &= \rho_s, & z < z_s\end{aligned}\tag{9.1}$$
$$\rho_s = \rho_u > \rho_l,$$

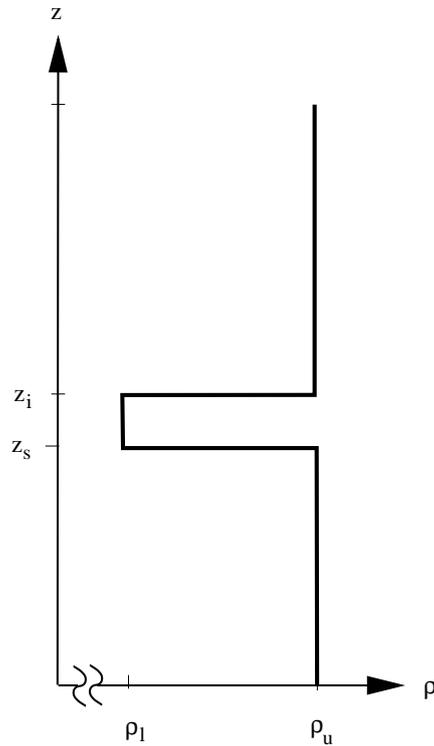


Figure 9.1: Diagram illustrating the initial density profile.

by this definition of stability is globally stable, and is studied in detail in this chapter. In the Boussinesq limit, the alternative configuration $\rho_s = \rho_l < \rho_u$, with ρ_s above ρ_u is dynamically equivalent, but practical considerations favoured the first configuration. This mixed unstable-stable system is rich with interesting features and has previously been explored experimentally and theoretically in [Jacobs & Dalziel \(2005\)](#). This chapter extends this work with molecular mixing experiments using the technique outlined in chapter 5, and compares with analogous synthetic diagnostics from ILES simulations.

9.2 Previous theoretical work

The same approach to an energy budget analysis as §4.2.4 was taken in [Jacobs & Dalziel \(2005\)](#) to obtain a late-time growth law for the current case. One would expect flow outside the mixing region to be irrotational, and at early times, when the height of the Rayleigh-Taylor instability growth is small compared with the distance to the stable interface, the velocity potential in the irrotational region should

be independent of the density jump across the stable interface. Hence the stable interface should have negligible influence on the growth of the instability at early times. However, at late times this is unlikely to be the case, and the resulting scaling is re-iterated below.

The energy balance equation 4.17, as before, can be written in the form

$$g \int (\overline{\rho_0} - \overline{\rho}) z dz = \frac{1}{2} \rho_b \int \overline{\mathbf{u}^2} dz, \quad (9.2)$$

where $\rho_0 = \rho(z, 0)$. For an initial configuration as equation 9.1, the expression on the left can be evaluated as a definite integral in a box of infinite depth,

$$\int_{-\infty}^{\infty} (\rho_{s,u} - \overline{\rho}) z dz - \int_{z_s}^{z_i} (\rho_u - \rho_l) z dz. \quad (9.3)$$

Provided the density profile is self-similar then there is a natural scaling for the first integral,

$$\begin{aligned} \rho_{s,u} - \overline{\rho} &= \hat{\rho} r(\zeta) \\ \overline{u^2} &= \hat{u}^2 s(\zeta), \end{aligned} \quad (9.4)$$

except in this instance the similarity variable ζ must account for the gradual shift of the centroid z_c of the density and kinetic energy distributions that occur as the stable interface becomes influential and the Rayleigh-Taylor instability evolves asymmetrically. The definition

$$\zeta = \frac{z - z_c(t)}{h(t)} \quad (9.5)$$

is used. By considering mass conservation, the second integral in expression 9.3 can also be expressed in terms of $r(\zeta)$. The equality

$$\int_{-\infty}^{\infty} (\rho_{s,u} - \overline{\rho}) dz = \int_{-\infty}^{\infty} (\rho_{s,u} - \rho_0) dz \quad (9.6)$$

is simply a re-statement of mass conservation, and noting that

$$\begin{aligned} \int_{-\infty}^{\infty} (\rho_{s,u} - \overline{\rho}) dz &= \hat{\rho} h \int_{-\infty}^{\infty} r(\zeta) d\zeta \\ &= (\rho_u - \rho_l) \Delta z, \end{aligned} \quad (9.7)$$

and

$$\int_{z_s}^{z_i} (\rho_u - \rho_l) z dz = \frac{1}{2} (\rho_u - \rho_l) \Delta z^2, \quad (9.8)$$

where $\Delta z = z_i - z_s$ and the integrals are evaluated in a transformed reference frame $\hat{z} = z - z_i$, the left-hand side of equation 9.2 becomes

$$\hat{\rho}gh \int_{-\infty}^{\infty} r(\zeta) \left(\zeta + \frac{z_c}{h} - \frac{\Delta z}{2h} \right) d\zeta. \quad (9.9)$$

We expect the ratio $\frac{z_c}{h}$ to be constant if the density and kinetic energy distributions are indeed self-similar, and at late time $\frac{\Delta z}{h} \rightarrow 0$, so the growth rate scaling is identical to equation 4.23, namely

$$\hat{\rho}gh \sim \frac{1}{2} (\rho_u + \rho_l) \left(\frac{dh}{dt} \right)^2. \quad (9.10)$$

However, from equation 9.7, we know that

$$\hat{\rho}h \int r(\zeta) d\zeta = (\rho_u - \rho_l) \Delta z, \quad (9.11)$$

so the scaling, instead of integrating to a t^2 profile, becomes linear in time,

$$\frac{\rho_u - \rho_l}{\rho_u + \rho_l} g \Delta z \sim \left(\frac{dh}{dt} \right)^2 \quad (9.12)$$

$$h \sim \sqrt{Ag \Delta z t}. \quad (9.13)$$

Since passive scalars will also conform to this growth profile, a separation of space and time of the form

$$\bar{\phi} = \hat{\phi}q(\zeta) \quad (9.14)$$

where ϕ represents dye-tagged fluid of density ρ_l , gives

$$\hat{\phi}h \int_{-\infty}^{\infty} q(\zeta) d\zeta = \Delta z. \quad (9.15)$$

Substituting for h , the maximum concentration $\hat{\phi}$ as a function of t follows

$$\hat{\phi} \sim \frac{\Delta z}{\sqrt{Ag \Delta z t}}. \quad (9.16)$$

The above analysis is only valid if the turbulent development of the instability produces self-similar density, passive scalar and kinetic energy profiles. Many previous studies have shown that classic two-layer Rayleigh-Taylor turbulence is indeed self-similar once it is well developed, but since no restrictions are made on the functional forms $r(\zeta)$ and $s(\zeta)$, in isolation there is no information in particular about the penetration across the stable interface to close the system. There are numerous

studies of grid-generated turbulence impinging on a stable interface (see [Fernando \(1991\)](#) for a review) and the penetration rate u_e is known to depend on Richardson number, and nominally has the form

$$\frac{u_e}{u_{turb}} \sim Ri^{-n}, \quad (9.17)$$

though the exponent n is in turn a function of Schmidt number. Given the qualitative similarity between grid and Rayleigh-Taylor generated turbulence above a stable interface, [Jacobs & Dalziel \(2005\)](#) went on to explore the connection. In the present context, the length and velocity scales l_{turb} and u_{turb} that constitute Ri can be derived from the above scaling, so that

$$Ri = \frac{\Delta\rho g l_{turb}}{\rho_b u_{turb}^2} \approx \frac{\hat{\rho} g h}{\rho_b \left(\frac{dh}{dt}\right)^2} \approx const. \quad (9.18)$$

More robust to measure experimentally than u_e is its time integral h_e , which we would thus expect to behave as

$$h_e \sim h Ri^{-n}, \quad (9.19)$$

and it is this equation which one would aim to correlate with experimental data.

9.3 Comparison of scalar transport and molecular mixing

9.3.1 Initial observations

A Fluorescein dye PLIF experiment was conducted, firstly to validate current technique against the previously published experiments of [Jacobs & Dalziel \(2005\)](#), and secondly to provide all the initialisation information for MOBILE simulations that could be fairly compared with experiment. The incident light sheet in the following experiments shines upwards from the bottom of the tank, and the camera orientation in this case is such that the barrier withdrawal can be observed, illustrated in the second image of the sequence in figure [9.2](#). Withdrawal is sufficiently rapid that the spatial variation in the development of the Rayleigh-Taylor instability is unimportant. Additionally, the stable interface provides a buoyancy force opposing any

barrier-withdrawal induced perturbations of the fluid, and because the interface sits just below the barrier, bulk overturning is prevented. Unlike the classic two-layer case, the perturbation imposed by the barrier remains confined to a small region at the extreme edge of the tank, except for low amplitude gravity waves that are initiated by the impulse from upper layer fluid reaching the stable interface.

The unstable interface grows initially according to $h = \alpha A g t^2$ until the velocity potential is significantly influenced by the stable interface. Thereafter, the instability growth is expected to slow down, since the supply of buoyant fluid to drive the instability is, unlike the classic case, finite. Energy is also lost from the system by doing mixing across the stable interface. Due to the inefficiency of doing work against buoyancy, penetration across the stable interface is modest, though momentum transport generated by the instability between the upper two layers grows to become comparable with the buoyancy forcing across the stable interface, and thus the interface itself eventually becomes highly distorted.

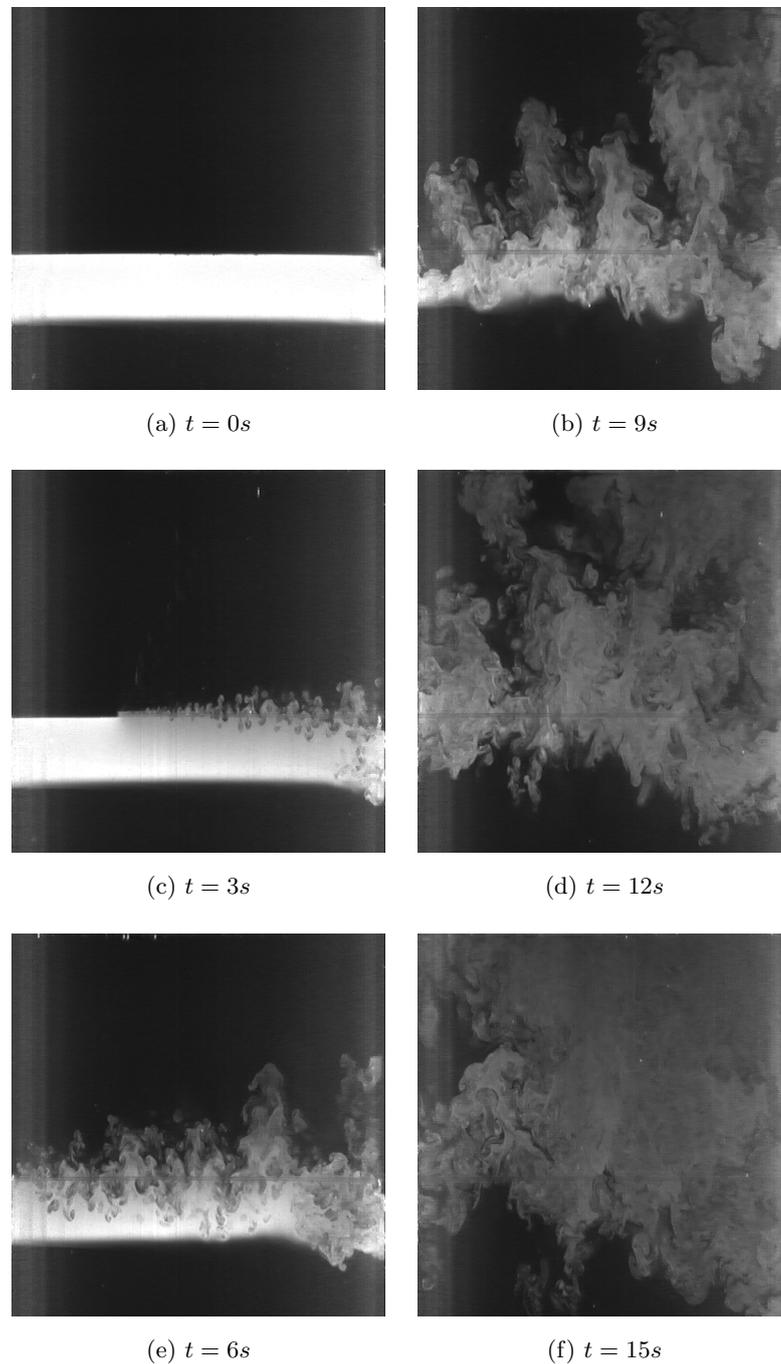


Figure 9.2: Middle layer scalar transport in the stable interface Rayleigh-Taylor problem. The field of view does not include the very bottom of the lower layer. The unstable interface Atwood number is $A = 2 \times 10^{-3}$, and stable interface Richardson number sits in the range $1.66 < Ri < 2.44$.

9.3.2 Profile self-similarity

The analysis in §9.2 is only strictly valid for self-similar density and kinetic energy profiles, so confirming that these are reasonable assumptions is essential for further comparison with the theory. The middle (relatively less dense) layer was impregnated with dye in the Fluorescein PLIF experiment. Naturally, as the Rayleigh-Taylor instability evolves, the dye migrates upwards into previously un-dyed fluid. Once penetration across the stable interface is under way, dye migrates downwards too. The efficiency of the transport mechanisms is obviously unequal, and the mean direction of transport is upwards, dominated by the stable layer. The self-similar form sought must eliminate this vertical time-dependency. Figure 9.3 shows horizontally averaged dye concentration as a function of height, normalised by maximum dye concentration and shifted by the vertical position of the concentration profile centroid. The curves are taken from the region after the lower extent of the mixing region has reached the stable interface but before the upper extent has reached the top of the tank.

The experimental curves have some noise contamination, which, as shown in [Jacobs & Dalziel \(2005\)](#), reduces significantly when averaged over larger ensemble of experiments, but the profiles clearly display self-similar collapse. The simulation does not suffer from such noise, also collapses well, and appears to predict both the shape and position of profiles seen in the experiment. One interesting discrepancy is the upper extent of dyed fluid (albeit at low concentration) seen in the experiment exceeds that in the simulation. This is most likely to have arisen from the well-known under-prediction of α in numerical simulation (see §5.3.2), though also perhaps through additional energy being supplied to fluid in the experiment by removal of the barrier. The scalar profile across the stable interface is, somewhat surprisingly, very well predicted by the numerical simulation. Direct visual examination of the numerical image sequence indicates that the detailed structure is not well captured - eddies tend to be longer-lived than in reality, but the macro-scale trends are correct.

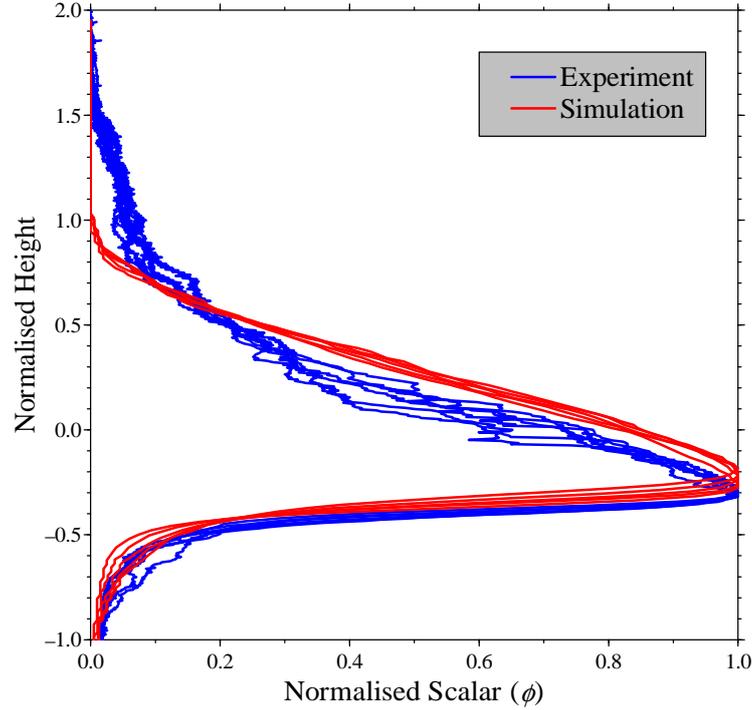


Figure 9.3: Self-similar collapse of density profiles at various times, experimental and numerical comparison.

