

THE UNIVERSITY OF CHICAGO

DATA DRIVEN MODELING OF THE LOW-ATWOOD SINGLE-MODE  
RAYLEIGH-TAYLOR INSTABILITY

A DISSERTATION SUBMITTED TO  
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES  
IN CANDIDACY FOR THE DEGREE OF  
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY  
MAXWELL HUTCHINSON

CHICAGO, ILLINOIS

DECEMBER 2016

Copyright © 2016 by Maxwell Hutchinson

All Rights Reserved

Dedicated to Tracey Ziev

## TABLE OF CONTENTS

LIST OF FIGURES . . . . .	vii
LIST OF TABLES . . . . .	xi
ACKNOWLEDGMENTS . . . . .	xii
ABSTRACT . . . . .	xiii
1 INTRODUCTION . . . . .	1
1.1 Formal definition . . . . .	1
1.2 Instances and motivation . . . . .	2
1.3 Terminology . . . . .	3
1.4 Note on the structure of the dissertation . . . . .	9
2 BACKGROUND . . . . .	10
2.1 Linear and weakly non-linear models . . . . .	10
2.1.1 Lord Rayleigh's linear model . . . . .	10
2.1.2 Viscous and diffusive linear models . . . . .	11
2.1.3 Weakly nonlinear expansions . . . . .	12
2.2 Potential flow models . . . . .	13
2.2.1 Layzer's unit Atwood model . . . . .	13
2.2.2 Goncharov's high Atwood model . . . . .	14
2.3 Buoyancy-drag models . . . . .	15
2.3.1 Bubble model of Davies and Taylor . . . . .	16
2.3.2 Tube model of Dimonte and Schneider . . . . .	16
2.3.3 Self-similar model of Oron . . . . .	17
2.4 Problems with single mode Rayleigh-Taylor modeling . . . . .	17
2.4.1 Pressure in the single-mode RTI . . . . .	18
2.4.2 Departure from potential flow . . . . .	18
2.4.3 Historical inconsistency in buoyancy-drag models . . . . .	19
2.5 Related work aimed at modeling stagnation and re-acceleration . . . . .	19
2.5.1 Vortex ring correction of Ramaprabhu . . . . .	20
2.6 Vorticity and viscosity in potential flow . . . . .	20
3 DATA-DRIVEN MODELING OF THE LOW-ATWOOD SINGLE-MODE RAYLEIGH-TAYLOR INSTABILITY . . . . .	21
3.1 Abstract . . . . .	21
3.2 Introduction . . . . .	21
3.3 Simple model . . . . .	24
3.3.1 Dynamics . . . . .	24
3.3.2 Mixing . . . . .	26
3.3.3 Coefficient constraints . . . . .	28

3.3.4	Coefficient estimation . . . . .	30
3.4	Numerical experiments . . . . .	33
3.4.1	Observables . . . . .	35
3.5	Growth stages through late times . . . . .	37
3.5.1	Exponential growth . . . . .	39
3.5.2	Saturation regime . . . . .	40
3.5.3	Viscous regime . . . . .	41
3.5.4	Diffusive regime . . . . .	41
3.5.5	Stagnation and re-acceleration . . . . .	45
3.6	Model fit coefficients . . . . .	46
3.6.1	Fitting . . . . .	46
3.6.2	Scope of the models . . . . .	49
3.6.3	Accuracy of the models . . . . .	50
3.6.4	Fit coefficients . . . . .	52
3.7	Conclusions . . . . .	60
3.8	Acknowledgements . . . . .	63
4	DIRECT NUMERICAL SIMULATION OF SINGLE MODE THREE-DIMENSIONAL RAYLEIGH-TAYLOR EXPERIMENTS . . . . .	64
4.1	Abstract . . . . .	64
4.2	Introduction . . . . .	65
4.3	Numerical methods . . . . .	67
4.3.1	Spectral element formulation . . . . .	67
4.3.2	Simulation setup . . . . .	69
4.3.3	Post-processing . . . . .	71
4.4	Results . . . . .	71
4.4.1	Validation . . . . .	71
4.4.2	Extension . . . . .	75
4.5	Conclusions . . . . .	82
4.6	Acknowledgements . . . . .	83
5	EFFICIENCY OF HIGH ORDER SPECTRAL ELEMENT METHODS ON PETASCALE ARCHITECTURES . . . . .	84
5.1	Abstract . . . . .	84
5.2	Introduction . . . . .	85
5.2.1	Outline . . . . .	86
5.3	Nek's Computational Core . . . . .	87
5.3.1	Governing equations and time-splitting . . . . .	87
5.3.2	Spectral element method . . . . .	88
5.3.3	Computational profile . . . . .	89
5.3.4	Order-dependent kernels . . . . .	90
5.4	Kernel Analysis and Optimization . . . . .	91
5.4.1	Small Matrix Multiplications . . . . .	91
5.4.2	Enhancing Element Update Performance by Streaming Stores . . . . .	94

5.4.3	Discussion of Performance Reproducers . . . . .	94
5.5	Scenarios and Performance . . . . .	97
5.5.1	Architectures . . . . .	97
5.5.2	Single mode Rayleigh-Taylor instability . . . . .	98
5.5.3	Time to accuracy . . . . .	103
5.5.4	Whole application performance . . . . .	105
5.6	Conclusion . . . . .	106
6	CONCLUSIONS . . . . .	109
6.1	New model for low-Atwood single mode . . . . .	109
6.2	Validation of DNS . . . . .	110
6.3	Convergence and performance . . . . .	112
6.4	Flow regimes . . . . .	113
6.5	Model coefficients . . . . .	113
6.6	Open problems . . . . .	114
	REFERENCES . . . . .	116

## LIST OF FIGURES

3.1 Slices of the scalar and vertical component of the velocity at early times and high Grashof number. The arrows indicate the dependence of the model terms on different span-wise length scales, and are identical in both figures. $C_1$ is related to the maximum cross sectional diameter of the bubble in the velocity field. $C_2$ is related to the nominal side-wall diameter of the bubble in the velocity field. $C_5$ is related to the nominal side-wall diameter of the bubble in the scalar field. $\delta$ is related to the interface thickness in the scalar field. The slice is in the plane $x = y$ , which passes through only bubble centers. . . . .	31
3.2 Comparison of height metrics at $\text{Gr} = 4.8 \times 10^4$ and $\text{Sc} = 1$ . . . . .	36
3.3 Bubble Froude number vs. nondimensional bubble height for $\text{Ra} = 10^{5.75}$ , $\text{Sc} = 4$ , simulation vs. model with successive terms enabled: first $C_4, C_6$ and $C_8$ , then $C_3$ , then $C_1$ , then $C_2$ , and finally $C_5$ and $C_7$ . The dashed vertical lines divide the trajectory into three regimes: linear growth for $H/\lambda < 0.05$ , saturation until $H/\lambda \approx 8$ , and viscosity beyond that. The stagnation and re-acceleration transients are seen beginning at $H/\lambda = 0.5$ and ending by $H/\lambda = 1.5$ . . . . .	38
3.4 Bubble Froude number vs. nondimensional bubble height for $\text{Ra} = 10^{4.5}$ , $\text{Sc} = 8$ , simulation vs. model with successive terms enabled: first $C_4, C_6$ and $C_8$ , then $C_3$ , then $C_1$ , then $C_2$ , and finally $C_5$ and $C_7$ . Dashed vertical lines divide the trajectory into four regimes: linear growth for $H/\lambda < 0.05$ , saturation until $H/\lambda \approx 0.1$ , viscosity until $H/\lambda \approx 1.3$ , and diffusion beyond that. The stagnation and re-acceleration transients, seen in Figure 3.3, are suppressed by the viscous regime, which onsets before the transient begins at $H/\lambda = 0.5$ . . . . .	38
3.5 Penetration depth, nondimensionalized, vs. the Rayleigh and Schmidt numbers. The dashed line separates completed from incomplete trajectories, which are clipped at $h/\lambda \approx 23$ . . . . .	42
3.6 Penetration depth, nondimensionalized, vs. the Rayleigh. The dashed line is a best fit line $h/\lambda = 2.41 \times 10^{-4}\text{Ra} - 0.36$ . . . . .	43
3.7 Example fit of the bubble height, top, and mixed volume, bottom, at $\text{Ra} =$ and $\text{Sc} = 8$ . The dashed line in the center plot is the stagnation velocity from potential flow models. The Froude number, which is a nondimensional velocity, is not directly fit but shows agreement between the model and simulation. . . .	47
3.8 Relative model errors versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain. . . . .	51
3.9 Best fit for $C_1$ versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain. . . . .	52

3.10 Best fit for $C_2$ versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain. . . . .	54
3.11 Best fit for $C_3$ versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain. . . . .	56
3.12 Best fit for $C_5$ versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain. . . . .	57
3.13 Best fit for $C_7$ versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain. . . . .	59
4.1 Initial condition, $\phi(x, y, 0, 0)$ , for 4.5 mode simulation. The bubbles are logically separated by the solid grid and each unique bubble is labeled. . . . .	68
4.2 Froude number vs. height, nondimensionalized by the wavelength. Lines are the derivative of cubic splines through simulation outputs. Points are from experiment via direct measurement of the bubble velocity [31]. The dotted horizontal line is positioned at Goncharov's theoretical value of $\pi^{-1/2}$ [19]. . . . .	73
4.3 Bubble height vs. time, nondimensionalized by the wavelength and linear growth rate, respectively. Lines are cubic splines through simulation outputs. Points are from experiment [31]. The points are shifted in time to minimize the square deviation from the spline summed over the plotted points. . . . .	74
4.4 Froude number as a function of height, nondimensionalized by the wavelength, by bubble in the 4.5 mode simulation. Solid line is from the height defined as the maximum taken over the entire span-wise domain. Dotted line is the periodic reference calculation. . . . .	76
4.5 Ratio of wall-bounded bubble height to periodic bubble height in the 4.5 mode simulation. . . . .	76
4.6 Spatial distribution of bubble heights at three times, with red denoting greater bubble heights and blue denoting lesser bubble heights. The first time is when the ratio of the wall-bounded to periodic bubble height is at a minimum, which is during the transition from linear to non-linear growth; a long wavelength mode is seen across the diagonal. The second time is taken from the stagnation phase; the long wavelength mode is still present. The third time is taken from the end of the simulation; the central bubbles are being squeeze significantly ahead of the bubbles near the wall. . . . .	77

4.7	Bubble velocity and bubble height vs. time, nondimensionalized by the wavelength and linear growth rate, for 4.5 mode simulations and experiment. Lines are from simulation output, one case with the same vertical extent as the simulation and in the other with twice that vertical extent. Points are from experiment via direct measurement of the bubble velocity and bubble height. The dotted horizontal line is positioned at Goncharov's theoretical value of $\pi^{-1/2}$ [19]. The solid vertical line marks the greatest bubble height reached in any of the experiments by Wilkinson and Jacobs [62].	79
4.8	Froude number as a function of height, nondimensionalized by the wavelength, by bubble in the 4.5 mode simulation with extended vertical extent. Solid line is from the height defined as the maximum taken over the entire span-wise domain. Dotted line is the periodic reference calculation. The dotted horizontal line is positioned at Goncharov's theoretical value of $\pi^{-1/2}$ [19]. The solid vertical line marks the greatest bubble height reached in any of the experiments by Wilkinson and Jacobs [62].	80
4.9	Scalar field in the horizontal mid-plane for 2.5, 3.5, and 4.5 mode simulations.	81
4.10	Secondary flow in the horizontal mid-plane. Background color is the vertical component of the vorticity. Contours are lines of constant pressure.	81
4.11	Dynamic pressure in the horizontal mid-plane.	81
5.1	Performance of the Helmholtz reproducer running on a single node of Shaheen for different implementation of small matrix multiplications. NTS denotes the usage of the non-temporal store optimized module.	96
5.2	Performance of the Helmholtz operator reproducer running on a single node of Mira for different implementation of small matrix multiplications.	97
5.3	Performance of the basis transformation reproducers using different implementation for the small matrix multiplications. NTS denotes the usage of the aforementioned non-temporal store optimized module. The top plot shows the diagonalization in the local Poisson operator, the middle one the prolongation and the bottom one the restriction case.	98
5.4	Scalar, vertical component of the velocity, vorticity component out of the plane, and pressure fields at end of simulation. The color scales are dimensionless, linear, and centered at the mean value over the image.	101
5.5	Weak scaling of bandwidth on Shaheen and Mira. In (a), Circles and crosses indicate memory bandwidth per core on Shaheen and Mira, respectively, vs. the problem size labeled by element size. In (b), the ratio of the bandwidths are shown vs. element size for common discretizations. The solid line indicates ratio of STREAM memory bandwidth.	102
5.6	Error with respect to bubble height, Equation 5.15, vs. the computational cost, in processor hours, on Shaheen (a) and Mira (b). Points are labeled as $(p+1, e)$ pairs, where $p$ is the order, $p+1$ is the element size, and $e$ is the number of elements in one dimension. More runs are present on Mira due to the smaller BGQ nodes evenly dividing more problem sizes.	104

5.7 Strong scaling (left) and weak scaling (right) on Shaheen on up to 131,072 cores (2/3) of the full 7 PFLOPS machine using an element size of 32. To avoid log plots we show per-core performance. . . . . 105

## LIST OF TABLES

3.1	Simulation conditions, fit coefficients, and relative errors. . . . .	45
4.1	Parameters of simulations. The aspect ratio is defined with respect to the quiescent amplitude, $a_0$ , from Equation 4.8. The last row is the simulation that extends the tank by a factor of two in the vertical direction. . . . .	69
4.2	Growth rate: linear theory vs. simulation. Theoretical values are calculated as in Equation 4.1. Simulation values are calculated as in Equation 4.10. The aspect ratio is shown for the second sample, $h_1/\lambda$ . Note the difference in Schmidt number between the two 4.5 mode cases. . . . .	72

## ACKNOWLEDGMENTS

I would like to thank Robert Rosner for allowing and supporting me to explore a variety of paths through this thesis and my doctoral research more broadly. I thank Aleksandr Obabko and Paul Fischer for supporting Nek5000 and working with me through the numerics. I also thank Oana Marin and Michel Schanen for their excitement and Elia Merzari and Ron Rahaman for their support. I thank Alexander Heinecke and Scott Parker for their help tuning the code to run on supercomputers. I thank Elizabeth Hicks for helpful comments and advice.

For computer time, this research partially used the resources of the Supercomputing Laboratory at King Abdullah University of Science & Technology (KAUST) in Thuwal, Saudi Arabia. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. I acknowledge support from the Department of Energy Computational Science graduate fellowship.

## ABSTRACT

The Rayleigh-Taylor instability is one of the most common and well studied phenomena in fluid dynamics. Despite research dating to the late 19th century, the non-linear dynamics of the interfacial instability are still not fully understood, particularly in the case when the two fluids have nearly the same density. It was recently demonstrated in this, the low-Atwood regime, that the idealized single-mode problem departs from established potential flow models in the form of a re-acceleration beyond the predicted terminal interface velocity. This thesis is an attempt to model that re-acceleration and, more broadly, the late time dynamics of the single-mode low-Atwood Rayleigh-Taylor instability.

The approach taken here is based on buoyancy-drag models, which express a force balance between buoyancy and parasitic drag. The dynamical buoyancy-drag model is supplemented with a mixing model that dilutes the buoyant force over time. These models are written deliberately generally, with 8 unique coefficients. Three of these coefficients are solved for by equating the early time behavior with that of well established linear theories. The remaining 5 coefficients are estimated by relating them to drag coefficients, friction factors, and geometric ratios in the interface shape.

To evaluate the model and compute the 5 unknown coefficients more precisely, a set of direct numerical simulations are performed over the relevant parameter space. These simulations are first validated against experimental data. Then they are shown to converge and their resolutions are chosen such as to minimize computational cost given the accuracy scale of the model. The 5 coefficients are fit to the resulting data set, and the model achieves better than 2% error in the bubble height and 4% error in the volume of mixed fluid. Three coefficients are nominally independent of the parameterization of the problem, while two are shown to vary with the Rayleigh number and the diffusivity.

# CHAPTER 1

## INTRODUCTION

### 1.1 Formal definition

The Rayleigh-Taylor instability occurs when the pressure and density gradients are in opposition [50]:

$$(\nabla P)(\nabla \rho) < 0 \quad (1.1)$$

The canonical example of the Rayleigh-Taylor instability is a heavy fluid superposed over a lighter one in a gravitational field. The standard terminology is based on this case.

Consider a horizontal planar interface at  $z = 0$  between two fluids of densities  $\rho_h > \rho_l$ , corresponding to the density of the heavier and lighter fluid, respectively. The denser fluid is at  $z > 0$  and the lighter at  $z < 0$ , and a gravitational acceleration  $-g\hat{z}$ . In equilibrium, the pressure must balance the gravitational force:

$$P(x, y, z) = \begin{cases} -\rho_h g z + C & \text{if } z > 0 \\ -\rho_l g z + C & \text{else} \end{cases} \quad (1.2)$$

A small perturbation is introduced, moving some heavy fluid below the interface and some light fluid above it but preserving the horizontal stratification of the pressure. The forcing on the heavier fluid below the interface is then:

$$\sum F = -\nabla P + F_g = \rho_l g - \rho_h g < 0, \quad (1.3)$$

so the heavier fluid is forced downwards through the lighter fluid, forming structures called spikes. Conversely, the forcing on the lighter fluid above the interface is:

$$\sum F = -\nabla P + F_g = \rho_l g - \rho_h g > 0, \quad (1.4)$$

so the lighter fluid is forced upwards through the heavier fluid, forming structures called bubbles. Perturbations in the interface grow, so the configuration is unstable.

## 1.2 Instances and motivation

The Rayleigh-Taylor instability is present in both natural and constructed systems at many scales. This section describes a few such systems of particular importance.

**Type Ia Supernovae** In Type Ia supernovae (SNe Ia), a white dwarf spontaneously ignites as its mass crosses the Chandrasekhar mass due to accretion from another source. The thermonuclear flame originates near the center of the star and burns outward [43]. The hot ash trailing the flame is lighter than the fuel, creating a Rayleigh-Taylor (RT) unstable flame front [63]. The primary RT instability and secondary Kelvin-Helmholtz instabilities wrinkle the flame front, enhancing mixing, burn rate, and flame speed [22]. These properties affect the rate of energy release and ejecta velocity, which can be observed. The use of SNe Ia as standard candles underlines interest in their description.

**Salt fingers and the thermohaline staircase** Salt fingers are an instance of the Rayleigh-Taylor instability set up by a difference in the diffusivity of two mass carriers [54, 34]. Consider a fluid in which the density perturbation is a linear combination of two fields. Let one of the fields have a stabilizing gradient and the other a destabilizing one. The combined system is unstable if the destabilizing field is less diffusive than the stabilizing one. A parcel of fluid perturbed from its equilibrium height will equilibrate with the stabilizing field before the destabilizing one, resulting in a buoyant force that pushes the parcel further from vertical equilibrium.

In the oceanic case, the stabilizing field is temperature and the destabilizing field is salinity. Near the surface, evaporation perturbs the salinity field creating parcels of salty dense fluid. As they sink, the parcels cool more quickly than they diffuse salt, further

increasing their density. This flow drives tall vertical convective cells, called salt fingers, that mix the oceans.

The same doubly diffusive buoyant process drives short broad convective cells at greater depths. Observed as thermohaline staircase [55], a series of sharp steps in the salinity and temperature versus depth, these cells occur at depths greater than a kilometer and extend for thousands of square miles.

**Inertial confinement fusion** Inertial confinement fusion (ICF) is a fusion technique that confines hot dense plasma temporarily via an implosion rather than magnetic field lines. In ICF, cryogenic deuterium is coated with a plastic ablator forming very small hollow spheres, or microcapsules. The microcapsules are positioned in a cylindrical *hohlraum*. The hohlraum interior is illuminated with a high energy burst of laser light and radiates x-rays into its interior. This x-ray radiation is absorbed by the plastic ablator causing it to rapidly expand and blow off the microcapsule, creating an implosion in the hydrogen fuel [41].

The implosion is not perfectly uniform; the x-ray radiation is not isotropic and there are asymmetries in the microcapsule. The resulting perturbations are Rayleigh-Taylor unstable: the dense plastic ablator is being accelerated through lighter hydrogen fuel. The carbon in the ablator is much better at radiating energy than the hydrogen, so Rayleigh-Taylor mixing cool the fuel, preventing ignition [18].

### 1.3 Terminology

In this section, we introduce common Rayleigh-Taylor terminology that will be used throughout the thesis.

**Atwood number** The Atwood number characterizes the density contrast:

$$A = \frac{\rho_1 - \rho_2}{\rho_1 + \rho_2} \in (-1, 1), \quad (1.5)$$

3

where  $\rho_1 > \rho_2$  corresponds to positive Atwood number. Negative Atwood numbers result in stable oscillating interfaces. There are three distinct regimes of unstable Atwood numbers: high, low, and moderate.

At high Atwood number, e.g. air and water, the flow is approximated in the limit of unit Atwood number. The internal dynamics of the light fluid are decoupled from the heavy fluid, and the transfer of momentum through the interface into the dense fluid can be neglected. In this regime, the flow of the dense fluid is nearly irrotational and potential flow models are reasonably accurate. The high Atwood regime requires the Atwood number to be at least one half.

At low Atwood number, e.g. salt water and fresh water, the flow is approximated in the limit of zero Atwood number. The governing equations can be simplified via the Boussinesq approximation: only terms with the product of the Atwood number and acceleration remain. The behavior of the rising fluid, i.e. bubbles, and falling fluid, i.e. spikes, are symmetric given symmetric initial conditions. It is often possible to treat the true two-phase problem as a single phase with an active scalar. The valid range of Atwood numbers for the Boussinesq approximation is problem dependent [40], but typically taken to be less than  $A = 0.10$ .

At moderate Atwood number, e.g. oil and water, not only is enough vorticity is generated to undermine potential flow models but also the density difference breaks the Boussinesq approximation and bubble-spike symmetry. These are truly two-phase problems in that the fluids are strongly coupled and have different governing parameters, e.g. viscosity.

**Boussinesq approximation** When the density contrast,  $\Delta\rho$ , is small compared to the average density  $\bar{\rho}$ , then the contrast can be neglected except when multiplied by the acceleration, as in  $Ag$ . This is known as the Boussinesq approximation and has the primary effect of neglecting differences in the inertia of the two fluids.

**Uniform constant property** In the spirit of the Boussinesq approximation, which neglects differences in the inertia of the two fluids, one can further assume the two fluids share all other material properties, such as viscosity. When the flow is both Boussinesq and has uniform constant properties, the density can be modeled as an active scalar with a buoyant forcing term.

**Miscible interface** In most cases where the Boussinesq and uniform constant property approximations are valid, the interface between the two fluids is miscible. For example, low-concentration solutions in a common solvent have interfaces governed by a diffusion coefficient,  $D$ , and small temperature gradients are governed by a thermal diffusivity  $\alpha$ . In the limit where  $D$  or  $\alpha$  goes to zero, the interface is immiscible but there is no surface tension. If there is no surface tension, the active scalar is governed by an advection-diffusion equation.

**Incompressible** The problem is said to be incompressible when each of the two fluids are. Formally, this means the flow is divergence free,  $\nabla \cdot u = 0$ . In simulations, incompressibility has the effect of integrating out acoustic waves, significantly reducing the need to resolve small time-scales in the flow.

**Single-mode** The single mode Rayleigh-Taylor instability constrains the initial perturbation of the interface to a single pure spatial frequency. In 3D, this typically implies two orthogonal wavevectors with the same wavelength. For the miscible uniform constant property Boussinesq Rayleigh-Taylor instability, the single mode initial condition is:

$$\phi(x, y, z, t = 0) = \operatorname{erf} \left( \frac{z + a_0 \cos(2\pi x/\lambda) \cos(2\pi y/\lambda)}{\delta} \right), \quad (1.6)$$

where  $a_0$  is the perturbation height,  $\lambda$  is the wavelength, and  $\delta$  is the interface thickness. Typically, the perturbation height and interface thickness are chosen to initially be much

smaller than the wavelength,  $a_0, \delta \ll \lambda$ .

**Multi-mode** The multi-mode Rayleigh-Taylor instability generically refers to the presence of more than one, and often many, wavelengths in the initial condition. Experimentally, the multi-mode initial condition is usually ‘natural’, in the sense that it is not deliberately perturbed and instead is due to natural noise sources. Computationally, the multi-mode condition is written as a sum of single modes with spectra  $z_0(k) \sim |k|^{-p}$  for  $p = 2, 3/2$ , or other rational factors. These modes are randomized within the spectral envelope.

**Governing equations** Given a miscible interface between two Boussinesq, uniform constant property incompressible fluids, the governing equations can be recast in terms of an active scalar:

$$\frac{du}{dt} + u \cdot \nabla u = \nu \nabla^2 u - \nabla P - Ag\phi\hat{z}, \quad (1.7)$$

$$\frac{d\phi}{dt} + u \cdot \nabla \phi = D \nabla^2 \phi, \quad (1.8)$$

$$\nabla \cdot u = 0, \quad (1.9)$$

where  $u$  is the fluid velocity,  $\phi$  is the active scalar that represents small mass differences, and the acceleration is aligned vertically with  $\hat{z}$ .

These equations have three governing parameters:  $(Ag)$ ,  $\nu$ , and  $D$ . In the single-mode case, the dominant length scale is  $\lambda$ , which supports the construction of two dimensionless numbers:

$$\text{Grashof} = \frac{Ag\lambda^3}{\nu^2}, \quad (1.10)$$

$$\text{Schmidt} = \frac{\nu}{D}. \quad (1.11)$$

If the initial and boundary conditions have no other length scales, e.g. the initial interface is a thin low-amplitude perturbation, i.e.  $a_0, \delta \ll 1$ , then these two parameters uniquely

identify the system.

It is convenient to introduce another dimensionless number, the Rayleigh number:

$$\text{Rayleigh} = \text{Grashof} \cdot \text{Schmidt} = \frac{Ag\lambda^3}{\nu D}, \quad (1.12)$$

with which many mixing-related quantities scale. The three dimensionless numbers are abbreviated Gr, Sc, and Ra, respectively.

It is difficult to define a Reynolds number in the absence of a velocity parameter. The square root of the Grashof number, which has the same scaling with the viscosity, is used instead. It is sometimes called the perturbation Reynolds number,  $\text{Re}_p = \sqrt{\text{Gr}}$ .

**Bubble height and spike depth** The bubble height refers to the distance beyond the initial interface that the bubble front has traveled, as a function of time. The standard experimental definition projects the density onto a line normal to the initial interface and defines the front interface as the point at which the density is at its 99th or 95th percentile. More formally:

$$H_p[\epsilon] = \sup \left\{ z : \int \tilde{\rho}(x, y, z) dx dy < (1 - \epsilon) \int \tilde{\rho}(x, y, \infty) dx dy \right\}, \quad (1.13)$$

where  $\tilde{\rho}$  is the deviation from the mean density,  $\tilde{\rho} = \rho - \bar{\rho}$ , and  $\epsilon$  is either 0.05 or 0.01 for the 95th and 99th percentile definitions, respectively.

If the fluids are miscible, this definition depends on the rate of diffusion across the interface. To avoid dependence on diffusion across the interface, we can base the bubble height on a measurement of the equi-molar interface, which is stationary under diffusion. To pick out the interface, we take a span-wise maximum instead of a span-wise sum:

$$H_m = \sup \left\{ z : \max_{x, y} \tilde{\rho}(x, y, z) < 0 \right\}. \quad (1.14)$$

However, diffusive mixing across the sides of an elongated bubble dilute as it grows. This dilution, which is observed as a linear profile in the span-wise maximum instead of an error function profile, can dip below  $\tilde{\rho} = 0$ , at which point the growth of  $H_m$  is also influenced by mixing.

In the absence of this affect,  $H_m$  tracks not only the equi-molar surface but also the inflection point in the profile  $\max_{x,y} \tilde{\rho}$ . This inflection point closely tracks the equi-molar surface at low diffusivity but is robust to diffusion across the bubble, remaining in the center of the error function region of the maximum density profile. Formally, this definition is:

$$H_i = \sup \left\{ z : \frac{d^2}{dz^2} \max_{x,y} \tilde{\rho}(x, y, z) = 0 \right\}. \quad (1.15)$$

Equivalent definitions can be given for the spike depth. If the flow is Boussinesq and the initial condition is symmetric, the bubble height and spike depth can be averaged.

**Mixing width** In some application, the interpenetration of the two fluids is secondary to their mixing. To measure the mixed-ness of the flow, we map the pure fluids to unity, mixed fluids to zero, and integrate. First, we define a normalized scalar measure of the density as:

$$\phi = \frac{2\rho - 2\bar{\rho}}{\rho_1 - \rho_2} \in [-1, 1], \quad (1.16)$$

then the purity of a fluid volume is  $|\phi|$ . The mix volume is simply the integral:

$$\Theta(t) = \int (1 - |\phi(x, y, z, t)|) dV, \quad (1.17)$$

and the mixing width  $\theta = \Theta/\mathcal{A}$ , where  $\mathcal{A}$  is the span-wise extent of the volume integral.

## 1.4 Note on the structure of the dissertation

This thesis contains three standalone papers, which are reproduced following a review of the state of modeling the Rayleigh-Taylor instability. They are logically ordered from the top down.

Chapter 3 contains a new buoyancy-drag model developed to describe the late time behavior in the low-Atwood, moderate Grashof case. It contains model coefficients that are fit to a data set of direct numerical simulations.

Chapter 4 validates those direct numerical simulations against experimental data from Wilkinson and Jacobs [62]. This establishes that the simulations contain the physical processes responsible for re-acceleration and other unexplained late-time phenomena. The paper takes advantage of the generality of numerical data to explore features of the flow that were not available to the original experiment, namely pressure driven secondary flows in the mid-plane and the interaction of the bubbles with one another.

Chapter 5 [27] contains a performance and convergence study of the numerical method and simulation software with the single mode Rayleigh-Taylor problem as a benchmark. The NekBox code, which was specialized specifically to this project, is shown to be an efficient tool for performing these calculations. Furthermore, the resolution is selected such that the simulation error is an order smaller than the expected model error, which ensures that the flow is sufficiently but not overly resolved. Maxwell Hutchinson authored all of sections 2 and 4 and the majority of sections 1 and 5.

## CHAPTER 2

# BACKGROUND

The study of the Rayleigh-Taylor instability (RTI) is primarily interested in evolution of the interface, that is the rate of penetration of the light fluid into the dense one and vice versa. While the volumetric mixing rate is relevant in some contexts, most flows have relatively low diffusivity, i.e., high Prandtl and Schmidt number, so mixing is dominated by transport rather than diffusion. In this section, we review approaches to modeling the evolution of the interface and experimental efforts to validate those models.

### 2.1 Linear and weakly non-linear models

The earliest models of the Rayleigh-Taylor instability were based on a linearization of the governing equations around small perturbations in the interface. Recently, with the aid of computational algebra, it has become possible to retain higher order terms in the expansion, demonstrating mode coupling and saturation amplitudes. However, even high order expansions fail as the interface loses analyticity.

#### 2.1.1 *Lord Rayleigh's linear model*

Lord Rayleigh considered a sinusoidal perturbation of an incompressible, inviscid, immiscible, quiescent stratified interface [35]. When the amplitude is small compared to the wavelength, the continuity and momentum equations can be linearised:

$$(\bar{\rho} + \tilde{\rho}) \left[ \frac{du}{dt} + u \nabla u \right] = -\nabla \tilde{P} + g(\bar{\rho} + \tilde{\rho}), \quad (2.1)$$

$$\nabla \cdot u = 0, \quad (2.2)$$

where  $u, \tilde{\rho}, \tilde{P}$  are the perturbation velocity, density, and pressure. They are small and their products neglected. The solution is sinusoidal with an exponential exponential with a growth rate:

$$w \sim e^{ikx} e^{-kz} e^{\gamma t}, \quad \gamma^2 = Agk, \quad (2.3)$$

where  $w$  is the vertical component of the velocity,  $g$  is the acceleration experienced by the fluid, and  $k = 2\pi/\lambda$  is the wave-number of the perturbation. Positive Atwood numbers correspond to unstable density stratifications, which grow exponentially. Negative Atwood numbers correspond to stable density stratifications, which oscillate.

### 2.1.2 Viscous and diffusive linear models

Chandrasekhar [7] and Hide [23] generalized the linear theory to viscous fluids by including an isotropic incompressible Newtonian shear stress. Chandrasekhar worked out the uniform constant property case,  $\mu_1 = \mu_2$ , and Hide included an approximate combination of distinct viscosities. Here, we are concerned with the simpler uniform constant property case:

$$\gamma = \sqrt{Agk + \nu^2 k^4} - \nu k^2, \quad (2.4)$$

where  $\nu$  is the kinematic viscosity. Note that the viscous growth rate has a fastest growing mode at finite wavenumber and that all wavenumbers have positive growth rates.