9.3.3 Time evolution of concentration field

Having established that self-similarity of the scalar field is a reasonable working assumption, some predictions of §9.2 are examined. Provided the scalar field is self-similar, the horizontally averaged scalar field can be decomposed into $\bar{\phi} = \hat{\phi}(\zeta)$ where $\hat{\phi} \sim \frac{\Delta z}{\sqrt{Ag\Delta z t}}$. Figure 9.4 shows $\frac{1}{\hat{\phi}}$ plotted against t for experiment and simulation, with the theoretical prediction superimposed. A virtual time-origin for the theoretical curve is admissible, since integration constants in the evaluation of $h(t)$ have been neglected, and this is physically consistent since self-similarity is not established until the Rayleigh-Taylor instability has reached the stable interface.

The maximum value of concentration remains close to unity until the instability reaches the stable interface ($t = 3.5s$) then gradually a self-similar form for the profile is established and $\hat{\phi} \sim \frac{1}{t}$. During the self-similar stage, experiment and simulation are entirely consistent, and consistent with the theoretical prediction of linearity. The gradient of the line is $\frac{\sqrt{Ag\Delta z}}{\Delta z}$, and arbitrary α -like constants have not

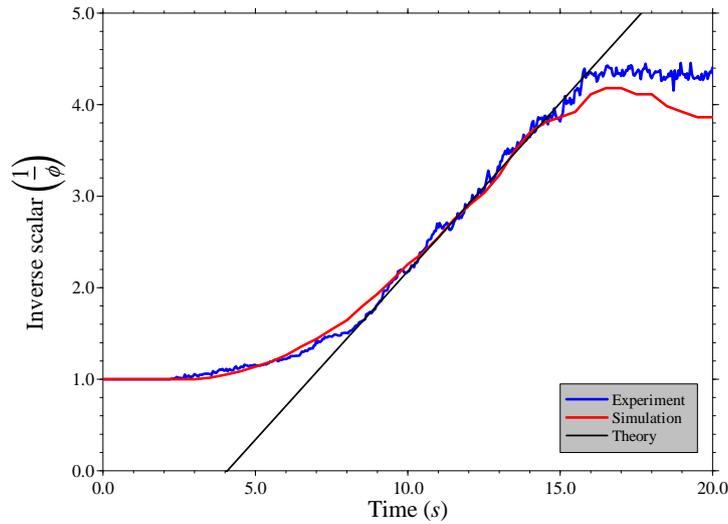


Figure 9.4: Decay of maximum concentration \hat{C} , plotted as an inverse quantity to compare analytical prediction with experiment and simulation.

been used to improve the fit. When the instability reaches the top of the box, the structure of the scalar profile changes again, and continuity requires that eventually $\hat{\phi}$ tends towards a constant as remaining available energy is spent. This appears to happen more quickly in the experiment than the simulation, and the time-evolution of the density profile suggests that the simulation reaches an end-state that is less well mixed than the experiment. After the instability reaches the top of the tank, the energy in the system decays, and the ILES method seems to under-predict the mixing in such instances. Lighter middle layer fluid congregates towards the top of the tank leaving a stably stratified end-state, which inevitably has a higher value of $\hat{\phi}$.

9.3.4 Interfacial growth

Further experiments were conducted using the Acridine-based RLIF, motivated by an interest in identifying any qualitative differences between the structure of mixing across the unstable interface and mixing across the stable interface. The light sheet comes once again from the tank bottom, but in the normal transverse orientation. An image sequence from the experiment identifying mixing between the upper (acidic) and middle (Acridine) layer is shown in figure 9.5, and qualitatively the structure of

the mixing is similar to classical Rayleigh-Taylor instability, with well-defined bubble and spike structures, modal interaction and progressive growth of the dominant length-scale. The fluid with the highest optical intensity (just over the critical volume fraction threshold) remains biased towards the bottom of the combined upper-middle region, which is consistent with the density profiles from the earlier experiments. This also indicates that the mixing is highly efficient, since relatively little dense fluid needs to fall to drive the mixing process, and that the proportion of dense fluid from the lower, non-acidic layer entrained into the upper-middle region is fairly small.

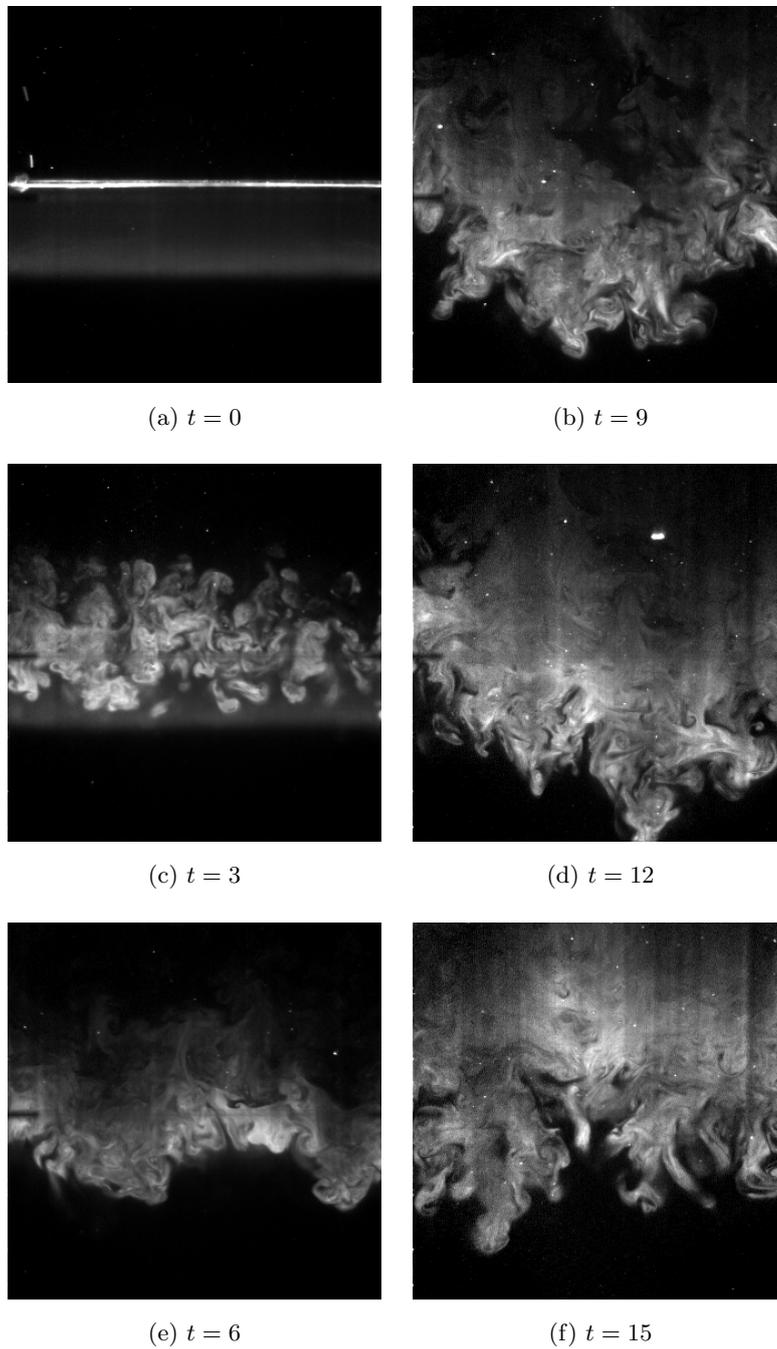


Figure 9.5: Visualisation of molecular mixing across the unstable interface in Rayleigh-Taylor instability confined asymmetrically by a stable interface. The initial stratification and the times at which images are shown are identical to those in figures 9.2 and 9.6.

To quantify the entrainment across the stable interface, another experiment was performed, this time with the lower layer acidic and the upper layer non-acidic. The middle layer again was impregnated with Acridine. An image sequence from this experiment is shown in figure 9.6. Immediately apparent is the layer of mixed fluid which exists before the experiment is initiated. Whereas in the unstable case the reagents are separated by the barrier until its removal, in the stable case mixing occurs during the stratifying process. A sponge and polystyrene float was used to minimise the thickness of the interface. Any possible dynamical influence of the non-zero thickness was a concern, since vortical structures impinging on the dense lower region might decelerate through a relatively thick interface and reduce the rate of mixing. A series of numerical tests were performed with a variety of thicknesses and profile shapes, each chosen carefully to maintain the same initial available potential energy, but there was no material change to the entrainment rate.

One interesting feature of this experiment is the obvious contrast in the turbulent structure that is visualised. Figure 9.6(c) clearly shows isolated Rayleigh-Taylor spike structures impinging on the interface and penetrating through. Subsequent images show how they drag lower layer fluid into the upper-middle region, in thin filamental structures. Consistent with experiment to visualise unstable interface mixing, very little lower layer fluid is entrained, but unlike the earlier experiment, the dominant visible length-scale does not grow significantly with time, and the mixed fluid remains filamental until viscous effects smear the gradients. At later time there is a pronounced vertical variation in the structure, with ‘newly’ mixed fluid just above the stable interface still filamental, while ‘older’ mixed fluid further above the interface is more homogeneously distributed.

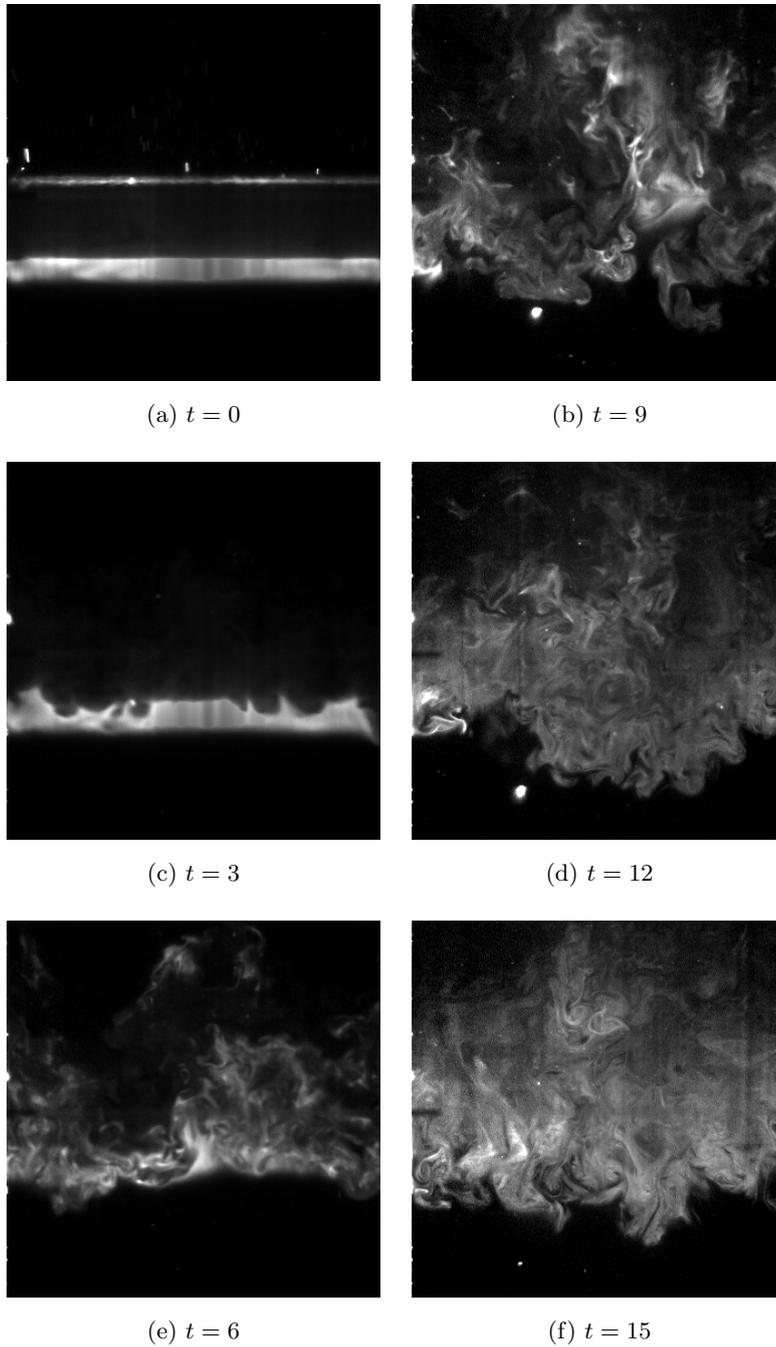


Figure 9.6: Visualisation of molecular mixing across the stable interface in Rayleigh-Taylor instability confined asymmetrically by this stable interface. Note that the tank-filling process induced mixing prior to release of the unstable interface and is thus identified by the diagnostic. The initial stratification and the times at which images are shown are identical to those in figures 9.2 and 9.5.

The erosion of stable interfaces by Rayleigh-Taylor instability has in [Jacobs & Dalziel \(2005\)](#) been compared to that by grid generated turbulence, and since the RLIF experiments yield the width of the interfaces directly, this diagnostic was used to evaluate the performance of MOBILE on this mixing problem. To permit fair comparison, the simulations were processed to provide a synthetic diagnostic. The experimental light sheet was modelled by an exponential function along each light ray

$$\frac{\partial I}{\partial s} = -\eta(\phi)I, \quad (9.20)$$

where I is the light intensity field, ϕ is the concentration field, $\eta(\phi)$ is a decay function (taken to be linear) controlling the rate of adsorption of light, and s is an ordinate along each ray, for convenience assumed to be aligned with the vertical. The local light intensity $I(y, z)$ and the calibration charts of chapter 5 were then used to calculate the Acridine spectral response, from which a green-filtered image was constructed. This matched the experiments that were recorded using a monochrome UNIQ camera with a green dichroic filter over the lens. Both experimental and simulation image sequences were then processed identically to obtain the interfacial thicknesses as functions of time.

The relationship between the growths of the interfaces is expected to be governed by the entrainment relationship $u_e \sim uRi^{-n}$ with a suitably defined Richardson number. The $h(t)$ and $h_e(t)$ profiles (more robustly available from noisy data sets than $u(t)$ and $u_e(t)$) are plotted in figure 9.7, and they show good agreement between simulation and experiment. Given the well-established issues with ILES mis-predicting α , a virtual time-origin was chosen so that onset impact time for the Rayleigh-Taylor instability reaching the stable interface was correlated with the experiment. The early discrepancy in the stable interface case is an experimental artefact, caused by the light sheet being reflected off the barrier, and therefore highly illuminated unmixed Acridine inadvertently contributes to the integral measure of $h_e(t)$ at first. This effect is of course not modelled in the simulation; the non-zero initial value in the simulation is due to the thickness of the diffuse interface, initialised to match the experiment, and indeed when the barrier is removed this is the level to which the integral measure reverts.

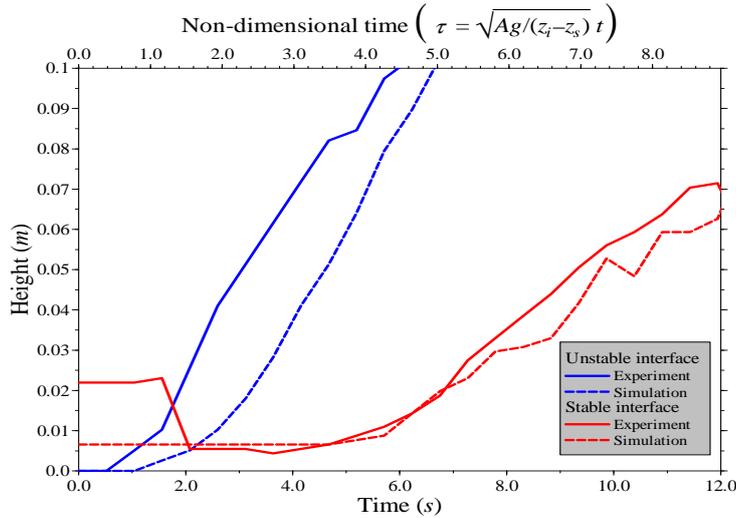


Figure 9.7: Integral measure of interfacial growth, comparison of experiment and simulation.

Clearly the growth rates of both interfaces are very well predicted by MOBILE, and after the early growth phase both are approximately linear. Linearity of $h(t)$ is expected from equation 9.13, but this makes no prediction about $h_e(t)$. That both functions are linear is encouraging confirmation that the entrainment relationship, equation 9.17, holds true in the Rayleigh-Taylor case, since for linear $h(t)$ the Richardson number is approximately constant. The Richardson number obtained by direct evaluation of equation 9.18 is $Ri = 3.78$, for a suitably chosen $\hat{\rho}$ and h , though this tacitly assumes that the turbulent length-scale is equal to (rather than simply scales with) the unstable interface height h . A more robust measure is to use equation 9.17, which yields $1.66 < Ri < 2.44$ for the range $1.75 > n > 1.00$, the accepted range from experiments discussed in Fernando (1991).

9.4 Summary

This chapter has re-examined the theoretical and experimental work of Jacobs & Dalziel (2005) and compared it to MOBILE simulations. The scaling relationships for $h(t)$ given by the theory require self-similarity of density and kinetic energy profiles is required. The evolution of the horizontally averaged middle (less dense) layer concentration was shown (in both experimental and numerical tests) to have self-

similar properties and thus is consistent with the requirements of the theory. A testable prediction of the theory has been investigated, namely the evolution of the maximum concentration, which the theory suggests follows $\hat{\phi} \sim \frac{1}{t}$, and the agreement is excellent. Using the new RLIF diagnostic, the growth of both stable and unstable interfaces (in separate experiments) was visualised, revealing qualitative differences between the spatial characters of the mixed fluid in each case. The unstable experiment showed length-scales growing with the Rayleigh-Taylor instability, as one would expect; in the stable experiment the mixed fluid near the stable interface has a filamental structure which does not grow in length-scale with time. Quantitatively, an integral measure of the interfacial growth was devised so that robust comparison of the growth rates was possible. The theory predicts linear late-time growth for the unstable case, and by invoking an entrainment hypothesis lifted from studies into grid generated turbulence, it was predicted in turn that the stable interface will grow linearly too if the Richardson number is constant. Experimental and numerical evidence appear to support both of these assertions. Despite the well-known uncertainties in predicting the rate of growth of Rayleigh-Taylor unstable flows with numerical simulation, MOBILE appears to perform remarkably well on this challenging test case.

Chapter 10

Conclusions

10.1 Review

In chapter 1 a historical overview of research into Rayleigh-Taylor instability introduced the topic and the three main avenues of research: modelling, experiments, and numerical simulation, which have contributed to the field. Chapters 2, 3 and 4 provide some detail on how existing techniques in these three avenues have been developed in this thesis for application to confined Rayleigh-Taylor instability.

Chapter 5 discussed a new experimental diagnostic technique called RLIF, which directly visualises molecular mixing in Rayleigh-Taylor instability. Chapters 6 and 7 studied Rayleigh-Taylor instability confined by geometry in a high-aspect-ratio domain, and this work was extended to confinement by stable linear stratification. Chapter 8 considers the Rayleigh-Taylor instability in a low-aspect-ratio domain confined on both sides by stable linear stratifications and focussed on energy transport and mixing efficiency. Chapter 9 examined asymmetric Rayleigh-Taylor instability, where it is confined on one side by a stable density interface. The remainder of this chapter summarises the main ideas and results, and identifies some interesting avenues for future work.

10.2 Themes and ideas

This thesis has examined the interaction of miscible fluids of different densities in various contexts where potential energy is unlocked from a system by Rayleigh-Taylor instability. The purest form of the instability, dense fluid overlying light fluid in free space with infinitesimal perturbations to the interface between the fluids, is not a configuration that can be created experimentally, so historically we have relied on a combination of technically feasible experiments and various forms of modelling to make progress. Most previous work has focussed on understanding how the instability grows, but this has been hampered because we only have a limited grasp of some of the sub-processes that control the instability. One poorly-understood process is molecular mixing, which is particularly important in miscible fluids. The motivation for this thesis has been to enhance our understanding both of the physics of molecular mixing, and the ways in which our models represent the physics.

Perhaps the best way to begin studying a process is to find a way of visualising it. In chapter 5 a new PLIF diagnostic technique, which employs a chemical indicator of mixedness, was developed and used to examine molecular mixing in Rayleigh-Taylor instability. Measures of molecular mixing used previously in the literature have been found not to be robust, and an alternative measure exploiting the properties of the new diagnostic has been proposed. Iso-surfaces of volume fraction are readily identified using the technique, and some new insight into the form of self-similarity of Rayleigh-Taylor instability has been uncovered. These developments led to a desire to understand individual components of the mixing process in more detail and explore them using other methods.

To begin decomposing the problem into more tractible components, ways to simplify the dynamics of the instability were investigated. The one-dimensional analogue of the pure instability is the high-aspect-ratio case of chapter 6, and this has the convenient property that the eddies are laterally confined. The eddies have $O(1)$ aspect ratio (approximately circular) and this constrains the rate of growth of the instability. The asymptotic behaviour of the mixing region height h follows $h \sim t^{\frac{2}{5}}$, compared with the $h \sim t^2$ one expects from the pure case. On any reasonable mea-

sure, the high-aspect-ratio ILES simulations exactly match the experiment, and this is surprising since numerical simulations have historically and consistently reported significantly lower ($O(50\%)$) values of the proportionality constant α in the pure instability relation $h = \alpha A g t^2$. The fixed size of the eddies and slower growth rate combine to adequately separate the turbulent eddy turnover time- and length-scales from the mixing region growth time- and length-scales. Thus the details of the mixing process, which are probably not captured accurately in the simulation, become unimportant to the overall instability growth. In contrast, the pure instability eddy turnover times and the range of eddy sizes scale with h , so any errors in the modelling of eddy interactions have an $O(1)$ effect on the estimate of instability growth.

Having ascertained that the MOBILE simulation tool performed well on the basic high-aspect-ratio case, the geometry was extended to include a bottom reservoir, sufficiently large that it could be regarded as an unbounded domain with fluid of constant density. This configuration was conceived to permit study in chapter 7 of Rayleigh-Taylor instability penetrating into an initially stably stratified layer, a state that cannot be realised experimentally in an overturned tank. As a stepping-stone, experiments of the homogenous-layer case were compared with a hierarchy of modelling tools. The analytical model, which yields $h \sim t^{\frac{2}{5}}$ in the homogenous case, assumes that the density profile remains self-similar, and that turbulent processes in the mixing zone have bulk properties that resemble a diffusion equation. Thus the density flux across the mid-plane is given by a simple relation $h\dot{h} \sim \kappa$, which can be closed with Prandtl's mixing length argument. Using the same assumptions, a general form for the local density flux can be obtained, and this relation can be integrated numerically. This low degree-of-freedom representation of the system sits intermediate in complexity between the analytical, 'zero'-dimensional model and the three-dimensional MOBILE simulation. It allows the upper boundary condition imposed by the top of the tank to be represented satisfactorily, and thus moves beyond the similarity model in its predictive capability. In particular the functional form of $\rho(z, t)$ is calculated with extraordinary accuracy.

The behaviour of the stratified high-aspect-ratio Rayleigh-Taylor problem was initially surprising. Whereas in the homogenous case the upper edge of the mixing zone

(h) was diffuse and ill-defined, the edge in the stratified case was extremely sharp. Analytical modelling of the instability growth was not especially successful since the second length-scale associated with the stratification introduces problems with models that require self-similarity. The one-dimensional numerical model gave extremely reliable predictions of the overall growth, and correctly predicted the sharpness of the mixing zone edge, but some resolution-dependent artefacts associated with the sharp edge remained. In attempting to refine the model to capture the interface dynamics, an energy transport equation was added into the system, which could decouple the location of production and dissipation of kinetic energy. Interestingly, allowing the model this freedom caused the predictions of $h(t)$ (for non-trivial parameter values) to diverge markedly from the experimentally obtained profile. It was deduced that production and dissipation in this system must therefore be balanced locally in time and space.

The work on high-aspect-ratio domains demonstrated the utility of simple analytical models, more general forms thereof, and ILES numerical simulations for investigating Rayleigh-Taylor instability. The rationale for studying this laterally confined case was to establish a benchmark problem tractable with a range of modelling tools, and then re-introduce complexity to explore the limitations of these models. In particular, ILES simulations are generally thought to perform poorly on mixing problems that include stable density gradients, and chapter 8 attempts unravel these issues. The configuration chosen was a Rayleigh-Taylor unstable interface sandwiched between two stable linear stratifications of the same gradient, in a low-aspect-ratio domain.

The instability grows as $h \sim t^2$ initially, before decelerating as the supply of potential energy is exhausted, and asymptotes towards a steady state. Using the exchange of potential and kinetic energy, a scaling for $h(t)$ was found analytically. However, the model embeds an assumption that the mixing region is uniformly well-mixed, and by construction this sets a well-defined end height $h(\infty)$. Intuition and simulation suggest the mixing region is not well mixed, and experiments confirm that $h_{expt}(\infty) > h_{model}(\infty)$ by a substantial margin. Further investigation showed that the height to which the instability reaches in this system is a function of the

‘aggregate mixing efficiency’, a parameter which describes how spent energy leaves the system at a micro-scale. Energy either leaves as heat by dissipation, or by changing the structure of the density field by doing mixing, and thereby changing its base state, or ‘background’ potential energy.

It is well established that the mixing efficiency in Rayleigh-Taylor flows is very close to 50%, *i.e.* equipartition of energy flux to heat and to mixing. In pure Rayleigh-Taylor instability this is the maximum value that is energetically possible, and indeed no other fluid process is known to generate a higher mixing efficiency. The unusual feature of the stratification-confined Rayleigh-Taylor instability is that it is energetically possible to have a higher mixing efficiency (up to 75%), and yet experimental evidence shows that the system still relaxes to 50%. Nothing phenomenological suggests equipartition of spent energy is unlikely, and it appears that this may be a fundamental feature of variable density turbulent flow. This of course raises the question of justifying the lower mixing efficiencies reported from other variable density fluid processes, but these are processes where turbulence in the system is not predominantly localised with density gradients. For instance, a system where turbulence is generated by an oscillating grid located beneath a stable density interface has a very low mixing efficiency since only a very small proportion of the turbulence reaches the interface to do mixing between the two fluids.

The ILES simulations in chapter 8 did not behave like the real fluid, but instead relaxed towards the maximum mixing efficiency solution of 75%. Mass and velocity diffusion are not explicitly calculated with the ILES method, so the apparent mixing and dissipation are created solely by the numerical scheme. The scheme is common to both mass and velocity advection so the numerical Schmidt number is $O(1)$. The real fluid has a much higher Schmidt number ($O(1000)$), and so the scales at which spent energy is dissipated are much larger than those associated with molecular diffusion of scalars. Despite this huge separation of scales, the real fluid has an equipartition of spent energy, so it is especially interesting that energy transport in ILES, where by construction there is no such separation of scales, is incorrectly biased towards scalar diffusion rather than dissipation.

The final method of confinement, examined in chapter 9 was a stable density

interface sitting beneath a Rayleigh-Taylor unstable interface. Initially uninhibited by the lower interface, the instability grows as $h \sim t^2$ but at late time this decelerates to $h \sim t$, since there is a finite volume of buoyant fluid driving the instability. While mixing and scalar entrainment across the stable interface in this instance is induced by Rayleigh-Taylor instability, the underlying behaviour is common to the well-studied problem where the turbulence is generated by an oscillating grid. The theoretical relations that underlie this problem are shown in the Rayleigh-Taylor context to continue to work well. Experiments were conducted both using standard and chemically reacting LIF diagnostic techniques. By using dye calibration data, ILES simulations were processed to resemble the experimental videos, and these fair comparisons indicate that ILES is a surprisingly useful tool for understanding this fluid system, despite the widely-held doubts about the method's performance in stably stratified flows.

10.3 Remarks on future directions

This thesis has unearthed several fascinating challenges to our understanding of variable density mixing processes, and has attempted to answer some, firstly by pruning the complexity of the problem down to tractable analogues, and then progressively re-introducing complexity to establish the boundaries of our modelling capability. Inevitably, many questions remain ignored, or are incompletely answered, so the following is a merely a list of avenues which presently appear ripe for further investigation.

At the heart of the scientific method is recurrent benchmarking against reality, and the level of detail available from experiments in many ways determines the scope for progress. The reactive LIF diagnostic technique has been hampered by inadequate levels of incident light (of the correct wavelength for excitation), and the related problem of noise thresholds in the camera equipment. The power of the technique is that it simultaneously yields full range scalar concentration information and a low-noise demarcation of an individual contour, yet with existing equipment it was not possible to exploit this capability. A logical next step is to deploy laser technology for illumination, and upgrade the camera. Only recently has high res-

olution, high speed, high sensitivity video technology become relatively affordable. This would permit much more detailed micro-scale investigation of the structures present in the mixing process and provide some confirmation (or otherwise) of accepted wisdom which we have hitherto only obtained from numerical simulation and phenomenological arguments.

At a more mundane level, the measurement and production of linear stratifications particularly important in chapter 8, was fraught with practical problems. Despite prior thought and immense care, consistent linear stratifications could not be produced, due to shallow layer mixing at an early stage of filling, mixing induced by shutting the barrier, fatigue of the flexible pipe in the peristaltic pump altering flow rates, tank leakage through the barrier seal, and minor issues with overfilling to offset any future leakage. Technology now exists to electronically control the stratifying process, and it would be perfectly possible to calibrate out the influence of these undesirable artefacts.

Measurement of the initial stratification in particular was extremely difficult, since the motorised conductivity probe had been designed not to traverse the full height of a tank: at the top the probe could only be started from a position part-way submerged, both to maintain electrical conductivity around its insulating tip and to siphon water through its interior. At the other end the probe could easily be damaged by being driven into a hard surface, and with every traverse required precautionary recalibration. A minor redesign of the probe, and a programmable linear stepper motor to drive it would resolve these issues.

While none of the above would materially change the results reported herein, these are obvious refinements that one would insist upon if undertaking similar work in future. A much less straightforward refinement is a replacement tank for the low-aspect-ratio Rayleigh-Taylor instability. The inherent flaw with the current design is the interaction between barrier removal and the instability, which leads to large-scale bulk motion that is not predominantly Rayleigh-Taylor driven. A long tank, with a barrier removed electrically at a steady velocity, would provide adequate separation of the instability and barrier removal timescales, would not exhibit the overturning circulation, and barrier end effects would become insignificant. The

Rayleigh-Taylor instability would therefore be predominantly spatially rather than temporally evolving, and one would envisage that this statistically steady scenario would greatly improve the reliability of the experimental measurements. Funding for this design concept has been sought.