LeLevier et al. [33] generalized the linear theory to continuous density gradients, specifically exponentially smoothed profiles for the form  $\bar{\rho} \pm e^{\mp Kz} \delta \rho$ :

$$\gamma = \sqrt{\frac{AgkK}{k + K}}. \quad (2.5)$$

Duff et al. [15] generalized the linear theory to miscible interfaces and incorporated Chandrasekhar and Hide's viscous theories, producing an combined expression for the growth

rate:

$$\gamma = \sqrt{\frac{Agk}{\psi(A, k\delta)} + \nu^2 k^4 - (\nu + D)k^2}, \quad (2.6)$$

where  $\delta$  is the instantaneous interface thickness,  $D$  is the diffusivity, and  $\psi$  is a function of the Atwood number and the product of the wavenumber and the interface thickness. For  $k\delta \ll 1$  and  $A \ll 1$ ,  $\psi \approx 1 + k\delta/\sqrt{\pi}$ . Note that for  $D > 0$  there is a wavenumber cutoff above which the growth rate is negative, i.e., the perturbation decays.

In the uniform constant property case,  $\delta = 2\sqrt{Dt}$ , introducing a time-dependence on the linear stability:

$$\gamma = \sqrt{\frac{Agk}{1 + \frac{2k}{\sqrt{\pi}}\sqrt{D(t + t_0)}} + \nu^2 k^4 - (\nu + D)k^2}, \quad (2.7)$$

where  $t_0$  is defined by the initial interface thickness:

$$t_0 = \frac{\delta_0^2}{4D}, \quad (2.8)$$

where  $\delta_0$  is the initial interface thickness.

### 2.1.3 Weakly nonlinear expansions

Jacobs and Catton provide a third order weakly non-linear theory for the inviscid unit Atwood Rayleigh-Taylor instability [30]. Their weakly non-linear theory is primarily used to compare linear growth rates across a variety of perturbation symmetries in 3D. Hexagonal and axi-symmetric perturbations are found to grow faster than rectangular perturbations.

Berning and Rubenchik extend the theory to arbitrary Atwood immiscible flows at any higher order, but analyze only the third order expansion [5]. They perform a similar geometric comparison to Jacobs and Catton, but also use the harmonic couplings to characterize linear saturation.

The perturbation expansion has been taken to at least the 10th order by Liu et al. [59].

However, there is limited progress to be made with such expansions, as singularities with branching point structures develop at moderate bubble displacements [5]. Put another way, the interface and velocity potentials are not analytic in the span-wise position, e.g., when the interface rolls up.

## 2.2 Potential flow models

The next class of models to be applied to the Rayleigh-Taylor instability are potential flow models. These models assume that little vorticity is generated and that it is confined to the interface, which is true at high Atwood numbers. At moderate and low Atwood numbers, there is significant generation and transport of vorticity via, for example, the Kelvin-Helmholtz instability, so these models break down.

### 2.2.1 Layzer's unit Atwood model

One of the first potential flow models is due to Layzer [32]. Layzer's model is of a bubble with  $\rho = 0$  rising in a fluid of density  $\rho = 1$ , which is unit Atwood number. The bubble and fluid are assumed to be incompressible and inviscid. The flow begins at rest, so there is no initial vorticity. Layzer claims the flow will therefore continue to be irrotational, because the viscous generation term of the vorticity equation is zeroed for inviscid, incompressible flows.

Since the flow is inviscid and irrotational, Layzer uses the potential flow technique, writing the velocity as the gradient of a scalar potential:

$$v = \nabla\Phi, \quad (2.9)$$

where  $v$  is the velocity and  $\Phi$  is the scalar potential. Incompressibility zeroes the Laplacian of the potential:

$$\nabla^2\Phi = 0. \quad (2.10)$$

A Bernoulli equation is used model the interface:

$$f(t) = \frac{d\Phi}{dt}(\eta(r, t), r, t) - \frac{1}{2} \left( \left( \frac{\partial \Phi}{\partial z} \right)^2(\eta(r, t), r, t) + \left( \frac{\partial \Phi}{\partial r} \right)^2(\eta(r, t), r, t) \right) - g\eta(r, t), \quad (2.11)$$

where  $\eta(r, t)$  is the height of the interface,  $g$  is the gravitational acceleration, and  $f(t)$  is an arbitrary function of time but not space. The flow is axially symmetric with a vanishing radial component at transverse walls and vanishing vertical component far away from the bubble:

$$\frac{\partial \Phi}{\partial r}(z, R, t) = 0, \quad \frac{\partial \Phi}{\partial z}(\pm\infty, r, t) = 0. \quad (2.12)$$

Finally, the fluid advects the interface:

$$\frac{d\eta}{dt}(r, t) = \frac{\partial \Phi}{\partial z}(\eta(r, t), r, t) - \frac{\partial \Phi}{\partial r}(\eta(r, t), r, t) \frac{d\eta}{dr}(r, t), \quad (2.13)$$

The bubble accelerates to a terminal velocity. That velocity, in two and three dimensions, is:

$$V_{2d} = \frac{1}{\sqrt{3}} \sqrt{\frac{gR}{\pi}}, \quad V_{cyl} = \sqrt{\frac{gR}{\beta_1}}, \quad (2.14)$$

where  $\beta_1$  is the first root of the first order Bessel function of the first kind:  $J_1(\beta_1) = 0$ . This velocity agrees with experimental results that were available to Layzer, e.g. those by Davies and Taylor [8].

### 2.2.2 Goncharov's high Atwood model

Goncharov extends the Layzer model to include two fluids of arbitrary density difference. In doing so, he makes a different choice of simplifying approximation for the Bernoulli equation [19]. The consideration of a second fluid with non-zero density turns the Bernoulli

equation Equation 2.11 into a difference:

$$\begin{aligned}
f(t) = & \rho_1 \frac{d\Phi_1}{dt}(\eta(r, t), r, t) - \rho_2 \frac{d\Phi_2}{dt}(\eta(r, t), r, t) \\
& - \rho_1 \frac{1}{2} \left( \left( \frac{\partial \Phi_1}{\partial z} \right)^2 (\eta(r, t), r, t) + \left( \frac{\partial \Phi_1}{\partial r} \right)^2 (\eta(r, t), r, t) \right) \\
& + \rho_2 \frac{1}{2} \left( \left( \frac{\partial \Phi_2}{\partial z} \right)^2 (\eta(r, t), r, t) + \left( \frac{\partial \Phi_2}{\partial r} \right)^2 (\eta(r, t), r, t) \right) \\
& - g\rho_1\eta(r, t) + g\rho_2\eta(r, t).
\end{aligned} \tag{2.15}$$

The Goncharov model keeps the free-slip boundary condition between the two fluids, which is exact only for  $A = 1$  and a reasonable approximation for  $\rho_1/\rho_2 \gg 1$ . In this respect, Goncharov's model should be reasonable for high-Atwood, nearly inviscid flows. The terminal velocity predicted is:

$$V = 1.02 \sqrt{\frac{2A}{1+A} \frac{g}{k}} = \frac{1.02}{\sqrt{\pi}} \sqrt{\frac{Ag\lambda}{1+A}}. \tag{2.16}$$

Similar potential flow models were introduced by Sohn [52] and Abarzhi et al. [2], with similar results.

### 2.3 Buoyancy-drag models

Buoyancy-drag models were developed concurrently with potential flow models, in part to provide a physical interpretation for their results. They balance buoyant and parasitic forces related to the geometry of a model bubble. Historically, buoyancy-drag models have had only 1 or 2 adjustable parameters, so they are evaluated more on their ability to reproduce specific features of the flow, e.g., the terminal velocity, rather than the full time-history. Here, we focus on models applicable to single-mode non-interacting bubbles.

### 2.3.1 Bubble model of Davies and Taylor

Early experiments on the Rayleigh-Taylor instability by Davies and Taylor [8] were performed by measuring the dynamics of large bubbles of gas rising through a dense liquid. In their analysis, they relate the terminal velocity of the bubble to a drag coefficient, implicitly defining a buoyancy-drag model of the form:

$$\dot{v}\rho\mathcal{V} = \rho g\mathcal{V} - C_D\pi d^2 \frac{1}{2}\rho v^2, \quad (2.17)$$

where  $v$  is the gas bubble velocity,  $\rho$  is the density of the liquid,  $g$  is the gravitational acceleration,  $\mathcal{V}$  is the bubble volume,  $C_d$  is a drag coefficient, and  $d$  is the bubble diameter. The coefficient  $C_d$  was found to take values between 0.52 and 1.37.

### 2.3.2 Tube model of Dimonte and Schneider

Dimonte and Schneider develop a buoyancy-drag model for tube-shaped bubbles [12, 11] based on Davies and Taylor's model, Equation 2.17. They let the ratio of the area to the volume go with the inverse bubble height,  $\mathcal{A}/\mathcal{V} \sim 1/h$ . They also add a rescaling of the buoyant term by  $\beta$ , attributed to Youngs:

$$\dot{v}_b = \beta Ag - C_d \frac{v_b^2}{h_b}, \quad (2.18)$$

where  $v_b$  is the bubble velocity,  $\beta < 1$  accounts for the relatively smaller buoyant portion of the bubble due to entrainment,  $C_d$  is a drag coefficient, and  $h_b$  is the bubble height.  $\beta$  and  $C_d$  depend on the Atwood number, but Dimonte proposes  $\beta = 1/2$  and  $C_d = 2$  for  $A \ll 1$  [13]. However, the model is stated to apply to self-similar bubble fronts, in which  $h_b \sim D$ .

### 2.3.3 Self-similar model of Oron

A model by Oron et al. also rescales the bubble mass [44]:

$$(\rho_1 + C_a \rho_2) \ddot{h} = (\rho_2 - \rho_1) g - \frac{C_d}{\lambda} \dot{h}^2 \rho_2 \mathcal{A}, \quad (2.19)$$

where  $\rho_2 > \rho_1$  are the densities of the two fluids,  $C_a$  is an added mass coefficient,  $h$  is the height of the bubble,  $g$  is the gravitational acceleration,  $C_d$  is a drag-like coefficient, and  $\lambda$  is a characteristic length. The use of  $\lambda$ , which is time-independent, implies the model is directed at self-similar flow. The values of  $C_a$  and  $C_d$  are assumed to be Atwood independent and set to agree with Layzer's theory:

$$C_a = 1, \quad C_d = 2\pi. \quad (2.20)$$

## 2.4 Problems with single mode Rayleigh-Taylor modeling

Simulations by Ramaprabhu et al. [48] have shown that, after stagnating at a constant velocity in agreement with the potential flow models, low-Atwood bubbles re-accelerate to velocities nearly twice the potential flow limit. The stagnation and re-acceleration phenomena were confirmed experimentally by Wilkinson and Jacobs [62]. Modeling the stagnation and re-acceleration phases is the primary open problem in the low Atwood single-mode Rayleigh-Taylor instability. The desire to describe this re-acceleration process, quantitatively, is the motivation for this thesis.

First, though, I will give three qualitative explanations for why re-acceleration should have been expected: one based on the pressure balance, one based on the assumptions of potential flow models, and another by identifying a historical inconsistency in the development of buoyancy-drag models.

### 2.4.1 Pressure in the single-mode RTI

If there is a terminal velocity regime, can it be due to form drag? In other words, can we have terminal velocity without viscosity? Let  $\nu \rightarrow 0$  and consider a fluid element lying on the axis of a bubble or spike at terminal velocity. By symmetry, it will only have a  $z$ -component of the velocity. The  $z$ -forces must balance:

$$-\phi g \hat{z} = \frac{\partial P}{\partial z}, \quad (2.21)$$

where  $\phi$  represents the mass. In the falling spike, the pressure would be decreasing with  $z$ . In the rising bubble, the pressure would be increasing with  $z$ . There would necessarily be a pressure gradient between the head of the bubble and the tail of the spike, and vice versa. As the bubble aspect ratio exceeded unity, the span-wise pressure gradient would exceed the gravitational forcing. The resulting span-wise flow would rapidly mix the two fluids, destroying the bubble and spike.

The form of the bubble and spike require the pressure to be reasonably homogeneous span-wise. Only at the bubble and spike tips do we observe a span-wise flow: the displacement of stationary fluid by the tip. In other words, the pressure drag is highly localized to the bubble and spike tips but cannot affect the flow in the stems of the bubbles and spikes. If the flow is terminal, then it must be attenuated predominately by viscous drag, which can act along the sidewalls, that is the stem, of the bubbles and spikes.

### 2.4.2 Departure from potential flow

The assumption that the flow is irrotational applies only at high Atwood number. At moderate and low Atwood numbers, the interface between the light and the dense fluid is a shear layer that generates vorticity. If the viscosity is low enough, secondary Kelvin-Helmholtz instabilities develop in the shear layer and transport vorticity into the center of the bubble.

While it is not obvious that vorticity should cause re-acceleration, it is clear that the flow is not irrotational, even away from the fluid interface, and therefore cannot be accurately modeled by potential flow.

#### *2.4.3 Historical inconsistency in buoyancy-drag models*

Buoyancy-drag models contain a buoyant term that goes with the bubble's volume and a form drag term that goes with the bubble's span-wise area. The model was originally developed to describe multi-mode self-similar flow, in which there is only one length scale, the dominant wavelength  $\lambda$ . Consequently, the ratio of the volume to the surface area is  $\lambda^{-1}$ , yielding a terminal velocity as a function of  $\lambda$ .

However, the single-mode RTI has two length scales: in addition to the wavelength  $\lambda$  there is the bubble height,  $h$ . In other words, single-mode RTI bubbles are cylindrical instead of spherical with an axis length that goes with the bubble height. The ratio of the volume to surface area is  $h^{-1}$ , not  $\lambda^{-1}$ , so force balance occurs when  $\dot{h} \sim \sqrt{h}$ , which is not terminal. Only by introducing a drag term that goes with the height  $h$ , such as skin drag, can a terminal velocity be recovered. This terminal velocity would be a function of the viscosity, and therefore cannot be described by potential flow.

### **2.5 Related work aimed at modeling stagnation and re-acceleration**

Since re-acceleration was observed experimentally, multiple attempts have been made to capture re-acceleration in the models.

### 2.5.1 Vortex ring correction of Ramaprabhu

Ramaprabhu et al. [47] attribute the reacceleration to the formation of a vortex ring at the bubble tip. They add a term to their buoyancy-drag model representing the centrifugal force per unit volume:

$$(\rho_2 g - \rho_1 g) + \rho_1 \frac{\omega_0^2 R}{4} = \frac{C_d \rho_2 v^2}{\lambda}, \quad (2.22)$$

where  $\omega_0$  is the average vorticity in the bubble tip. The model does not provide an evolution equation for  $\omega_0$ ; it is measured from simulations ad hoc making the model descriptive but not predictive. The model agrees qualitatively, but not quantitatively, from the onset of stagnation at bubble height  $h/\lambda \approx 0.5$  through re-acceleration at  $h/\lambda \approx 2.0$ , but doesn't capture linear growth at early times or the dynamics at late times. Furthermore, they compare the vortex ring model to simulations using two different codes; the two codes disagree quantitatively over the re-acceleration regime and disagree qualitatively over what follows it.

## 2.6 Vorticity and viscosity in potential flow

Banerjee et al. attempt to describe re-acceleration by adding viscous and vortical effects to a potential flow model [3]. Similar to the vortex ring correction to buoyancy-drag, the vorticity is an input to the potential flow model. Instead of using data from simulations, Banerjee et al. write the vorticity as an analytic function of time independent of the Atwood number. The resulting dynamics have a single re-acceleration phase before reaching an asymptotic terminal velocity.

# CHAPTER 3

## DATA-DRIVEN MODELING OF THE LOW-ATWOOD SINGLE-MODE RAYLEIGH-TAYLOR INSTABILITY

### 3.1 Abstract

The Rayleigh-Taylor instability (RTI) pervades classical fluid dynamics and is essential to a diversity of phenomena, e.g., salt fingers, thermonuclear flames, and inertial confinement fusion, but remains poorly understood in dissipative systems. Recently, the single-mode RTI has shown experimentally and numerically to deviate from established potential flow models when the Atwood number was less than 1/2. Attempts to explain the deviation, termed re-acceleration, have been ad hoc and hindered by a dearth of data at late times and high aspect ratios. This paper present buoyancy-drag and mixing models that include dissipative terms and match the linear theory. To inform the model, a numerical experiment is performed, simulating a range of Grashof and Schmidt numbers and reaching bubble heights  $34 \times$  the bubble width. The model coefficients are estimated by physical argument and then fit to the numerical results. The model error is less than 2% for the bubble height and 4% for the volume of mixed fluid. An attempt is made to interpret variations in the fit parameters with the Rayleigh and Schmidt numbers, where present, but it is hindered by many of the simulations interacting with the boundaries. Simulations in higher aspect ratio domains would improve the model.

### 3.2 Introduction

The Rayleigh-Taylor instability has been the subject of considerable study since its characterization by Lord Rayleigh in the 19th century [35]. Despite this, many aspects of the non-linear growth remain poorly understood. Analytic models based on potential flow have been

reasonably effective for flows with a large density jump, i.e. an Atwood number,  $\Delta\rho/\sum\rho$ , near unity [32, 32]. However, recent experiments at low Atwood number have demonstrated a significant departure from the potential flow limit: Rayleigh-Taylor bubbles are seen accelerating past the terminal velocity predicted by potential flow models [62]. The acceleration persists beyond times which are experimentally accessible, so numerous efforts have been made to compute the late time flow numerically [47, 61]. Ultimately, the goal is to use these numerical experiments to inform a simple model that captures the key features and mechanisms of the flow, as potential flow models do for high Atwood number flows.

This study concerns the dynamics of the single-mode Rayleigh-Taylor instability (sRTI), where the interface between a heavy fluid and a light fluid is perturbed with a single wavelength  $\lambda$  and corresponding wavenumber  $k$ . If the Atwood number is low, then at early times the interface grows exponentially with a rate given by a linear approximation [15]:

$$\gamma = \sqrt{\frac{Agk}{1 + \pi^{-1/2}k\delta} + \nu^2k^4 - (\nu + D)k^2}, \quad (3.1)$$

where  $A$  is the Atwood number,  $g$  is the local acceleration,  $\delta$  is the interface thickness,  $\nu$  is the kinematic viscosity, and  $D$  is the diffusivity. As the amplitude approaches the wavelength, the linear growth saturates. At unit Atwood number, the non-linear regime is described by potential flow, in which Layzer [32] found the interface approaches a terminal velocity proportional to the root of the acceleration and wavelength:  $v_L \sim \sqrt{g\lambda}$ : Potential flow models have been extended to  $A < 1$ , with the most successful model by Goncharov [19]:

$$v_G = \frac{1}{\sqrt{\pi}} \sqrt{\frac{Ag\lambda}{1 + A}} \quad (3.2)$$

Experiments by Wilkinson and Jacobs [62] show that after reaching the velocity given by Equation 3.2, low Atwood bubbles unexpectedly accelerate a second time. This is termed ‘re-acceleration’, with the terminal velocity replaced by a ‘stagnation velocity’. Re-acceleration

was not present in popular potential flow and buoyancy-drag models, and attempts to capture re-acceleration have been the emphasis of recent model development in the low Atwood number regime.

Experiments have thus far been unable to observe more than the onset of the re-acceleration phase. Therefore, the community has turned to numerical studies to compute late-time dynamics. At least two such efforts have been undertaken, one by Ramaprabhu et al. [47] and one by Wei and Livescu [61], leading to slightly different conclusions. In the study by Ramaprabhu, the flow accelerates to around twice the stagnation velocity and then decelerates back to the stagnation velocity, indicating that re-acceleration may be a transient. In the study by Wei and Livescu, the flow re-accelerates and then breaks up into many small pockets of buoyant fluid, which themselves continue to accelerate at nearly a fixed rate.

Buoyancy-drag models have been proposed for the multi-mode and single-mode dynamics of the bubble front. They balance the buoyant force of the bubble with a drag force to predict the front velocity as a function of the characteristic length. Ramaprabhu proposes an additional forcing term based on the development of a vortex ring at the bubble tip, but relies on observation of the mean vorticity experimentally or in numerical simulations. There is a need for a predictive model for re-acceleration that relies only on the initial conditions.

In Section 3.3, we propose a simple buoyancy-drag model for the dynamics of the smRTI. In Section 3.4, we present a battery of numerical trials. In Section 3.5, we describe the regimes present in late time bubble trajectories. In Section 3.6, we fit the unconstrained model coefficients to the numerical results.. Finally, in Section 3.7, we assess the current state of smRTI and highlight outstanding questions.

### 3.3 Simple model

We base our model on the buoyancy-drag models of [44]:

$$(\rho_1 + \rho_2)\mathcal{V}\ddot{h} = (\rho_2 - \rho_1)g\mathcal{V} - C\dot{h}^2\rho\mathcal{A}, \quad (3.3)$$

where  $\rho_1$  and  $\rho_2$  are the densities of the light and heavy fluid,  $\mathcal{V}$  is the characteristic volume of the bubble  $g$  is the acceleration,  $C$  is a drag-like coefficient, and  $\mathcal{A}$  is the characteristic cross sectional area of the bubble. Making the Boussinesq approximation,  $\rho_1 \approx \rho_2$  yields:

$$\ddot{h} = Ag - \frac{C}{2}\dot{h}^2\frac{\mathcal{A}}{\mathcal{V}}. \quad (3.4)$$

In the self-similar regime there is only one length scale, so  $\mathcal{A}/\mathcal{V} \sim 1/\lambda$ . However, in the single-mode regime that is the focus of this study, the bubbles are elongated, producing two length scales: a span-wise scale  $\lambda$  and a stream-wise scale  $h$ . Therefore, for the smRTI  $\mathcal{A}/\mathcal{V} \sim \frac{1}{h}$  and the model of Oron et al. yields unbounded velocities:

$$\ddot{h} = Ag - \frac{C}{2}\frac{\dot{h}^2}{h}. \quad (3.5)$$

Because the strength of the form drag relative to buoyancy decreases at high aspect ratio, we must consider other drag terms, such as skin drag, that grow at least linearly with  $h$ .

#### 3.3.1 Dynamics

We begin by listing the external forces the bubble experiences. The first is the buoyant force:

$$F_b = C_0Ag\lambda^2h, \quad (3.6)$$

where  $C_0$  is an unknown coefficient. The next is the form drag:

$$F_f = C_1 \lambda^2 \dot{h}^2, \quad (3.7)$$

where  $C_1$  is similar to a drag coefficient. The next is the viscous, or skin, drag:

$$F_s = C_2 \nu h \dot{h}, \quad (3.8)$$

where  $C_2$  is another unknown coefficient and  $\nu$  is the kinematic viscosity.

To complete the dynamic equation, we must characterize the inertia of the bubble. The bubble is roughly cylindrical with a height  $h$ , so we expect an inertial term of the form  $\lambda^2 h$ . However, consider the limit of  $h \rightarrow 0$ . Here, streamlines must extend from bubble to spike, which has a characteristic separation  $\lambda$  for an inertial term of the form  $\lambda^3$ . Therefore, we expect the inertia to be a mix of a term that goes as  $\lambda^2 h$  and one that goes as  $\lambda^3$ :

$$I = C_3 \lambda^2 h + C_4 \lambda^3, \quad (3.9)$$

where  $C_3$  and  $C_4$  are two more unknown coefficients.

The complete dynamic equation is:

$$\ddot{h} = \frac{C_0 A g \lambda^2 h - C_1 \lambda^2 \dot{h}^2 - C_2 \nu h \dot{h}}{C_3 \lambda^2 h + C_4 \lambda^3}. \quad (3.10)$$

Without loss of generality, we can let  $C_0 = 1$  and simplify:

$$\ddot{h} = \frac{A g h - C_1 \dot{h}^2 - C_2 \nu (h/\lambda^2) \dot{h}}{C_3 h + C_4 \lambda}. \quad (3.11)$$

We can nondimensionalize by defining a dimensionless length and time:

$$z = \frac{h}{\lambda}, \quad \tau = \sqrt{\frac{Ag}{\lambda}} t, \quad (3.12)$$

which simplifies:

$$\ddot{z} = \frac{z - C_1 \dot{z}^2 - C_2 \text{Gr}^{-1/2} z \dot{z}}{C_3 z + C_4}, \quad (3.13)$$

where the derivative is with respect to  $\tau$  and  $\text{Gr} = A_0 g \lambda^3 \nu^{-2}$  is the Grashof number.

### 3.3.2 Mixing

As the bubble height grows, the velocity approaches a terminal value specified by the balance between buoyancy and skin drag. At terminal velocity, the flux of pure fluid into the bubble is bounded. However, the interfacial mixing continues to grow with the interfacial area, which grows with  $h$ . Therefore, for any finite diffusivity, the bubble will ultimately diffuse away. For this reason, we must include the effects of interfacial mixing, at least to the first order.

The quantity of mixed fluid,  $m$  is computed directly from the time, bubble height, diffusivity, and initial interface thickness. The quantity of mixed fluid is defined as the integral of the absolute value of the scalar:

$$m(t) = \int (1 - \text{abs} [\phi(x, y, z, t)]) dV, \quad (3.14)$$

where we assume the mean scalar is zero,  $\int \phi dV = 0$ .

We approximate the volume integral by a 1D integral across the interface multiplied by the surface area:

$$m(t) \approx S \int (1 - \text{abs} [\phi_1(r)]) dr, \quad (3.15)$$

where  $S$  is the surface area and  $\phi_1$  is a model 1D scalar profile:

$$\phi_1(r) = \frac{1}{2} \left( \operatorname{erf} \left[ \frac{r}{\delta} \right] - \operatorname{erf} \left[ \frac{r-d}{\delta} \right] \right), \quad (3.16)$$

where  $\delta$  is the interface width and  $d$  is the diameter of the bubble.

The surface area has contributions from the bubble tip and side walls:

$$S = \left( C_6 \lambda^2 + C_5 \lambda h \right), \quad (3.17)$$

where  $C_5$  and  $C_6$  are unknown coefficients.  $C_5$  scales the perimeter of span-wise slices of the bubble while  $C_6$  rescales bubble tip.

To the first order, the diameter is half the wavelength:  $d \approx \lambda/2$ . However, the cylindrical bubbles do not always fill the span-wise domain. This can be seen by values of  $C_5$  that are below 4, the value corresponding to space-filling rectangular bubbles. Therefore, we adjust the diameter using  $C_5$ :

$$d = \frac{\lambda}{2} \frac{C_5}{4}. \quad (3.18)$$

The interface width is modeled by simple 1D diffusion:

$$\delta(t) = 2\sqrt{D(t+t_0)}, \quad (3.19)$$

where  $t_0$  is chosen to match  $\delta(0)$  to the initial condition.

We perform the integral through the bubble:

$$\begin{aligned} \int_{-d/2}^{d/2} (1 - |\phi_1(r)|) dr &= \frac{\delta}{\sqrt{\pi}} \left( 1 - \exp \left[ -\frac{d^2}{\delta^2} \right] \right) \\ &\quad + d \left( 1 - \operatorname{erf} \left[ \frac{d}{\delta} \right] \right). \end{aligned} \quad (3.20)$$

The mixed mass must still be connected to the dynamics equation via the Atwood number.

ber:

$$A = A_0 \left(1 - \frac{m}{V}\right), \quad (3.21)$$

where  $A_0$  is the pure Atwood number and  $V$  is the volume of the bubble. As in the dynamics equation, we define the volume as a mixture of  $\lambda^3$  and  $\lambda^2 h$  terms:

$$V = \left(C_8\lambda^3 + C_7\lambda^2 h\right), \quad (3.22)$$

where  $C_7$  and  $C_8$  scale the volume analogously to  $C_5$  and  $C_6$ .

The volume of mixed fluid,  $m(t)$ , can be measured directly in the simulations. This gives meaning to the value of  $m(t)$  independent of the ratio  $m(t)/V$ . Therefore, unlike in the dynamics, where a coefficient could be discarded without loss of generality, all four of  $C_5, C_6, C_7$  and  $C_8$  are necessary. The overall scale factor cannot be removed if we want to compare to mixed volume measurements.

### 3.3.3 Coefficient constraints

First, consider the limit where  $D = 0$ ,  $\nu = 0$ , and  $h \rightarrow 0$ . The dynamical equation becomes

$$\ddot{h} = \frac{Ag}{C_4\lambda}h, \quad (3.23)$$

which matches Rayleigh's original linear stability analysis if

$$C_4 = 1/(2\pi). \quad (3.24)$$

When  $\nu > 0$ , the growth rate  $\ddot{h}$  is given by Duff's linear theory:

$$\ddot{h} = \left(\sqrt{Agk + \nu^2 k^4} - \nu k^2\right)^2 h, \quad (3.25)$$

where  $k = 2\pi/\lambda$  is the wavenumber. Setting this equal to Equation 3.23 yields:

$$C_4 = \frac{1 + 2x \left( \sqrt{1+x^2} + x \right)}{2\pi}, \quad (3.26)$$

where,

$$x = \sqrt{\frac{8\pi^3\nu^2}{Ag\lambda^3}} = \sqrt{\frac{(2\pi)^3}{\text{Gr}}}, \quad (3.27)$$

and  $\text{Gr}$  is the Grashof number.

Next, consider the initial quantity of mixed mass for small sharp interfaces,  $\delta(0), a_0 \rightarrow 0$ .

We assume the initial condition is an error function profile:

$$M(t=0) = \lambda^2 \int_{-\infty}^{\infty} \text{erf} \left[ \frac{z}{\delta} \right] = \frac{2\lambda^2\delta}{\sqrt{\pi}}. \quad (3.28)$$

Equating this to the product of Equation 3.17 and Equation 3.20:

$$\frac{2\lambda^2\delta}{\sqrt{\pi}} = \frac{C_6\lambda^2\delta}{\sqrt{\pi}}, \quad (3.29)$$

which implies that  $C_6 = 2$ .

Next, consider the limit when  $\delta \rightarrow 0$  and  $h \rightarrow 0$ . In the linear theory, the Atwood number is rescaled:

$$A = \frac{A_0}{1 + \pi^{-1/2}k\delta} = A_0 \left( 1 - \frac{k\delta}{\sqrt{\pi} + k\delta} \right). \quad (3.30)$$

We equate this to Equation 3.21:

$$\frac{2\pi\delta}{\lambda(\sqrt{\pi} + 2\pi\delta/\lambda)} = \frac{C_6}{C_8} \frac{\delta}{\lambda\sqrt{\pi}}, \quad (3.31)$$

or

$$\frac{C_6}{C_8} = \frac{2\pi\sqrt{\pi}}{\sqrt{\pi} + 2\pi\delta/\lambda}. \quad (3.32)$$

The variable  $\delta$  is associated with the mixing model, not the dynamics, so it would be convenient to have  $C_8$  independent of  $\delta$ . We've defined  $C_6 = 2$  in the limit of  $\delta(0), a_0 \rightarrow 0$ , so we can add a term that goes to zero at  $\delta = 0$ :

$$C_6 = \frac{2}{1 + 2\sqrt{\pi}\delta/\lambda}, \quad (3.33)$$

which constrains  $C_8$  to be:

$$C_8 = \frac{1}{2\pi}, \quad (3.34)$$

which is the same as  $C_4$  in the inviscid case.

### 3.3.4 Coefficient estimation

The parameter  $C_1$  scales the form drag and serves as a drag coefficient. Because we have let  $C_0 = 1$ , the force balance is really aggregated over two rising bubbles and two falling spikes, each with diameter  $\lambda/2$ . Therefore, we multiply the force on a single bubble of diameter  $\lambda/2$  by 4. Now, we relate  $C_1$  to the drag coefficient  $C_d$  in the drag equation:

$$C_1 \lambda^2 \dot{h}^2 = 2C_d \mathcal{A} \dot{h}^2, \quad (3.35)$$

where  $\mathcal{A}$  is the cross sectional area, so  $C_1$  can be estimated using drag coefficients of similar objects:

$$C_1 = 2C_d \frac{\mathcal{A}}{\lambda^2}, \quad (3.36)$$

where  $\mathcal{A} \approx (\lambda/2)^2$ . Initially, the bubble tip is a flat plate, which has  $C_d = 1.28$ . At late times, the bubble is closer to an elongated cylinder, which has  $C_d = 0.82$ , but with a somewhat streamlined tip, which further reduces drag. We expect  $C_1 \approx 0.64$ , but possibly much smaller if the bubble takes a streamlined shape. However, if the bubble spreads to have a diameter greater than  $\lambda/2$ ,  $C_1$  could be greater than 0.64.