All of the work in this thesis considers scalar transport in Rayleigh-Taylor driven flows, but our intrinsic understanding of variable density mixing processes will only move forward substantially when we have techniques that provide simultaneous velocity and scalar information. There are several ways one might achieve this aim - particle image velocimetry combined with PLIF is an obvious first step. Stereoscopic imaging is a natural progression to obtain out-of-plane velocity information; conceivably PLIF could be extended into the third dimension too by oscillating the position of the light sheet plane. Perhaps a more revealing approach might be to derive the velocity information from the evolution of a scalar field. There are existing methods in the computer graphics literature that use scalar gradient information to infer a velocity field (*e.g.* used in digital speed cameras), but the problem is inherently ill-posed and a regularisation suitable for fluid mechanics applications is non-trivial. Kalman filtering, a technique used to reconcile numerical weather predictions with observational data is one popular method. Another perspective is to regard the search for a velocity field evolution consistent with a given scalar field evolution as a search for the initial condition of the system, and tackling this inverse problem could lead to major leaps in our understanding of mixing. Funding has also been sought to pursue these ideas.

However important a contribution experiments make to our knowledge-base, modelling is the route by which that knowledge develops into understanding. Numerical simulation is the pinnacle of our modelling capability, yet frequently it makes incorrect predictions that we cannot explain. Chapter 8 revealed from macro-scale measurements what appears to be a fundamental truth about variable density mixing at a micro-scale, namely the equipartition of spent energy into mixing and dissipation, yet the numerical simulations do not replicate this behaviour. While a complete understanding of energy transport in these flows remains elusive, it would seem a potentially fruitful avenue of study to build such fundamental laws into numerical

codes and evaluate whether they then better represent the macro-scale behaviour.

In particular the ILES method embeds some questionable assumptions about the micro-scale dynamics. Energy leaves the numerical system most rapidly when velocity gradients in the flow direction are steep, but in the physical system, transverse shear is primarily responsible for energy loss. In a flow whose macroscopic behaviour is close to isotropic, these discrepancies may be masked, which might explain the relative success of ILES for Rayleigh-Taylor flows and its relatively poor performance in some other cases. It would be a worthwhile exercise to manipulate the numerical scheme to respond isotropically to steep gradients and compare macro-scale behaviour.

In ILES flows, fluids mix by cell averaging at each timestep. Momentum is treated in a similar way, and therefore the effective Schmidt number is $O(1)$. In this special case, both dissipation and mixing occur at the grid scale. If a small explicit viscosity were introduced to raise the Schmidt number, a distinct Batchelor scale would be created and it is not immediately clear whether the computed mixing efficiency would change. This might provide some insight into the (errant) behaviour of the ILES method in the doubly stratified case discussed in chapter 8.

The MOBILE code used throughout this thesis has exclusively used the ILES method, but in light of the above suggestions, several ways in which the numerical algorithm should be modified have been identified. From a physical standpoint, simulating the long tank experiment requires a larger computer (exploiting the existing parallelisation) and it would be desirable to have an adaptive mesh refinement capability to optimally allocate computational resource.

One piece of work in progress is an algorithmic modification that incorporates a solid phase with a ‘Volume-of-Fluid’ approach, allowing withdrawal of the barrier, and other arbitrary moving solid bodies to be explicitly modelled. This would have improved the correlation between experiment and simulation in the present work, particularly in the classic Rayleigh-Taylor case since the bulk overturning of the fluid forms a significant component of the system development. Ensuring net volume conservation with more than two scalars is a well-known issue for non-linear numerical schemes, and in an incompressible pressure projection code, this is an especially

important property to maintain. Some work is yet required to achieve this in the general case.

10.4 Final thoughts

The above remarks on possible avenues for further work are in some sense pedagogical developments of work already completed. Technology improves and becomes more affordable, so incremental (that is not to say insubstantial) progress can be made in clearing a backlog of open questions. There are however wider scientific questions that are inspired by, rather than following from the work herein. Do the same relationships hold in mixing between non-Newtonian fluids? What happens to the equipartition of spent energy when surface energy begins to play a role? Does an emulsion have the same energetic behaviour at a macro-scale as a miscible mixture? If a water column is stably stratified in salt but unstably stratified in temperature, is there sufficient non-linearity in the mixing to self-sustain Rayleigh-Taylor instability? If we could resolve the Batchelor scale, would our cartoon-like intuition of turbulent structures still hold? Since equipartition of spent energy appears to exist at non-unity Schmidt numbers, does it follow that available energy distribution is scale-invariant in the inertial range? These are questions that are easy to pose, but some may take a more than a lifetime to answer.

Appendix A

Validation exercises on MOBILE

A.1 Introduction

This appendix contains details of several verification and validation test problems for the code used elsewhere in this thesis. The aim of these tests is firstly to verify that the errors in the numerical approximation to the equations of motion reduce with increasing resolution, and provide evidence that any convergence with resolution tends towards a ‘real’ physical solution of the equations of motion. Secondly, these tests are designed to demonstrate that MOBILE can be applied to a variety of standard fluid problems for which there is ample existing experimental or numerical data and theoretical understanding in the literature, and shown to compare well with this material. The cases considered herein are

- Single-mode Rayleigh-Taylor instability
- High-aspect-ratio Rayleigh-Taylor instability
- Kelvin-Helmholtz instability
- Lock-exchange gravity currents

A.2 Grid convergence verification

A.2.1 Single-mode Rayleigh-Taylor instability

Single-mode Rayleigh-Taylor instability is an idealised test case, which is easy to define numerically, but very challenging to achieve experimentally. Some theoretical results exist, indeed as discussed in §4.2.2, the potential flow approach has been widely used, and provides the basis for many models of multi-mode interaction. The calculations presented here are intended to compare against Ramaprabhu *et al.* (2006), which explores the limitation of potential flow theory to single-mode Rayleigh-Taylor instability, using a number of numerical codes from various research institutions. The problem specification has been recycled and forms ‘Test Problem One’ announced for the 11th International Workshop on the Physics of Compressible Turbulent Mixing (Santa Fe, July 2008), which should eventually accumulate a community-wide repository of directly comparable simulation data on this and other easily defined, if far from simple problems.

Buoyancy-drag models predict saturation of individual bubble or spike structures towards a terminal velocity,

$$u_{\infty} = Fr \sqrt{\frac{Ag\lambda}{1+A}}, \quad (\text{A.1})$$

re-expressed in terms of a body Froude number Fr from equation 4.15, where λ is the wavelength of the structure. However, Ramaprabhu *et al.* (2006) discovers that the single-mode instability undergoes the expected exponential initial growth (though numerical evidence for this is not conclusive), reaches a terminal velocity as predicted by potential flow theory ($Fr = 0.56$), but subsequently re-accelerates. Figure A.1 shows MOBILE prediction of bubble rise in a $32 \times 32 \times 256$ doubly periodic domain, where the perturbation is aligned so the spike penetrates downwards in the middle of the box. These images are reconstructed from two half-slices out of phase by π in both of the periodic directions. The Atwood number chosen for this study was $A = 0.25$, to demonstrate the non-Boussinesq capability of MOBILE (which was not required when comparing simulations with laboratory experiments in previous chapters).

The most obvious parameter on which to demonstrate convergence with increas-

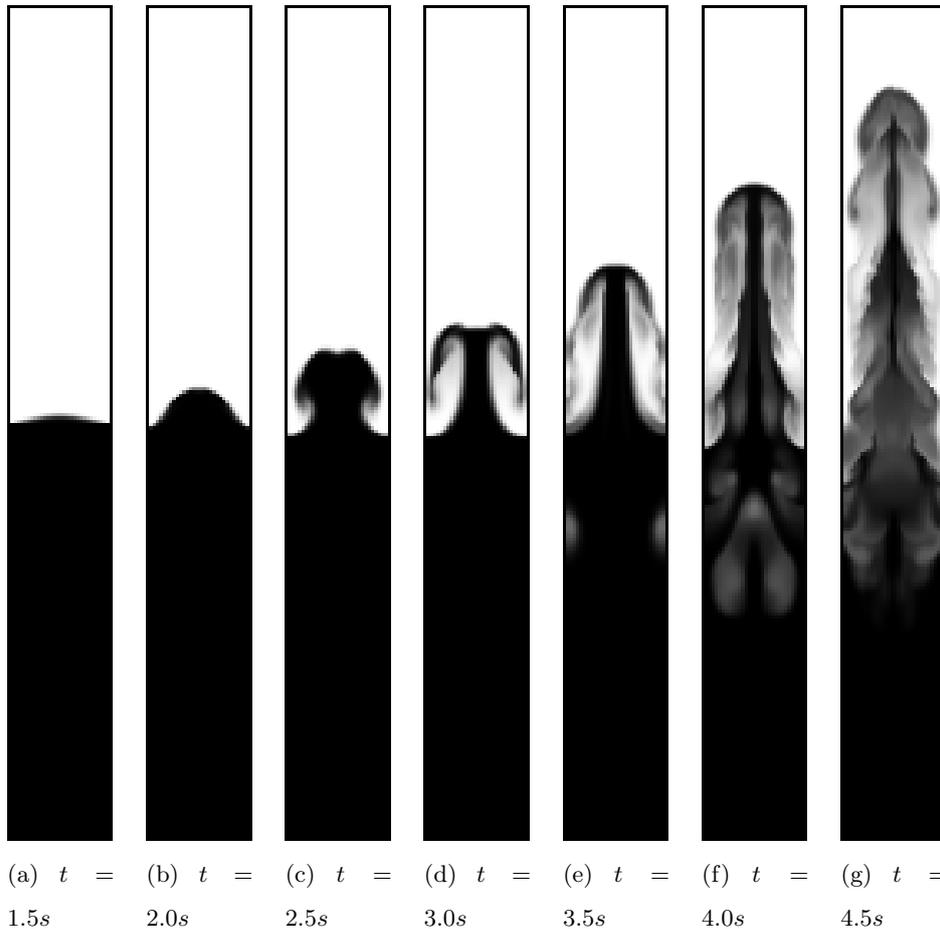


Figure A.1: Single-mode bubble growth with Atwood number $A = 0.25$. Scalar concentration is represented in greyscale, and spatial resolution is $32 \times 32 \times 256$.

ing resolution, is the non-dimensional bubble and spike height as a function of non-dimensional time. Figure A.3 shows these results ranging in resolution from 4×4 to 64×64 cross-section. The 4×4 simulation is very poorly resolved so one would not expect it to capture many of the features of the instability, and indeed the 8, 16 and 32 simulations appear to establish a trend for the growth of the bubbles and spikes. The non-Boussinesq Atwood number gives rise to more quickly growing spikes than bubbles, and this accords with predictions from potential flow theory. Spikes tend to have smaller frontal area so their terminal velocity is expected to be higher.

The 64×64 simulation is somewhat anomalous in this regard, and diverges from the trend established by lower resolution simulations. Examination of the data sug-

gests that with increasing resolution (decreasing numerical dissipation), symmetry is broken progressively earlier, and by 64×64 , symmetry is lost while the bubble is still forming. The strong planar symmetry present initially can be observed in the image sequence in figure A.2.

The deviation of bubble and spike trajectories from the trend established by 8×8 , 16×16 and 32×32 cases corresponds to the point where the net vorticity associated with the loss of symmetry causes the bubble to break down. The reason symmetry breaks in the first place is a sensitivity in the numerical algorithm to changes in direction of upwinding. Currently the code uses conditional branching to decide which direction in which to acquire upwind information for the advection stencil. Code such as

```
if(u_face>0.0) { u_upw=u_left; }
else           { u_upw=u_right; }
```

is sensitive to numerical noise at the axis of symmetry where horizontal velocities should be exactly zero, and it appears that rapid changes in upwinding decisions can lead to exponential growth of asymmetry. One possible solution is to use a blending function to interpolate between the current upwind scheme and a new centred scheme as $\pm u \rightarrow 0$, but this would greatly complicate and slow down execution of the advection routine, and may risk violation of monotonicity. In almost all real-world problems, exact preservation of an unstable symmetry is of negligible importance, and the effort required to develop such a scheme purely for this pathological case could not be justified. It should be noted that the RTI-3D code results in Ramaprabhu *et al.* (2006) also exhibit an anomalous trend in its 64×64 test (see figure 11(a) in that paper) and, in common with RTI-3D, the advection algorithm for MOBILE is based on the principles outlined in Andrews (1995).

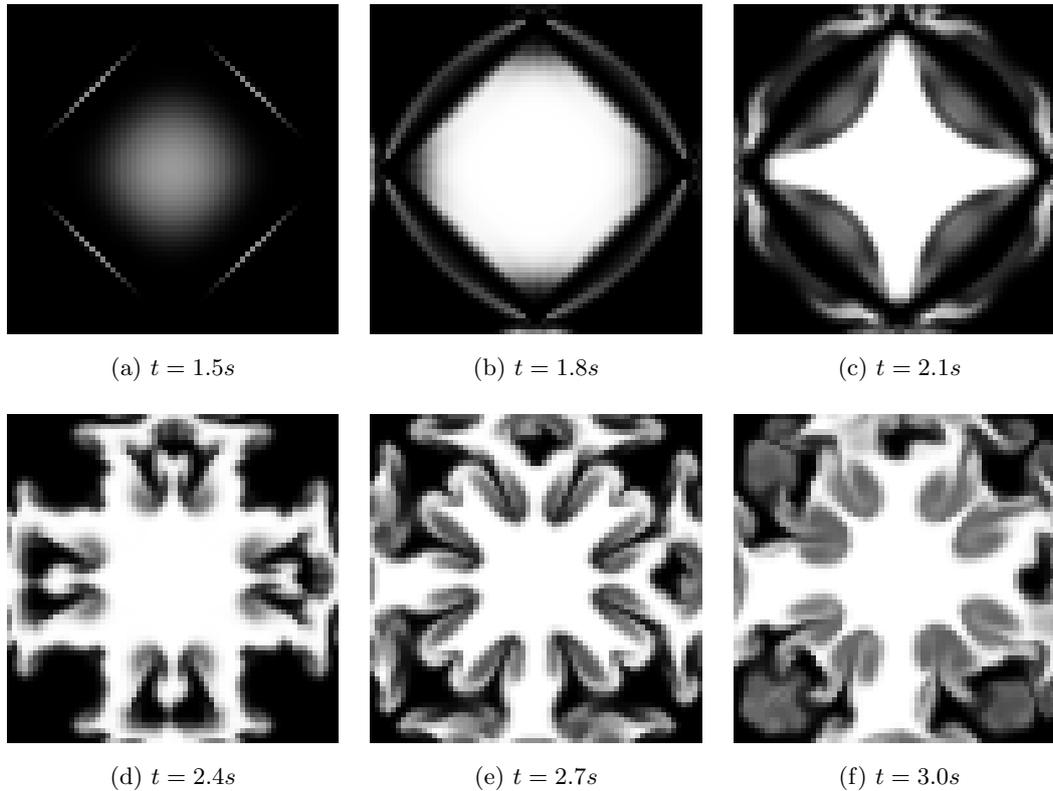


Figure A.2: Horizontal section of scalar concentration at mid-height plane, showing patterns of increasing complexity with time. The Atwood number is $A = 0.25$.

In the presentation of subsequent data, the 64×64 case is discarded for the reasons outlined above, and the 4×4 case is discarded for aesthetic reasons since both its trend and the spatial quantisation of the data make further comparison unproductive. The behaviour of bubble and spike velocity is shown in figure A.4, normalised for Atwood number and gravity in the same manner as a terminal velocity with a Froude number of unity. Clearly good convergence is exhibited, and the re-acceleration reported in figure 3(b) of Ramaprabhu *et al.* (2006), is observed.

A more stringent test is to compare non-dimensional velocity against the non-dimensional height, since this gives a better indication that the bubble and spike structures have the correct buoyancy and drag when they have reached a given aspect-ratio. Bubble data from Ramaprabhu *et al.* (2006) made available for ‘Test Problem One’, is cross-plotted in black in figure A.5. This sits reassuringly closely over the MOBILE 32×32 data through the initial growth, deceleration/terminal

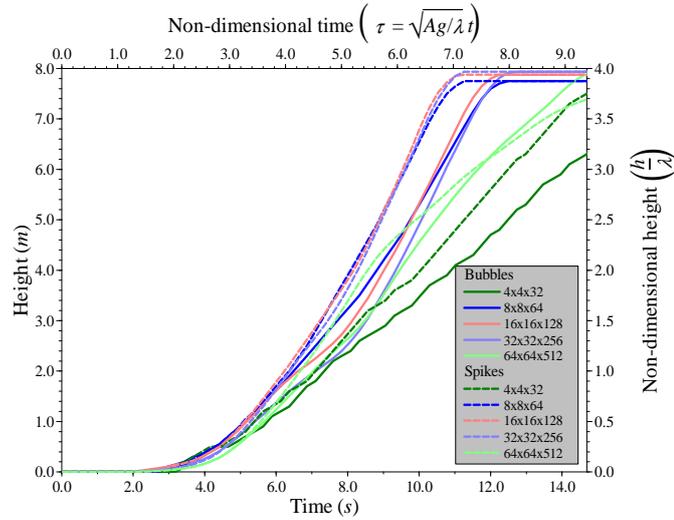


Figure A.3: Convergence of single-mode instability growth evolution with increasing resolution.

velocity, and re-acceleration phases, though it is less clear why codes using similar methods should mis-match at later time.

Taken together, the MOBILE results in themselves demonstrate convergence on a solution with increasing resolution until breakdown of symmetry, and performance relative to other codes is comparable, even to the extent that problems appear at the same points.

A.2.2 High-aspect-ratio Rayleigh-Taylor instability

A second grid convergence study was performed on high-aspect-ratio Rayleigh-Taylor instability, since it provides contrast in time- and velocity-scales, and the problem setup is more representative of real-world mixing applications where the code is intended to be used. Chapter 6 examines in some detail the comparison between simulation, theory and experiment. Here the same initial density profile is used, but the focus is on the convergence of solution with increasing resolution, and confirmation that simulations match previously published experimental data. The feasible limit on resolution for a dual core 2GB single memory computer was $64 \times 64 \times 2560$. However, as already remarked in §6.2.3, the data storage requirements to retain for presentation satisfactory time-resolution of the simulation exceeded the available

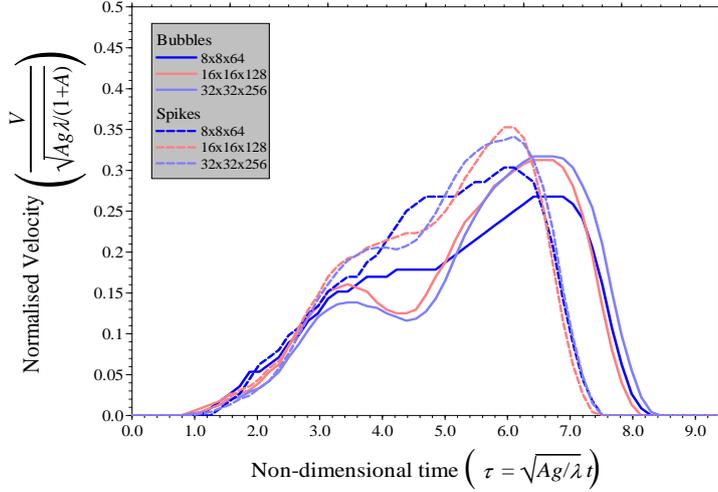


Figure A.4: Convergence of single-mode instability growth velocity with increasing resolution.

hard drive capacity, so data for this case is only available at 20s intervals, instead of every 2s on lower resolution cases.

The main parameter on which to demonstrate grid convergence is the envelope growth. The 4×4 and 8×8 cases are inaccurate and grossly overpredict the initial growth. The 16×16 case and thereafter appear to approximate well the analytically predicted $h \sim t^{\frac{2}{5}}$ functional form; an experiment matching this form is cross-plotted in black in figure A.6. This data is from Dalziel *et al.* (2008), but because of the experimental overturning leading to a significant region of mixing around the interface height, a virtual origin correction has been applied to visually match the computational data. The scalar concentration threshold taken to define the instability front is 0.98, as chosen in chapter 6. There appears to be some oscillatory component to the grid convergence, since 16×16 and 64×64 cases appear to correspond closely, with the 32×32 undershooting somewhat. The computational cost of doing simulations at yet higher resolution to confirm this was prohibitive, and since the 64×64 case is poorly under-resolved in time, it is difficult to be certain exactly how the trend lies relative to the other cases.

The experimental results from Dalziel *et al.* (2008) show that the horizontally averaged vertical profile of density has a self-similar form after the initial transients

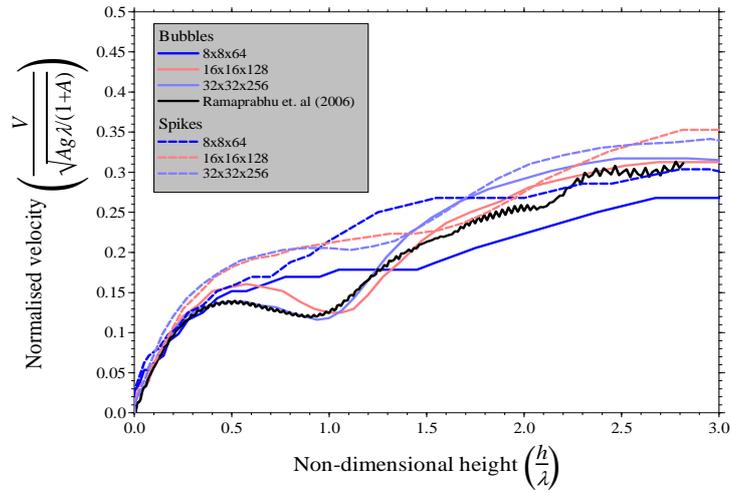


Figure A.5: Plot of instability growth against instability velocity.

have died away. In figure A.7 the profiles for the adequately resolved simulations (16, 32 and 64 cases) are cross-plotted with one another and the experimental data, on normalised scales of height and scalar concentration. Normalised height is defined as the range between the 2% and 98% contour, which explains why the profiles extend beyond this range. This figure establishes that the experimental and simulation vertical profiles have a similar structure, and the time-variation of the structure is insignificant. However, it is difficult to extract any direct confirmation of grid convergence because the density of information on the plot.

Figure A.8 resolves this by making the (very reasonable) assumption based on the evidence in figure A.7 that the functional form is invariant with time, so the time-mean of the self-similar form is plotted for each resolution. The results clearly demonstrate not only that the simulations are grid convergent on this fairly exacting metric, but also that they compare well with the experimental results once adequately resolved.

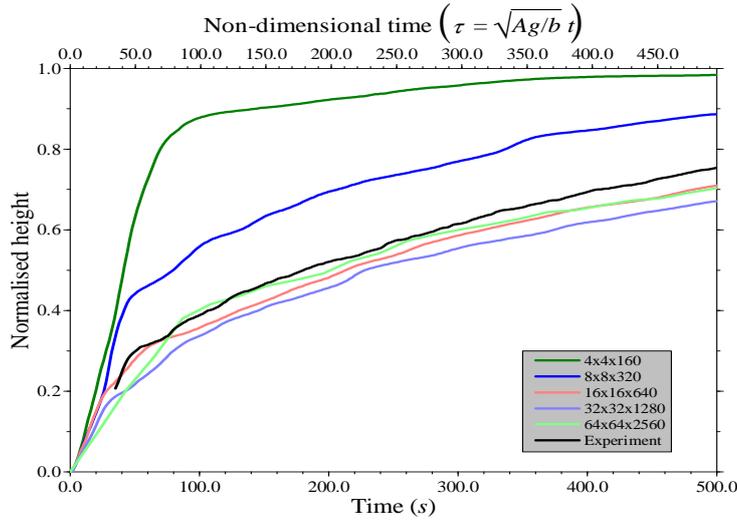


Figure A.6: Convergence of instability growth evolution with increasing resolution, and compared with experiment.

A.3 Validation on other mixing problems

A.3.1 Stratified Kelvin-Helmholtz instability

Shear across an interface gives rise to a very beautiful instability known as Kelvin-Helmholtz, where small disturbances on the interfacial surface grow exponentially in amplitude at first, then at later time roll up into vortical structures.

In the ideal case, the shear interface between opposing flows is an unbounded two-dimensional vortex sheet in unstable equilibrium; by the Biot-Savart law each point on the vortex sheet is experiencing induced velocities applied by all other points on the sheet, and these integrate to zero. However, any distortion of the interface modifies the distance between the mis-positioned points and neighbouring vorticity, such that the net induced velocity is non-zero. Since this velocity is directed perpendicular to the sheet, the system is unstable and exponential growth can be expected. Non-linear development occurs where the vortex sheet self-advects perturbations towards one-another. These perturbed regions tend to clump together and form mutual orbits, inducing yet more vorticity into them and reducing the vorticity density elsewhere. The vortical regions organise themselves in to recognisable finite-sized billow-like features.

In the laboratory, a tank with an initially horizontal stable density stratification

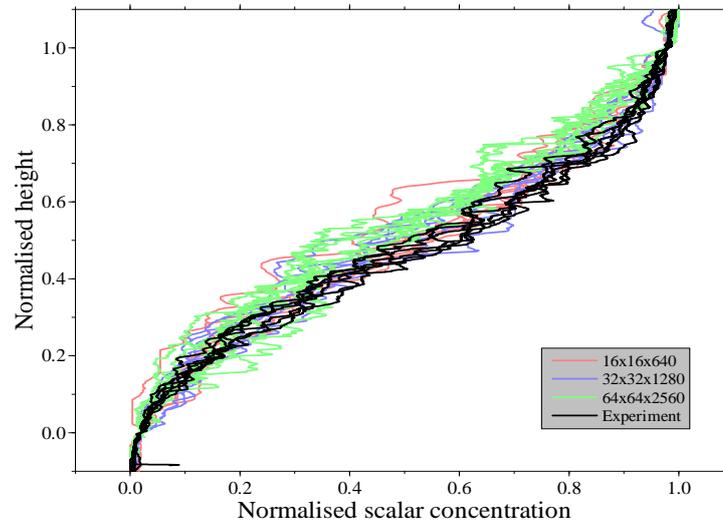


Figure A.7: Demonstration of self-similarity of scalar concentration profile from MOBILE simulation at three resolutions, and compared with experiment.

is tilted to a small angle. The dense fluid accelerates downwards, and the light fluid rises. Except at the tank ends, the interface initially remains aligned with the tank after tilting. In the case where the interface sits at half-height in the tank, the fluid on both sides of the interface accelerate equally in opposite directions and, subject to a Richardson number stability criterion, any initial perturbations on the interface grow without being displaced in either direction. Fluid reaching the tank ends has nowhere to escape, and encroaches on the other density layer. As the Froude number increases, the fluid building back forms a hydraulic jump. Fortunately, the fluid ahead of the travelling jump is unaware in advance of its arrival and, although the region exhibiting the instability diminishes, the instability itself is unaffected. For these reasons, the rotation rate, final tilt angle and upper and lower densities were selected such that the instability could fully develop before the hydraulic jumps meet in the middle.

Previous research on the instability provided a benchmark for expected outcomes, but the precise development of the instability is (naturally) a strong function of the initial conditions, and to obtain a high quality match between experiment and MOBILE simulation, a characterisation of an experimental setup was required. As an illustration of the influence initial conditions have on the instability development,

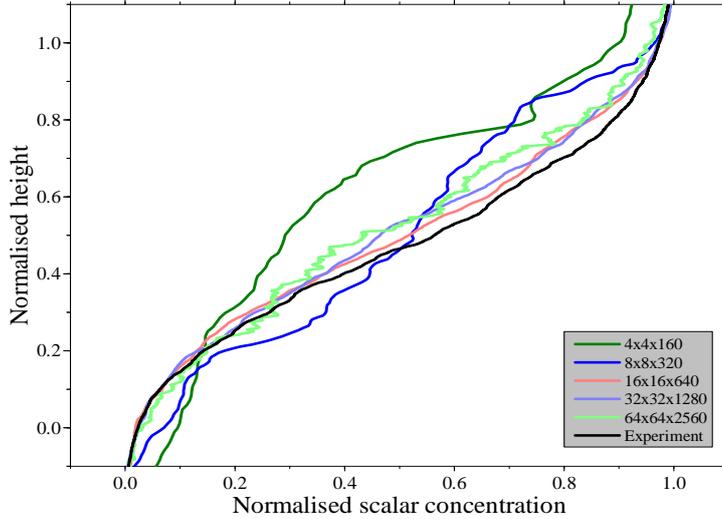


Figure A.8: Convergence of mean normalised scalar concentration profile with increasing resolution. MOBILE simulations are compared with experiment.

two simulations are shown in figure A.9, with the upper image in each pair showing the instability growth where there is no explicit perturbation to the interface, and the lower image where the perturbation was chosen to match an experiment.

Clearly, and unsurprisingly, the hydraulic jump velocities are unaffected by the perturbation; this suggests that the perturbation does not affect the acceleration of fluid in the middle region during the initial, pre-onset phase. Since floating-point error will inevitably lead to some small perturbation of the interface even in the nominally unperturbed case, Kelvin-Helmholtz instability has the potential to develop. The MOBILE code uses a multi-grid convergence acceleration, described in some detail in §3.3.1, for the pressure correction calculation, and in these simulations is converged according to a global residual criterion

$$\left| \frac{res_n}{res_0} \right| < 1e^{-4}, \quad (\text{A.2})$$

on the finest grid level (where n is the iteration counter). The coarsest grid level is also forced to satisfy this criterion, to ensure that low wavenumber (potentially larger amplitude) pressure corrections are applied quickly, minimising the number of V-cycle sweeps that are required. Since only global measures of error are available as convergence criteria, there is ample freedom to satisfy these constraints but develop, from numerical noise, an arbitrarily large interfacial perturbation. Indeed Kelvin-

Helmholtz billows emerge in the initially unperturbed case with a most-unstable mode period which equates with the coarsest grid scale, 32 cell-widths.

In attempting to match the experiment, care was taken to model the whole process of tilting the tank, obviously by rotating the gravitational acceleration component, but also by modifying the velocities with source terms to account for the change of reference frame, *i.e.* the Coriolis term ($2\boldsymbol{\Omega} \times \mathbf{x}$) centripetal acceleration ($\boldsymbol{\Omega} \times \boldsymbol{\Omega} \times \mathbf{x}$) and angular acceleration ($\frac{\partial \boldsymbol{\Omega}}{\partial t} \times \mathbf{x}$). Additionally, the interface was observed to have a slow gravity wave residual from the filling and stratifying process, and this was found to be instrumental in setting the size and period of the billows, which varied along the tank. Theoretical considerations (Thorpe (1985)) highlight the importance of the initial (non-zero) interface thickness in determining the most unstable mode. Unfortunately the thickness cannot be reliably estimated from the experiment because the gradients of refractive indices in the interfacial region act such that the apparent dye gradient is steepened, hence the interface thickness is underestimated. Trial and error was required to find a thickness for the simulation interface that favoured the growth of modes observed in the experiment. The chosen profile was given by

$$\rho = \rho_l + \frac{1}{2}(\rho_u - \rho_l) \left(1 + \operatorname{erf} \left(\frac{2(z - z_i)}{25.4\sqrt{\pi H}} \right) \right) \quad (\text{A.3})$$

where H is the tank height. The interfacial height variation that emerged in the experiment was measured, scaled in amplitude and used as a representative initial condition for the simulation. The image sequences in figure A.10 show a direct comparison of the simulation and experiment in the restricted domain that was captured on video during the experiment. For clarity the tilt angle of the tank is shown.