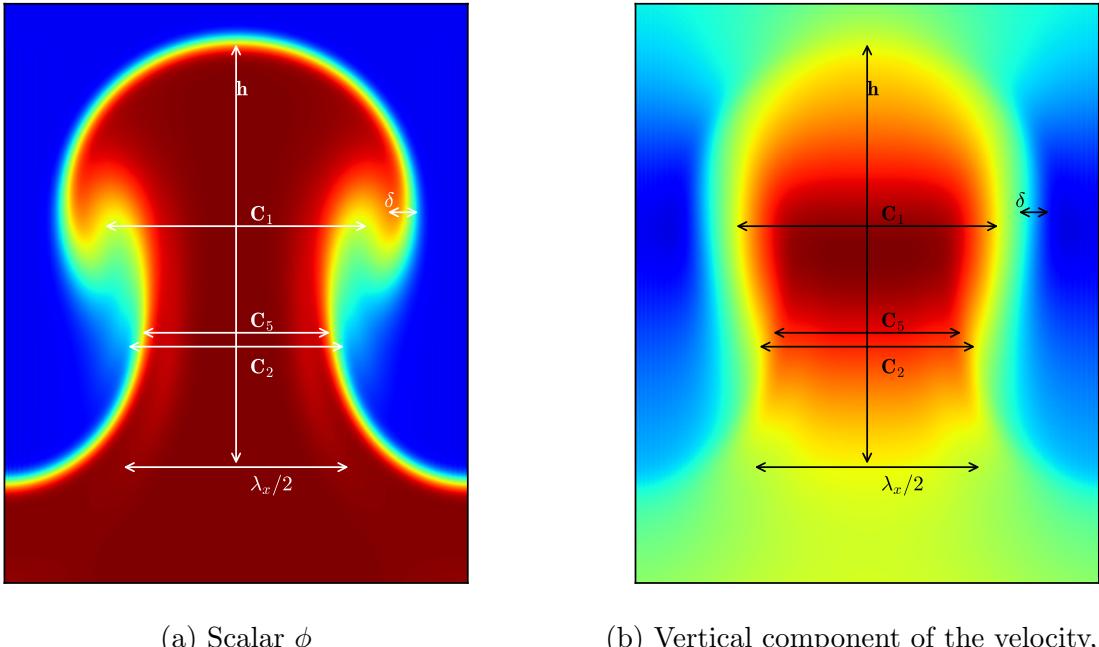


Figure 3.1: Slices of the scalar and vertical component of the velocity at early times and high Grashof number. The arrows indicate the dependence of the model terms on different span-wise length scales, and are identical in both figures.  $C_1$  is related to the maximum cross sectional diameter of the bubble in the velocity field.  $C_2$  is related to the nominal side-wall diameter of the bubble in the velocity field.  $C_5$  is related to the nominal side-wall diameter of the bubble in the scalar field.  $\delta$  is related to the interface thickness in the scalar field. The slice is in the plane  $x = y$ , which passes through only bubble centers.

Next, consider the limit when  $h \rightarrow \infty$  and  $D = 0$ : The dynamical equation becomes

$$\ddot{h} = \frac{A_0 g - C_2 \nu (1/\lambda) \dot{h}}{C_3}, \quad (3.37)$$

which leads to a terminal velocity of

$$\dot{h} = \frac{A_0 g \lambda^2}{C_2 \nu}, \quad (3.38)$$

or a nondimensional velocity, i.e., Froude number,

$$\text{Fr} = \frac{dz}{d\tau} = \frac{\sqrt{\text{Gr}}}{C_2}. \quad (3.39)$$

The case of extended bubbles and spikes affected only by viscous drag is highly analogous to flow through a square duct. The pressure drop,  $\Delta p$ , along a duct is given by the Darcy-Weisbach formula:

$$\Delta p = \frac{f_D}{2} \frac{v^2 L}{d}, \quad (3.40)$$

where  $L$  is the length of the duct,  $v$  is the mean velocity,  $d$  is the hydraulic diameter, and  $f_D$  is the Darcy friction factor. In our case,  $L = h$ ,  $v = \dot{h}$ , and  $\Delta p = A_0 g h$ , so

$$A g = \frac{f_D}{2} \frac{\dot{h}^2}{d}. \quad (3.41)$$

For laminar flows in circular pipes,  $f_D = \bar{f}_D = 64/\text{Re}$ , so

$$A g = \frac{f_D}{\bar{f}_D} 32 \nu \frac{\dot{h}}{d^2}. \quad (3.42)$$

The hydraulic diameter  $d = \lambda/2$ , so

$$\dot{h} = \frac{\bar{f}_D}{f_D} \frac{A_0 g \lambda^2}{128 \nu}. \quad (3.43)$$

This gives an estimate for  $C_2$ :

$$C_2 \approx 128 \frac{f_D}{\bar{f}_D}, \quad (3.44)$$

where the ratio  $f_D/\bar{f}_D$  is affected by the geometry and departure from laminar flow. For example, for square ducts  $f_D/\bar{f}_D \approx 0.889$ , so  $C_2 \approx 114$  [17].

The product of the coefficient  $C_5$  and  $\lambda h$  gives the interfacial area of the side of the bubble. Therefore,  $C_5$  captures information both about the bubble shape and the bubble diameter. If the bubbles had diameter  $\lambda/2$  and were smooth and rectangular, then  $C_5 \approx 4$ . If the bubble has a lower surface area shape, e.g., cylindrical, or is thinner, then  $C_5 < 4$ .

These diameters, along with the relevant span-wise length scales in the preceding coefficients, are sketched in Figure 3.1. The diameter associated with  $C_1$  is defined as the entrainment width at the bubble tip, in contrast to the width at the bubble center used for  $C_2$ . The diameter associated with  $C_5$  is defined similarly to  $C_2$ , but with respect to the scalar interface. The interface width  $\delta$  also depends on the scalar representation of the bubble.

### 3.4 Numerical experiments

To evaluate the simple model of Section 3.3, we conduct a battery of direct numerical simulations. The novel components of the simple model, i.e., viscous drag and mixing, are most pronounced at late times. We primarily direct our effort at simulating higher aspect ratio domains to allow the bubble to reach a dissipative flow.

The numerical experiments simulate the incompressible Navier-Stokes equations with the Boussinesq approximation:

$$\frac{\partial u}{\partial t} + u \cdot \nabla u = \nu \nabla^2 u - \nabla P + Ag\phi, \quad (3.45)$$

$$\frac{\partial \phi}{\partial t} + u \cdot \nabla \phi = D \nabla^2 \phi, \quad (3.46)$$

where  $u$  is the velocity,  $\nu$  is the kinematic viscosity,  $P$  is the pressure,  $\phi$  the nondimensional density, and  $D$  is the diffusivity of  $\phi$ .

The initial conditions are quiescent, with a horizontal interface perturbed by a product of cosine functions and smeared by an error function:

$$\phi(x, y, z, t = 0) = \operatorname{erf} \left( \frac{z + a_0 \cos(2\pi(x/\lambda)) \cos(2\pi(y/\lambda))}{\delta} \right), \quad (3.47)$$

where  $a_0$  is the initial amplitude and  $\delta$  is the initial interface thickness. Both  $a_0$  and  $\delta$  are taken to be small enough to minimize their effects on the solution, 0.01 and 1/128, respectively. The governing equations and initial condition have four dimensional parameters:  $\nu$ ,  $D$ ,  $Ag$ ,  $\lambda$ . These are combined into 2 nondimensional numbers, the Grashof number and the Schmidt number:

$$\operatorname{Gr} = \frac{A_0 g \lambda^3}{\nu^2}, \quad \operatorname{Sc} = \frac{\nu}{D}. \quad (3.48)$$

The Grashof number serves the role of a Reynolds number for instability problems without a consistent characteristic velocity. For this reason, the root of the Grashof number is sometimes called the *perturbation Reynolds number* [61]:

$$\operatorname{Re}_p = \sqrt{\frac{A_0 g \lambda^3}{\nu^2}}. \quad (3.49)$$

The domain is  $[0.5, 0.5, 64]$  and rotated 45 degrees in the span-wise plane to model  $\lambda = \sqrt{2}$ , transforming the initial condition to:

$$\phi(x, y, z, t = 0) = \operatorname{erf} \left( \frac{z + a_0 \cos(\pi(x + y)) \cos(\pi(x - y))}{\delta} \right). \quad (3.50)$$

This is done so the span-wise boundaries at  $x = \{0, 0.5\}$  and  $y = \{0, 0.5\}$  are symmetric. The length of the domain is  $64/\sqrt{2} \approx 45.2$  wavelengths with no-slip walls at the top and bottom. Based on a previous validation of the smRTI with no-slip boundaries, we expect

the bubble to be unaffected by the top and bottom walls until it reaches 75% of the height, or about  $17\lambda$ . This provides significantly more data than the  $h < 4\lambda$  results of Ramaprabhu et al. [47].

The model introduced in Section 3.3 assumes the bubbles and spikes are coherent structures, that is they travel at some velocity and have a well defined interface. As the Grashof number increases and the bubbles and spikes break up, they depart from the assumptions of the model. On the other hand, at low Grashof number and finite diffusivity, diffusion moves the  $\phi = 0$  interface, as opposed to simply transporting the scalar across it, which also departs from the model assumptions. For these reasons, we restrict our study to an intermediate range of Grashof numbers: those which are large enough to sustain bubble dynamics while not being so large as to break the bubbles apart. This range has been identified empirically to be approximately  $6 \times 10^2 \leq \text{Gr} \leq 6 \times 10^5$  for Schmidt numbers greater than 1.

The number of spatial samples needed to resolve the advection-diffusion equation for the scalar goes with the Peclet number to the third power. It is prohibitively expensive to perform calculations at high Schmidt numbers and high Grashof numbers.

Simulations are performed with the NekBox version [26] of the Nek5000 code [16], which has been previously validated against single-mode Rayleigh-Taylor experiments [24, 62]. The spectral element method implemented by NekBox has purely dispersive errors and converges exponentially with the spectral order [10]. The resolution parameters, the number of spectral elements, the order of the spectral elements, and the time step were chosen to achieve an accuracy of  $10^{-4}$  in the bubble aspect ratio [27].

### 3.4.1 Observables

**Bubble height** For miscible RTI, the shape of the scalar field is due to a combination of advection in the bubble and diffusion across the interface. We can assume an error function-like profile across the interface at the bubble tip, but diffusion across the bubble sidewalls

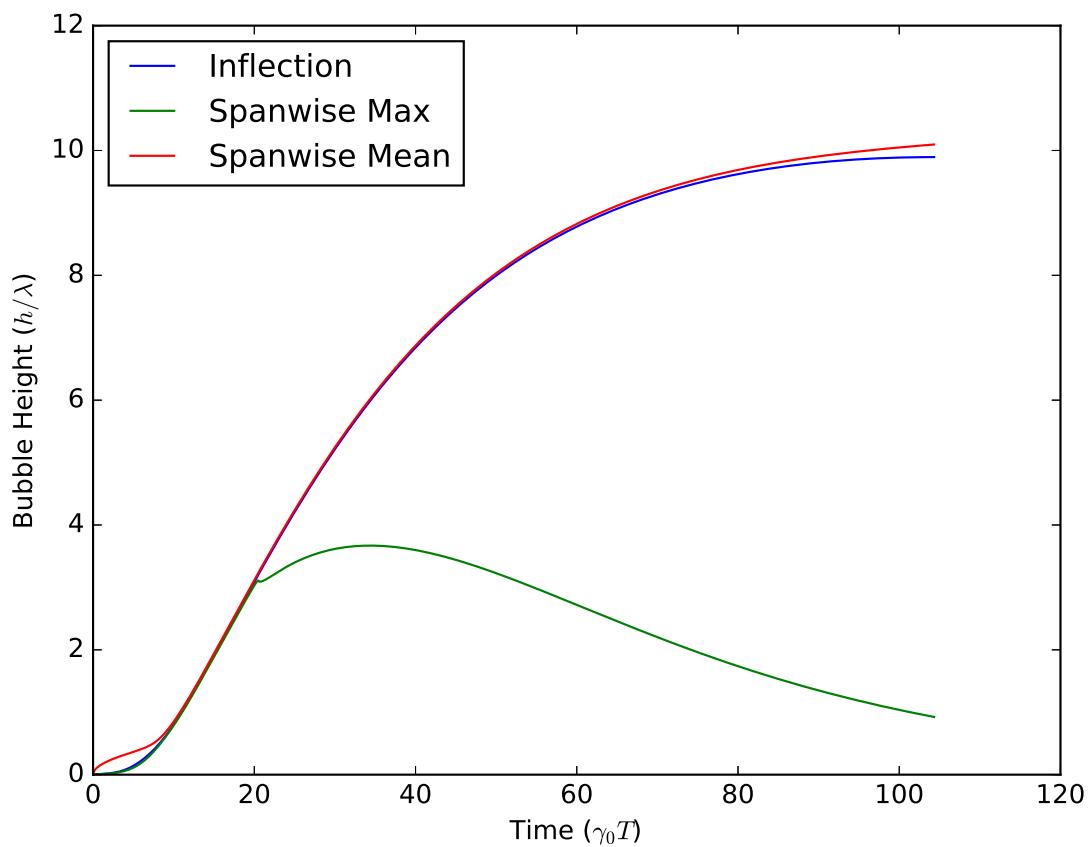


Figure 3.2: Comparison of height metrics at  $\text{Gr} = 4.8 \times 10^4$  and  $\text{Sc} = 1$ .

results in a linear profile in both the span-wise mean and maximum of the scalar. To separate the definition of the bubble tip from sidewall mixing, which is incorporated by the decreasing effective Atwood number, we introduce a new definition: the bubble tip is defined as the inflection point in the span-wise maximum scalar profile. For the symmetric case, this span-wise maximum of the scalar is equivalent to the value along the bubble axis. While mixing leads to a linear decay behind the bubble tip, the profile remains sharp near the bubble tip, decoupling the position of the inflection point from the sidewall mixing.

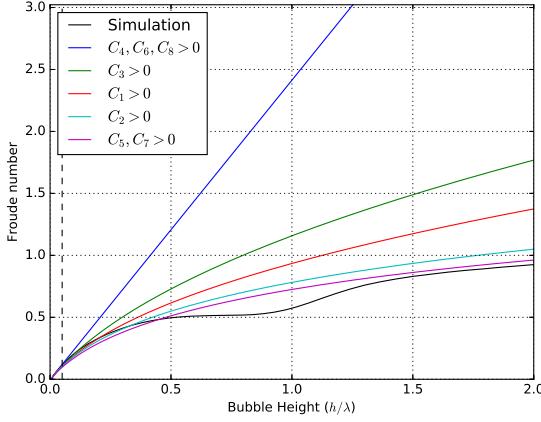
This definition of the bubble height is compared to two more traditional definitions, based on a cutoff in the mean or maximum profiles, in Figure 3.2. At early times, the definition based on the mean profile grows diffusively. At late times, the definition based on the max profile kinks as the linear part of the profile crosses zero and then stagnates. The definition based on the inflection avoids both breakdowns while agreeing with the two traditional definitions within each of their valid ranges.

**Mixed volume** The scalar is normalized such that  $\phi \in [-1, 1]$  and the average  $\bar{\phi} = 0$ . The purity of the fluid is therefore  $|\phi|$  and the volume of mixed fluid is given by a simple integral:

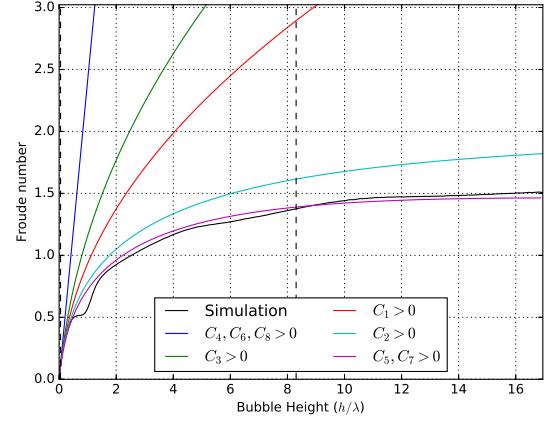
$$M(t) = \int (1 - |\phi(x, y, z)|) dV \quad (3.51)$$

### 3.5 Growth stages through late times

Single mode experiments have been limited to bubble heights of  $1.8\lambda$  [62]. Simulations have reached in  $4\lambda$  in 3D [47] and  $9\lambda$  in 2D [61]. Here, we present trajectories that continue up to  $17\lambda$ , for example Figure 3.3. However, the more dissipative bubbles stop rising at lower aspect ratios, for example Figure 3.4. Plotted with the nondimensional bubble height,  $h/\lambda$ , and nondimensional mixed volume,  $M/\lambda^3$  is the nondimensional velocity, or Froude number,

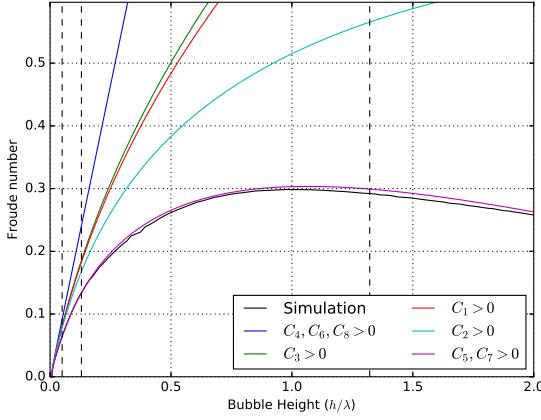


(a) Early times

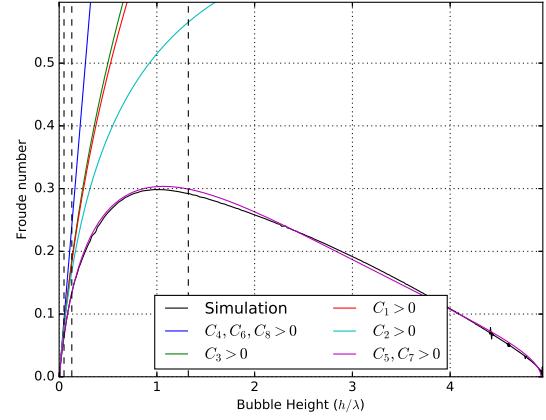


(b) Late times

Figure 3.3: Bubble Froude number vs. nondimensional bubble height for  $\text{Ra} = 10^{5.75}$ ,  $\text{Sc} = 4$ , simulation vs. model with successive terms enabled: first  $C_4, C_6$  and  $C_8$ , then  $C_3$ , then  $C_1$ , then  $C_2$ , and finally  $C_5$  and  $C_7$ . The dashed vertical lines divide the trajectory into three regimes: linear growth for  $H/\lambda < 0.05$ , saturation until  $H/\lambda \approx 8$ , and viscosity beyond that. The stagnation and re-acceleration transients are seen beginning at  $H/\lambda = 0.5$  and ending by  $H/\lambda = 1.5$ .



(a) Early times



(b) Late times

Figure 3.4: Bubble Froude number vs. nondimensional bubble height for  $\text{Ra} = 10^{4.5}$ ,  $\text{Sc} = 8$ , simulation vs. model with successive terms enabled: first  $C_4, C_6$  and  $C_8$ , then  $C_3$ , then  $C_1$ , then  $C_2$ , and finally  $C_5$  and  $C_7$ . Dashed vertical lines divide the trajectory into four regimes: linear growth for  $H/\lambda < 0.05$ , saturation until  $H/\lambda \approx 0.1$ , viscosity until  $H/\lambda \approx 1.3$ , and diffusion beyond that. The stagnation and re-acceleration transients, seen in Figure 3.3, are suppressed by the viscous regime, which onsets before the transient begins at  $H/\lambda = 0.5$ .

which is defined as:

$$\text{Fr} = \frac{v}{\sqrt{Ag\lambda}}. \quad (3.52)$$

The dissipation in the bubble is characterized by the Rayleigh number:

$$\text{Ra} = \text{GrSc} = \frac{A_0 g \lambda^3}{\nu D}. \quad (3.53)$$

In the spirit of recent analyses [47, 61], we try to identify distinct growth regimes. To do so, we consider the behavior of the simple model with coefficients set to zero. Initially, only  $C_4$ ,  $C_6$ , and  $C_8$  affect bubble dynamics. The growth is exponential, with a rate in agreement with the linear theory, so we term this the linear regime. As the bubble grows, the  $C_3$  term reduces the growth rate to a limiting value of  $Ag/C_3$ . Because this represents the transition from exponential growth to free-fall, we term it the saturation regime. The  $C_1$  term has the same qualitative effect as the  $C_3$  term, so it doesn't distinguish a unique regime. The  $C_2$  term does impose a limiting velocity scale, so it departs significantly from the  $C_3$  dynamics in the viscous regime. Finally, the  $C_5$  term, balanced by the  $C_7$  term, mixes the fluid and reduces the effective Atwood number in the diffusive regime. Because the relative onset of the viscous and diffusive regimes depends on the Schmidt number, they are grouped together into the dissipative regime. Ultimately, the bubble stops rising. The bubble height at this point, which is also the maximum bubble height, is called the penetration depth.

### 3.5.1 Exponential growth

When the amplitude is small and the interface is thin, the linear theory and numerical results identify exponential growth:  $\ddot{h} = \gamma^2 h$ . Similarly, in the limit  $h, \delta \rightarrow 0$ , the simple model yields a growth rate:

$$\gamma = \sqrt{\frac{A_0 g k}{2\pi C_4(1 + C_8/C_6 \pi^{1/2} k \delta)}}, \quad (3.54)$$

which differs from the linear theory by the absence of a  $-Dk^2$  term. This is the only term that can cause unstable interfaces to decay, i.e.,  $\gamma < 0$  for positive Atwood numbers. In the simple model, all unmixed bubbles grow while in the linear theory highly diffusive bubbles decay.

The other terms, i.e. , those scaled by  $C_1, C_2, C_3, C_5$  and  $C_7$ , can be omitted while the description of the exponential growth is unaffected.

### 3.5.2 Saturation regime

The exponential growth saturates as the bubble height increases. In the simple model, this captured by the  $C_3$  term:

$$\ddot{h} = \frac{Ag}{(C_3 h + C_4 \lambda)(1 + C_8/C_6 \pi^{1/2} k \delta)}, \quad (3.55)$$

which becomes significant when  $h \approx C_4 \lambda / C_3$ . We will find in the following section that  $C_3$  takes a value of about 1, so, omitting viscous corrections, saturation halves the growth rate at  $h \approx \lambda/(2\pi)$ .

The definition of the start of the saturation regime is somewhat arbitrary. Here, we propose the definition  $h/\lambda < 0.05$  as it is where the exponential and saturated plots of the Froude number versus height visually deviate. The important thing is that this threshold is independent of  $Ag$ , the Grashof number, and the Rayleigh number. Within the saturation regime, the acceleration takes a limiting value of  $Ag$ , defining a saturation velocity:

$$v_s \sim \sqrt{Agh}. \quad (3.56)$$

### 3.5.3 Viscous regime

As the bubble grows, so does its surface area. The viscous drag, which scales linearly with the bubble height and bubble velocity, ultimately balances the buoyancy to limit the velocity, as in Equation 3.38. The scaling of the onset of this regime can be found by equating the saturation velocity and the viscous velocity:

$$\sqrt{Ag h} \sim \frac{Ag\lambda^2}{\nu}. \quad (3.57)$$

Solving for  $h/\lambda$  yields:

$$\frac{h_\nu}{\lambda} \sim \frac{Ag\lambda^3}{\nu^2} = \text{Gr}, \quad (3.58)$$

where  $h_\nu$  is the onset of the viscous regime.

As in the saturation case, the particular definition of the onset is arbitrary. Here, we will define  $h_\nu$  such that when  $h_\nu/\lambda = 0.5$ , the viscous velocity has the potential flow value:  $v_\nu = \pi^{-1/2} \sqrt{Ag\lambda}$ . This works out to be:

$$\frac{h_\nu}{\lambda} = \frac{\pi}{4C_2^2} \text{Gr} \approx 4.7 \times 10^{-5} \text{Gr}, \quad (3.59)$$

where we've let  $C_2 = 128$ , the nominal value from Poiseuille flow.

### 3.5.4 Diffusive regime

The rate of diffusion across the surface of the bubble also scales with the bubble height. When the viscous drag limits the bubble's velocity, the flux of pure fluid into the bubble, which goes as the velocity, is unable to match the flux of mixed fluid through the interface. Mixing dilutes the buoyant fluid, reducing the effective Atwood number and therefore the bubble velocity. Ultimately, the effective Atwood number reaches zero and eventually the bubble stops rising.

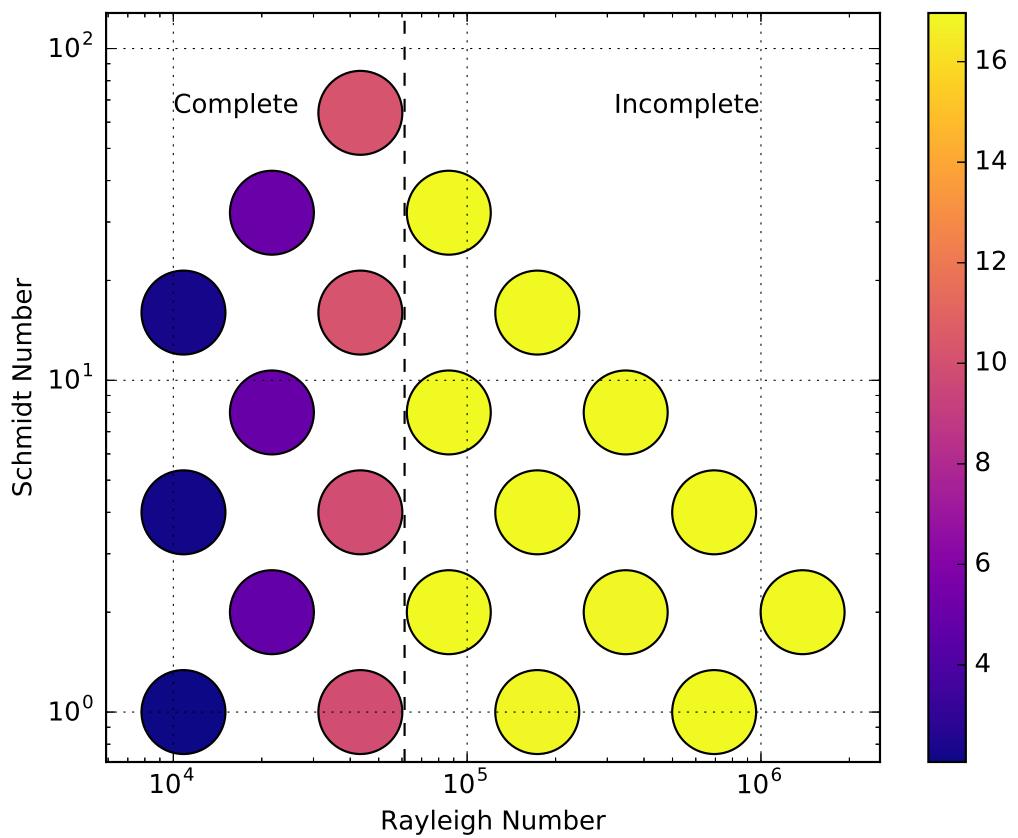


Figure 3.5: Penetration depth, nondimensionalized, vs. the Rayleigh and Schmidt numbers. The dashed line separates completed from incomplete trajectories, which are clipped at  $h/\lambda \approx 23$ .

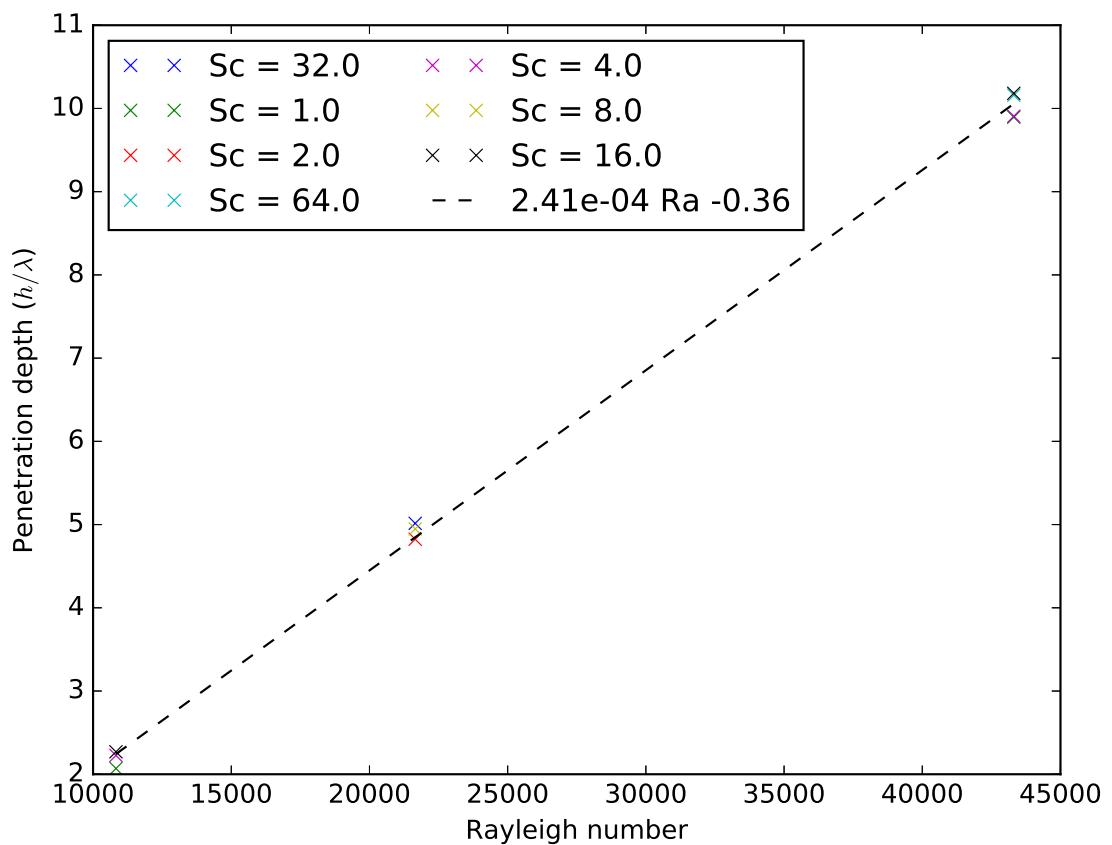


Figure 3.6: Penetration depth, nondimensionalized, vs. the Rayleigh. The dashed line is a best fit line  $h/\lambda = 2.41 \times 10^{-4}Ra - 0.36$ .

The penetration depth is the maximum height of the bubble, which is the height of the bubble at the stopping condition of this analysis. We can estimate the scaling of the penetration depth as the product of a characteristic velocity with a characteristic time-scale. The late-time velocity is the viscous velocity,  $v_\nu$ , while the time scale is given by diffusion:

$$\tau_D = \frac{\lambda^2}{D}. \quad (3.60)$$

Combining and nondimensionalizing yields:

$$\frac{h(\infty)}{\lambda} = \frac{Ag\lambda^3}{\nu D} = \text{Ra}, \quad (3.61)$$

so the penetration depth should go linearly with the Rayleigh number.

The penetration depth is plotted as a function of Rayleigh and Schmidt number in Figure 3.5, and, indeed, depends strongly on the Rayleigh number before being clipped by the top walls. Furthermore, the relationship to the Rayleigh number is linear over the cases shown here, as shown in Figure 3.6.

We can use a similar analysis to define the onset of the diffusive regime when the interface width is a quarter wavelength,  $\delta = \lambda/8$ , which is a quarter of the nominal bubble diameter. This results in a bubble height:

$$\frac{h_D}{\lambda} = \frac{1}{128C_2} \text{Ra} \approx 6.1 \times 10^{-5} \text{Ra}, \quad (3.62)$$

where we've again let  $C_2 = 128$ .

The ratio of the diffusive to viscous heights, Equation 3.62 and Equation 3.59, respectively, is linear with the Schmidt number, or about  $1.3\text{Sc}$ . For  $\text{Sc} > 1$ , the portion of the trajectory that is governed by viscosity alone increases with the Schmidt number. However, for  $\text{Sc} < 1$  the diffusive regime dominates the viscous regime entirely.

Gr	Ra	Sc	$C_1$	$C_2$	$C_3$	$C_7$	$C_5$	$E_d$	$E_m$
$6.9 \cdot 10^5$	$1.4 \cdot 10^6$	2	0.46	136.3	1.00	1.05	3.18	0.74%	2.04%
$6.9 \cdot 10^5$	$6.9 \cdot 10^5$	1	0.45	146.7	1.00	1.14	2.48	0.49%	1.33%
$1.7 \cdot 10^5$	$6.9 \cdot 10^5$	4	0.47	201.9	1.00	3.16	3.29	0.43%	1.02%
$1.7 \cdot 10^5$	$3.5 \cdot 10^5$	2	0.41	197.3	1.00	2.48	2.70	0.34%	0.55%
$1.7 \cdot 10^5$	$1.7 \cdot 10^5$	1	0.38	179.8	1.00	1.73	2.39	0.30%	0.24%
$4.3 \cdot 10^4$	$3.5 \cdot 10^5$	8	0.30	161.1	1.00	2.49	3.86	0.29%	0.24%
$4.3 \cdot 10^4$	$1.7 \cdot 10^5$	4	0.28	147.9	1.00	2.01	3.14	0.26%	0.18%
$4.3 \cdot 10^4$	$8.7 \cdot 10^4$	2	0.39	99.0	1.00	1.28	2.57	0.25%	1.42%
$4.3 \cdot 10^4$	$4.3 \cdot 10^4$	1	0.00	96.5	1.54	1.02	2.27	0.14%	1.22%
$1.1 \cdot 10^4$	$1.7 \cdot 10^5$	16	0.00	158.8	1.00	2.65	3.63	0.15%	1.54%
$1.1 \cdot 10^4$	$8.7 \cdot 10^4$	8	0.46	111.0	1.00	1.44	2.72	0.37%	3.46%
$1.1 \cdot 10^4$	$4.3 \cdot 10^4$	4	0.54	79.4	1.00	1.10	2.40	0.33%	3.68%
$1.1 \cdot 10^4$	$2.2 \cdot 10^4$	2	0.00	82.4	2.03	1.04	2.36	0.09%	1.50%
$1.1 \cdot 10^4$	$1.1 \cdot 10^4$	1	1.05	119.2	2.94	1.00	2.26	2.87%	1.24%
$2.7 \cdot 10^3$	$8.7 \cdot 10^4$	32	1.22	85.0	1.02	1.42	2.83	0.38%	2.77%
$2.7 \cdot 10^3$	$4.3 \cdot 10^4$	16	0.56	72.7	1.20	1.08	2.44	0.75%	3.67%
$2.7 \cdot 10^3$	$2.2 \cdot 10^4$	8	0.19	69.9	2.41	1.07	2.45	0.25%	2.32%
$2.7 \cdot 10^3$	$1.1 \cdot 10^4$	4	0.91	68.2	3.23	1.00	2.40	1.12%	1.12%
$6.8 \cdot 10^2$	$4.3 \cdot 10^4$	64	1.63	65.7	2.44	1.12	2.57	0.40%	2.91%
$6.8 \cdot 10^2$	$2.2 \cdot 10^4$	32	0.52	66.9	3.94	1.07	2.49	0.55%	2.46%
$6.8 \cdot 10^2$	$1.1 \cdot 10^4$	16	0.00	69.5	9.10	1.00	2.47	0.22%	0.95%

Table 3.1: Simulation conditions, fit coefficients, and relative errors.

### 3.5.5 Stagnation and re-acceleration

Absent from the previous discussion is the stagnation and re-acceleration of the bubble around  $h/\lambda = 1$ , as seen in Figure 3.3. There are no terms in the buoyancy-drag model capable of producing an inflection point in the Froude number vs. bubble height, so stagnation and re-acceleration cannot be controlled by turning a model coefficient on or off. This suggests that the buoyancy-drag model is missing a term.

However, the buoyancy-drag model does have a limiting velocity, the viscous velocity  $v_\nu$ . When the viscous velocity is near or below the stagnation velocity,  $\text{Fr} \approx \pi^{-1/2}$ , saturation and re-acceleration are suppressed. Otherwise, stagnation and re-acceleration temporarily interrupt the saturation regime. The stagnation and re-acceleration is a transient regime that only occurs at sufficiently high Grashof numbers.

## 3.6 Model fit coefficients

### 3.6.1 Fitting

The mixing model defines the quantity of mixed fluid,  $m(t)$ , as an analytic function:

$$\{H(t), \delta(0), \lambda, D, C_5\} \rightarrow m(t), \quad (3.63)$$

where  $H(t)$  is the bubble height,  $\delta(0)$  is the initial interface thickness,  $\lambda$  is the wavelength,  $D$  is the diffusivity, and  $C_5$  is a mixing coefficient. The values of  $H(t)$ ,  $\delta(0)$ ,  $\lambda$ , and  $D$  are taken from the numerical experiments, allowing for the definition of an mixing error:

$$E_m = \|m[C_5] - M(t)\|_2, \quad (3.64)$$

where  $M(t)$  is the reference value from the numerical experiments. To compute  $C_5$ , the error is minimized under the constraint  $C_5 > 0$ . The fitting problem is non-linear but 1D dimensional, so it can be solved with a sequential least squares minimizer, which finds local minima, wrapped with a basin hopping scheme, which samples across the local minima [58].

The dynamics model defines the bubble height,  $h(t)$ , as the solution to an ordinary differential equation. The dynamics model is integrated using the variable-coefficient ordinary differential equation (VODE) solver for stiff systems, creating a map from the dynamics coefficients to the height:

$$\{h(0), \delta(0), \lambda, \nu, D, \{C\}\} \rightarrow h(t), \quad (3.65)$$

where  $h(0)$  is the initial bubble height,  $\delta(0)$  is the initial interface thickness,  $\lambda$  is the wavelength,  $\nu$  is the kinematic viscosity,  $D$  is the diffusivity, and  $\{C\}$  are the model coefficients.  $\delta(0)$ ,  $\lambda$ ,  $\nu$ , and  $D$  are taken from the simulation while  $C_5$  is taken from independently fitting

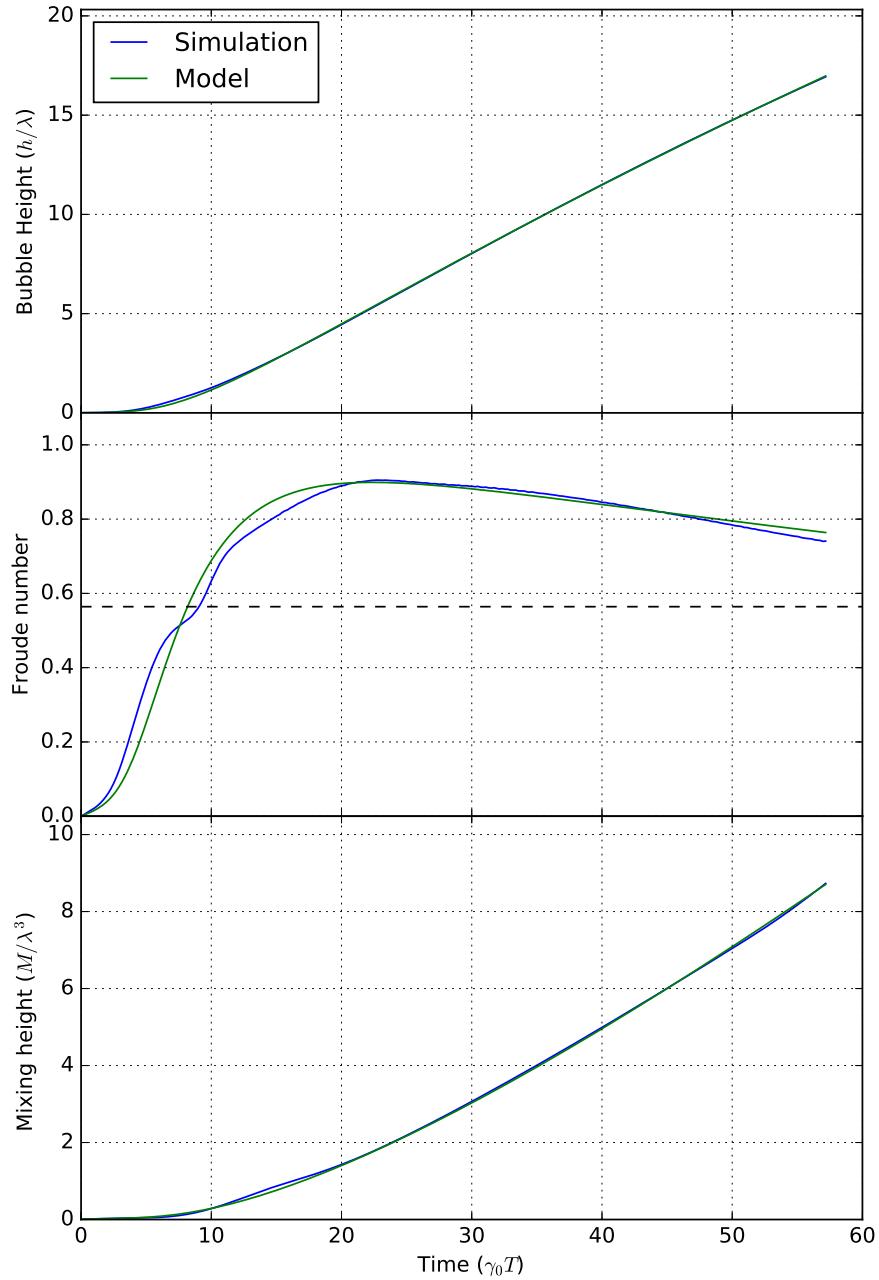


Figure 3.7: Example fit of the bubble height, top, and mixed volume, bottom, at  $\text{Ra} =$  and  $\text{Sc} = 8$ . The dashed line in the center plot is the stagnation velocity from potential flow models. The Froude number, which is a nondimensional velocity, is not directly fit but shows agreement between the model and simulation.

the mixing model, allowing for the definition of a dynamics error:

$$E_d = \|h[C_1, C_2, C_3, C_7] - H(t)\|_2, \quad (3.66)$$

where  $H(t)$  is the reference value from the numerical experiments. The non-linear global minimization problem is solved with the covariance matrix adaptation evolution strategy (CMA-ES). CMA-ES iteratively refines a sample distribution that evolves towards the global optimum. Computing the model error is very inexpensive compared to the simulations, so we choose a broad initial distribution with a large population size. The stochastic solution is polished with a sequential least-squares local minimization.

However, the high Rayleigh trajectories are incomplete, in that the trajectory is truncated when the bubble gets close to the top wall, leading to under-constrained systems. Therefore, we regularize the fit by adding a term to the model error:

$$R = \beta \left\| \frac{\mathbf{C} - \bar{\mathbf{C}}}{\bar{\mathbf{C}}} \right\|_2, \quad (3.67)$$

where  $\beta$  is the regularization parameter,  $\mathbf{C}$  is the vector of model coefficients, and  $\bar{\mathbf{C}}$  are the coefficient estimates. Although the buoyancy-drag model is non-linear, this L2 regularization can be thought of as a Tikhonov regularization [49] or ridge regression [38]. The regularization parameter is chosen to be an order smaller than the model error,  $\beta = 0.1E_d$ . If the regularized dynamic error ends up lower than the unregularized error, the unregularized problem must not have converged to the local minima. In those cases, the unregularized fit is repeated with the regularized coefficient values as a starting seed. Then the regularized fit is repeated with the updated definition. In this way, the two types of fits are iterated until consistency is reached.

The fitting process defines a mapping from the Grashof and Schmidt numbers to the

model coefficients and model errors:

$$(\text{Gr}, \text{Sc}) \rightarrow (C_1, C_2, C_3, C_5, C_7, E_m, E_d) \quad (3.68)$$

The rest of this section explores the relationships in this mapping.

### 3.6.2 Scope of the models

The proposed models aim to describe the mixing and dynamics in rising Rayleigh-Taylor bubbles and falling Rayleigh-Taylor spikes. The symmetry of the governing equations equates the spike behavior to that of the bubbles, so we will omit spikes from the following discussions. In highly viscous and diffusive cases, the bubbles may not reach late time highly non-linear dynamics. For this analysis, a bubble is considered to be covered by the model only if its height exceeds its wavelength before it stops rising. Experiments which do not meet that condition are discarded.

The bubble grows until mixing dilutes its buoyancy sufficiently for it to stop rising. After this point, it slowly recedes due to diffusing across the bubble tip. The model does not account for this diffusive effect, which moves the center of the interface rather than just broadened it, so bubble trajectories are clipped beyond the point at which the bubble velocity is zero. The height at that point in the trajectory is maximal and is called the *penetration depth*.

The penetration depth increases with Rayleigh number. For high Rayleigh number cases, the bubble continues to grow until it begins to interact with the top boundary, given the finite computational domain. Based on previous validation studies [24], we clip the trajectory when the bubble height reaches 75% of the domain height. Experiments in which this clipping occurs are incomplete. Those cases should be re-simulated with a larger computational domain, at greater computational cost, to collect trajectories which reach their penetration

depth.

However, there is still information in the incomplete experiments. In the following sections, incomplete experiments will be marked as such, but the trend in the complete experiments is often seen to continue smoothly into incomplete ones. Though not definitive, those suggest that the data present in the incomplete experiments is sufficient to constrain the corresponding characterization of the flow. Conversely, in some cases the behavior in the incomplete experiments departs from that in the complete ones. In those cases, it is difficult to differentiate between Rayleigh-dependent behavior and the side-effects of underconstrained fitting.

It should be noted that the computational cost of a complete trajectory goes as:

$$\text{Gr}^4 \text{Ra}^2 \max(1, \text{Sc}^4). \quad (3.69)$$

The fourth power of the Grashof and Schmidt numbers come from stability constraints: three from the spectral constraint and 1 from the explicit time-stepping constraint. The square of the Rayleigh number comes from the penetration depth, which both increases the length of the domain and the number of time-steps taken. For  $\text{Sc} \geq 1$ , the cost simplifies to  $\text{Ra}^6$ . For the runs in this study, the cheapest incomplete trajectories will cost  $64\times$  more than most expensive completed ones. The most expensive incomplete trajectory, at  $\text{Ra} \approx 1.4 \times 10^6$ , will cost over  $10^9\times$  more than cheapest completed one.

### 3.6.3 Accuracy of the models

The accuracy is characterized by the mixing and dynamics model errors relative to the maximum mix volume and bubble height, respectively. The relative model errors are plotted versus the Rayleigh and Schmidt numbers in Figure 3.8. In both models, the nominal error is less than 5%, but the two errors have differing structures.

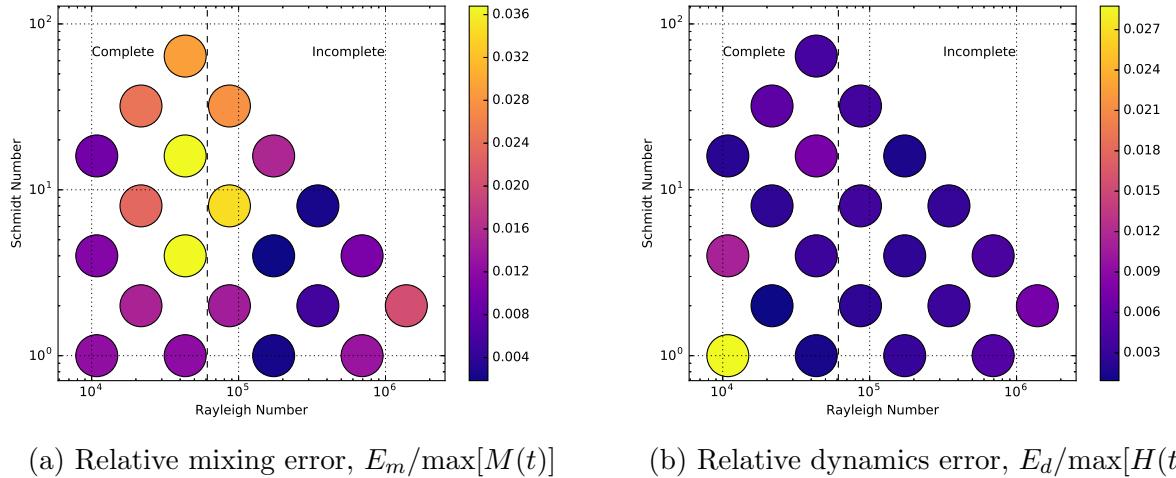


Figure 3.8: Relative model errors versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain.

The relative mixing error is about 3% for the complete experiments, with generally greater relative error at greater Rayleigh and Schmidt numbers. Among the incomplete experiments, relative error decreases with Rayleigh number before increasing again. This is likely due to the incompleteness; the greater error above  $\text{Ra} = 10^6$  suggests that the overall trend is increasing. Overall, the relationship between the accuracy of the mixing model, the Rayleigh, and Schmidt numbers is not fully determined.

The relative dynamics error is about 1%, with the exception of the unit Schmidt  $\text{Ra} \approx 10^4$  case. The mixing error for the outlying case is typical, so the error cannot be attributed to the treatment of mixing. This case will be considered more in the following sections. Outliers aside, the relative error decreases with Schmidt number and Rayleigh number, both for complete and incomplete trajectories. This indicates the dynamics model is most accurate, at least relative to the mixing height, when there is less mixing and drag. Another factor is re-acceleration, which adds some relatively constant error that is amortized more when the penetration depth is greater.

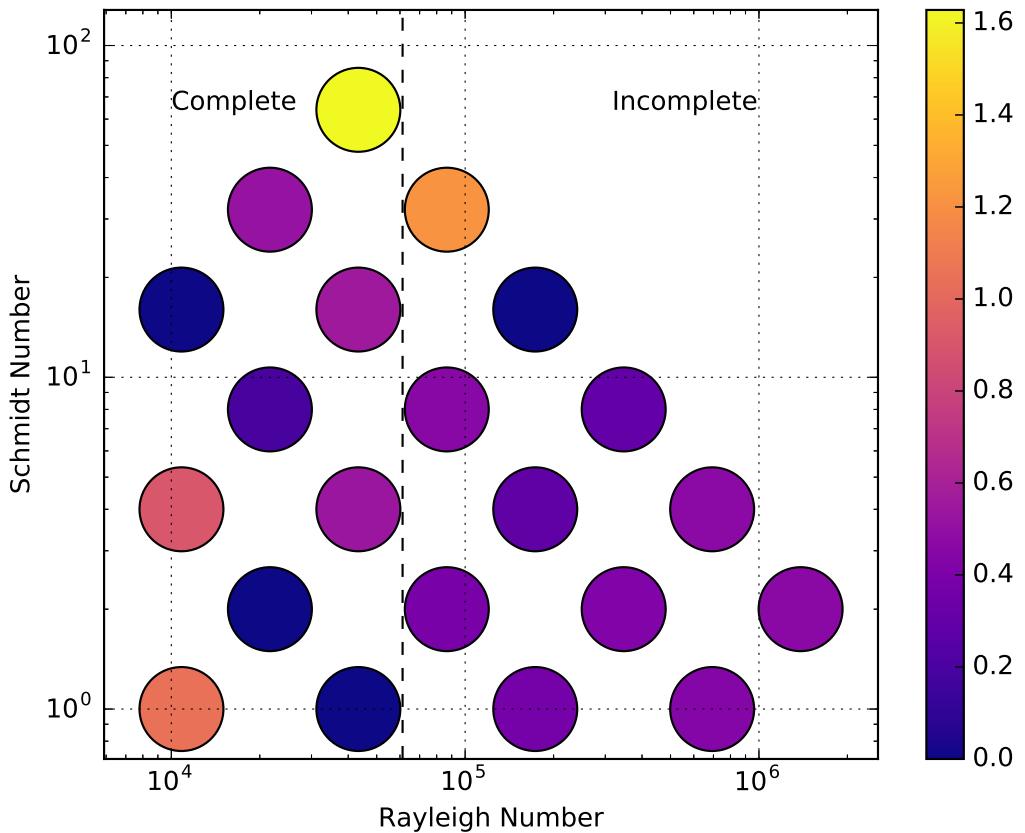


Figure 3.9: Best fit for  $C_1$  versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain.

### 3.6.4 Fit coefficients

The proposed model has 5 undetermined parameters. In each case, we estimate the value a priori by physical arguments. Then, the estimates are used as the starting point for fitting, i.e., minimization of the model error over the scope of the model.

## Form drag coefficient, $C_1$

The coefficient  $C_1$  is related to the drag coefficient  $C_d$  by Equation 3.36. Therefore we expect it to take a value around or less than 0.64, which corresponds to the drag coefficient of a flat plate and a surface area  $A = \lambda^2$ .

The  $C_1$  term is most influential early in the flow, so we expect it to be well constrained even in the incomplete trajectories. However, the  $C_1$  term is qualitatively redundant with the  $C_3$  inertial term, as they both bound the acceleration but not the velocity. Similarly, viscous drag is weak but still present at early times. It is possible that late-time effects that influence  $C_2$  could have an indirect effect on the value of  $C_1$ .

The fit values of  $C_1$  are plotted versus the Rayleigh and Schmidt numbers in Figure 3.9. The majority of trajectories are closely grouped between 0.3 and 0.6, which fits the drag coefficient rationale. However, there are outliers both at  $C_1 = 0$  and  $C_1 > 0.9$ . These outliers are troubling because each type occurs at both high and low Schmidt number and at low to moderate Rayleigh numbers. There is no clear pattern, but there are no outliers at the high Rayleigh number.

## Skin drag coefficient, $C_2$

The coefficient  $C_2$  scales the viscous drag. As with  $C_1$ ,  $C_2$  can be related to a standard measure of drag, in this case the Darcy friction factor, which provides an estimate of 113. The  $C_2$  term is linear with  $h$  and  $\dot{h}$ , so its influence is greatest at moderate to late times. Therefore, values of  $C_2$  taken from incomplete trajectories should be taken with a grain of salt.

The fit values of  $C_2$  are plotted versus the Rayleigh and Schmidt numbers in Figure 3.10. For the completed trajectories,  $C_2$  is about 60 and grows slightly with the Grashof number while being nearly independent of the diffusivity. At higher Rayleigh numbers the incomplete trajectories also show  $C_2$  growing with Grashof number, but the effect is much stronger.

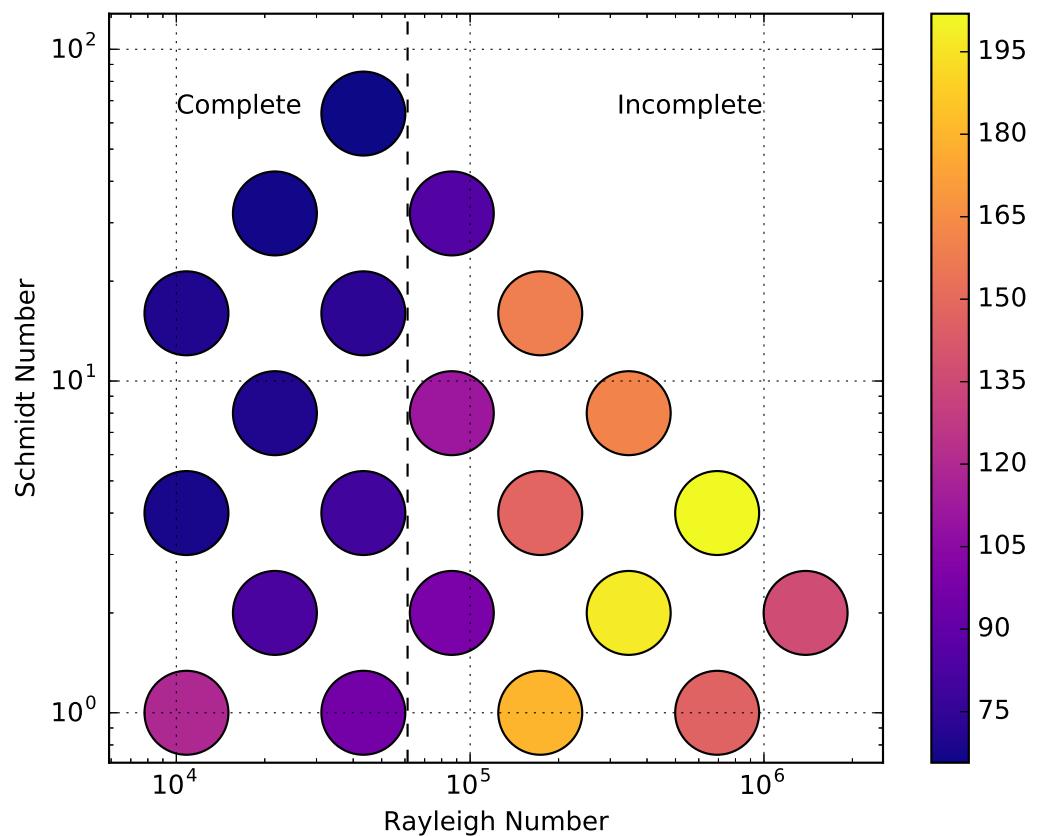


Figure 3.10: Best fit for  $C_2$  versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain.

There is weaker growth as the diffusivity decreases. Finally, at the highest Grashof numbers the  $C_2$  coefficient moderates again.

The presence of a positive relationship between  $C_2$  and the Grashof number in both the complete and incomplete trajectories suggests it is a real effect, though its strength is unclear. The relationship between  $C_2$  and the diffusivity is much weaker and could disappear as the high Rayleigh trajectories are completed.

One possible mechanism for increasing  $C_2$  with increasing Grashof number is the development of small amplitude Kelvin-Helmholtz structures on the bubble surface. These structures are suppressed at low Grashof number and grow more rapidly at higher Grashof number. They would enhance the transport of momentum across the bubble interface, thereby increasing the viscous drag coefficient.

### Inertial coefficient, $C_3$

The coefficient  $C_3$  gives the ratio of the inertial height to the buoyant height. For  $C_3 = 1$ , the maximum bubble acceleration is  $Ag$  while  $C_3 > 1$  represents the entrainment of neutrally or anti-buoyant fluid that contributes to the inertia but not the forcing. Mixing, which also reduces the ratio of the forcing to the inertia, is accounted for explicitly with the  $C_5$  coefficient and fit independently to the mixed volume observable, which prevents it from compensating for entrainment. The  $C_3$  term is linear with the height, so its influence is most pronounced at greater values of the height.

The fit values of  $C_3$  are plotted versus the Rayleigh and Schmidt numbers in Figure 3.11. The majority of trajectories have  $C_3 = 1$ , indicating that entrainment is not significant. For completed low Rayleigh high Schmidt flows  $C_3$  increases to a value of 9.1. However, the nominal value of 1 is recovered within the completed trajectories. Because the  $C_3$  term depends on the height, it is relatively underconstrained at lower Rayleigh numbers. If the model values of  $C_4$ ,  $C_6$ , or  $C_8$  are incorrect, the  $C_3$  term has the greatest ability to

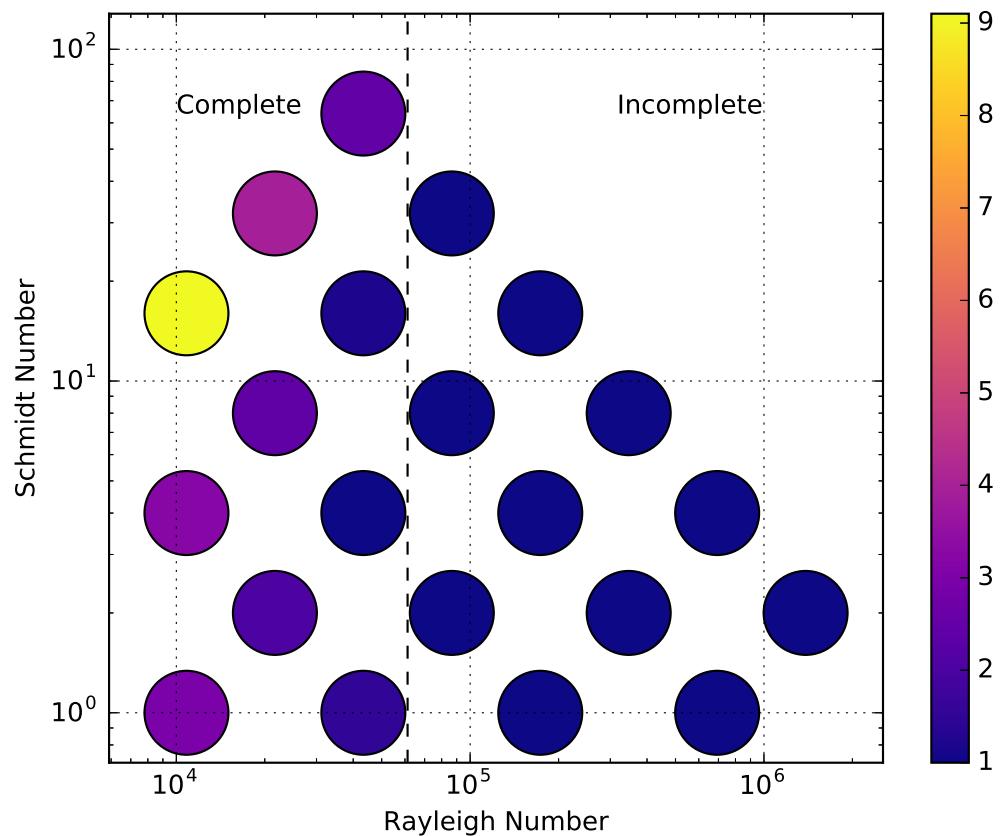


Figure 3.11: Best fit for  $C_3$  versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain.

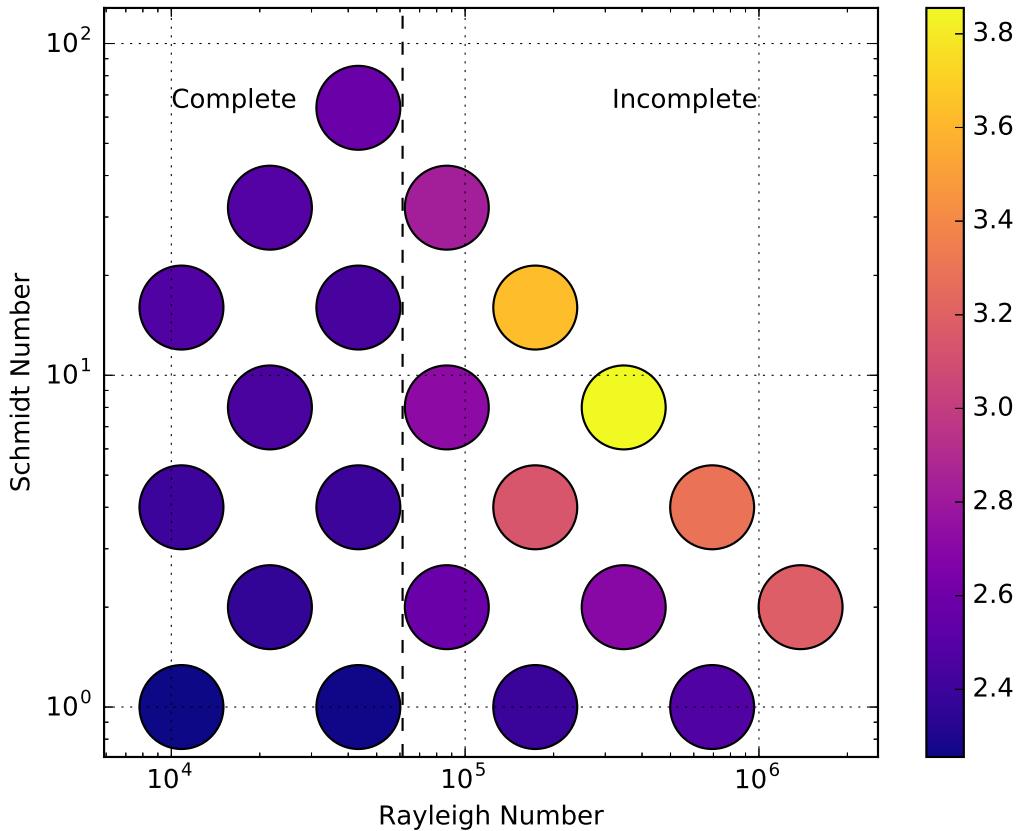


Figure 3.12: Best fit for  $C_5$  versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain.

compensate for the error when the height and velocity are small. However, the height is still small, so  $C_3$  would have to change significantly. The author believes that  $C_3 = 1$  is therefore the nominal value, but there is a breakdown in the model at low Rayleigh numbers that is influencing the fit of  $C_3$ . Identifying and correcting this model breakdown would be expected to recover  $C_3$  at low Rayleigh number.

## Interfacial area coefficient, $C_5$

The parameter  $C_5$  gives the ratio of the span-wise circumference of the scalar interface to the wavelength. If the bubble were rectangular in cross section with diameter  $\lambda/2$ , then  $C_5 = 4$ . If the bubble had a circular cross section with diameter  $\lambda/2$ , then  $C_5 \approx \pi$ . If the bubble diameter is less than a half-wavelength, then  $C_5 < \pi$ . The mixing width  $\delta$  increases with time, so the  $C_5$  term is most influential at late times. Therefore, the values of  $C_5$  in the incomplete trajectories should be taken with a grain of salt.

The fit values of  $C_5$  are plotted versus the Rayleigh and Schmidt numbers in Figure 3.12. Among the completed trajectories,  $C_5$  is a much stronger function of the Schmidt number than the Rayleigh number, increasing in both. The values are between 2 and 3, indicate a thinning of the bubble that decreases the mix rate by reducing surface area and total quantity of pure light fluid transported into the dense fluid.