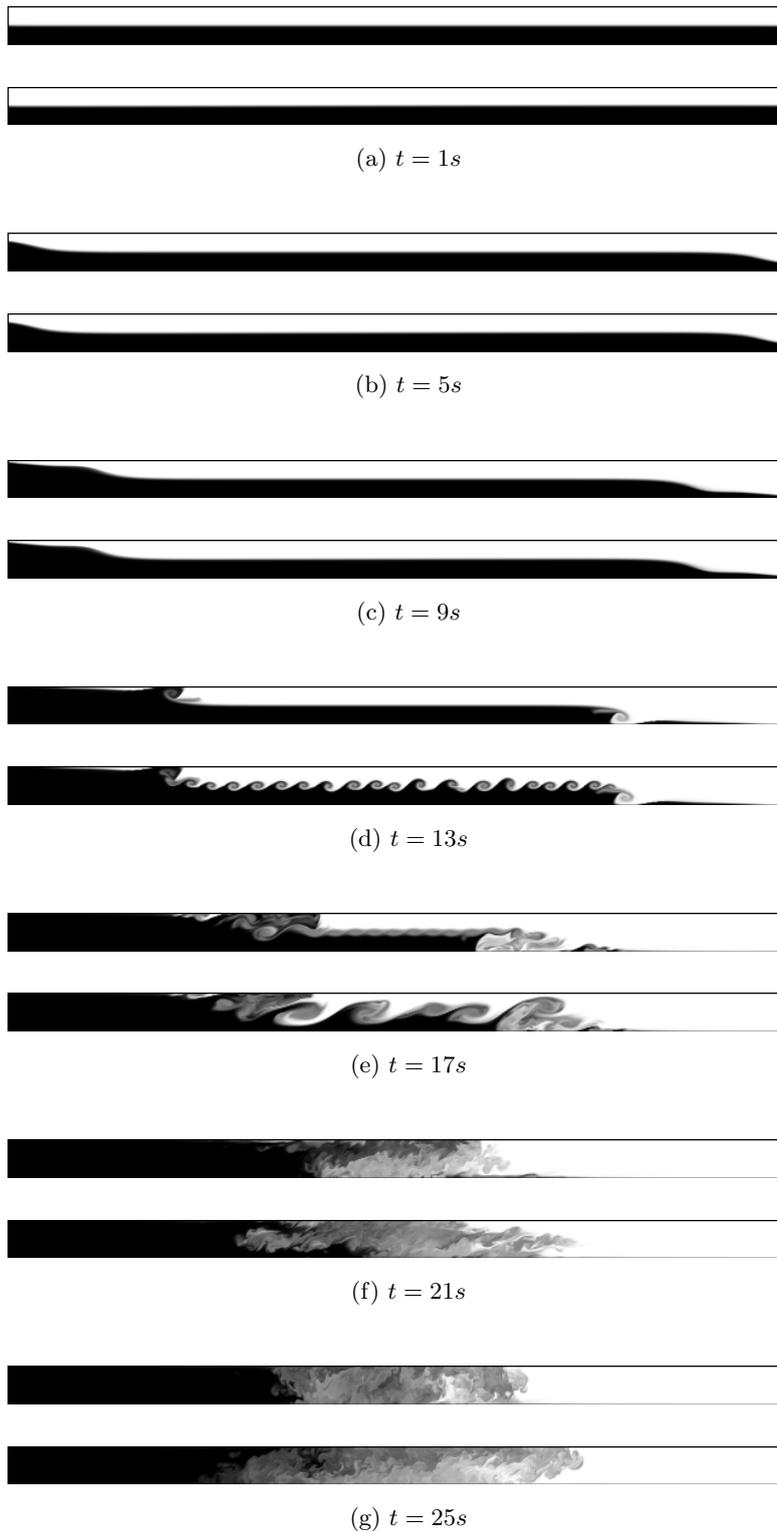


Figure A.9: MOBILE simulations of Kelvin-Helmholtz instability comparing no interfacial perturbation (upper image of each pair) with one (lower image of each pair) perturbed to closely match an experiment. For ease of presentation the tank inclination is not shown.

The features that the simulation captures well are the diameter and period of the billows in the imaging region, and their apparent growth rate. It was found, through numerous trial simulations, that the spatial variation of the size of the billows is due to a gravity wave oscillating in the tank before the experiment was initiated. Towards the left of the imaging region the interface sits approximately at the tank half-height, towards the right, it is substantially lower than half-height, and this seems to promote growth of larger billows. The larger wavelength appears to arise particularly where there is a dominant perturbation among weaker neighbours, which leads the weaker vortices by mutual induction to be engulfed into the larger ones. In the third pair of images in the sequence, such a merging event can be seen in both experiment and simulation. Predicting exactly which vortices will become dominant and which will become engulfed is a poorly conditioned problem, and since the simulation initial condition only approximates the observed initial conditions, achieving a precise match between simulation and experiment is challenging. A rigorous, systematic reverse-engineering of the experimental outcome was beyond the scope of this validation exercise.

Qualitatively however, the smaller billows on the left break down into incoherent mixed fluid at an earlier time than those on the right - as can be seen in the fourth experimental image of the sequence. The larger vortices protrude further into the ‘freestream’ flow and their extremities are advected, giving a distinctive elongated ‘cats-eye’ shape. The simulated vortices preserve their coherence for longer than the experiment, and there are two contributing causes. Firstly, the maximum affordable resolution was $32 \times 64 \times 1280$ and this gives a somewhat lower ‘effective’ Reynolds number than the experiment, and secondly, the experimental apparatus was retro-fitted with a turbulator mesh along its base, since in a carefully conducted experiment it is entirely possible for interfacial perturbations to be too small to fully develop Kelvin-Helmholtz instability in the required time. The turbulator adds high wavenumber velocity disturbances to the flow which help initiate the instability, but may also be hastening its breakdown. To reduce the discrepancy between simulation and experiment, low amplitude noise of random phase was added to the density field initial condition at the interface, but this has been rather less effective than hoped

at assisting vortex breakdown. Interestingly, because the vorticity density is higher in the simulation than the experiment, a secondary instability has time to develop before the hydraulic jumps arrive. Given appropriate conditions, such secondary instabilities have previously been observed in experiments ([Thorpe \(1985\)](#)). The arrival time of the hydraulic jumps (as seen at the edges of the final image pair, figure [A.10\(g\)](#)) is also correctly predicted.

The exchange of energy in the tilting tank is the key to understanding the mixing taking place in Kelvin-Helmholtz billows and previous work uses the concept of decomposing potential energy into two components, available and background potential energy. As discussed in [§8.3](#), changes in background potential energy are uniquely associated with mixing events since adiabatic rearrangement of an unmixed density field cannot produce the same stratification (and hence potential energy) as one in which some mixing has occurred. The graphs of figure [A.11](#) show how energy is exchanged in the tilted tank simulation. The measurements of potential energy are made in the reference frame of gravity once the tank is stationary in its tilted state, so while the tank is rotating, there is an error of $O(1 - \cos(\theta_0 - \theta))$. Kinetic energy is measured in the tank reference frame (for convenience), so during unsteady rotation a force is applied by the tank to the fluid, and as shown in the figure, this manifests itself as an increase in the fluid kinetic energy, after rotation has stopped.

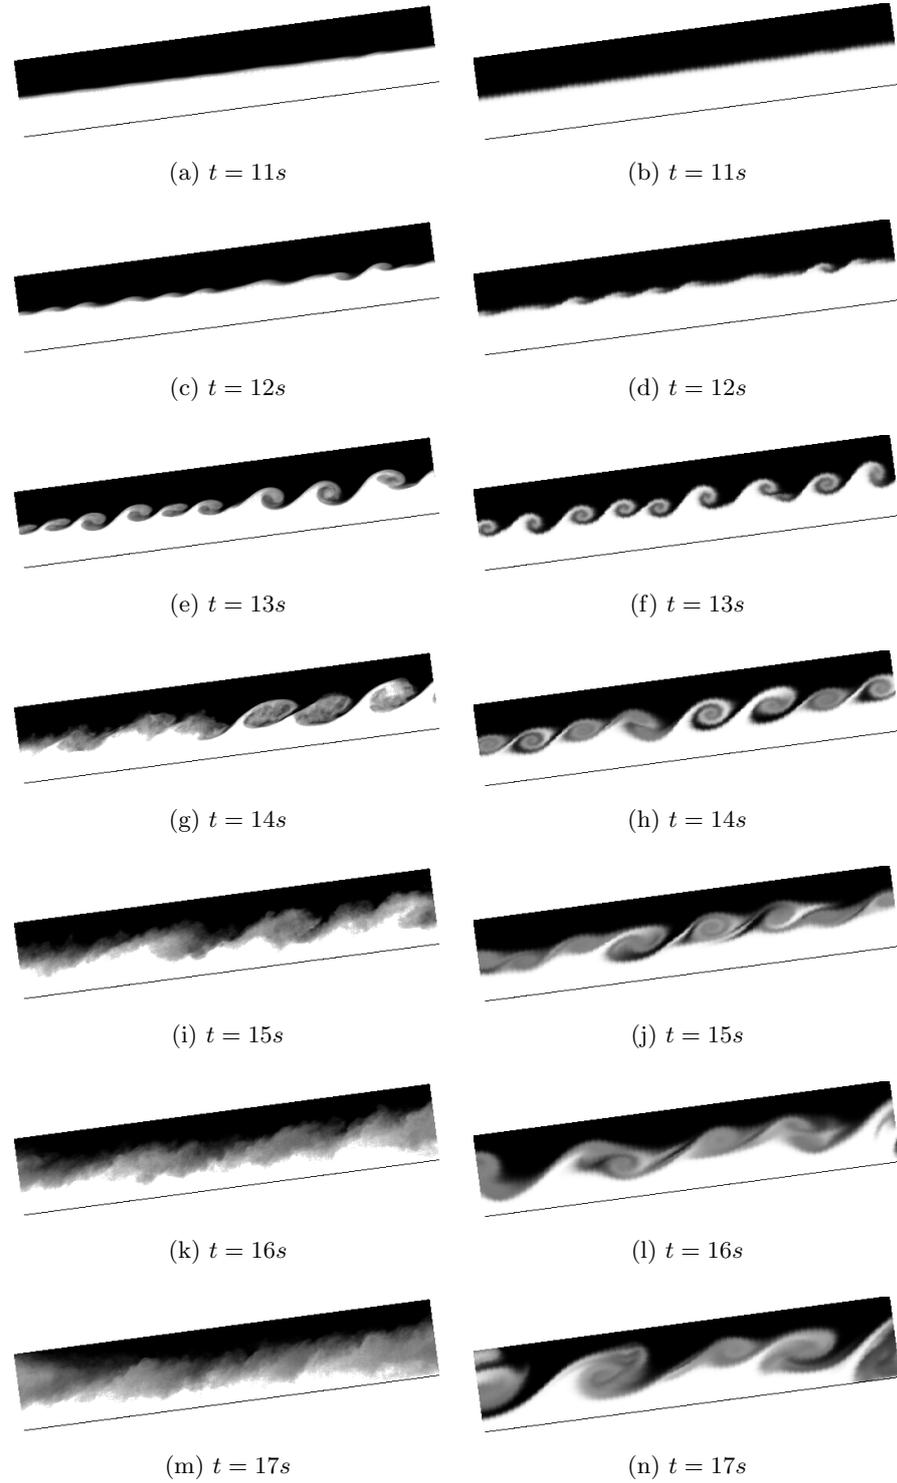


Figure A.10: Comparison of experiment and MOBILE simulation of Kelvin-Helmholtz instability. The tank tilt angle is 7.58° and the upper and lower layer densities in both cases are $\rho_u = 998.2 \text{ kgm}^{-3}$, $\rho_l = 1012.5 \text{ kgm}^{-3}$. The experimental sequence runs down the left column, and simulation down the right.

The non-dimensional time-scale is based on freestream velocity at onset, U_{onset} , and the density-layer depth, $\frac{H}{2}$. All the energies have been normalised by the potential energies of the initial state in its tilted configuration, and the minimum energy final state - an unmixed stable stratification. The total available energy (the sum of kinetic and available potential) shown in black rises above 1 at the point where the tank rotation ceases, since this additional energy arises from work done by an externally imparted force. Elsewhere, as one would expect, the total available energy monotonically decreases as energy either leaves the system by dissipation or by contributing to raising of the background potential energy by doing mixing.

Until the onset of Kelvin-Helmholtz instability, the kinetic energy in the system increases as $E_k \sim t^4$. Available potential and kinetic energies are simply being exchanged as the dense layer falls and lighter layer rises. The system is approximately in freefall with a gravitational component $g \sin \theta$ and dissipative mechanisms are insignificant at this stage, so the total available energy is approximately constant. After the onset of Kelvin-Helmholtz instability, kinetic energy is being removed from the freestream to fuel the instability growth and lift dense fluid upwards. This reduces the rate of exchange of potential with kinetic energy, and as dissipative mechanisms become more important, the supply of available energy declines, and in consort the background potential energy increases. Unfortunately, since both Kelvin-Helmholtz billows and the hydraulic jumps give rise to mixing and dissipation, it is difficult with this analysis to attribute increases in background potential energy to either feature individually.

[Patterson *et al.* \(2006\)](#) isolated the Kelvin-Helmholtz instability by examining billows in a small intermediate region of the tank and computing the total, background and available potential energies in the reference frame of the tank. While this analysis neglects the fluxes of energy into and out of the small region, increases in the background potential energy can be exclusively attributed to the mixing caused by the Kelvin-Helmholtz billows.

As discussed in detail in §8.3, the mixing efficiency is a measure of the relative flux of energy lost to heat by dissipation relative to the flux that does mixing and raises the background potential energy. Unfortunately, it is far from straightforward

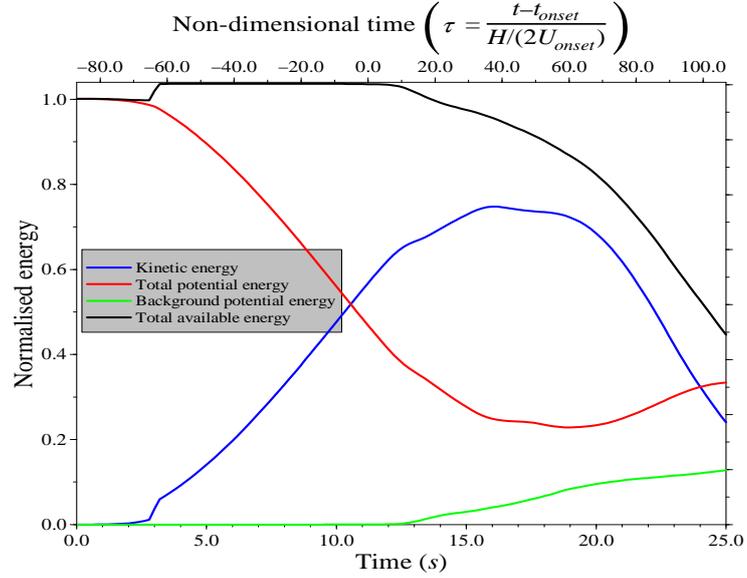


Figure A.11: Time-evolution of total potential, background potential, kinetic and available energy in the tilting tank experiment, measured in the reference frame of gravity once the tank is stationary in its tilted state.

to calculate a Kelvin-Helmholtz mixing efficiency from a tilted tank experiment, because the instantaneous measure requires knowledge of the velocity field and the density field simultaneously, and the aggregate measure includes mixing due to other artefacts such as hydraulic jumps. However, the relative behaviour of the total potential and the background energies gives some insight into when mixing takes place in a billow's evolution. For the purposes of code validation, existing data from [Patterson *et al.* \(2006\)](#) is plotted against the experiment and matching simulation shown earlier in figure [A.10](#). To maintain consistency between studies, the potential energies have been evaluated in the reference frame of the (tilted) tank, and no correction for energy fluxed into and out of the imaging region has been made. The results are shown in figure [A.12](#), where the energies are normalised arbitrarily.