The incomplete trajectories contain richer behavior, with a local maximum at  $\text{Ra} = 10^{5.5}$  and  $\text{Sc} = 8$ . In aggregate, higher Rayleigh number trajectories have increasing  $C_5$  with decreasing diffusivity. The dependence on the Grashof number is peaked at  $\text{Gr} \approx 4.3 \times 10^4$ . The increasing  $C_5$  with increasing Schmidt number in the completed trajectories is consistent with these two effects, which can be seen the smoothness of Figure 3.12.

A possible mechanism for increasing  $C_5$  with decreasing diffusivity is the development of structures on the scalar interface the increase the effective surface area. This is related to the Kelvin-Helmholtz structures proposed to explain increasing  $C_2$  with Grashof number, but with additional dependence on the diffusivity which can otherwise smear out the structures.

Another possible mechanism is a change in the bubble diameter. Higher Grashof number bubbles have thinner momentum boundaries, allowing more buoyant fluid to flow freely through the stem of the bubble. This sustains the bubble diameter at late times, while more viscous bubbles can thin. This would explain an increase of  $C_5$  with the Grashof number.

The author has no mechanism by which to explain the local maximum value and expect

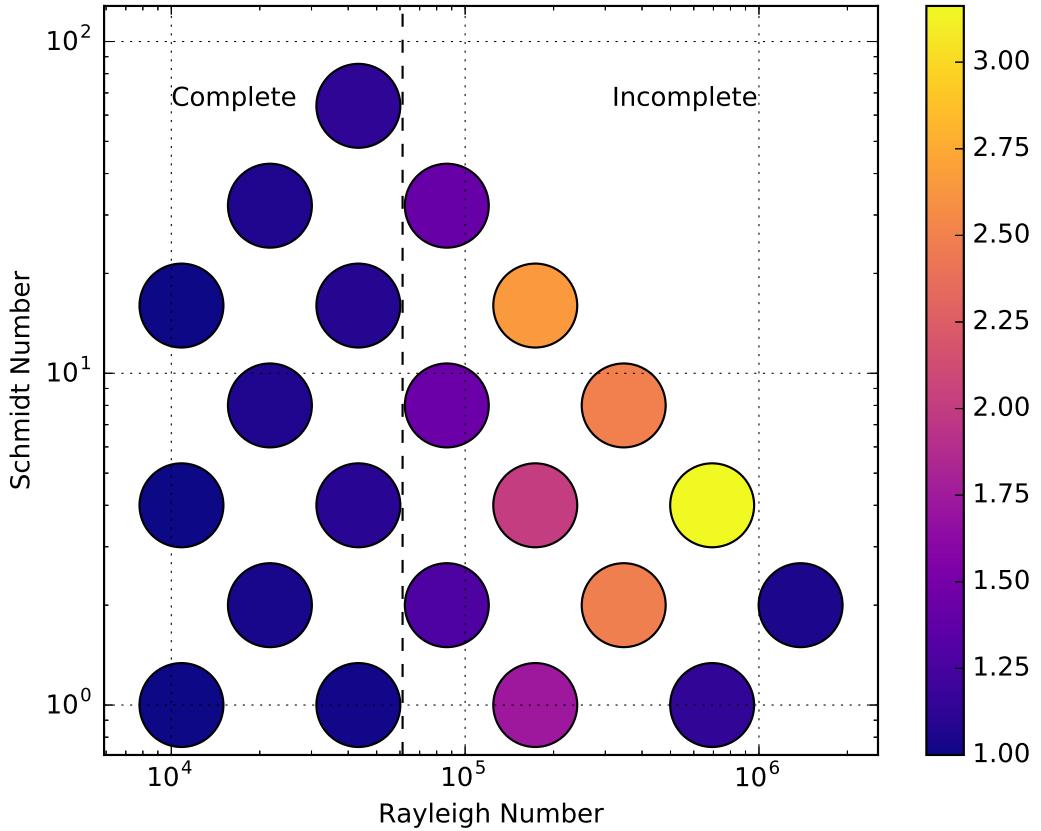


Figure 3.13: Best fit for  $C_7$  versus Rayleigh and Schmidt numbers. Experiments on the left side of the dashed line completed when the bubble stopped rising. Experiments on the right side of the dashed line are incomplete, having approached the vertical boundaries of the simulated domain.

it to disappear with trajectory completion by default. It would be very interesting if it remained, and further motivates completing the high Rayleigh number trajectories.

### Pure fluid coefficient, $C_7$

Similar to  $C_3$ ,  $C_7$  gives the ratio of the buoyant volume to the maximal mixed volume. A value of  $C_7 = 1$  implies the effective Atwood number is zeroed when  $M(t) = \lambda^2 h$ . Values greater than one imply that some portion of the positively buoyant fluid that would be in bubble has become entrained into the neighboring spike, allowing  $M(t) > \lambda^2 h$  while retaining

net buoyancy.

The fit values of  $C_7$  are plotted versus the Rayleigh and Schmidt numbers in Figure 3.13. The complete trajectories all have  $C_7 \approx 1$ , while the incomplete trajectories at higher Rayleigh numbers have increasing  $C_7$ . It is possible that at higher Rayleigh numbers some volume of light fluid detaches from the bubble and is transported into the spike. However, it is more likely that in the incomplete cases  $C_7$ , which influences the dynamics most at high mixing volumes, is underconstrained. In that case, we would expect  $C_7 \approx 1$  with no dependence on the Grashof, Rayleigh, or Schmidt numbers when the trajectories are completed.

### 3.7 Conclusions

We have proposed a simple ODE model for the growth of single mode Rayleigh-Taylor bubbles and spikes at low Atwood number. The model targets an intermediate range of Grashof numbers and high Rayleigh numbers, in which the single mode perturbation grows into an array of coherent bubbles and spikes. The dynamics of the bubbles are described in terms of buoyancy, viscous drag, and form drag. The buoyant force is scaled by a mixing factor related to the volume fraction of mixed fluid within the bubble, which is modeled by diffusion across the bubble's interface.

We have presented high fidelity spectral element simulations that reach later times and higher aspect ratios than previously available. The trajectory of the bubble can be roughly divided into regimes based on which terms in the model must be included. The first regime is exponential growth, which requires only the  $C_4$ ,  $C_6$ , and  $C_8$  terms, each of which is set by the linear theory. Next is the saturation regime, which adds the  $C_3$  inertial term and onsets around  $h = 0.05\lambda$ . The  $C_1$  form drag term can be added for better agreement, but doesn't change the dynamics qualitatively. Next is the viscous regime, which adds the  $C_2$  skin drag term and onsets around  $10^{-4}\text{Gr}$ . The last is the diffusive regime, which adds the  $C_5$  and  $C_7$

mixing terms and onsets around  $10^{-4}\text{Ra}$ .

The model proposed here is unable to describe stagnation and re-acceleration seen at higher Grashof numbers. However, the model is very accurate both before and after stagnation and re-acceleration, respectively. This demonstrates stagnation and re-acceleration to be transients in the flow, occurring around  $h = \lambda$ , suggesting that the physical processes that lead to them are initially absent, grow to some critical extent, and then decay relative to the magnitude of other processes. The bubble tip vortex ring is a strong candidate, as suggested by others [3, 47]. The model could be extended to account for the build-up of vorticity at the bubble tip.

The viscous drag, which is absent in other buoyancy-drag models, is essential to recovering terminal behavior at high aspect ratios and low diffusivity. Without viscous drag or mixing, the buoyant force grows with the aspect ratio while the form drag does not, leading to unbounded bubble velocity. In practice, as the velocity grows at high Reynolds number, the bubble interface breaks up leading to enhanced turbulent mixing. At moderate Reynolds number, viscosity bounds the bubble velocity, generally above the bound given by potential flow theories.

For any non-zero diffusivity, mixing reduces the buoyant force and the bubble ultimately stops rising. The penetration depth, i.e. the height of the bubble when it stops, scales linearly with the Rayleigh number. The relation implicitly defines a critical Rayleigh number below which the bubbles do not rise:  $\text{Ra}_c \approx 1500$ .

The proposed model has 8 descriptive parameters, 3 of which are constrained by the linear theory. These three are the  $C_4\lambda^3$  term in the inertia, the  $C_6\lambda^2$  term in the surface area, and the  $C_8\lambda^3$  term in the bubble volume. The presence of the  $C_4$  term demonstrates that the volume of fluid that begins to circulate at early times is independent of the bubble height.  $C_4$  is an increasing function of the viscosity, indicating that the viscous entrainment increases this volume, resulting in the reduction in growth rate predicted by the linear theory.

The physical interpretation of the five remaining parameters provides a prior for their value. Those parameters are similar to a drag coefficient, a friction factor, and three geometric ratios. The  $C_1$  term is estimated by relation to the drag coefficient of a flat plate. The  $C_2$  term is estimated by relation to the Darcy friction factor in a square duct. The  $C_3$  term is estimated as unity such that the inviscid immiscible acceleration is  $Ag$ . The  $C_7$  term is estimated as unity such that the fully mixed Atwood number is zero. The  $C_5$  term is estimated as  $\pi$ , which corresponds to cylindrical bubbles with diameter  $\lambda/2$ .

To calculate the 5 unconstrained model parameters, we fit the model to a battery of direct numerical simulations at moderate Grashof number, high Rayleigh number, and high aspect ratio. The simulations provide trajectories for the bubble height and volume of mixed fluid. The single mixing parameter is fit directly to mixed fluid measurements from numerical simulations. The remaining four parameters are fit with L2 regularization around the prior estimates. The resulting model reproduces simulated trajectories with relative errors in the bubble height less than 2% and in the volume of mixed fluid less than 4%.

The  $C_3$  and  $C_7$  coefficients, which scale the height in the denominator of Equation 3.11 and Equation 3.21, respectively, take values very near unity except for the lowest Rayleigh numbers, in the case of  $C_3$  and the incomplete trajectories, in the case of  $C_7$ . The  $C_1$  drag-type coefficient is typically between 0.3 and 0.6, with outliers at zero and above 0.9. The author has no direct explanation for the outliers and suggest they are indirectly caused by other early-time breakdowns in the model. It is conceivable that adding a vortical term, which would be most pronounced at early times, would align these cases with nominal range of values. The  $C_2$  friction factor-type coefficient is an strongly increasing function of the Grashof number and weakly increasing function of the Schmidt number, suggesting that shear instabilities could be enhancing the transport of momentum and consequently the drag. Similarly, the  $C_5$  mixing area coefficient is decreasing with diffusivity, suggesting the development of structure on the interface is smoothed in the diffusive cases. The  $C_5$

coefficient has a peaked dependence on the Grashof number, with a local maximum internal to the simulated trajectories. The author has no explanation.

While the simple model is sufficient to describe coherent, steady bubbles and spikes at low Atwood numbers and high Peclet numbers, few real flows fall within this regime. In this regard, the simple model proposed here is just one example of a general approach to defining models as a general force balance with coefficients based on limiting cases and simulation data.

The data presented here is a relatively sparse sample of Rayleigh-Schmidt space, intended to explore the efficacy and sensitivity of the model and its coefficients on the governing parameters of the problem. The model is predictive, in that it can predict the trajectory of the bubble from the initial condition. However, the model is sensitive to variations in the model coefficients, particularly  $C_2$ , and  $C_5$ , that themselves depend on the Rayleigh and Schmidt numbers.. To predict a bubble trajectory for a case within the convex hull of the parameter space explored here, but not at one of those points, the coefficients would need to be interpolated. For accurate interpolation, the parameter space should be more thoroughly sampled and the interpolation cross-validated.

### 3.8 Acknowledgements

M. H. acknowledges helpful conversations with Robert Rosner, Aleksandr Obabko, and especially Elizabeth Hicks, and the support of a Department of Energy Computational Science graduate fellowship.

For computer time, this research partially used the resources of the Supercomputing Laboratory at King Abdullah University of Science & Technology (KAUST) in Thuwal, Saudi Arabia. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

# CHAPTER 4

## DIRECT NUMERICAL SIMULATION OF SINGLE MODE THREE-DIMENSIONAL RAYLEIGH-TAYLOR EXPERIMENTS

### 4.1 Abstract

The single-mode Rayleigh-Taylor instability (smRTI) is well defined, poorly understood, and applicable to many fluid flows directly and through its relationship to multi-mode Rayleigh-Taylor models. This study reproduces three low-Atwood smRTI experimental runs (Wilkinson and Jacobs, 2007) in a specialized version of the Nek5000 spectral element code. The simulations use the initial amplitude, wavelength, acceleration, Atwood number, and viscosity from the three specific experiments and impose no-slip and no-flux boundaries on the velocity and scalar, respectively. The simulations are shown to reproduce the linear, saturation, stagnation, and re-acceleration phases of the smRTI seen in the experiments. Additionally, access to the full velocity and scalar fields demonstrates three different finite size effects: wall drag, wall lift, and a long wavelength mode along the diagonal. One of the simulations is extended by a factor of two in the vertical direction and the resulting late-time dynamics reach Froude numbers around 1.8, higher than previously reported. Finally, inspection of the span-wise flow reveals secondary flows of the first kind that transport the scalar from the bubble-spike interfaces into the bubble and spike centers. The agreement between simulations and experiments inspires confidence in the spectral element method for studying the Rayleigh-Taylor instability.

## 4.2 Introduction

The Rayleigh-Taylor instability occurs when a denser fluid is supported by a lighter one. Low-amplitude perturbations in the interface between the two fluids grow exponentially with a rate that is well modeled by linear stability analysis [15]:

$$\gamma = \sqrt{\frac{Agk}{\psi} + \nu^2 k^4 - (\nu + D)k^2}, \quad (4.1)$$

where  $g$  is the local acceleration,  $k$  is the wave-number,  $\nu$  is the kinematic viscosity,  $D$  is the diffusivity,  $A$  is the Atwood number, which characterizes the density difference:

$$A = \frac{\rho_h - \rho_l}{\rho_h + \rho_l}, \quad (4.2)$$

and  $\psi$  describes the effect of the interface thickness and is a function of  $A$ ,  $k$ , and the thickness  $\delta$ . In the low Atwood number limit:

$$\psi = 1 + \frac{k\delta}{\sqrt{\pi}}. \quad (4.3)$$

At larger amplitudes, the perturbations grow non-linearly with the light fluid rising through the heavier fluid in ‘bubbles’ and the heavy fluid falling through the lighter fluid in ‘spikes.’ Early experiments by Davies and Taylor [8] and potential flow models by Layzer [32] for  $A \approx 1$  suggested that the bubbles reach a terminal velocity, and later experiments by Dimonte and Schneider [12], also at  $A \approx 1$ , showed that dense spikes free-fall. On the other hand, recent experiments by Wilkinson and Jacobs [62] and simulations by Ramaprabhu et al. [48, 47], Wei and Livescu [61], and others [53] show that, at Atwood numbers less than one half, the constant velocity regime is followed by a re-acceleration regime in which the velocity doubles. The dynamics beyond re-acceleration have not been established, with Ramaprabhu et al. observing a return to the velocity of potential flow [19] while Wei and

Livescu report continued constant acceleration.

Here, we consider the low-Atwood number limit. The Boussinesq approximation, which ignores density differences that don't multiply the gravitational acceleration, simplifies the governing equations to a single incompressible phase with an active scalar representing the buoyancy:

$$\begin{aligned}\frac{D}{Dt}u &= -\nabla P + \nu \nabla^2 u - A\vec{g}\phi, \\ \frac{D}{Dt}\phi &= D\nabla^2\phi,\end{aligned}\tag{4.4}$$

where  $u$  is the velocity,  $P$  is the pressure, and  $\phi$  is the scalar that controls the density gradient. Without loss of generality, we define  $-1 \leq \phi \leq 1$ . These equations have a symmetry under inversion of the scalar and acceleration,  $\phi \rightarrow -\phi, \hat{g} \rightarrow -\hat{g}$ , so we know the bubbles and spikes have the same dynamics given the same initial conditions. The parameter space of the equations are described by two nondimensional numbers: the Grashof number,

$$\text{Gr} = \frac{Ag\lambda^3}{\nu^2},\tag{4.5}$$

where  $\lambda$  is a characteristic length, and the Schmidt number,

$$\text{Sc} = \frac{\nu}{D}.\tag{4.6}$$

We approximate the governing equations numerically using the spectral element method (SEM) [10]. The SEM converges exponentially with respect to spectral order and has purely dispersive errors, making it a natural method for direct numerical simulations of mixing problems. Unlike pseudo-spectral methods, it handles no-slip boundaries and can evenly sample the interior of the domain. We use a specialized version of the Nek5000 community code, NekBox [26], customized specifically to study the low Atwood number Rayleigh-Taylor

instability. NekBox restricts the domain to a tensor product of orthogonal bases and employs fast spectral coarse preconditioners for the pressure Poisson equation. It is roughly an order of magnitude faster than the more general Nek5000.

It is the goal of this study to validate direct numerical simulations of the low-Atwood single mode Rayleigh-Taylor instability in NekBox against the best available experimental data, that from Wilkinson and Jacobs [62]. In the future, we will apply the same numerical methods to the broader question of the late-time dynamics of the low Atwood smRTI.

In Section 4.3, we review the spectral element method and describe the numerical parameters of the simulations. In Section 4.4, we compare the numerical and experimental results, extend the domain in extent and time to reach higher aspect ratios, and introduce new behavior in the span-wise flow. In Section 4.5, we discuss the validity of our methods for simulating Rayleigh-Taylor flows, the limits of wall-bounded single mode experiments, and the implications of secondary flows to mixing.

## 4.3 Numerical methods

### 4.3.1 Spectral element formulation

The governing equations, Equation 4.4, are discretized using the spectral element method [10]. The domain is first divided into cubic elements. Each element is represented by a tensor product of  $n^3$  Gauss-Lobatto-Legendre (GLL) quadrature points of order  $p = n - 1$ . The elements are coupled by continuity at the shared points on their faces.

Time is discretized with a 3rd order backwards difference formula (BDF3). The linear and non-linear terms are split, and the non-linear convection operator is extrapolated with a 3rd order scheme that sets the 3rd derivative at the extrapolated point to zero (EX3). The full discrete system is 3rd order in time,  $p$ th order in space, and has purely dispersive errors.

Initially, no filtering is applied. Filtering is applied when the scalar field  $\phi$  exhibits

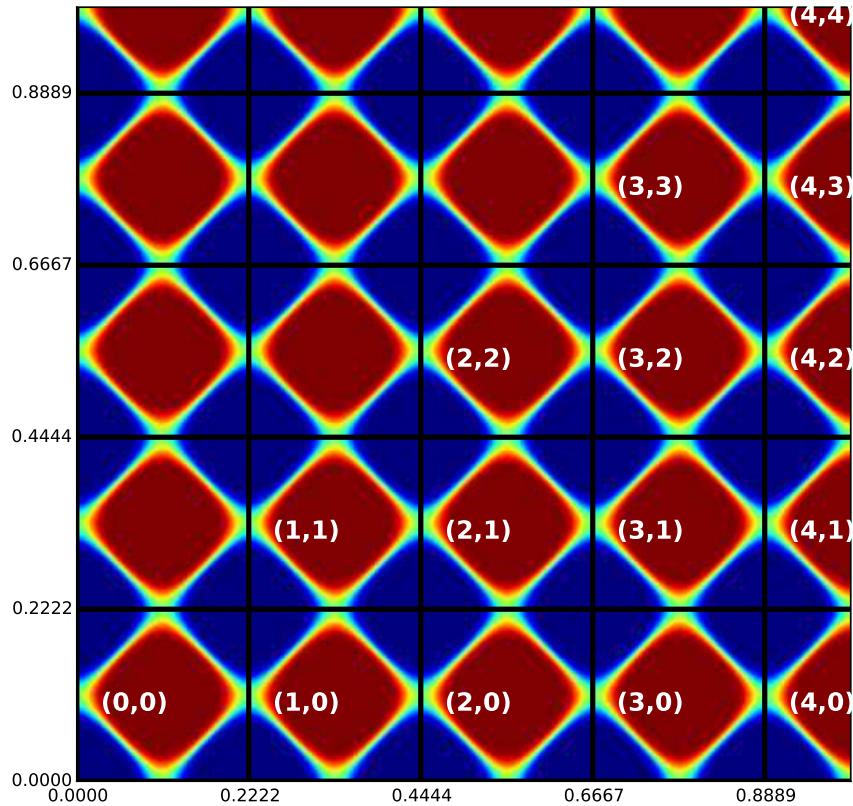


Figure 4.1: Initial condition,  $\phi(x, y, 0, 0)$ , for 4.5 mode simulation. The bubbles are logically separated by the solid grid and each unique bubble is labeled.

small scale oscillations on what should otherwise be smooth steep boundaries. For Schmidt numbers greater than or equal to unity, the scalar field is less numerically stable than the velocity, so instability in the scalar precedes instability in the velocity. The filter attenuates the highest frequency mode of the velocity and scalar fields within each element by 5% at each time-step.

Modes	Grashof	Schmidt	Aspect
2.5	$5.12 \times 10^6$	7.0	0.037
3.5	$3.98 \times 10^6$	3.5	0.025
4.5	$4.66 \times 10^5$	1.0	0.041
4.5	$4.66 \times 10^5$	3.5	0.041

Table 4.1: Parameters of simulations. The aspect ratio is defined with respect to the quiescent amplitude,  $a_0$ , from Equation 4.8. The last row is the simulation that extends the tank by a factor of two in the vertical direction.

### 4.3.2 Simulation setup

Three simulations were conducted to reproduce experiments with 2.5, 3.5, and 4.5 modes across the diagonal. Then, an extension of the 4.5 mode experiment with twice the vertical extent was performed. Finally, a 4.5 mode calculation with periodic boundary conditions and unit Schmidt number was performed as a reference for comparing the growth of the bubbles in the wall-bounded flows.

The boundary conditions for the non-periodic simulations are no-slip for the velocity and no-flux (insulating) for the scalar. The initial conditions for the velocity are quiescent. The initial conditions for the scalar are the product of two cosines smeared by an error function in the z-direction:

$$\phi(x, y, z, 0) = \text{erf} \left[ \frac{z + a_0 \cos(k_x x) \cos(k_y y)}{\delta} \right], \quad (4.7)$$

where  $a_0$  is the initial amplitude,  $k_x$  and  $k_y$  are the wave-numbers in the x and y directions, respectively, and  $\delta$  is the interface thickness. In our case,  $k_x = k_y$ . An example initial condition is plotted in Figure 4.1. By symmetry, we know that the average scalar is zero,  $H_{i,j} = H_{j,i}$ , and that the spike dynamics are identical to the bubbles.

The Atwood number, local acceleration, and initial amplitude were taken from experimental measurements by Wilkinson and Jacobs [62, 31]. In the experiment, the interface has a non-zero initial velocity and the bulk flow is not measured. Instead of trying to model

the bulk flow, we used the linear theory, Equation 4.1, to transform the flow to quiescence. Specifically, the system

$$H = a_0 \cosh(\gamma(t - t_0)) \quad (4.8)$$

$$V = a_0 \gamma \cosh(\gamma(t - t_0))$$

where  $H$  and  $V$  were the experimental measures of the initial interface height and velocity, respectively, is solved for  $a_0$  and  $t_0$ , the quiescent amplitude and time. The quiescent amplitude is used for the simulations. The experiments to reproduce were selected such that solutions to Equation 4.8 existed and the bubble reached the greatest heights.

The viscosity is matched to the experimental fluids, but the diffusivity is not. The Schmidt number in the experiments is in excess of 1000. The computational cost of the simulation goes with the Schmidt number to the 4th power, so directly simulating such high Schmidt number was not possible. Instead, the Schmidt number was varied in the range of 1 to 7 to gain a qualitative understanding of its influence on the flow.

It should be noted that, while the Atwood number was used to scale the local acceleration, the Boussinesq approximation implies that the Atwood number has been taken to zero while keeping the product  $Ag$  fixed. The generally accepted Boussinesq limit is  $A = 0.05$ , which is three times smaller than the Atwood number in this case. Previous simulations directed at single-mode re-acceleration have been performed at  $A \geq 0.15$ , so it is not known whether re-acceleration persists in the limit as  $A \rightarrow 0$ .

The three simulations reproducing experimental runs are conducted on the Mira supercomputer at Argonne Leadership Computing Facility (ALCF). The resolution and time-step was chosen to ensure numerical stability: a  $256 \times 256 \times 512$  mesh of 7th order elements for 11,509,170,176 degrees of freedom. The simulation was distributed over 524,288 cores and 1,048,576 MPI processes. 64 outputs were written to disk, each 6/8ths of a TiB. The number of elements and degrees of freedom are doubled for the extension to twice the vertical extent.

### 4.3.3 Post-processing

The simulation outputs the velocity, pressure, and scalar fields at the Gauss-Lobatto-Legendre points in double precision. They are post-processed into low-dimensional observables: the bubble height and two-dimensional slices of the velocity, vorticity, pressure, and scalar through the horizontal mid-plane and vertical diagonal. Post-processing is performed using the ‘nek-analyze’ post-processing framework, which implements a MapReduce-like [9] backend for parallel, out-of-core analysis.

The bubble height is defined as:

$$H = \sup \left\{ z : \min_{x,y} \phi(x, y, z) < 0 \right\}, \quad (4.9)$$

which avoids measuring the diffusive growth by tracking the center of the interface profile instead of the ends. The bubble velocity is found by fitting a cubic spline to  $H(t)$  and differentiating.

The height of individual bubbles,  $H_{i,j}$ , is found by restricting  $\min_{x,y}$  in Equation 4.9 to the span-wise square of diagonal length  $\lambda$  centered on the bubble in the  $i$ -by- $j$ -th position. The bubble domains and labels are shown in Figure 4.1.

## 4.4 Results

### 4.4.1 Validation

#### Linear growth rate

We compute the growth rate from the bubble height at the first simulation output time:

$$\gamma \approx \frac{1}{t_1} \text{acosh} \left( \frac{h(t_1)}{h(0)} \right), \quad (4.10)$$

Modes	Theory	Simulation	Aspect ratio
2.5	12.94	12.96	0.089
3.5	22.24	22.24	0.098
4.5	12.10	11.05	0.124
4.5	12.31	11.52	0.128

Table 4.2: Growth rate: linear theory vs. simulation. Theoretical values are calculated as in Equation 4.1. Simulation values are calculated as in Equation 4.10. The aspect ratio is shown for the second sample,  $h_1/\lambda$ . Note the difference in Schmidt number between the two 4.5 mode cases.

and collect the results in Table 4.2. Because the simulations targeted the non-linear regime, the height is not available until characteristic time  $\tau = t\gamma \sim 2$  and aspect ratio  $h/\lambda \approx 0.1$ . Therefore, we expect the simulation value to be below the theoretical value given by Equation 4.1 due to saturation.

The agreement is good, with only the lower Grashof, higher aspect ratio 4.5 mode calculation deviating more than a part in one hundred. The 2.5 mode simulation outperforms the theory slightly. This could be due to the long-wavelength finite size effect discussed later, which is stronger for fewer modes in the finite domain.

## Froude number

The experiments observe only the diagonal plane and measure the height with respect to the most internal bubble. Therefore, we plot in Figure 4.2 the nondimensional velocity, i.e., the Froude number, of the central bubble alone:

$$\text{Fr} = \frac{v}{\sqrt{Ag\lambda}} \quad (4.11)$$

In each case, the simulation exhibits the same qualitative behavior as the experiment: exponential growth saturating to a stagnation velocity around Goncharov's theoretical value of  $\text{Fr} = \pi^{-1/2}$  [19]. In cases where the simulation and experimental data extend in time, the

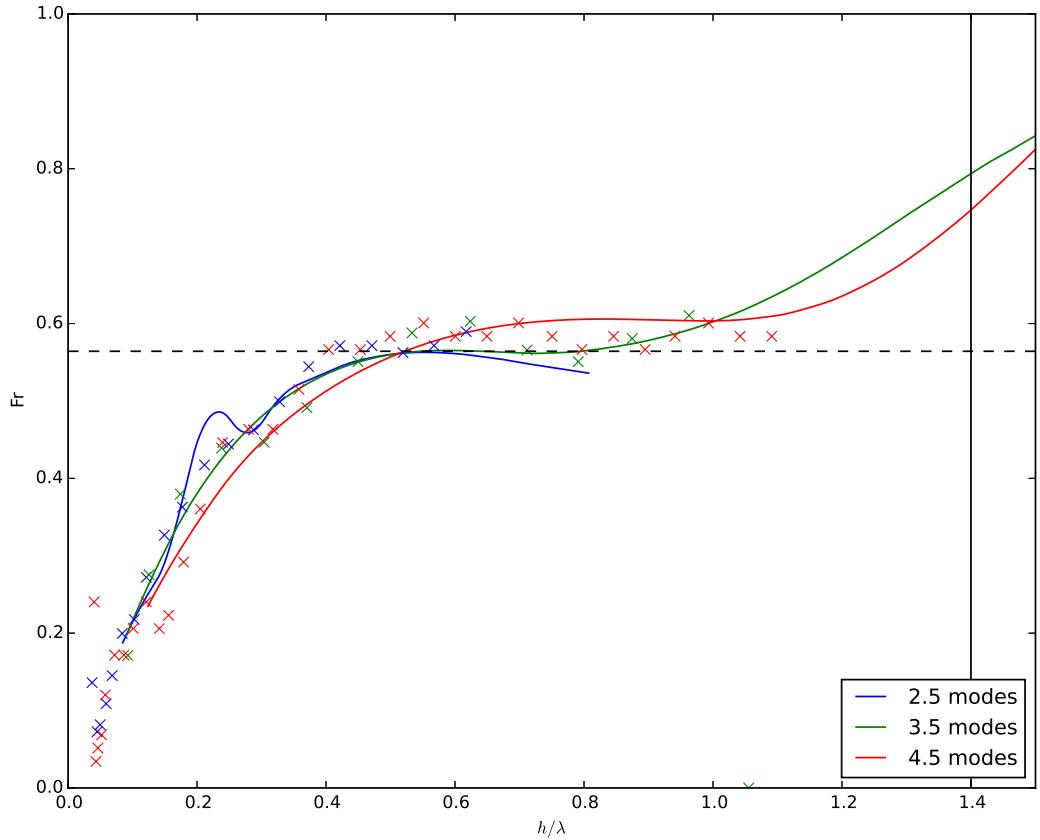


Figure 4.2: Froude number vs. height, nondimensionalized by the wavelength. Lines are the derivative of cubic splines through simulation outputs. Points are from experiment via direct measurement of the bubble velocity [31]. The dotted horizontal line is positioned at Goncharov's theoretical value of  $\pi^{-1/2}$  [19].

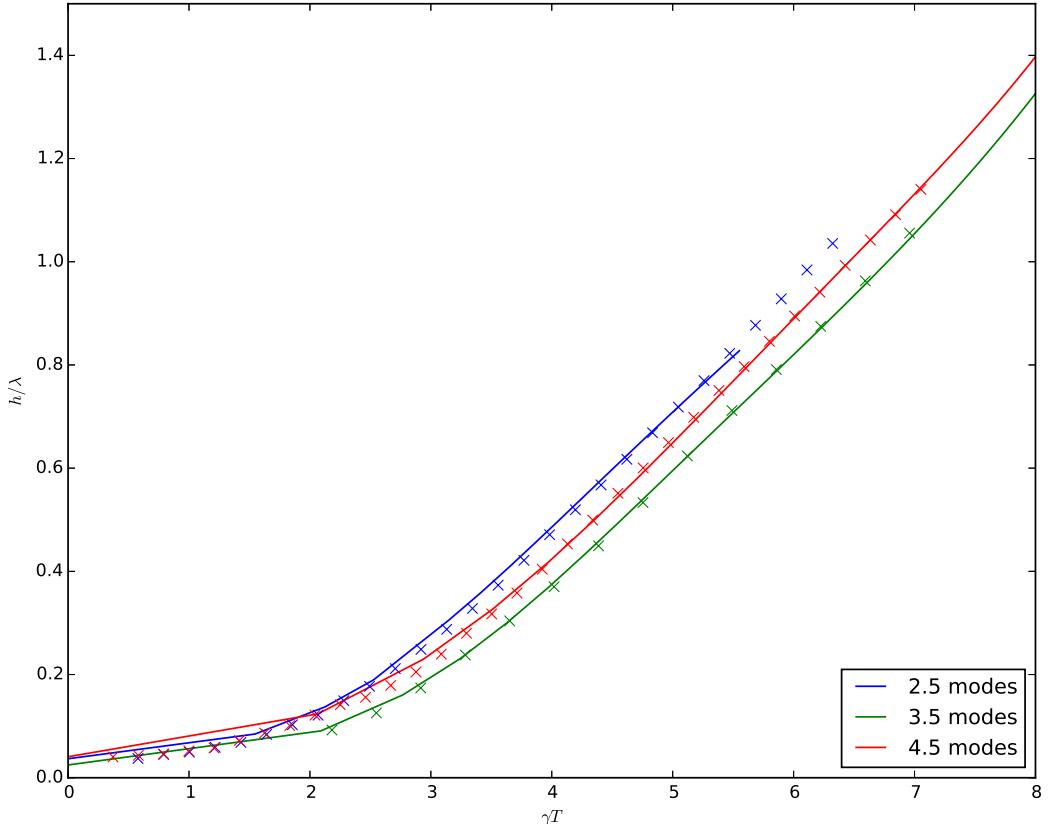


Figure 4.3: Bubble height vs. time, nondimensionalized by the wavelength and linear growth rate, respectively. Lines are cubic splines through simulation outputs. Points are from experiment [31]. The points are shifted in time to minimize the square deviation from the spline summed over the plotted points.

beginnings of re-acceleration are also seen.

The 2.5 mode case exhibits peculiar behavior, with two inflection points around aspect ratio  $h/\lambda = 0.2$ . This is not an issue with the splines, as can be seen in Figure 4.3, which plots the nondimensional height vs. nondimensional time without splines. The 2.5 mode case, which has the highest Grashof and Schmidt numbers, went unstable. Upon inspection, the 5% filtering was unable to suppress the oscillations in the scalar field. This simulation could be retried with greater resolution, but, given the stability of the 3.5 mode case, the computational cost would be up to  $43\times$  greater, which is why it was not repeated in this study.

#### 4.4.2 Extension

Given the agreement between the simulations and the experiment, we can use the simulations to explore the flow in ways that are not readily accessible experimentally. In this study, we extend the subject cases and analysis in three ways. First, we calculate the height of each bubble in the tank individually and use the results to study finite size effects. Second, we extend the 4.5 mode experiment by a factor of two in the vertical extent of the domain and simulation time to delay interaction with the top boundary and reach later times and larger bubble heights. Finally, we consider the span-wise, vs. stream-wise, flow by taking slices of the midplane and observing pressure driven secondary flows. These three extensions are a small sample of the types of observations that are available numerically.

#### Wall effects

The individual bubble heights for the 4.5 mode case are computed as in Equation 4.9, but restricted to the span-wise domain nearest to the bubble tip, which is marked in Figure 4.1. The bubble Froude numbers are plotted in Figure 4.4, alongside the aggregate and purely periodic values.

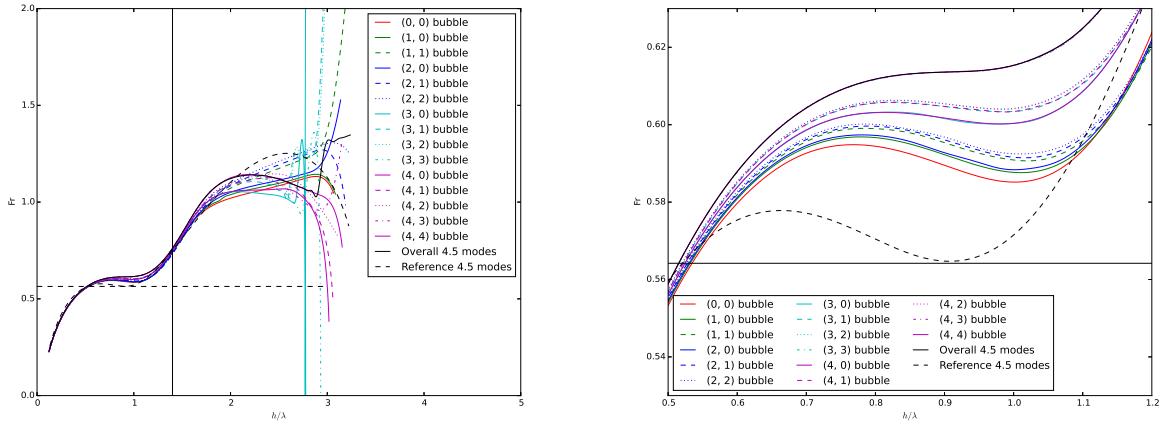


Figure 4.4: Froude number as a function of height, nondimensionalized by the wavelength, by bubble in the 4.5 mode simulation. Solid line is from the height defined as the maximum taken over the entire span-wise domain. Dotted line is the periodic reference calculation.

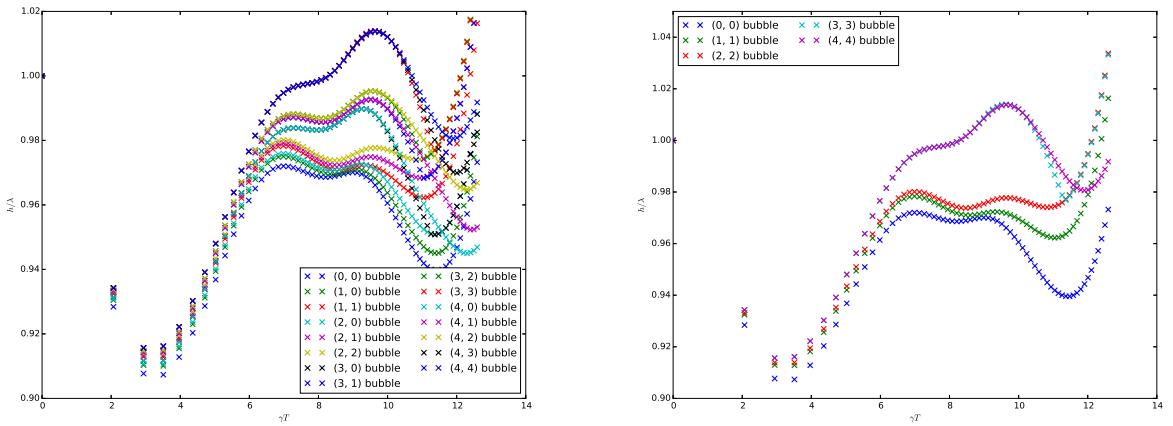


Figure 4.5: Ratio of wall-bounded bubble height to periodic bubble height in the 4.5 mode simulation.

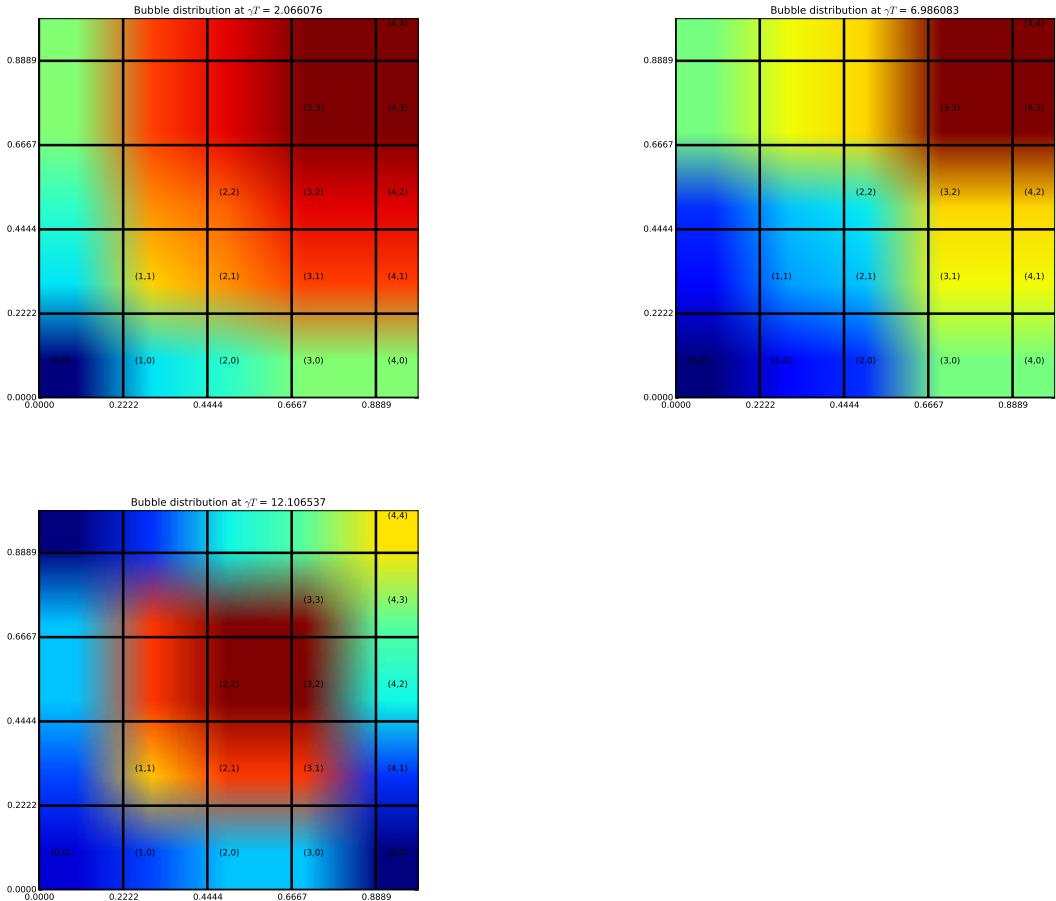


Figure 4.6: Spatial distribution of bubble heights at three times, with red denoting greater bubble heights and blue denoting lesser bubble heights. The first time is when the ratio of the wall-bounded to periodic bubble height is at a minimum, which is during the transition from linear to non-linear growth; a long wavelength mode is seen across the diagonal. The second time is taken from the stagnation phase; the long wavelength mode is still present. The third time is taken from the end of the simulation; the central bubbles are being squeeze significantly ahead of the bubbles near the wall.

There are at least three mechanisms by which the walls divert the flow from its fully-periodic preference. The first is that the no-slip boundaries that pass along the diagonal of the boundary bubbles and spikes create a boundary layer that viscously damps vertical flow. The second is that the pressure gradient from the boundary layer pushes the boundary bubbles and spikes towards the interior of the domain. These two effects have been studied independently in the context of multi-phase bubbles rising near walls [56], and are characterized as wall drag and lift forces. Finally, the finite nature of the span-wise lattice of bubbles and spikes, coupled with the vertical symmetry condition that the total bubble volume must equal the total spike volume, breaks one of the 4-fold symmetries of the infinite cubic lattice causing a local aggregation of spikes in the  $(0, 0)$  corner and of bubbles in the  $(4, 4)$  corner. The local aggregation sets up one low-amplitude long-wavelength mode across the diagonal.

These three mechanisms promote and penalize the growth of different bubbles in the finite lattice, allowing us to infer the relative magnitudes of the effects based on the performance of the bubbles compared to their periodic counterparts. The wall drag penalizes the growth of bubbles that contain a boundary: the bubbles in the 4th column. Because the effect also penalizes the growth of the spikes that contain boundaries, it should encourage the growth of the bubbles adjacent to those spikes: those in the 0th row. The effect should alternate and diminish towards the interior of the domain. The wall lift pushes bubbles and spikes at the boundary towards the interior. This reduces the form drag on the interior bubbles by increasing the pressure on their trailing edges. When adjacent bubbles actually touch, skin drag is also reduced. Overall, wall lift promotes the growth of the interior bubbles. Finally, the long-wavelength mode promotes growth of bubbles in the bubble heavy corner and penalizes growth of bubbles in the spike heavy corner.

The spatial distribution of the bubble heights can be seen in Figure 4.6 and the heights relative to the periodic bubble can be seen in Figure 4.5. At moderate times, the  $(4, 4)$  bubble leads and the  $(0, 0)$  bubble trails, indicating that the long-wavelength mode due to symmetry

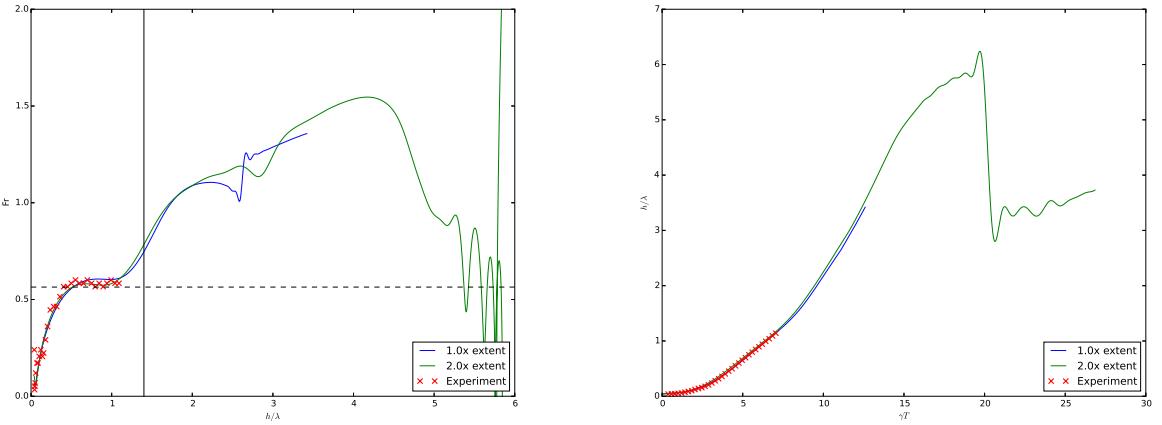


Figure 4.7: Bubble velocity and bubble height vs. time, nondimensionalized by the wavelength and linear growth rate, for 4.5 mode simulations and experiment. Lines are from simulation output, one case with the same vertical extent as the simulation and in the other with twice that vertical extent. Points are from experiment via direct measurement of the bubble velocity and bubble height. The dotted horizontal line is positioned at Goncharov's theoretical value of  $\pi^{-1/2}$  [19]. The solid vertical line marks the greatest bubble height reached in any of the experiments by Wilkinson and Jacobs [62].

breaking is the dominant effect. Additionally, all of the bubbles under-perform their periodic counterpart, indicating that the wall drag has damped the overall flow but with less spatial dependence than the long-wavelength mode. At late times, the central  $(2, 2)$ ,  $(2, 1)$  bubbles are accelerated while the edge bubbles break down, indicating the growing importance of the wall lift effect. The wall lift ultimately leads to bubble collisions that destroy the bubble lattice, enhance mixing, and break down the flow.

### Late-time behavior

The 4.5 mode simulation was repeated with twice the vertical extent and simulation time. Additionally, the Schmidt number was increased from 1 to 3.5 to reduce late-time mixing not present in high Schmidt number experiments. Figure 4.7 compares the short unit-Schmidt and long moderate-Schmidt trajectories, which are widely in agreement. The reduction in bubble acceleration around aspect ratio  $h/\lambda = 2$  is present in both the short and the long

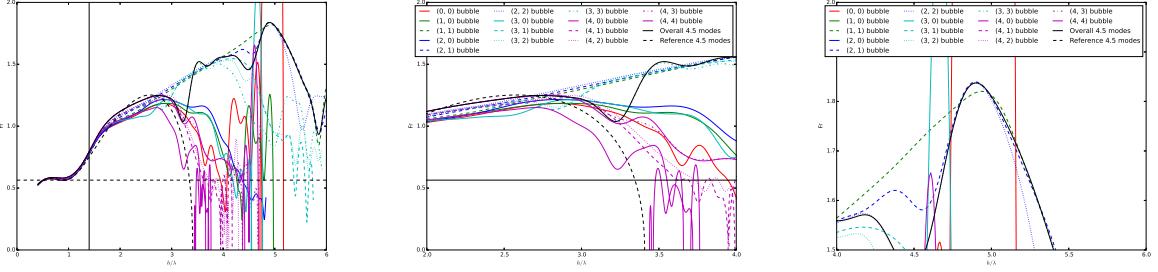


Figure 4.8: Froude number as a function of height, nondimensionalized by the wavelength, by bubble in the 4.5 mode simulation with extended vertical extent. Solid line is from the height defined as the maximum taken over the entire span-wise domain. Dotted line is the periodic reference calculation. The dotted horizontal line is positioned at Goncharov's theoretical value of  $\pi^{-1/2}$  [19]. The solid vertical line marks the greatest bubble height reached in any of the experiments by Wilkinson and Jacobs [62].

simulations, so it is unlikely to be due to the vertical domain boundaries. It is not, however, present in the periodic calculation, so it could be a wall effect.

Figure 4.8 mimics Figure 4.4 but for the late-time case. The periodic reference trajectory, calculated with the original domain size, rapidly decays after reaching a maximum around aspect ratio  $h/\lambda = 2.5$  due to interactions with the top of the domain. Its maximum Froude number is around 1.2, consistent with previous calculations. The central bubbles in the extended late-time run continue to experience constant acceleration past aspect ratio 3 and Froude number 1.2, with the (1, 1) bubble continuing to aspect ratio 5 and Froude number 1.8.

The decay of the velocity of the periodic reference bubble around  $h/\lambda = 2.5$  suggests the late-time simulations would interact with the top boundary around  $h/\lambda = 5$ . In fact, this is exactly when the (1, 1) and (2, 2) bubbles begin to decay. However, the bubbles closer to the boundaries break down much earlier. The (4, 4) bubble, for example, reaching maximum Froude number around  $h/\lambda = 2.75$ . We can infer that the wall lift that drives the boundary bubbles into the interior bubbles destroys the periodic ordering. It is not clear if the decay of the interior bubbles at  $h/\lambda = 5$  is due to the top boundary or the wall lift destroying the

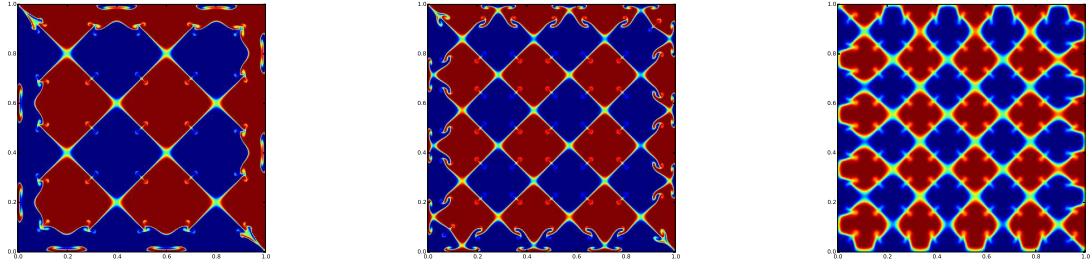


Figure 4.9: Scalar field in the horizontal mid-plane for 2.5, 3.5, and 4.5 mode simulations.

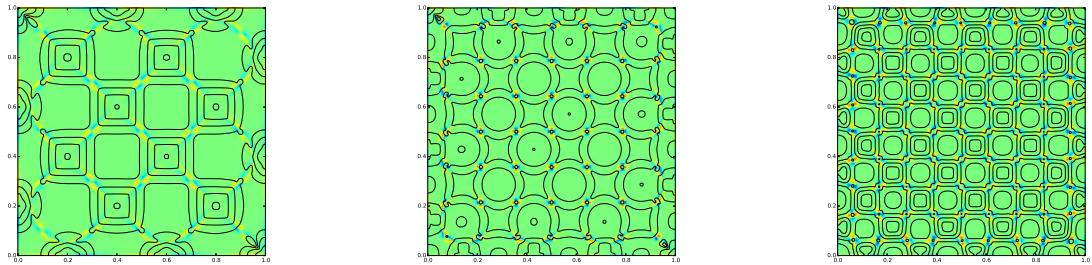


Figure 4.10: Secondary flow in the horizontal mid-plane. Background color is the vertical component of the vorticity. Contours are lines of constant pressure.

periodic ordering.

## Secondary flow

In addition to the bubble height and vertical slices in the diagonal plane, we observe the horizontal mid-plane. The span-wise scalar field,  $\phi(x, y, 0)$ , exhibits plume structures that

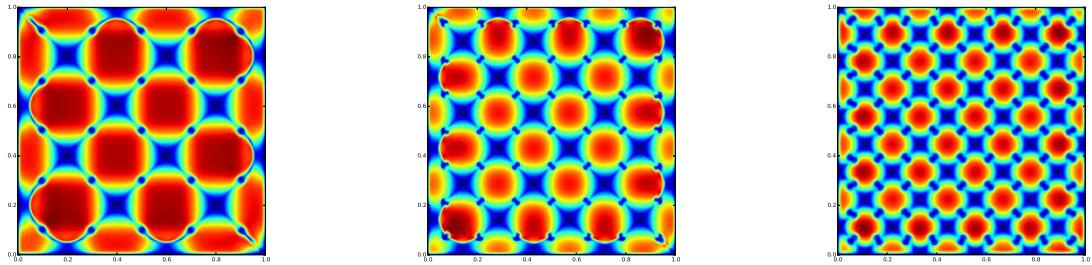


Figure 4.11: Dynamic pressure in the horizontal mid-plane.

penetrate the bubble faces, as seen in Figure 4.9. The cause of these plumes is advection by secondary flows, as seen in Figure 4.10. Overlaying the pressure with the span-wise velocity reveals it to be secondary flow of the first kind: secondary flow due to span-wise pressure gradients. The span-wise pressure gradient comes from the dynamic pressure of the rising and falling bubbles and spikes, in contrast to the stationary points at their interfaces.

The secondary flow advects mixed fluid from the interface into the centers of the bubbles and spikes. Enhanced mixing reduces the effective Atwood number of the bubbles and spikes, but the magnitude of this effect is not clear. As a secondary flow of the first kind, this mixing mode is present even at low Reynolds numbers.

## 4.5 Conclusions

The simulations described here reproduce the growth rate, stagnation velocity, and reacceleration of the low-Atwood single mode Rayleigh Taylor instability for three experimental runs by Wilkinson and Jacobs. These reproductions inspire confidence not only in the NekBox code, but also in the Boussinesq approximation for  $A = 0.15$  and the low-Schmidt approximation.

In wall-bounded flows, the bubbles and spikes nearest to the no-slip boundaries experience lift and drag forces that slow their non-linear growth and push them towards their inner neighbors. There is an additional effect due to the finite domain breaking one of the 4-fold symmetries from the purely periodic problem. In the wall-bounded initial condition, one corner of the domain has an excess of bubbles while the opposite has an excess of spikes. This sets up a long-wavelength mode across the diagonal that encourages bubble growth in one corner and discourages it in the other.

Ultimately, the bubble-bubble and spike-spike collisions may destroy the single-mode ordering of the flow at aspect ratio 5, but the onset of velocity decay may alternatively be due to the upper boundary. If the decay is due to collisions, it would limit the use of

wall-bounded flows as proxies for periodic flows to moderate aspect ratios. The ability of wall-bounded flows to approximate periodic ones at high aspect ratio warrants further study.

In addition to causing collisions, the growing boundary layer squeezes the flow in the span-wise direction, accelerating it past its fully-periodic trajectory in the stagnation and re-acceleration phases.

The inner bubbles experience near constant acceleration from aspect ratio 2 to aspect ratio 5, with a maximum Froude number of 1.8. This contrasts results by Ramaprabhu et al. [47] that show saturation post-reacceleration at  $\text{Fr} \approx 1$ . The saturation could be explained by excess mixing or the finite size of the domain, which extends to  $h/\lambda = 6$  in their case and 9 in ours. Alternatively, the acceleration could be artificially sustained by the wall lift force pushing the boundary bubbles into the interior ones.

Single-mode Rayleigh-Taylor flows develop span-wise pressure gradients with local minima in bubble and spike centers and local maxima in bubble and spike corners. The pressure drives secondary flows of the first kind in the form of vortex quads centered on bubble-spike interface centers. These span-wise flows mix the fluid across otherwise laminar interfaces, perturbing the scalar profiles.

## 4.6 Acknowledgements

M. H. is grateful for useful conversations with Jeffrey Jacobs, Robert Roser, Aleksandr Obabko, Elia Merzari, Oana Marin, and Elizabeth Hicks.

M. H. acknowledges support from the Department of Energy Computational Science graduate fellowship. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

# CHAPTER 5

## EFFICIENCY OF HIGH ORDER SPECTRAL ELEMENT METHODS ON PETASCALE ARCHITECTURES

MAXWELL HUTCHINSON, ALEXANDER HEINECKE, HANS PABST, GREG HENRY, MATTEO PARSANI, AND DAVID KEYES

### 5.1 Abstract

High order methods for the solution of PDEs expose a trade-off between computational cost and accuracy on a per degree of freedom basis. In many cases, the cost increases due to higher arithmetic intensity while affecting data movement minimally. As architectures tend towards wider vector instructions and expect higher arithmetic intensities, the best order for a particular simulation may change.

This study highlights preferred orders by identifying the high order efficiency frontier of the spectral element method implemented in Nek5000 and NekBox: the set of orders and meshes that minimize computational cost at fixed accuracy. First, we extract Nek's order-dependent computational kernels and demonstrate exceptional hardware utilization by hardware-aware implementations. Then, we perform production-scale calculations of the nonlinear single mode Rayleigh-Taylor instability on BlueGene/Q and Cray XC40-based supercomputers to highlight the influence of the architecture. Accuracy is defined with respect to physical observables, and computational costs are measured by the core-hour charge of the entire application. The total number of grid points needed to achieve a given accuracy is reduced by increasing the polynomial order. On the XC40 and BlueGene/Q, polynomial orders as high as 31 and 15 come at no marginal cost per timestep, respectively. Taken together, these observations lead to a strong preference for high order discretizations that use fewer degrees of freedom. From a performance point of view, we demonstrate up

to 60% full application bandwidth utilization at scale and achieve  $\approx 1$  PFlop/s of compute performance in Nek's most flop-intense methods.

## 5.2 Introduction

The solution of partial differential equations (PDEs) is a core problem in HPC, with particular application to computational materials science and fluid dynamics. PDEs are solved by discrete approximation: space and time are sampled and the PDEs is translated into a relation on those samples. From a mathematical point of view, these approximations are characterized by stability conditions and convergence rates. Schemes which do not satisfy stability conditions usually fail catastrophically with values that diverge to infinity. The convergence rate describes the relationship between the resolution and the error. For a characteristic inter-sample spacing  $h$ , a method is of order  $p$  if the error goes as  $h^p$ . High order methods are schemes with convergence rates higher than third order [60], many of which expose the order as a user input.

From a computational point of view, the approximations are characterized by sparsity, locality, and arithmetic intensity. As the order increases, the sparsity and locality typically decrease while the arithmetic intensity increases. The improved convergence rates are ‘paid for’ with more floating point operations (FLOP), on a per sample basis, while, for a given error tolerance, the number of samples can be decreased. The relationship between these computational characteristics and computational cost is complicated by features common to modern architectures: vector instructions, deep caches, and arithmetic-to-data movement imbalance.

Here, we explore the relationship between order, accuracy, cost, and architecture. We identify the user-facing properties of high order methods: the accuracy in observables, time to solution, resource usage, and required scale. We also identify the user-defined inputs: the physical problem, the order, the total number of samples, the number of processors, and the

computer architectures. To make the study more practical, we focus on the specific task of optimizing a study of the single-mode Rayleigh-Taylor instability (smRTI) as a parameter sweep over Grashof and Prandtl numbers. This is a high throughput use-case, where the relevant cost is resource usage and scale is fixed with respect to the size of the problem and assumed to not be a limitation. This leaves us with the accuracy and resource usage versus the order, number of samples, and computer architectures.

We select the NekBox version of the Nek5000 code (together: Nek), which implements the spectral element method (SEM) [46] with tunable order, is known to scale to a million ranks [42], and has been used for Rayleigh-Taylor problems in the past [25]. NekBox takes advantage of static, uniform meshes to solve the coarse part of the preconditioner with FFTs or DCTs, improving efficiency and scalability. We extract representative order-dependent kernels from Nek and analyze their performance on BlueGene/Q and Cray XC40 supercomputers.

We also conduct a set of application benchmarks to measure the cost and accuracy. The cost is computed in core-hours, in the same way most users are charged. The accuracy is computed with respect to the smRTI’s bubble height and mix volume, which are the most common observables studied in the smRTI community. The benchmarks vary the order and total number of samples, and are conducted on the Mira and Shaheen XC40 supercomputers at Argonne Leadership Computing Facility (ALCF) and KAUST Supercomputing Laboratory (KSL), respectively.

### 5.2.1 *Outline*

In Section 5.3, we review the SEM as implemented in Nek. In Section 5.4, we introduce LIBXSMM for hardware-aware implementation of Nek’s performance critical kernels, and demonstrate their performance in isolation. In Section 5.5, we perform a convergence/performance study of SEM discretizations for the smRTI problem and present full-application

performance at scale. Section 5.6 concludes with a discussion of preferred orders on the BlueGene/Q and Cray XC40 supercomputers.

### 5.3 Nek's Computational Core

#### 5.3.1 Governing equations and time-splitting

Nek5000 and NekBox solve the incompressible Navier–Stokes and continuity equations:

$$\frac{\partial u}{\partial t} + u \cdot \nabla u = -\frac{1}{\rho} \nabla p + \nu \nabla^2 u + f \quad \nabla \cdot u = 0, \quad (5.1)$$

where  $u$  is the flow velocity,  $\rho$  is the fluid density,  $p$  is the pressure,  $\nu$  is the kinematic viscosity, and  $f$  consists of user-defined forcing terms. Additionally, Nek can solve advection-diffusion equations for scalars, such as the temperature or mass fraction:

$$\frac{\partial \phi_i}{\partial t} + u \cdot \nabla \phi_i = \alpha_i \nabla^2 \phi_i + q_i, \quad (5.2)$$

where  $\phi_i$  is the scalar value,  $\alpha_i$  is the diffusivity, and  $q_i$  is a user-defined source term, each for the  $i$ th scalar.

The time derivative is discretized with a backward difference formula (BDF), within which the nonlinear and forcing terms are extrapolated (EX):

$$\sum_{j=0}^k \frac{\beta_i}{\Delta t} M u_i^{n-j} = -\frac{1}{\rho} D_i p^n + \nu K u_i^n + \sum_{j=1}^n a_j \left[ M f_i^{n-j} - (C u_i)^{n-j} \right], \quad (5.3)$$

where  $M$  is the mass matrix,  $C$  is the convection matrix,  $K$  is the stiffness matrix,  $D$  is the gradient matrix,  $i \in \{1, 2, 3\}$  are the spatial dimension indexes,  $n$  is the time level index, and  $k$  is the formal order of accuracy of the BDF/EX scheme. The pressure is decoupled

from the new velocity,  $u^n$ , by taking the divergence:

$$Kp^n = D_i \sum_{j=1}^n a_j F_i^{n-j}, \quad (5.4)$$

where  $F_i^n = Mf_i^n - (Cu_i)^n$ , which results in the Poisson pressure equation. Finally, the pressure is incorporated back into Equation 5.3:

$$\left[ \nu K + \frac{b_0}{\Delta t} M \right] u_i^n = -D_i \frac{p^n}{\rho} + \sum_{j=1}^k \left[ a_j F_i^{n-j} + \frac{b_j}{\Delta t} M u^{n-j} \right], \quad (5.5)$$

which results in three Helmholtz velocity equations.

These steps are the core of Nek5000 and NekBox: the explicit calculation of right-hand sides, a Poisson solver for the pressure, Equation 5.4, and a Helmholtz solver for the three components of the velocity, Equation 5.5.

### 5.3.2 Spectral element method

Nek5000 and NekBox implement SEM: a two-level discretization constructed from tensor products of Gauss-Lobatto-Legendre (GLL) quadrature points within elements and continuity across elements, forming a mesh. Fields are represented as

$$u(x, y, z) = \sum_{i=0}^p \sum_{j=0}^p \sum_{k=0}^p \tilde{u}_{i,j,k,e} h_i(x) h_j(y) h_k(z), \quad (5.6)$$

where  $p$  is the polynomial order of the method,  $e(x, y, z)$  is the index of the element in the mesh, and  $h_i(x)$  is the  $i$ th Lagrange polynomial through the GLL points of element  $e$ . The choice of Lagrange polynomials leads to diagonal mass matrices and related geometric factors. The spectral basis within each element enjoys exponential convergence with respect to the polynomial order. GLL points do not sample space uniformly, so concatenating elements is more effective at reducing grid spacing than increasing spectral order. Many

small elements are also better able to match complex geometries than fewer larger ones. The spectral element method is able to satisfy both the demand for geometric flexibility with quasi-uniform coverage and spectral convergence, but the particular choice of the spectral order versus the number of elements can be difficult to optimize.

In SEM, operators are written as the product of a local operator and *direct stiffness summation*, which enforces continuity at the shared element boundaries. The local operators are decomposed into tensor products of 1D operators. The general form of an operator  $A$  is:

$$A = (A_x \times I_y \times I_z) + (I_x \times A_y \times I_z) + (I_x \times I_y \times A_z), \quad (5.7)$$

where  $A_x, A_y, A_z$  are 1D projections of the operator  $A$  and  $I$  is the identity matrix. In this way, linear operators from  $R^{N \times N \times N} \rightarrow R^{N \times N \times N}$  can be evaluated in  $O(N^4)$  operations instead of  $O(N^6)$  [57]. This reduces the arithmetic intensity of operator evaluation in SEM to  $O(p)$ .