The onset time is consistent between experiment and matching simulation, and the growth rates of the energies are comparable. The delay between growth of total potential energy and the corresponding background potential energy indicates that in both cases mixing is not the dominant process in the early stages of the Kelvin-Helmholtz instability, and develops only later, once the billow has overturned and dense fluid overlies light fluid. A delay in production of background potential

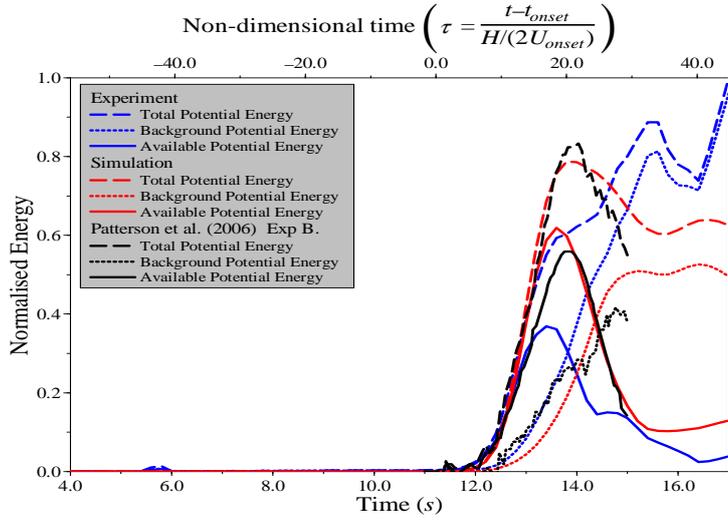


Figure A.12: Time-evolution of total potential, available potential and background potential energy in Kelvin-Helmholtz billows, measured in the reference frame of the tank.

energy is a feature common to both experiment *A* and *B* reported in [Patterson et al. \(2006\)](#). However, despite a qualitative visual similarity between all three experiments, and some degree of quantitative agreement in the energy statistics, many features are inconsistent. Mixing occurs much sooner after onset in experiment *B* than in experiment *A* (indeed the measure of background potential energy in *A* dips erroneously below the zero axis for a time), and the total potential energy increases monotonically in *A*, but peaks in *B*. Given that this variance occurred just in a small corner of the parameter space {density difference, tilt angle, interface thickness} and therefore in a very limited range of Reynolds numbers ($O(300)$), it is not so surprising that the experiment to which the simulation was nominally matched followed the characteristics of experiment *A*, but the simulation followed the characteristics of experiment *B*. Despite the obvious and incompletely explained discrepancies, the simulation statistics sit well within the range of experimental variation, and this offers reasonable confidence that MOBILE does indeed perform well on the Kelvin-Helmholtz problem.

A.3.2 Lock-exchange gravity currents

Gravity currents are a class of flow that occur when a density difference between two fluids under the influence of gravity gives rise to a predominantly horizontal motion along a boundary or interface. In nature these can be witnessed in powder snow avalanches, pyroclastic flows emanating from erupting volcanoes, turbidity currents over continental shelves in the deep oceans, the arrival of sea fog at the coast, and wind associated with the arrival of cold fronts.

In the laboratory it is most common to perform Boussinesq experiments with fresh and salt water, and to use a swiftly removed lock between the two fluids to initiate the gravity current. This approach has been used since O'Brien & Chernov (1934), and particularly in more recent work, *e.g.* Huppert & Simpson (1980); Shin *et al.* (2004). To determine the validity of MOBILE in this context, three gravity current cases have been selected for simulation and compared against our theoretical understanding from the literature. All the cases examined are so-called ‘full depth’ releases of buoyant fluid, where a lock divides the 2.85m tank into a more dense and a less dense segment. By moving the position of the simulated lock release, the eventual depth to which the gravity current settles is altered, and the ratio of these depths is an important parameter in predicting the form of the flow. The depth ratio (or equivalently, the ratio of lock length to tank length) of the sequence of images in figure A.13 is 50%. The images show vertical slices through the current, and the Kelvin-Helmholtz billows that arise from the velocity difference across the interfacial surface can be seen clearly.

It turns out that the details of the front and mixing induced across the interface plays only a secondary role in the development of the flow. The underlying system is well modelled by the one-dimensional shallow water approximation, and so Bernoulli’s principle and conservation of mass can be enforced. In a reference frame moving with the front, the front tip is a stagnation point, and by considering Bernoulli, it is reasonable to expect that the front should have a constant Froude number Fr . Conservation of mass can be written down for late times (once initial transients associated with the lock release are no longer important) if one assumes a self-similar form for the height $h(x, t) = h(t)f(x/L)$ of the gravity current in rela-

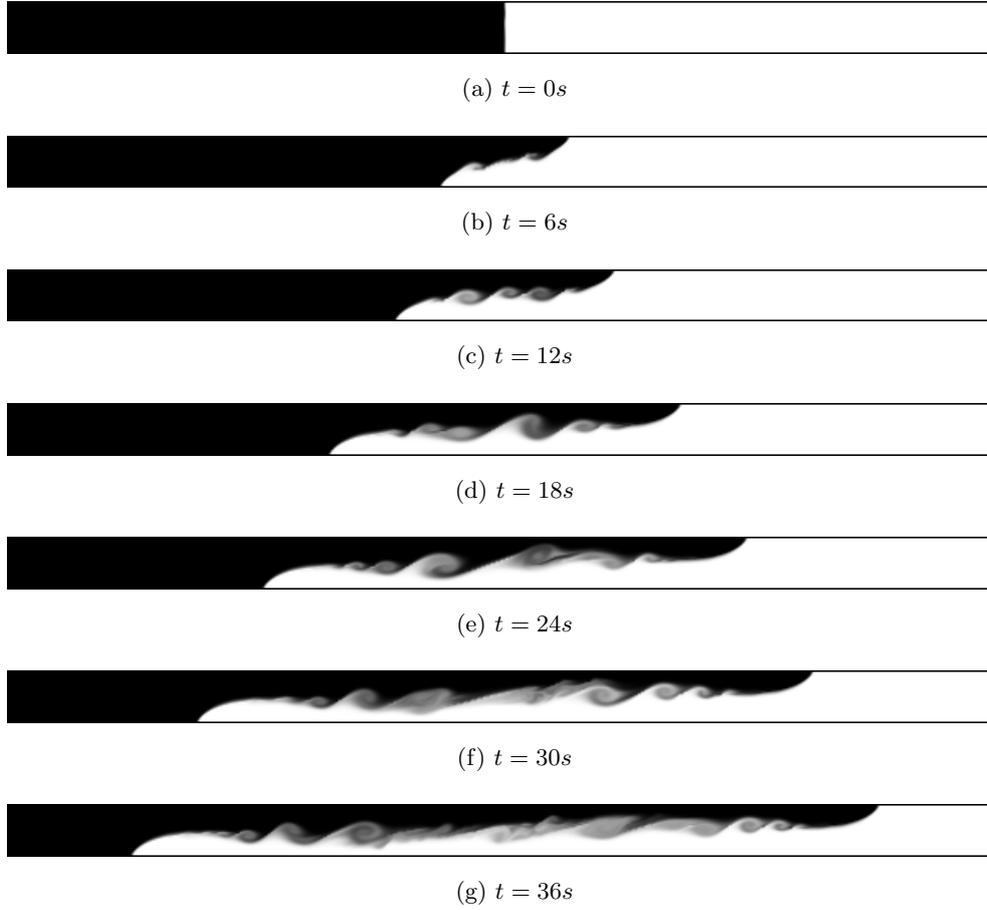


Figure A.13: Boussinesq half-depth gravity current evolution predicted by MOBILE simulation, $\rho_{\text{left}} = 1000 \text{kgm}^{-3}$, $\rho_{\text{right}} = 1003 \text{kgm}^{-3}$.

tion to its length L . Taking both these elements together, we can construct a simple box-model system for the release of a finite volume of relatively buoyant or dense fluid into an unbounded one-dimensional domain,

$$\frac{\partial L}{\partial t} = Fr \sqrt{g'h} = u_{\text{front}} \quad (\text{A.4})$$

$$L(t)h(t) = L(0)h(0) = V_0/b, \quad (\text{A.5})$$

where g' is a reduced gravity, V_0 is the initial volume and b is the tank width.

Eliminating h ,

$$\frac{\partial L}{\partial t} = Fr \left(g' \frac{L_0 h_0}{L} \right)^{\frac{1}{2}}, \quad (\text{A.6})$$

and integrating,

$$\frac{2}{3} L^{\frac{3}{2}} = Fr (g' L_0 h_0)^{\frac{1}{2}} t, \quad (\text{A.7})$$

we recover

$$L \sim t^{\frac{2}{3}}. \quad (\text{A.8})$$

This functional form is found at late time, when initial transients associated with the lock release have become negligible.

Early time behaviour can also be predicted. One might expect a perfect lock release between Boussinesq fluids to evolve with buoyant fluid occupying approximately the upper half of the depth of the enclosing tank near the lock, and the denser fluid to occupy the lower half, *i.e.* $h(0, \epsilon) = \frac{1}{2}H$ where H is the tank height. In a strictly energy conserving flow, this can indeed be shown to happen, though there are several caveats to this detailed in Benjamin (1968). Should the fluids indeed occupy the tank in this manner, then application of Bernoulli's principle to the flow with velocity u_{wake} in the wake just behind the head of one of the currents yields,

$$\frac{1}{2}\rho_b u_{wake}^2 = g(\rho_u - \rho_l)h, \quad (\text{A.9})$$

and by considering continuity for the wake flow u_{wake} compared with the freestream u_{free} well ahead of the front, in the reference frame travelling with one of the fronts we have,

$$\begin{aligned} u_{free} &= \frac{H-h}{H}u_{wake} \\ &= \frac{1}{2}u_{wake} \end{aligned} \quad (\text{A.10})$$

and substituting for u_{wake} from Bernoulli,

$$\begin{aligned} u_{free} &= \frac{1}{2}\sqrt{2g\frac{\rho_u - \rho_l}{\rho_b}h} \\ &= \frac{1}{\sqrt{2}}\sqrt{g'h} \end{aligned} \quad (\text{A.11})$$

so there is a clear prediction for the Froude number of the front under conditions of energy conservation. This can only be valid in the early stages of the flow, since dissipative mechanisms, including mixing, play a role at later time. However, it does provide an unambiguous estimate of the front velocity immediately after the lock release. This appears to contradict the earlier prediction of $L \sim t^{\frac{2}{3}}$, which by definition implies that the initial front travels infinitely quickly, a state that is clearly not physically realisable. A study into the transition between the early- and

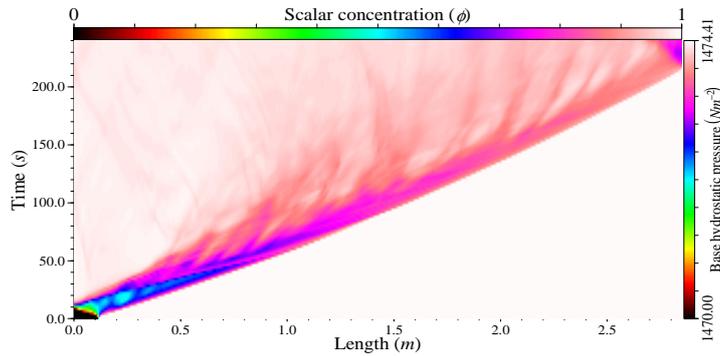


Figure A.14: Time-series image of an asymmetric lock release with 4% depth ratio, showing progress of the front and internal characteristics reflected off the endwall.

late-time behaviours was conducted by [Rottman & Simpson \(1983\)](#) who examined the lock exchange problem by considering the full two-layer shallow water equations. The initial lock release is an expansion fan centred on the lock position, but in cases where an endwall is present on one side of the lock, the wave reflects back into the flow and forms a backward facing hydraulic jump (bore) which travels more quickly than the front itself (which has $Fr = \frac{1}{\sqrt{2}} < 1$). When the jump reaches the front, the front can no longer travel at a constant velocity, and thereafter the late-time behaviour is a better estimate of the evolution. [Figure A.14](#) shows just such a case. The initial expansion and reflected jump are clearly visible in the colour plot - which represents both depth averaged pressure and scalar concentration, and the demarcation between early-time $L \sim t$ and late-time $L \sim t^{\frac{2}{3}}$ front behaviour is very clearly the arrival point of the jump. Wave-like streaks behind the front at later time are Kelvin-Helmholtz billows on the approximately horizontal portion of the interface behind the front, and they tend to move with the mean velocity of upper and lower layers.

The front position for this case has been extracted from the time-series image of [figure A.14](#), and is shown in green in [figure A.15](#) together with the early-time and late-time predictions. A virtual origin correction is necessary for the late-time behaviour since the early-time analysis predicts a finite initial velocity that persists for a substantial period and this is inconsistent with $L \sim t^{\frac{2}{3}}$. Two other depth ratios ($\frac{h(t=\infty)}{H}$) are also shown in this figure, tending towards the limiting cases. No virtual

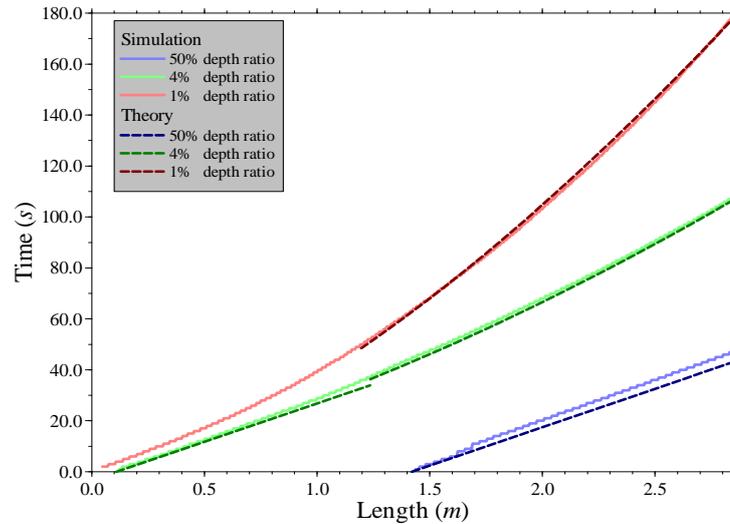


Figure A.15: Front velocity as a function of time for three depth ratios, and compared with appropriate theoretical predictions for early and late time.

origin correction has been made in the 1% depth ratio case, to illustrate how the initial transient behaviour becomes insignificant at late time. The 50% depth ratio case does not reach a state where endwall wave reflections can interfere with the front velocity, so the front remains close to its early-time prediction until it reaches the end of the simulated tank.

From the above results it is clear that MOBILE has accurately captured the features of gravity currents that are well established in the experimental and theoretical literature, not only the predictions made on energy conserving assumptions, such as initial front velocity, the expansion fan at the lock, the late-time $L \sim t^{\frac{2}{3}}$ behaviour, but also the transient, dissipative elements such as the hydraulic jump and interfacial Kelvin-Helmholtz billows.

A.4 Summary

This appendix contains validation and verification studies, which demonstrate that the MOBILE code used elsewhere in this thesis is grid convergent, and correctly predicts fluid behaviour in variable density flows for which there is a well-established knowledge-base in the literature. The single-mode Rayleigh-Taylor instability is a standard test-case for such codes, and the IWPCTM11 Test Problem 1 was per-

formed and compared against an available data-set. To demonstrate capability beyond the Boussinesq regime, the Atwood number was $A = 0.25$. Symmetry breakdown at high resolution was explained, and to provide supplementary evidence of grid convergent behaviour, the high-aspect-ratio Rayleigh-Taylor instability was investigated. This is a less pathological test case, and demonstrates grid convergence very clearly. A study of the stratified Kelvin-Helmholtz instability was performed to ascertain that MOBILE can perform well in other flows and configurations. To ensure that comparison between experiment and simulation was as fair as possible, a new experiment was conducted, and the initial conditions and interfacial perturbation carefully measured and implemented in the simulation. Qualitatively the two match reasonably well, but some inconsistencies remain. Quantitative comparison with published experiments using the same apparatus provided a benchmark. The variability between experiments is large, but the simulation falls well within the acceptable range, when considering instability onset time, total potential energy growth, and background potential energy growth. Lock-exchange gravity currents were also studied, and the front position compared very successfully with well-understood theoretical predictions for early- and late-time behaviour. The results taken together suggest strongly that MOBILE is an effective and reliable tool for modelling a wide variety of variable density, non-Boussinesq and rotating flows.