### 5.3.3 Computational profile

The spectral element method, as implemented in Nek5000 and NekBox, spends its time in three computational motifs: sparse communication, vector-vector, and matrix-matrix. The sparse communication comes from the direct stiffness summations and the coarse part of the pressure preconditioner. The vector-vector workload comes from inner products in the solvers and frequent rescaling by geometric factors, which are shaped like the diagonal mass matrix. The matrix-matrix workload comes from local operator evaluation.

The direct-stiffness summation is handled by a stand-alone library [29, 45]. In Nek, the pressure solve takes roughly 30% of the run-time, distributed between operator application, inner products, and the preconditioner. The preconditioner is multigrid with a local additive Schwarz part and the global coarse part [36]. In NekBox, the coarse part of the pressure

preconditioner is solved directly with FFTs or fast cosine transforms, and typically takes less than 5% of the total runtime. Local communication makes up a small portion of NekBox's run time at moderate numbers of points per processor, and Nek5000 and NekBox weak scale effectively to millions of ranks [25].

The efficiency of the vector-vector computation is generally left to the compiler, aided by aggressive loop merging in the solvers. For architectures that support them, the compiler needs help issuing non-temporal stores, which are performance optimal only if the working set is larger than the last level cache. These stores are used in parts of the solver and local element evaluation, and are discussed further in Section 5.4.

Matrix-matrix is the most accessible and performance critical portion of the workload. In particular, it is the only part of Nek that depends on the order, holding the total degrees of freedom (DOFs) fixed.

### 5.3.4 Order-dependent kernels

There are two matrix-matrix routines that sit inside of the iterative solvers: the Helmholtz operator and a basis transformation.

The Helmholtz operator is found on the left-hand side of Equation 5.4 and Equation 5.5:

$$Hu = (h_1 K + h_2 M)u,$$

where the special case of  $h_2 = 0$  is the Poisson operator.

```

1: procedure LOCAL HELMHOLTZ OPERATOR( $Hu, u, h_1, h_2$ )
2:    $(Hu)_{i,j,k} \leftarrow (G_x)_{i,j,k} * \sum_l (K_x)_{i,l} u_{l,j,k}$             $\triangleright$  matrix-multiply size  $(p^2, p, p)$ 
3:   for  $k = 0 \rightarrow p$  do
4:      $(Hu)_{i,j,k} += (G_y)_{i,j,k} * \sum_l (K_y)_{j,l} u_{i,l,k}$             $\triangleright$  matrix-multiply size  $(p, p, p)$ 
5:   end for
6:    $(Hu)_{i,j,k} += (G_z)_{i,j,k} * \sum_l (K_z)_{k,l} u_{i,j,l}$             $\triangleright$  matrix-multiply size  $(p, p^2, p)$ 

```

```
7:    $(Hu)_{i,j,k} += h_1(Hu)_{i,j,k} + h_2 M_{i,j,k} u_{i,j,k}$ 
```

```
8: end procedure
```

$G$  is a constant diagonal matrix derived from geometric terms and subscripts within parenthesis refer to spatial directions. Matrix sizes are given in BLAS notation: rows in result, columns in result, inner dimension.

The basis transformation is used to diagonalize the local Poisson operator in the overlapping Schwarz preconditioner, to restrict and interpolate the solution and residual in the multigrid preconditioner, and to dealias the convection operator.

```
1: procedure TRANSFORM( $v, u$ )  
2:    $f_{i,j,k} \leftarrow \sum_l (A_x)_{i,l} u_{l,j,k}$   $\triangleright$  matrix-multiply size  $(p^2, p, p)$   
3:   for  $k = 0 \rightarrow p$  do  
4:      $g_{i,j,k} \leftarrow \sum_l (A_y)_{j,l} f_{i,l,k}$   $\triangleright$  matrix-multiply size  $(p, p, p)$   
5:   end for  
6:    $v_{i,j,k} \leftarrow \sum_l (A_z)_{k,l} g_{i,j,l}$   $\triangleright$  matrix-multiply size  $(p, p^2, p)$   
7: end procedure
```

## 5.4 Kernel Analysis and Optimization

### 5.4.1 Small Matrix Multiplications

The implementation of fast matrix multiplications, i.e., the BLAS library's GEMM routines, and dense linear algebra more generally is one of computer science's best studied fields. However, large matrices [20] have been the primary focus and, as a result, vendor-tuned BLAS implementations do not provide optimal performance when used for the small GEMMs in Nek. Several BLAS libraries recently introduced so-called batched interfaces to speed-up series of independent and small multiplications by exploiting parallelism and amortizing calling overheads [28]. As Nek performs dependent GEMMs within each element, batched

execution would necessarily be inter-element, inhibiting important caching optimization and consuming significantly more memory bandwidth. Therefore, most of Nek’s computer science related work was devoted on speeding up small GEMMs [51]. Parts of Nek5000 and the related NekCEM codes have been independently ported to OpenACC [37, 45] to speed-up small GEMMs.

Today, Nek5000 and NekBox ship with a FORTRAN-based matrix-matrix implementation called `mxm_std`. By default, `mxm_std` explicitly defines multiple interfaces corresponding to values of the inner dimension  $k$ , and provides unrolled FORTRAN primitives to the compiler. For IBM’s BlueGene series, common sizes are manually implemented for best performance in FORTRAN assembly-intrinsics in `mxm_bgq`. Similarly, `mxm_std` features some special case optimizations targeting AMD’s Opteron processor, which is used in the United States’ largest system, Titan, at Oak Ridge National Laboratory.

In order to ensure the best possible performance on a range of modern Intel processors, featuring different versions of Advance Vector Extensions (AVX) instructions, we would need to conduct a long and complicated tuning effort of Nek’s `mxm_std` akin to the narrow customizations already present. Instead, we integrated an early prototype of the LIBXSMM library [1, 21] into NekBox. LIBXSMM provides highly-optimized single-threaded small matrix-multiplication routines tuned for all recent Intel processors. It is already successfully used in the quantum chemistry application CP2K and high-order finite element seismic wave equation solver SeisSol [6].

In contrast to `mxm_std`, LIBXSMM creates a specific kernel implementation for each small matrix multiplication size and optimizes that kernel specifically for each set of vector extensions. Each kernel is composed from a priori known and best-performing basic blocks. Remainder handling can be performed either explicitly by application-side padding or internally by slightly less efficient fill-in basic blocks. We rely on the latter in our integration of LIBXSMM into NekBox.

Listing 5.1: Integration of LIBXSMM into NekBox’s element-local Helmholtz operator. `xmm1`, `xmm2`, `xmm3` are persistent functions pointers to amortize LIBXSMM’s dispatching overhead. The `libxsmm_dispatch` call JITs the requested kernel and populates the persistent function pointers.

```

logical, save :: init = .false.
type(LIBXSMM_DMM_FUNCTION), save :: xmm1, xmm2, xmm3

! lazy initialization of function-private function pointers
! to eliminate dispatching overhead
if (.not. init) then
  call libxsmm_dispatch(xmm1, nx, ny*nz, nx, 1.0_dp, 0.0_dp)
  call libxsmm_dispatch(xmm2, nx, ny, ny, 1.0_dp, 0.0_dp)
  call libxsmm_dispatch(xmm3, nx*ny, nz, nz, 1.0_dp, 0.0_dp)
  init = .true.
endif

! element-local operation
call libxsmm_call(xmm1, C_LOC(wddx), C_LOC(u(1,1,1)), C_LOC(work1))
do iz=1,nz
  call libxsmm_call(xmm2, C_LOC(u(1,1,iz)), C_LOC(wddy), C_LOC(work2(1,1,iz)))
enddo
call libxsmm_call(xmm3, C_LOC(u(1,1,1)), C_LOC(wddz), C_LOC(work3))

! element update
au(:,:,:) = h1* ( work1*gx + work2*gy + work3*gz ) + h2*b*u

```

We leverage LIBXSMM’s experimental just-in-time (JIT) compilation feature to adapt at runtime to Nek’s spectral order. The JIT feature generates a small matrix multiplication when its size is requested for the first time and caches compiled code until the application process is terminated. Additionally, LIBXSMM can expose the function pointer to the application to bypass future dispatches when call patterns are simple.

As an example, we provide the integration of LIBXSMM into NekBox’s local Helmholtz kernel from Section 5.3.4 in Listing 5.1. This fragment is called within a loop over elements that is typically long enough to amortize overheads. When entering the element-local operator for the very first time, we request the required kernels from the LIBXSMM library, which JIT compiles them internally, and store the corresponding functions pointers into persistent variables to avoid dispatching in subsequent calls. Compared to the pseudo-code fragment, cf. 5.3.4, we use temporary buffers to separate matrix-matrix from vector-vector operations, which are performed in one step at the end of each element. The other common matrix-matrix motifs, basis transformation in particular, are optimized analogously.

### 5.4.2 Enhancing Element Update Performance by Streaming Stores

Caches in Intel processors are designed as write-back caches with read-for-ownership (RFO). Therefore, writing to a vector in main memory costs two operations: a load into the cache and the write. Nek performs many such element updates, cf. Listing 5.1, and long vector updates in linear solvers. Compiling the Helmholtz element update leads to 5 streams being explicitly read (`gx`, `gy`, `gz`, `b`, `u`), one RFO of `au` and one write of `au`. As we stream through all elements the RFOs are harmful for two reasons: a) they consume bandwidth and therefore can cause a  $\approx 16\%$  performance drop; and b) they unnecessarily occupy cache space and might evict useful data.

Since the SSE2 instruction set, the Intel architecture offers so-called non-temporal stores (NTS). These special instructions write data directly into main memory without generating RFOs and consuming cache. They operate best when being executed on vector-length aligned addresses, as cache-line splits are impossible. The compiler cannot fulfill the alignment requirement for all orders, because Nek stores field data compactly, which prohibits semi-automatic generation of NTS. Therefore, we implemented a FORTRAN interface module with a C-backend and x86 intrinsics that applies loop-peeling to leverage NTS for the majority of stores in long, potentially unaligned updates. This module covers the important kernels of Nek by offering NTS-enhanced primitives to: a) set an 1d-array to a fixed value b) copy an 1d-array c) multiply component-wise an 1d array, and d) perform the Helmholtz element update, including the special case of the Poisson operator,  $h_2 = 0$ . For case b), Listing 5.2 depicts Intel AVX2 code.

### 5.4.3 Discussion of Performance Reproducers

In order to analyze the performance of LIBXSMM integration and the NTS module, we have implemented standalone reproducers of the identified small matrix multiplication motifs. They are included in the LIBXSMM library as examples and performance tests. In contrast

Listing 5.2: Loop peeling approach including determining the middle section for which aligned NTS instructions can be used.

```

void stream_vector_copy( const double* i_a,
                        double*          io_c,
                        const int         i_length) {

    int l_n = 0;
    int l_trip_prolog = 0;
    int l_trip_stream = 0;

    /* init the trip counts to determine aligned middle section */
    stream_init( i_length, (size_t)io_c, &l_trip_prolog, &l_trip_stream );

    /* run the prologue */
    for ( ; l_n < l_trip_prolog; l_n++ ) {
        io_c[l_n] = i_a[l_n];
    }

    /* run the bulk, using streaming stores */
    for ( ; l_n < l_trip_stream; l_n+=8 ) {
        _mm256_stream_pd( &(io_c[l_n]), _mm256_loadu_pd(&(i_a[l_n])) );
        _mm256_stream_pd( &(io_c[l_n+4]), _mm256_loadu_pd(&(i_a[l_n+4])) );
    }

    /* run the epilogue */
    for ( ; l_n < i_length; l_n++ ) {
        io_c[l_n] = i_a[l_n];
    }
}

```

to NekBox, they are parallelized via OpenMP instead of MPI, but the performance agrees within 10% of a full NekBox execution at scale. We used a single node of the Cray XC40 and BlueGene/Q, cf. Section 5.5.1, for generating performance data in this section.

Figure 5.1 compares the performance of Intel MKL 11.2.1, Nek’s own `mxm_std`, and LIBXSMM with and without non-temporal stores. For all element sizes, LIBXSMM offers the best performance, but the difference for orders  $\leq 16$  are very small as the execution is heavily memory bandwidth bound. A significant boost is possible by leveraging NTS: we are able to sustain 100% of the STREAM triad bandwidth (101.6 GiB/s) up to an element size of 16. For larger problems, the small GEMM performance is more important. Here LIBXSMM is up to 2 $\times$  faster than `mxm_std` und up to 40% faster than Intel MKL.

In case of very low orders the benefit of NTS is greater than 16%, which we attribute to NTS avoiding cache pollution. For medium sized orders we exactly see the expected 16%, and large problems have additional bandwidth available such that RFOs are less harmful.

The performance numbers for the basis transformation on Shaheen are comparable to

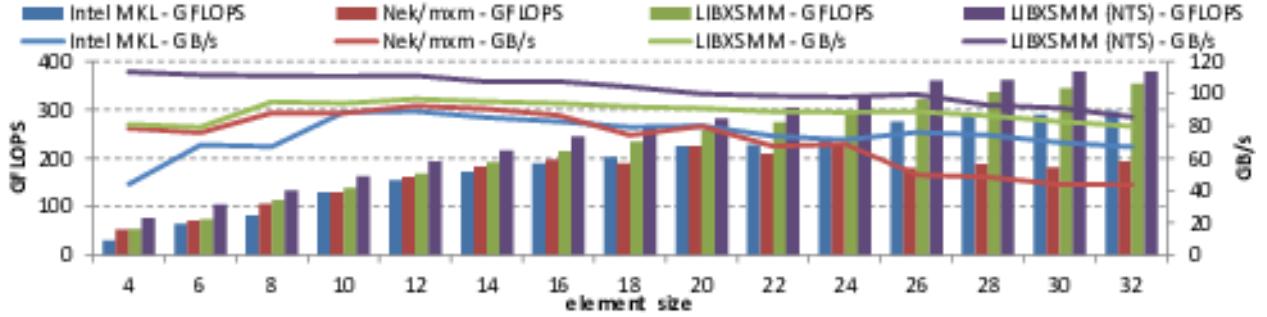


Figure 5.1: Performance of the Helmholtz reproducer running on a single node of Shaheen for different implementation of small matrix multiplications. NTS denotes the usage of the non-temporal store optimized module.

the Helmholtz operator and therefore not plotted. To summarize them, LIBXSMM-based GEMMs are the fastest and, due to higher computational demand, NTS are only important of for very small 1d sizes. LIBXSMM is able to achieve 50% of maximum floating-point performance for moderate orders. LIBXSMM ranges from 4× faster than `mxm_std` and Intel MKL at the smallest order to 40% faster at the largest.

The performance of the Helmholtz kernel is representative of the basis transformations kernel on Mira as well. To compare with Shaheen, Figure 5.2 repeats the Helmholtz operator reproducer experiment on a single node of Mira. IBM ESSL version 5.1.1 is used as the vendor library in place of Intel MKL. In place of LIBXSMM, `mxm_bgq`, which features QPX SIMD instructions, is used for the sizes that it supports. When no QPX implementation is available, `mxm_bgq` falls back to `mxm_std`. Up to element size 16, Nek’s `mxm_std` and `mxm_bgq` libraries are a better choice compared to IBM ESSL. For larger element sizes (except 22 and 24) the performance is comparable. However, the fraction of available bandwidth used is significantly worse than on Shaheen. Even at high element sizes, Shaheen is at 80% bandwidth utilization with LIBXSMM and 50% without, whereas Mira runs at 17%. The relative efficacy of `mxm_bgq` on Mira, where available, highlights the strength of LIBXSMM: the ability to automatically issue the best available vector instructions at any size.

Figure 5.3 depicts corresponding performance numbers for the basis transformation re-

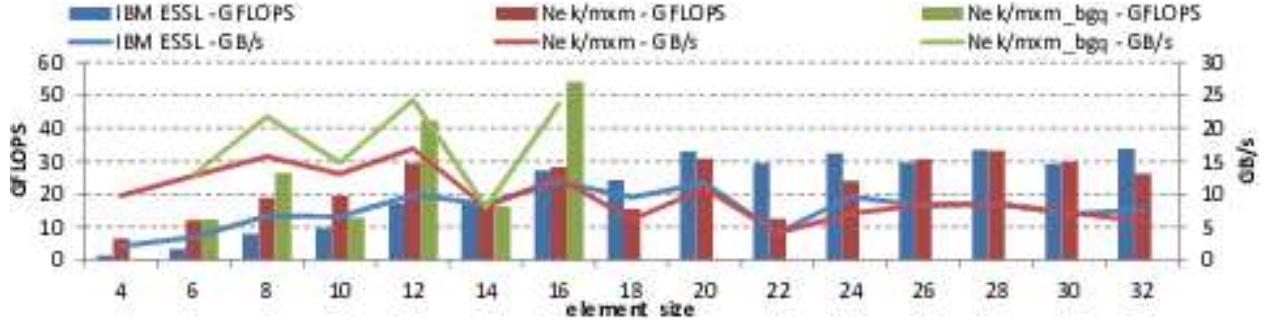


Figure 5.2: Performance of the Helmholtz operator reproducer running on a single node of Mira for different implementation of small matrix multiplications.

producer in three use cases: a) unitary transformation from element size to element size, b) prolongation/dealiasing from 1d size to  $(3/2)$  the element size, and c) restriction/aliasing from 1d size to  $(2/3)$  the element size. Note that the  $(3/2)$  factor implies some dimensions are significantly larger than the element size shown on the x-axis.

As with the Helmholtz reproducer, the LIBXSMM-based executions are the fastest and due to higher computational demand; NTS are only important for very small 1d sizes. LIBXSMM is able to achieve 50% of maximum floating-point performance for medium sized orders. In direct comparison to `mxm_std` and Intel MKL, the speed-up of LIBXSMM varies between close to  $4\times$  at very small order to roughly 40% at very large order.

## 5.5 Scenarios and Performance

### 5.5.1 Architectures

We run on two supercomputers: Mira at the ALCF and Shaheen XC40 at the KSL. Mira is a IBM BlueGene/Q with 49,152 nodes. Each node has 16 cores with 4 hardware threads per core and can support 204.8 GFLOPS and 30 GiB/s main memory bandwidth, measured by [39]. Shaheen is a Cray XC40 with 6144 nodes. Each node has two Intel<sup>®</sup> Xeon<sup>®</sup> E5-2698v3 (code-named Haswell) processors with 16 cores each and can support around 1177.6 GFLOPS and 101.6 GiB/s main memory bandwidth, measured by [39]. Shaheen's

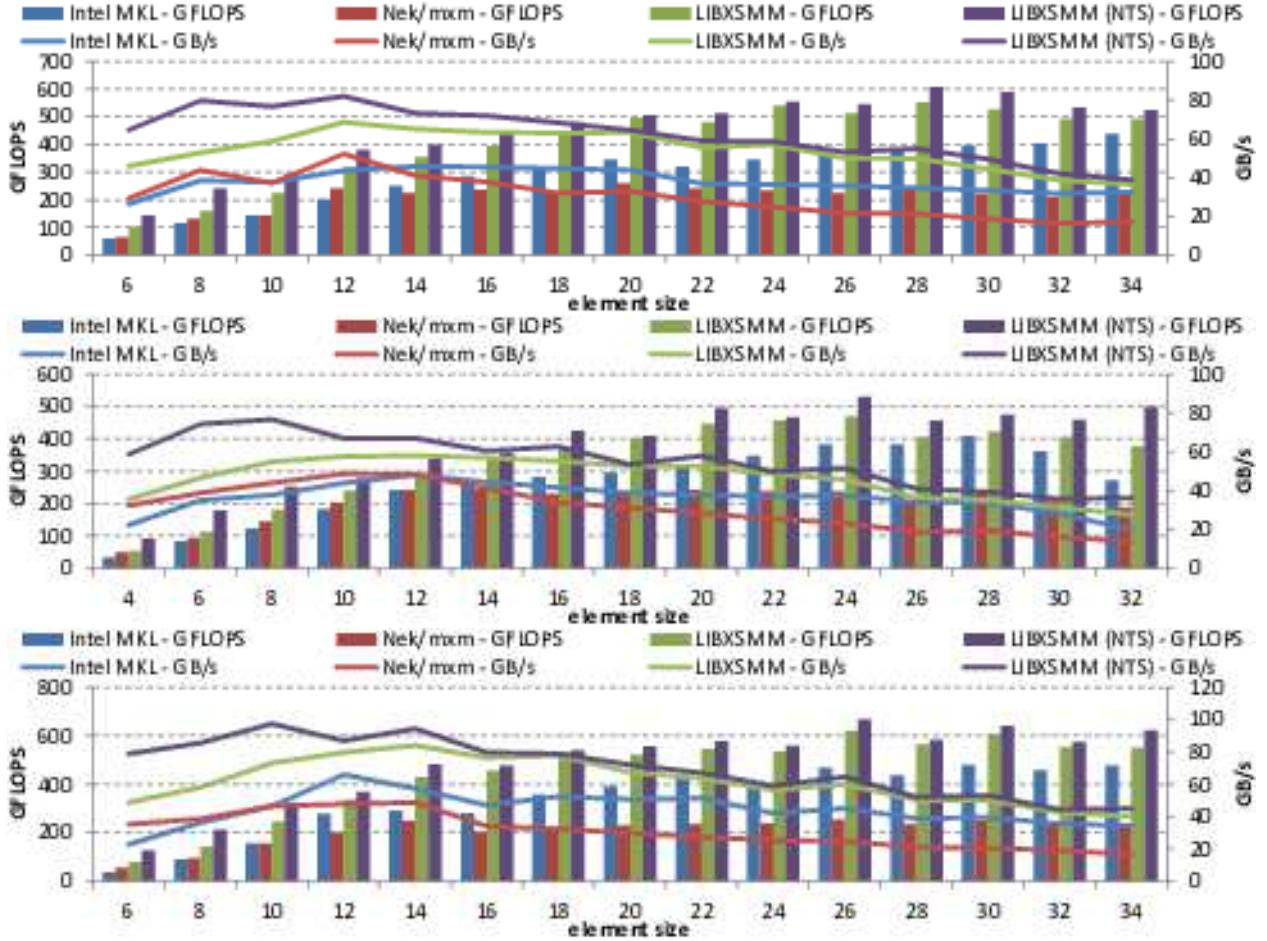


Figure 5.3: Performance of the basis transformation reproducers using different implementation for the small matrix multiplications. NTS denotes the usage of the aforementioned non-temporal store optimized module. The top plot shows the diagonalization in the local Poisson operator, the middle one the prolongation and the bottom one the restriction case.

cores therefore have  $2.9\times$  the floating point and  $1.7\times$  the memory bandwidth of Mira's BlueGene/Q cores.

### 5.5.2 Single mode Rayleigh-Taylor instability

The Rayleigh-Taylor instability (RTI) occurs when the pressure and density gradients point in opposite directions, as in the canonical case of a heavy fluid supported on top of a lighter fluid in a gravitational field. The Rayleigh-Taylor growth rate is an increasing function of the wave-number, up to a viscous cutoff, making the smallest scales grow fastest. Because

energy is pumped into the system at small scales, the RTI is notoriously difficult to model numerically [14].

The RTI describes how the dense fluid is pushed through and mixes with lighter fluid. This dynamic mixing process is essential to the behavior of flows found in exploding stars [4], the oceans and atmosphere [34], and inertial confinement fusion. In the latter case, dense plastic ablator is pushed into and mixed with the lighter hydrogen fuel. The carbon-laden ablator radiates energy much more quickly than the fuel, reducing hot-spot temperature and preventing ignition. The study of the RTI and related mixing is a priority research direction for inertial confinement fusion performance [18].

Nek5000 and NekBox [26] are used to model the incompressible Boussinesq equations, which approximate the RTI at low density contrasts:

$$\frac{\partial u}{\partial t} + u \cdot \nabla u = -\nabla p + \nu \nabla^2 u + \tilde{g}T \quad (5.8)$$

$$\frac{\partial T}{\partial t} + u \cdot \nabla T = \alpha \nabla^2 T \quad (5.9)$$

$$\nabla \cdot u = 0, \quad (5.10)$$

where  $T$  is a scalar that can be interpreted as a temperature, in which case  $\alpha$  is the thermal diffusivity and  $\tilde{g}$  is the product of the gravitational acceleration and the thermal expansion coefficient.

The single-mode Rayleigh-Taylor instability (smRTI) restricts the initial perturbation of the interface to be sinusoidal, and is generally considered in periodic span-wise boundary conditions:

$$T(x, y, z, 0) = A \cdot \text{erf} \left[ \frac{z + a_0 \cos(2\pi x/\lambda) \cos(2\pi y/\lambda)}{\delta} \right], \quad (5.11)$$

where  $A \in (0, 1]$  is the Atwood number,  $\lambda$  is the wavelength,  $a_0$  is the initial interface amplitude, and  $\delta$  is the initial interface width. This simplification allows the problem to be defined by only two dimensionless numbers in the limit of  $a_0, \delta \rightarrow 0$ , the Grashof number

(Gr) and the Prandtl number (Pr):

$$\text{Gr} = \frac{A\tilde{g}\lambda^3}{\nu^2}, \quad \text{Pr} = \frac{\nu}{\alpha}. \quad (5.12)$$

Even under these simplifications, the late-time behavior is not well understood. Experiments are prone to spurious low-wavelength modes that dominate the dynamics at late times, while the cost of direct numerical simulations is quadratic with the domain's aspect ratio.

It would be valuable to systematically sample the Grashof-Prandtl space with high fidelity simulations at late-time/high-aspect-ratio to better inform experimental design and model development. Such a study would be very expensive, so it is important to select a cost-minimizing strategy.

We take this problem, the selection of a cost-minimizing strategy for the late-time smRTI, as our motivation. In addition to the isolated reproducers discussed in Section 5.4, we present NekBox application benchmarks based on smRTI with typical Nek settings. The aim of these benchmarks is to identify minimum cost discretizations that attain a given accuracy.

The benchmarks are conducted for combinations of the element size taken from

$$\{4, 6, 8, 10, 12, 14, 16, 32\}, \quad (5.13)$$

and span-wise mesh size taken from

$$\{2, 4, 8, 12, 16, 24, 32, 48, 64, 96, 128\}. \quad (5.14)$$

The total number of points ranges from around 1 million to 4 billion. The problem is weak-scaled: the number of elements per rank is chosen as to consume approximately half of the available main memory, or around 16k and 262k points per rank on Mira and Shaheen, respectively. The problems are constrained to fill an integer number of nodes, which puts a

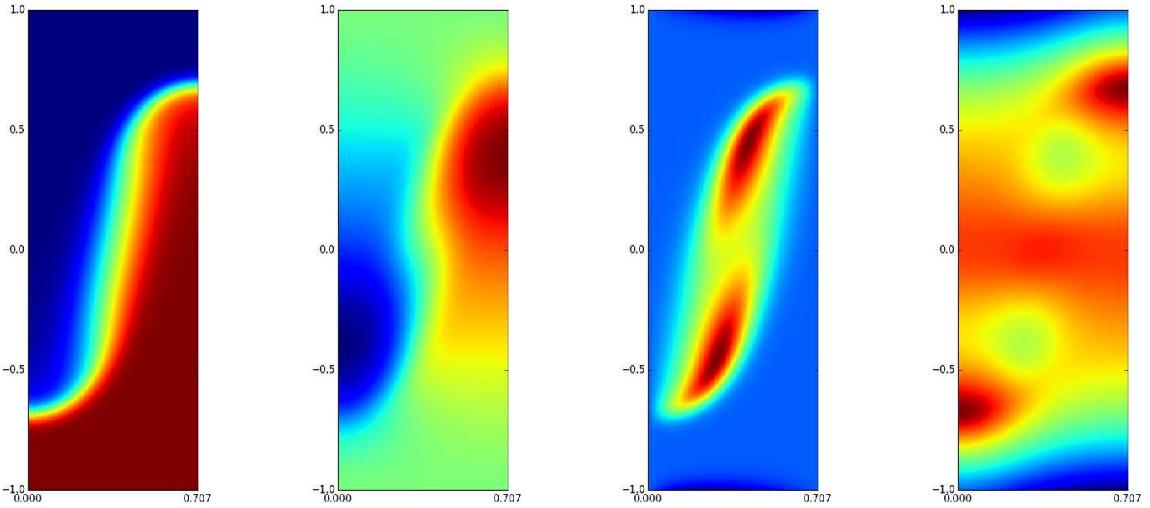
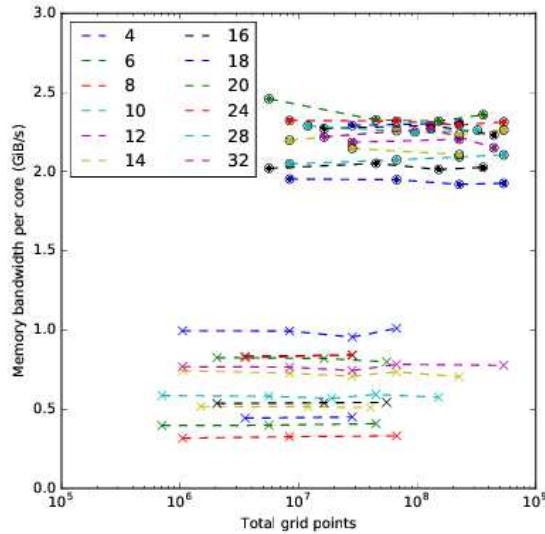


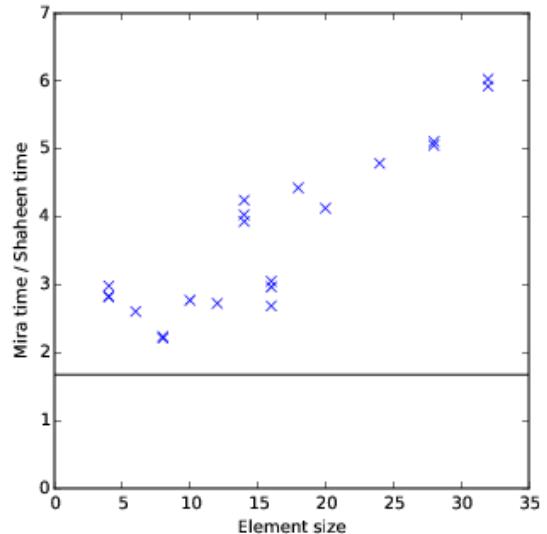
Figure 5.4: Scalar, vertical component of the velocity, vorticity component out of the plane, and pressure fields at end of simulation. The color scales are dimensionless, linear, and centered at the mean value over the image.

lower bound on the mesh size and excludes some cases that would partially fill nodes. The domain is a box with dimension  $[0, 0.5]^2 \times [-1, 1]$ , and the elements are cubic. The spanwise boundary conditions are symmetric and the vertical boundary conditions are no-slip in velocity and no-flux (insulating) in the scalar. The initial condition is stationary in velocity with a scalar given by Equation 5.11, the Grashof number is 17,324, and the Prandtl number is 1. The timestep is calculated based on a Courant number of 0.4, which scales linearly with the number of elements and quadratically with the size of the element due to the spacing of the GLL nodes. The Courant condition is defined only in a linear limit, so during the initial exponential growth regime the Courant number is computed using the stagnation velocity,  $\sqrt{A\tilde{g}/(\pi\lambda)}$ .

Outputs are written at regular intervals in simulated time, constant across problem sizes. Therefore, smaller problems perform a greater share of I/O, as is the common case in CFD. Nek5000 and NekBox write separate files for separate ranks. The number of ranks that participate in I/O is a fixed proportion of the total number of ranks.



(a) Bandwidth



(b) Ratio

Figure 5.5: Weak scaling of bandwidth on Shaheen and Mira. In (a), Circles and crosses indicate memory bandwidth per core on Shaheen and Mira, respectively, vs. the problem size labeled by element size. In (b), the ratio of the bandwidths are shown vs. element size for common discretizations. The solid line indicates ratio of STREAM memory bandwidth.

Slices of the end of the simulation are shown in Figure 5.4. Two observables are calculated in post-processing: the *bubble height* and the *mix volume*:

$$H = \sup \left\{ z : \min_{x,y} T(x,y,z) < T_0 \right\}, \quad \Theta = \int |T - T_0| dV, \quad (5.15)$$

where  $T_0$  is the volumetric average temperature. These two observables are common to smRTI models and lie at opposite ends of the locality spectrum: the bubble height is defined by the neighborhood of the bubble tip while the mix volume is an integral over the entire domain. The root mean square error in each observable is computed over all the outputs.

### 5.5.3 Time to accuracy

For each simulation, we compute the FLOP rate and aggregate memory bandwidth. NekBox includes explicit FLOP and memory operation counters and timers in the most performance critical regions of the code. Memory operations are counted assuming single-element intermediate data stays in cache, and therefore does not contribute to main memory bandwidth. These counters are consistent with those used in the reproducers. The whole application is not covered, so the counters can be considered lower bounds on the whole-application performance.

The attained memory bandwidth per core on Shaheen and Mira are plotted in Figure 5.5. On Shaheen, bandwidth is constant with respect to the number of elements and a weak function of the order, ranging from around 65 to 75% of peak. On Mira, bandwidth is still constant with respect to the scale, but varies more strongly with polynomial order, especially at orders greater than 16 and those not divisible by 4. It ranges from around 15 to 50% of peak. The `mxm_bgq` library, discussed in Section 5.4, is used, resulting in performance spikes at QPX-supported orders, e.g., 8.

The accuracy is plotted versus the computational cost for a variety of discretizations in Figure 5.6. The error in bubble height and mix volume are strongly correlated, so only the error in the height is plotted. As expected, doubling the spectral order while keeping the number of elements fixed, e.g.,  $(4, 32) \rightarrow (8, 32)$  and  $(8, 8) \rightarrow (16, 8)$ , significantly improves the accuracy, but also increases the cost by 16-32 $\times$ . The first 8 $\times$  is due to an increase in the number of degrees of freedom, the next 2 $\times$  is due to the shorter timestep, and, when compute-bound, the final 2 $\times$  is due to an increase in the floating point load. Doubling the spectral order while keeping the number of points fixed, e.g.,  $(16, 8) \rightarrow (32, 4)$  and  $(8, 8) \rightarrow (16, 4)$ , increases the cost by 2-4x, as expected, but also improves the accuracy. Doubling the spectral order while halving the number of points in each direction, e.g.,  $(8, 32) \rightarrow (16, 8)$  and  $(14, 16) \rightarrow (28, 4)$ , reduces the cost by 4-8 $\times$  while maintaining or slightly improving the

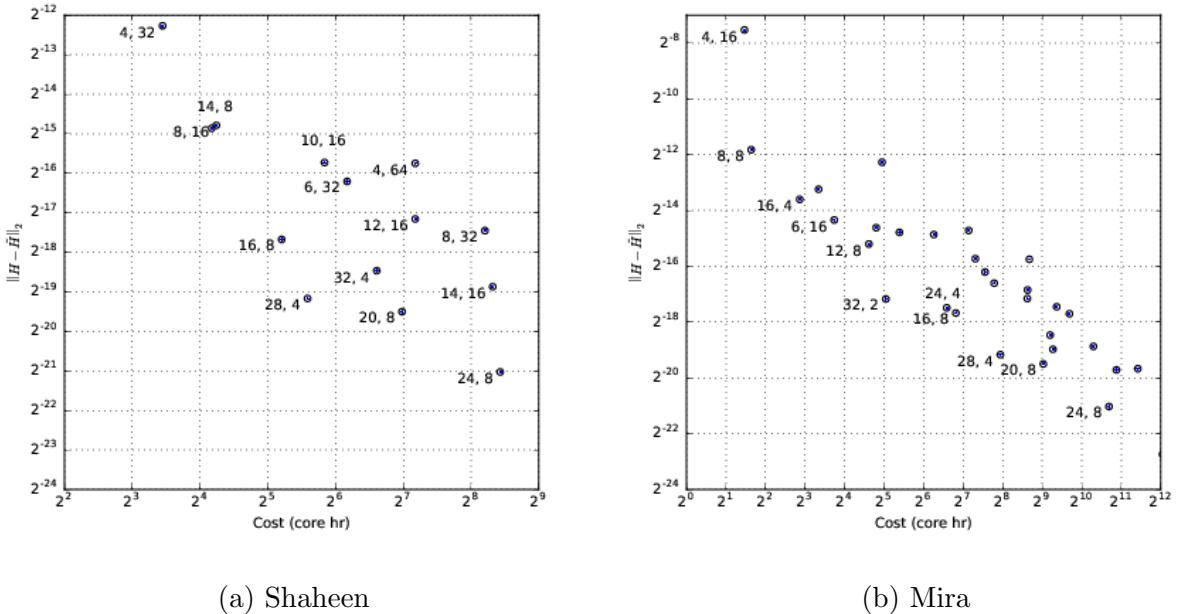


Figure 5.6: Error with respect to bubble height, Equation 5.15, vs. the computational cost, in processor hours, on Shaheen (a) and Mira (b). Points are labeled as  $(p+1, e)$  pairs, where  $p$  is the order,  $p+1$  is the element size, and  $e$  is the number of elements in one dimension. More runs are present on Mira due to the smaller BGQ nodes evenly dividing more problem sizes.

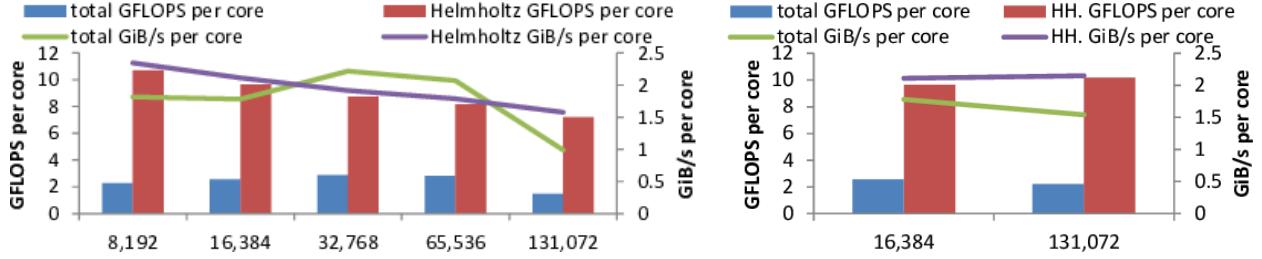


Figure 5.7: Strong scaling (left) and weak scaling (right) on Shaheen on up to 131,072 cores (2/3) of the full 7 PFLOPS machine using an element size of 32. To avoid log plots we show per-core performance.

accuracy.

We define the *efficiency frontier* as the set of discretizations that minimize computational cost for fixed accuracy or, equivalently, minimize error given fixed computational cost. The efficiency frontiers on Mira and Shaheen are comprised of discretizations with very high orders, given our constraints. The most efficient schemes are those with element size greater than 16, except for very low accuracy simulations.

#### 5.5.4 Whole application performance

To date, our largest calculation on Shaheen occupied 131,072 cores, as depicted in Fig. 5.7 for element size 32. NekBox achieved 197 TiB/s memory bandwidth and 290 TFLOPS in weak scaling. This corresponds to 47.8% of peak memory bandwidth sustained over the entire application at high order. In case of strong scaling these numbers are slightly lower with 130 TiB/s and 195 TFLOPS. However, the Helmholtz operator, as the most compute intense subroutine, is able to achieve up to 0.94 PFLOPS in strong and 1.33 PFLOPS in weak-scaling on 131,072 cores. We also consider 65,536 cores runs, occupying 1/3 of Shaheen. These runs achieved at least 135.6 TiB/s memory bandwidth and 184.9 TFLOPS. This corresponds to 67.5% of peak memory bandwidth sustained over the entire application at high order. Finally, extrapolating to full machine, NekBox would reach at least 406.8 TiB/s and 554.6 TFLOPS. At the same scale, a weak scaling of the Helmholtz operator would result into

1.9 PFLOPS out of 7 PFLOPS performance.

## 5.6 Conclusion

NekBox enhanced by LIBXSMM generated kernels on Shaheen XC40 executes the performance critical, order-dependent components of Nek above 80% of peak memory bandwidth. For comparison, compiled code on the BlueGene/Q architecture is only able to reach 50% of peak and for many polynomial orders operates around 30%. Therefore, despite only having  $1.7\times$  the memory bandwidth, Shaheen's cores outperform Mira's cores by  $3\text{-}6\times$ , with the greatest improvement at high order and for sizes that are not divisible by the vector width, 4 in this case. NekBox is able to scale 67.5% utilization rates to 65,536 cores on Shaheen.

For the smRTI, the efficiency frontier, i.e., the discretizations that minimize cost given accuracy or minimize error given cost, have polynomial orders between 15 and 31, higher than are typically used in spectral element schemes. The presence of high order schemes on the efficiency frontier can be understood by the combination of two effects. First, the increase in arithmetic intensity is hidden by the imbalance between floating point capabilities and memory bandwidth, providing high order at no marginal cost on a per time step basis. Second, higher order schemes with fewer degrees of freedom are more accurate than lower order schemes with more degrees of freedom. It is generally possible to maintain accuracy by increasing the order while decreasing the total degrees of freedom, and, consequently the total cost.

Generally the order should be chosen to be at least large enough to saturate the floating point capabilities of the architecture in the order-dependent kernels, because increasing the order to that point significantly improves accuracy at no marginal computational cost. On BlueGene/Q, this mark is polynomial order 15, while on the Cray XC40 it is 31.

For many problems and observables, the calculation may additionally benefit from increasing the order until just before single-element operations spill out of cache. The im-

provement in accuracy is exponential with the polynomial order, so the degrees of freedom needed to achieve a level of accuracy can decrease. The increase in the cost with respect to order for compute-bound orders is linear, so if the decrease in the number of degrees of freedom needed is super-linear, the net result is a less expensive calculation. Usage in this way, which exceeds the largest element sizes that we ran on Shaheen, warrants further study.

More generally, high order methods with high locality, the structured elements in SEM being only one example, are able to take advantage of wider vectors and higher compute to memory ratios to reach higher order at little to no marginal cost on a per-step basis. However, increases in cost can come in through coupling to the choice of time-step and an increase in iteration counts in the solvers. These increases can often be mitigated by reducing the total number of degrees of freedom, relative to an equivalent lower-order calculation.

The next generation will include supercomputers featuring the Xeon Phi processor code-named Knights Landing, e.g., Cori at NERSC with more than 20 PFLOPS. As the architecture continues to evolve, we can see that updated node-level optimizations and order-sensitivity studies are key to helping scientists continue to perform large scale, high efficiency simulations.

## Acknowledgment

This research used the resources of the Supercomputing Laboratory at the King Abdullah University of Science & Technology (KAUST) in Thuwal, Saudi Arabia. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

We acknowledge useful conversations with Paul Fischer, James Lottes, Aleksandr Obabko, Oana Marin, Michel Schanen, Scott Parker, Vitali Morozov, Matthew Otten, and Robert Rosner.

Optimization Notice: Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, compo-

nents, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>.

Intel, Xeon, and Intel Xeon Phi are trademarks of Intel Corporation in the U.S. and/or other countries.

# CHAPTER 6

## CONCLUSIONS

The following sections summarize the main contributions made by this thesis to the study of the Rayleigh-Taylor instability. The last section discusses questions that this thesis has been unable to address but could be fruitful avenues of future study.

### 6.1 New model for low-Atwood single mode

I have presented a simple model for the low-Atwood single mode Rayleigh-Taylor instability. The model has two components: an ordinary differential equation for the bubble height based on buoyancy and drag and an analytic expression for the volume of mixed fluid as a function of the bubble height, wavelength, time, and initial interface thickness. The model components are coupled by an analytic expression for the effective Atwood number as a function of the pure Atwood number, bubble height, wavelength, and volume of mixed fluid. The model components are separable in that changes to the mixing model formulation shouldn't significantly affect the dynamics model so long as the changed model is accurate, and vice versa.

The model has 8 unique coefficients. Three of the coefficients can be solved for by constraining the model to match the linear theory, i.e., the limit of small amplitude bubbles and thin interfaces. The remaining 5 parameters can be estimated by physical arguments. One is related to a form drag coefficient. Another is related to a friction factor. Two give the ratios of alternative definitions of bubble volumes, e.g., based on the scalar versus the velocity, and are nominally unity. The last scales the scalar interface surface area and is related to the shape of the bubble, e.g., cylindrical vs rectangular.

The estimation of the model coefficients is insufficient to evaluate the model's accuracy, both qualitatively in the representation of known features of the flow and quantitatively in its

predictiveness of the bubble height and mixed volume as a function of time. Neither does the model formulation establish the accuracy of the coefficient estimates or their independence of Grashof and Schmidt numbers. To answer these questions, we turn to direct numerical simulations to generate a dataset of low-Atwood single mode Rayleigh-Taylor experiments over a range of Grashof and Schmidt numbers.

## 6.2 Validation of DNS

Before performing the direct numerical simulations, their methodology is validated against experimental data from drop tower experiments of Wilkinson and Jacobs [62]. The simulations solve the incompressible Navier–Stokes equations for a single fluid. The density is modeled with the Boussinesq approximation, turning it into an active scalar coupled to the flow via the advection-diffusion equation. The incompressible Navier–Stokes and advection-diffusion equations are solved with the spectral element method (SEM). The method is implemented in the NekBox code, a descendant of Nek5000 specifically tuned for box geometries.

The simulations make three key approximations. First, the Boussinesq approximation, which is equivalent to the limit of zero Atwood number for the fixed product of the Atwood number and local acceleration,  $Ag$ . The experiments have Atwood numbers around 0.15, which is greater than the generally accepted Boussinesq range. Second, the experiments were miscible but with Schmidt numbers greater than 1000. For Schmidt numbers greater than unity, the cost of the simulation goes with the Schmidt number to the fourth power. Therefore, the simulations were performed with Schmidt numbers of 1, 3.5, and 7. Finally, the experiments established an initial perturbation by exciting a standing wave in a stable density stratification. However, only the initial interface amplitude and velocity were measured. These measurements were transformed into a quiescent initial condition by assuming the early time dynamics of the linear theory. If the single mode dynamics depend on the

initial condition, this could be a substantial approximation.

The simulations were in agreement with the experiments, both qualitatively with respect to the presence of stagnation and re-acceleration, and quantitatively with respect to the stagnation velocity and bubble height at the onset of re-acceleration. The highest Grashof number simulation was unstable, which helped characterize the resolution requirements for the simulation. The Boussinesq approximation, low Schmidt number approximation, and linear initial condition approximations preserve the re-acceleration phenomena. More generally no physical processes beyond the incompressible Navier-Stokes with an active scalar are necessary to produce stagnation and re-acceleration, so they can be studied with the spectral element method in NekBox.

In addition to validation, one of the simulations was repeated in a tank with twice the vertical extent to study late time dynamics and the interaction with the tank boundaries. Three distinct effects of the finite size tank were observed. First, the initial condition had a long-wavelength mode imposed along the diagonal that preferred the growth of bubbles in one corner and spikes in the opposite corner. Second, the bubbles and spikes adjacent to the walls experienced drag along the walls. Third, the bubbles and spikes lifted off the walls towards the center of the tank, which squeezed the interior bubbles forward.

These boundary interactions limit the bubble height that can be reached in drop-tank experiments to about one tank diameter. The technique used by Wilkinson and Jacobs to drive the initial standing wave suffers from attenuation at high wave-numbers and is limited to 4.5 modes across the diagonal. Together, these limit the bubble aspect ratio,  $h/\lambda$ , to 5 to 10, at which point the high Grashof bubbles in the experiment are still accelerating. To experimentally access the late-time dynamics of periodic arrays of bubbles requires techniques to drive initial perturbations with many more modes.

### 6.3 Convergence and performance

The calculations required to generate data to fit the model are very expensive. The cost of simulating a trajectory until the bubble stops rising goes with the Rayleigh number to the 6th power. Of the Rayleigh numbers within the scope of the model, less than half in log-space are practical to complete. Therefore, it is important not only to improve the performance of the method but also to avoid over-resolving the flow.

The spectral element method has two resolution parameters: the size of the mesh of elements and the polynomial order of each element. Those two parameters were optimized empirically. A test case was designed at high Rayleigh number and unit Schmidt number, but with a shorter vertical extent than the full problem. For each of many combinations of mesh resolution and polynomial order, the flow was evolved to a fixed time corresponding to a nondimensional bubble height  $h/\lambda \approx 1$ . The error in the bubble height and mixed volume was compared to the computational cost. It was found that very high spectral orders, between 19 and 31, provided the highest accuracy for a given computational cost or, conversely, were the lowest cost for fixed accuracy.

Production simulations were performed at polynomial order 31, with a mesh size dependent on the Grashof and Schmidt numbers. Based on the results of the simulations, we estimate the cost of a complete trajectory as:

$$f_T \approx 3.54 \times 10^{-23} \text{Ra}^6 \text{ core hours,} \quad (6.1)$$

where core hours are based on Shaheen XC40 at the King Abdullah University of Science and Technology (KAUST).

## 6.4 Flow regimes

The terms in the model naturally divide the flow into four regimes: the exponential growth regime, the saturation regime, the viscous regime, and the diffusive regime. The onset of each regime with respect to the bubble height can be estimated by scaling arguments. The exponential growth regime is governed by the linear theory and well studied. The linear growth saturates at a fixed aspect ratio, taken here to be  $h/\lambda = 0.05$ , independent of the Rayleigh and Schmidt numbers. The saturation limits the acceleration to  $Ag/C_3$ , with  $C_3 \approx 1$ . The acceleration is damped somewhat by the form drag, but still with constant acceleration. As the bubble elongates, viscous drag on the side walls imposes a terminal velocity scale  $Ag\lambda^2/C_2$ . The onset of the viscous regime scales with the Grashof number,  $h/\lambda \sim \text{Gr}$ . As the bubble slows, the inflow of pure fluid falls behind the mixing of fluid across the interface. The mixing reduces the effective Atwood number, slowing the bubble until it stops rising. This is the diffusive regime. The onset scales with the Rayleigh number,  $h/\lambda \sim \text{Ra}$ . Similarly, the maximum height of the bubble, i.e. the penetration depth, scales linearly with the Rayleigh number. The relationship is  $h(\infty)/\lambda = 2.41 \times 10^{-4}\text{Ra} - 0.36$ , defining a critical Rayleigh number,  $\text{Ra}_c = 1500$ , below which the bubble does not rise.

## 6.5 Model coefficients

The 5 unconstrained model parameters were fit to the numerical trajectories. First, the mixing coefficient  $C_5$  was fit directly to the mixed volume. Then, given  $C_5$ , the remaining four coefficients were fit with L2 regularization about the parameter estimates.

The  $C_3$  and  $C_7$  terms take a nominal value of 1, with outliers at the lowest Rayleigh numbers, in the case of  $C_3$ , and incomplete trajectories, in the case of  $C_7$ . The parameter dependence of  $C_3$  is postulated to be due to model breakdowns in highly dissipative cases. The parameter dependence of  $C_7$  is expected to disappear as the high Rayleigh trajectories

are completed.

The  $C_1$  drag-type coefficient takes nominal values between 0.3 and 0.6. There are outliers at zero and greater than 0.9, at low to moderate Rayleigh number and both low and high Schmidt number. The cause of these outliers is a mystery; the majority of trajectories are consistently within the nominal range.

The  $C_2$  friction factor-type coefficient clearly varies across the parameter space. It is a strongly increasing function of the Grashof number and a weakly increasing function of the Schmidt number. A possible explanation is the development of shear instabilities along the bubble sides that enhance the transport of momentum out of the bubble. These instabilities would be strongest at high Grashof number and could be stabilized somewhat by low Schmidt numbers.

Similarly, the  $C_5$  mixing coefficient varies with the diffusivity and Grashof number. It is a strongly decreasing function of diffusivity and is peaked with respect to the Grashof number. Similar to the development of momentum structures at high Grashof number, the low diffusivity cases could develop scalar structures on the interface that enhance mixing. The dependence on the Grashof number is mysterious; it is not clear what would cause peaked behavior. However, the mixing coefficient captures shape information that could depend on the viscosity and shear instabilities would transport the scalar just as they transport momentum.

## 6.6 Open problems

The stagnation and re-acceleration transient is not captured by the proposed buoyancy-drag model. Specifically, the model is unable to produce an inflection point at finite amplitude, while the stagnation and re-acceleration transient has two. This suggests the model is incomplete. Vorticity, specifically the development of a vortex ring at the bubble tip, isn't a part of the buoyancy-drag model and is a strong candidate for the missing term.

The majority of the simulations that have been performed are incomplete in that the bubble begins to interact with the top boundary before it stops rising. The truncation of the dynamics may affect the fit coefficients. In particular, the  $C_2$ ,  $C_5$ , and  $C_7$  terms are the most important to the dynamics at late times, so the fit values from clipped trajectories come with a grain of salt. These cases should be completed by simulation in domains with greater aspect ratios. The cost of completing each of the next set of simulations presented here is about 15 million core-hours on Shaheen XC40.

There are experiments and simulations that could be performed to isolate the effects of subsets of the model parameters. These experiments deviate from the standard low-amplitude thin interface single-model initial condition but should still be covered by the model. For example, the value of the  $C_3$  term is greatest at late times but the forces are nearly balanced by that point, weakening its influence. To isolate it, the pure Atwood number could be spontaneously doubled, throwing the numerator of the dynamics model out of balance and accelerating the bubble. The rate of that acceleration would depend strongly on  $C_3$ . The experimental equivalent of this procedure would be to rapidly increase the local acceleration.

Recent work on the single-mode instability has focused on cubic lattices of bubbles and spikes. Weakly non-linear theories have shown that a hexagonal lattice grows more rapidly, at least at early times. The lattice should affect the bubble geometry, so it would be interesting to compute the variation of the parameters across lattices and search for simple relations between them. It would also be interesting to see how the onset of the stagnation and reacceleration, around  $h/\lambda = 1/2$  and 1 for the cubic lattice, respectively, depends on the lattice type.

## REFERENCES

- [1] LIBXSMM v1.0.2, 2015.
- [2] S. Abarzhi, K. Nishihara, and J. Glimm. Rayleigh–Taylor and Richtmyer–Meshkov instabilities for fluids with a finite density ratio. *Physics Letters A*, 317(5-6), 2003.
- [3] R. Banerjee, L. Mandal, S. Roy, M. Khan, and M. R. Gupta. Combined effect of viscosity and vorticity on single mode Rayleigh-Taylor instability bubble growth. *Physics of Plasmas*, 18, 2011.
- [4] J. Bell, M. Day, C. Rendleman, S. Woosley, and M. Zingale. Direct numerical simulations of type ia supernovae flames. II. The Rayleigh-Taylor instability. *The Astrophysical Journal*, 608(2), 2004.
- [5] M. Berning and A. M. Rubenchik. A weakly nonlinear theory for the dynamical Rayleigh-Taylor instability. *Physics of Fluids*, 10(7), 1998.
- [6] A. Breuer, A. Heinecke, L. Rannabauer, and M. Bader. High-Order ADER-DG Minimizes Energy-and Time-to-Solution of SeisSol. In *International Conference on High Performance Computing*, 2015.
- [7] S. Chandrasekhar. The character of the equilibrium of an incompressible heavy viscous fluid of variable density. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51(1), 2008.
- [8] R. M. Davies and G. Taylor. The Mechanics of Large Bubbles Rising through Extended Liquids and through Liquids in Tubes. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 200(1062), 1950.
- [9] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1), 2008.
- [10] M. O. Deville, P. F. Fischer, and E. H. Mund. *High-order methods for incompressible fluid flow*, volume 9. Cambridge University Press, 2002.
- [11] G. Dimonte. Spanwise homogeneous buoyancy-drag model for Rayleigh-Taylor mixing and experimental evaluation. *Physics of Plasmas*, 7(2000), 2000.
- [12] G. Dimonte and M. Schneider. Turbulent Rayleigh-Taylor instability experiments with variable acceleration. *Physical Review E*, 54(4), 1996.
- [13] G. Dimonte and M. Schneider. Density ratio dependence of Rayleigh-Taylor mixing for sustained and impulsive acceleration histories. *Physics of Fluids*, 12(2), 2000.

- [14] G. Dimonte, D. L. Youngs, A. Dimits, S. Weber, M. Marinak, S. Wunsch, C. Garasi, A. Robinson, M. J. Andrews, P. Ramaprabhu, A. C. Calder, B. Fryxell, J. Biello, L. Dursi, P. MacNeice, K. Olson, P. Ricker, R. Rosner, F. Timmes, H. Tufo, Y.-N. Young, and M. Zingale. A comparative study of the turbulent Rayleigh–Taylor instability using high-resolution three-dimensional numerical simulations: The Alpha-Group collaboration. *Physics of Fluids*, 16(5), 2004.
- [15] R. E. Duff, F. H. Harlow, and C. W. Hirt. Effects of diffusion on interface instability between gases. *Physics of Fluids*, 5, 1962.
- [16] P. Fischer *et al.* *Nek5000 User Documentation*. Argonne National Laboratory, 2016.
- [17] S. M. Ghiaasiaan. *Convective heat and mass transfer*. Cambridge University Press, 2011.
- [18] V. Goncharov and O. A. Hurricane. Panel 3 Report : Implosion Hydrodynamics. *LLNL Report LLNLTR-562104*, 2012.
- [19] V. N. Goncharov. Analytical Model of Nonlinear, Single-Mode, Classical Rayleigh–Taylor Instability at Arbitrary Atwood Numbers. *Physical Review Letters*, 88(13), 2002.
- [20] K. Goto and R. A. Geijn. Anatomy of high-performance matrix multiplication. *ACM Transactions on Mathematical Software*, 34(3), 2008.
- [21] A. Heinecke, H. Pabst, and G. Henry. LIBXSMM: A high performance library for small matrix multiplications, 2015. Poster and extended abstract presented at SC’15.
- [22] E. P. Hicks. Rayleigh-Taylor Unstable Flames – Fast or Faster? *The Astrophysical Journal*, 803(2), 2015.
- [23] R. Hide. The character of the equilibrium of an incompressible heavy viscous fluid of variable density: an approximate theory. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51, 1955.
- [24] M. Hutchinson. Direct numerical simulation of single mode three-dimensional Rayleigh–Taylor experiments. *in preparation*.
- [25] M. Hutchinson. Direct numerical simulation of single mode three-dimensional Rayleigh–Taylor experiments. *arXiv preprint arXiv:1511.07254*, 2015.
- [26] M. Hutchinson. NekBox v2.0.0, 2015.
- [27] M. Hutchinson, A. Heinecke, H. Pabst, G. Henry, M. Parsani, and D. Keyes. Efficiency of high order spectral element methods on petascale architectures. In *International Conference on High Performance Computing*. Springer International Publishing, 2016.
- [28] Intel Corporation. Intel MKL 11.3 Release Notes, 2015. Introduced (S/D)GEMM\_BATCH and (C/Z)GEMM3M\_BATCH functions to perform multiple independent matrix-matrix multiply operations.

- [29] I. Ivanov, J. Gong, D. Akhmetova, I. B. Peng, S. Markidis, E. Laure, R. Machado, M. Rahn, V. Bartsch, A. Hart, and P. Fischer. Evaluation of parallel communication models in nekbone, a nek5000 mini-application. In *2015 IEEE International Conference on Cluster Computing*. IEEE, 2015.
- [30] J. W. Jacobs and I. Catton. Three-dimensional Rayleigh-Taylor instability Part 1. Weakly nonlinear theory. *J. Fluid Mech.*, 187, 1988.
- [31] J. Jacobs. personal communication.
- [32] D. Layzer. On the Instability of Superposed Fluids in a Gravitational Field. *The astrophysical journal*, 1955.
- [33] R. LeLevier, G. Lasher, and F. Bjorklund. Effect of a density gradient on Taylor Instability. Technical report, Lawrence Radiation Laboratory, 1955.
- [34] P. F. Linden. On the structure of salt fingers. *Deep Sea Research and Oceanographic Abstracts*, 20(November 1972), 1973.
- [35] Lord Rayleigh. Investigation of the character of the equilibrium of an incompressible heavy fluid of variable density, Scientific papers. *Proc. London Math. Soc.*, (39), 1883.
- [36] J. W. Lottes and P. F. Fischer. Hybrid multigrid/Schwarz algorithms for the spectral element method. *Journal of Scientific Computing*, 24(1), 2005.
- [37] S. Markidis, J. Gong, M. Schliephake, E. Laure, A. Hart, D. Henty, K. Heisey, and P. Fischer. Openacc acceleration of the nek5000 spectral element code. *International Journal of High Performance Computing Applications*, 29(3), 2015.
- [38] D. W. Marquardt and R. D. Snee. Ridge regression in practice. *The American Statistician*, 29(1), 1975.
- [39] J. D. McCalpin. STREAM: Sustainable memory bandwidth in high performance computers. Technical report, University of Virginia, Charlottesville, Virginia, 1991-2007. A continually updated technical report. <http://www.cs.virginia.edu/stream/>.
- [40] K. O. Mikaelian. Boussinesq approximation for rayleigh-taylor and richtmyer-meshkov instabilities. *Physics of Fluids*, 26(5), 2014.
- [41] National Research Council. *An Assessment of the Prospects for Inertial Fusion Energy*. The National Academies Press, Washington, DC, 2013.
- [42] N. Offermans, O. Marin, M. Schanen, J. Gong, P. Fischer, P. Schlatter, A. Obabko, A. Peplinski, M. Hutchinson, and E. Merzari. On the strong scaling of the spectral element solver nek5000 on petascale systems. In *Solving Software Challenges for Exascale*. Springer, 2016.

- [43] E. S. Oran and V. N. Gamezo. Origins of the deflagration-to-detonation transition in gas-phase combustion. *Combustion and Flame*, 148(1), 2007.
- [44] D. Oron, L. Arazi, D. Kartoon, A. Rikanati, U. Alon, and D. Shvarts. Dimensionality dependence of the RayleighTaylor and RichtmyerMeshkov instability late-time scaling laws. *Physics of Plasmas*, 8(6), 2001.
- [45] M. Otten, J. Gong, A. Mametjanov, A. Vose, J. Levesque, P. Fischer, and M. Min. An MPI/OpenACC implementation of a high order electromagnetics solver with GPUDirect communication. *International Journal of High Performance Computing*, 2016.
- [46] A. T. Patera. A spectral element method for fluid dynamics: Laminar flow in a channel expansion. *Journal of Computational Physics*, 54(3), 1984.
- [47] P. Ramaprabhu, G. Dimonte, P. Woodward, C. Fryer, G. Rockefeller, K. Muthuraman, P.-H. Lin, and J. Jayaraj. The late-time dynamics of the single-mode Rayleigh-Taylor instability. *Physics of Fluids*, 24(7), 2012.
- [48] P. Ramaprabhu, G. Dimonte, Y.-N. Young, A. C. Calder, and B. Fryxell. Limits of the potential flow approach to the single-mode Rayleigh-Taylor problem. *Physical Review E*, 74(6), 2006.
- [49] T. Roths, M. Marth, J. Weese, and J. Honerkamp. A generalized regularization method for nonlinear ill-posed problems enhanced for nonlinear regularization terms. *Computer physics communications*, 139(3), 2001.
- [50] D. Sharp. An overview of Rayleigh-Taylor instability. *Physica D: Nonlinear Phenomena*, 1984.
- [51] J. Shin, M. W. Hall, J. Chame, C. Chen, P. F. Fischer, and P. D. Hovland. Speeding up Nek5000 with autotuning and specialization. In *Proceedings of the 24th ACM International Conference on Supercomputing*. ACM, 2010.
- [52] S.-I. Sohn. Simple potential-flow model of Rayleigh-Taylor and Richtmyer-Meshkov instabilities for all density ratios. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 2003.
- [53] S.-I. Sohn. Late-Time Vortex Dynamics of Rayleigh-Taylor Instability. *Journal of the Physical Society of Japan*, 80, 2011.
- [54] M. E. Stern and J. S. Turner. Salt fingers and convecting layers. In *Deep Sea Research and Oceanographic Abstracts*, volume 16, 1969.
- [55] R. I. Tait and M. R. Howe. Thermohaline staircase. *Nature*, 231, 1971.
- [56] F. Takemura, S. Takagi, J. Magnaudet, and Y. Matsumoto. Drag and lift forces on a bubble rising near a vertical wall in a viscous liquid. *Journal of Fluid Mechanics*, 461(461), 2002.

- [57] H. M. Tufo et al. Terascale spectral element algorithms and implementations. In *Proceedings of the 1999 ACM/IEEE conference on Supercomputing*, 1999.
- [58] D. J. Wales and J. P. Doye. Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A*, 101(28), 1997.
- [59] L. F. Wang, W. H. Ye, Z. F. Fan, and Y. J. Li. Nonlinear saturation amplitude in the Rayleigh-Taylor instability at arbitrary Atwood numbers with continuous profiles. *EPL (Europhysics Letters)*, 90(1), 2010.
- [60] Z. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary, H. Deconinck, R. Hartmann, K. Hillewaert, H. Huynh, N. Kroll, G. May, P.-O. Persson, B. van Leer, and M. Visbal. High-order CFD methods: current status and perspective. *International Journal for Numerical Methods in Fluids*, 72(8), 2013.
- [61] T. Wei and D. Livescu. Late-time quadratic growth in single-mode Rayleigh-Taylor instability. *Physical Review E*, 86(4), 2012.
- [62] J. P. Wilkinson and J. W. Jacobs. Experimental study of the single-mode three-dimensional Rayleigh-Taylor instability. *Physics of Fluids*, 19(12), 2007.
- [63] M. Zingale, S. E. Woosley, C. A. Rendleman, M. S. Day, and J. B. Bell. Three-dimensional Numerical Simulations of Rayleigh-Taylor Unstable Flames in Type Ia Supernovae. *The Astrophysical Journal*, 632, 2005.