Bibliography

- ANDREWS, M. J. 1995 Accurate computation of convective transport in transient two-phase flow. *Int. J. Num. Meth. Fluids* **21-3**, 205–222.
- ANDREWS, M. J., KRAFT, W. N. & MUESCHKE, N. J. 2007 Progress with molecular mixing measurements and high Atwood number experiments at Texas A&M university. In *60th Annual Meeting of the Division of Fluid Dynamics*. American Physical Society, Salt Lake City, US.
- BENJAMIN, T. B. 1968 Gravity currents and related phenomena. *J. Fluid Mech.* **31**, 209–248.
- BICKLEY, W. G. 1937 The plane jet. *Phil. Mag.* **7:23**, 727–731.
- CABOT, W. H. & COOK, A. W. 2006 Reynolds number effects on Rayleigh-Taylor instability with implications for type 1a supernovae. *Nature Physics* **2**, 562.
- COLLELA, P. 1985 A direct Eulerian MUSCL scheme for gas dynamics. *SIAM J. on Computing* **6**, 104–117.
- COOK, A. W., CABOT, W. & MILLER, P. L. 2004 The mixing transition in Rayleigh-Taylor instability. *J. Fluid Mech* **511**, 333.
- DAHM, W. J. A. & DIMOTAKIS, P. E. 1987 Measures of entrainment and mixing in turbulent jets. *AIAA J.* **25:9**, 1216–1223.
- DALZIEL, S. B. 1993 Rayleigh-Taylor instability: Experiments with image analysis. *Dyn. Atmos. Oceans* **20**, 127–153.
- DALZIEL, S. B. 2001 Toy models for Rayleigh-Taylor instability. In *Proceedings of IWPCTM8*. Lawrence Livermore National Laboratory.

- DALZIEL, S. B., LINDEN, P. F. & YOUNGS, D. L. 1999 Self-similarity and internal structure of turbulence induced by Rayleigh-Taylor instability. *J. Fluid Mech.* **399**, 1–48.
- DALZIEL, S. B., PATTERSON, M. D., CAULFIELD, C. P. & COOMARASWAMY, I. A. 2008 Mixing efficiency in high-aspect-ratio Rayleigh-Taylor experiments. *Phys. Fluids* **2008:6**, 14.
- DAVIES, R. M. & TAYLOR, G. I. 1950 The mechanics of large bubbles rising through liquids in tubes. *Proc. Roy. Soc. A* **200**, 375.
- DEBACQ, M., FANGUET, V., HULIN, J.-P. & SALIN, D. 2001 Self-similar concentrations profiles in buoyant mixing of miscible fluids in a vertical tube. *Phys. Fluids* **13**, 3097.
- DEBACQ, M., HULIN, J.-P., SALIN, D., PERRIN, B. & HINCH, E. J. 2003 Buoyant mixing of miscible fluids of varying viscosities in vertical tubes. *Phys. Fluids* **15**, 3846.
- DIMONTE, G. & SCHNEIDER, M. 1996 Turbulent Rayleigh-Taylor instability experiments with variable acceleration. *Phys. Rev. E* **54**, 3740–3743.
- DIMONTE, G., YOUNGS, D. L., DIMITS, A., WEBER, S., MARINAK, M., WUNCH, S., GARASI, C., ROBINSON, A., ANDREWS, M. J., RAMAPRABHU, P., CALDER, A. C., FRYXELL, B., BIELLO, J., DURSI, L., MACNEICE, P., OLSON, K., RICKER, P., ROSNER, R., TIMMES, F., TUFO, H., YOUNG, Y. N. & ZINGALE, M. 2004 A comparative study of the turbulent Rayleigh-Taylor instability using high-resolution three-dimensional numerical simulations: The Alpha-group collaboration. *Phys. Fluids* **16-5**.
- DUFF, R. E., HARLOW, F. H. & HIRT, C. W. 1962 Effects of diffusion on interface instability between gasses. *Phys. Fluids* **5**, 417.
- EMMONS, H. W., CHANG, C. T. & WATSON, B. C. 1959 Taylor instability of finite surface waves. *Proc. Roy. Soc. A* **7-2**, 177.

- FERNANDO, H. J. S. 1991 Turbulent mixing in stratified fluids. *Annu. Rev. Fluid Mech.* **23**, 455–493.
- FORTUIN, J. M. H. 1960 Theory and application of two supplementary methods for constructing density gradient columns. *J. Polymer. Sci.* **44**, 505–515.
- GEDDES, C. D. 2001 Optical halide sensing using fluorescence quenching: Theory, simulations and applications - a review. *Meas. Sci. Tech.* **12**, 53–88.
- GILES, M. B. 1990 Non-reflecting boundary conditions for Euler equation calculation. *AIAA J.* **28**, 2050–2058.
- GODUNOV, S. K. 1959 A difference scheme for numerical solution of discontinuous solution of hydrodynamic equations. *Math. Sbornik* **47**, 271–306.
- HAHN, M. & DRIKAKIS, D. 2005 Large eddy simulation of compressible turbulence using high-resolution methods. *Int. J. Num. Meth. Fluids* **47**, 971–977.
- HIGGINSON, R. C., DALZIEL, S. B. & LINDEN, P. F. 2003 The drag on a vertically moving grid of bars in a linearly stratified fluid. *Expt. Fluids* **34**, 678–686.
- HOLFORD, J. M., DALZIEL, S. B. & YOUNGS, D. L. 2001 Rayleigh-Taylor instability at a tilted interface in laboratory experiments and compressible numerical simulations. In *Proceedings of IWPCTM8*. LLNL, USA.
- HOLFORD, J. M., DALZIEL, S. B. & YOUNGS, D. L. 2003 Rayleigh-Taylor instability at a tilted interface in laboratory experiments and numerical simulations. *Lasers & Particle Beams* **21**, 491.
- HU, H. & KOOCHESFAHANI, M. M. 2002 A novel method for instantaneous, quantitative measurement of molecular mixing in gaseous flows. *Exp. Fluids* **31**, 202–209.
- HUPPERT, H. E. & SIMPSON, J. E. 1980 The slumping of gravity currents. *J. Fluid Mech.* **99**, 785–799.
- HUPPERT, H. E. & TURNER, J. S. 1981 Double diffusive convection. *J. Fluid Mech.* **106**, 299–329.

- INOAMOV, N. A., OPARIN, A. M., DEM'YANOV, A. YU., DEMBITSKIOE, L. N. & KHOKHLOV, V. A. 2001 On stochastic mixing caused by the Rayleigh-Taylor instability. *J. Exp. Theor. Phys.* **92**, 714.
- JACOBS, J. W. & DALZIEL, S. B. 2005 Rayleigh-Taylor instability in complex stratifications. *J. Fluid Mech.* **542**, 251–279.
- KOOCHESFAHANI, M. M. & DIMOTAKIS, P. E. 1985 Mixing and chemical reactions in a turbulent liquid mixing layer. *J. Fluid Mech.* **170**, 83–112.
- LANE-SERFF, G. F. 1989 Heat flow and air movement in buildings. PhD thesis, DAMTP, University of Cambridge, UK.
- LAWRIE, A. G. W. & DALZIEL, S. B. 2006*a* Rayleigh-Taylor driven mixing in a multiply stratified environment. In *Sixth International Symposium on Stratified Flows*. University of Western Australia, Australia.
- LAWRIE, A. G. W. & DALZIEL, S. B. 2006*b* Rayleigh-Taylor driven mixing in a stratified environment. In *Proceedings of IWPCMT10* (ed. M. Legrand). CEA, France.
- LAYZER, D. 1955 On the instability of superposed fluids in a gravitational field. *Astroph. J.* **122-1**, 1.
- LESTER, K. S. & CLEMENS, N. T. 2003 The structure of fine-scale scalar mixing in gas-phase planar turbulent jets. *J. Fluid Mech.* **488**, 1–29.
- LEWIS, G. D. J. 1950 The instability of liquid surfaces when accelerated in a direction perpendicular to their planes. *Proc. Roy. Soc. A* **202**, 81.
- LINDEN, P. F. 1977 The flow of a stratified fluid in a rotating annulus. *J. Fluid Mech.* **79**, 435–447.
- LINDEN, P. F., REDONDO, J. M. & YOUNGS, D. L. 1994 Molecular mixing in Rayleigh-Taylor instability. *J. Fluid Mech.* **265**, 97–124.
- MARGOLIN, L. G., RIDER, W. J. & GRINSTEIN, F. F. 2006 Modeling turbulent flow with implicit LES. *J. Turbulence* **7**, 1–27.

- MAXWORTHY, T. 1987 The nonlinear growth of a gravitationally unstable interface in a Hele-Shaw cell. *J. Fluid Mech.* **177**, 207–232.
- MENEVEAU, C. & KATZ, J. 2000 Scale invariance in Large Eddy Simulation. *Annu. Rev. Fluid Mech.* **32**, 1–32.
- MIKAELIAN, K. O. 2008 Limitations and failures of the Layzer model for hydrodynamic instabilities. *Phys. Rev. E* **78**, 015303(R).
- NEVMERZHITSKY, N. K., MESHKOV, E. E., ABND I. G. ZHIDOV, A. G. IOOLEV, PYLEV, I. G. & SOKOLOV, S. S. 1994 Wave processes effect on the dynamics of turbulent mixing at liquid layer surface accelerated by a compressed gas. In *Proceedings of IWPCMT4* (ed. Youngs Linden & Dalziel). University of Cambridge, Cambridge, UK.
- O'BRIEN, M. P. & CHERNO, J. 1934 Model law for motion of salt water through fresh. *Trans. ASCE* **99**, 576–594.
- OFER, D., ALON, U., MCCRORY, D. SHVARTS R. L & VERDON, C. P. 1996 Modal model for the nonlinear multimode Rayleigh-Taylor instability. *Phys. Plasmas* **3-8**, 3073.
- PATTERSON, M. D., CAULFIELD, C. P., MCELWAINE, J. N. & DALZIEL, S. B. 2006 Time-dependent mixing in stratified Kelvin-Helmholtz billows. *Geophys. Res. Lett.* **33**, L15608.
- PRANDTL, L. Z. 1925 Bericht uber untersuchungen zur ausgebildeten turbulenz. *Z. angew. Math, Mech.* **5**, 136–139.
- PULLIN, D. I. 2000 A vortex-based model for the subgrid flux of a passive scalar. *Phys. Fluids* **12**, 2311.
- RAMAPRABHU, P., DIMONTE, G., YOUNG, Y.-N., CALDER, A. C. & FRYXELL, B. 2006 Limits of the potential flow approach to the single-mode Rayleigh-Taylor problem. *Phys. Rev. E* **76**, 066308.
- RAYLEIGH, LORD 1883 Investigation of the character of the equilibrium of an incompressible heavy fluid of variable density. *Proc. Lond. Math. Soc.* **XIV**, 70–177.

- READ, K. I. 1984 Experimental investigation of turbulent mixing by Rayleigh-Taylor instability. *Physica D* **12**, 45–58.
- RIDER, W. 2007 Effective subgrid modeling from the ILES simulation of compressible turbulence. *J. Fluids Eng.* **129**, 1493.
- RIKANATI, A., ORON, D., ALON, U. & SHVARTS, D. 2000 Statistical mechanics merger model for hydrodynamic instabilities. *Astroph. J. Suppl. Ser.* **127**, 451.
- ROBEY, H. F., KANE, J. O., REMINGTON, B. A., DRAKE, R. P., HURRICANE, O. A., LOUIS, H., WALLACE, R. J., KNAUER, J., KEITER, P., ARNETT, D. & RYUTOV, D. D. 2001 An experimental testbed for the study of hydrodynamic issues in supernovae. *Phys. Plasmas* **8-5**, 2446.
- ROTTMAN, J. W. & SIMPSON, J. E. 1983 Gravity currents produced by instantaneous releases of a heavy fluid in a rectangular channel. *J. Fluid Mech.* **135**, 95–110.
- SCASE, M. M. & DALZIEL, S. B. 2004 Internal wave fields and drag generated by a translating body in a stratified fluid. *J. Fluid Mech* **498**, 289–313.
- SCASE, M. M. & DALZIEL, S. B. 2006 Internal wave fields generated by a translating body in a stratified fluid: An experimental comparison. *J. Fluid Mech* **564**, 305–331.
- SCHMITT, R. W. 1995 The salt finger experiments of Jevons (1857) and Rayleigh (1880). *J. Phys. Ocean.* **25**, 8–17.
- SHIN, J. O., DALZIEL, S. B. & LINDEN, P. F. 2004 Gravity currents produced by lock exchange. *J. Fluid Mech.* **521**, 1–34.
- SMAGORINSKY, J. 1963 General circulation experiments with the primitive equations. *Month. Weather Rev.* **91**, 99–164.
- SMOLARKIEWICZ, P. K., MARGOLIN, L. G. & WYSZOGRODZKI, A. A. 2007 Implicit large eddy simulation in meteorology: from boundary layers to climate. *J. Fluid Eng.* **129**, 1533.

- SNIDER, D. M. & ANDREWS, M. J. 1994 Rayleigh-Taylor and shear driven mixing with an unstable thermal stratification. *Phys. Fluids* **6**, 3324–3334.
- STRANG, G. 1968 On the construction and comparison of difference schemes. *SIAM J. on Num. Anal.* **5:3**, 506–517.
- TAYLOR, G. I. 1950 The instability of liquid surfaces when accelerated in a direction perpendicular to their planes. *Proc. R. Soc. Lond.* **201**, 192–196.
- THORNBER, B., DRIKAKIS, D. & YOUNGS, D. 2008 Large eddy simulation of multi-component compressible turbulent flows using high-resolution methods. *Computers and Fluids* **37**, 867–876.
- THORPE, S. A. 1985 Laboratory observations of secondary structures in Kelvin-Helmholtz billows and consequences for ocean mixing. *Geophys. Astrophys. Fluid Dyn.* **34**, 175–199.
- TURNER, J. S. 1968 The influence of molecular diffusivity on turbulent entrainment across a density interface. *J. Fluid Mech.* **33**, 639–656.
- TURNER, J. S. 1974 Double-diffusive phenomena. *Ann. Rev. Fluid Mech.* **6**, 37–54.
- V VLIET, E., V BERGEN, S. M., DERKSEN, J. J., POERTELA, L. M. & V D AKKER, H. E. A. 2004 Time-resolved, 3d, laser-induced fluorescence measurements of fine-structure passive scalar mixing in a tubular reactor. *Expt. Fluids* **37**, 1–21.
- WADDELL, J. T., NIEDERHAUS, C. E. & JACOBS, J. W. 2001 Experimental study of Rayleigh-Taylor instability: Low Atwood number liquid systems with single-mode initial perturbations. *Phys. Fluids* **13**, 1263–1273.
- WILKINSON, J. P. & JACOBS, J. W. 2007 Experimental study of the single-mode three-dimensional Rayleigh-Taylor instability. *Phys. Fluids* **19**, 124102.
- WILSON, P. N. & ANDREWS, M. J. 2002 Spectral measurements of Rayleigh-Taylor mixing at small Atwood number. *Phys. Fluids* **14**, 938–945.
- WINTERS, K. B., LOMBARD, P. N., RILEY, J. J. & D’ASARO, E. A. 1995 Available potential energy and mixing in density-stratified fluids. *J. Fluid Mech.* **289**, 115–128.

YOUNGS, D. L. 1984*a* Accurate numerical methods for volume-fraction transport in one-dimensional multifluid flow. *Tech. Rep.*. AWRE report.

YOUNGS, D. L. 1984*b* Numerical simulation of turbulent mixing by Rayleigh-Taylor instability. *Physica D* **12**, 32–44.

ZUFIRIA, J. 1988 Bubble competition in Rayleigh-Taylor instability. *Phys. Fluids* **31-3**, 446.