Discontinuous Galerkin Methods and Phase-Field Models for Numerical Simulations of Compressible Interfacial Flows with Shocks and Vorticity

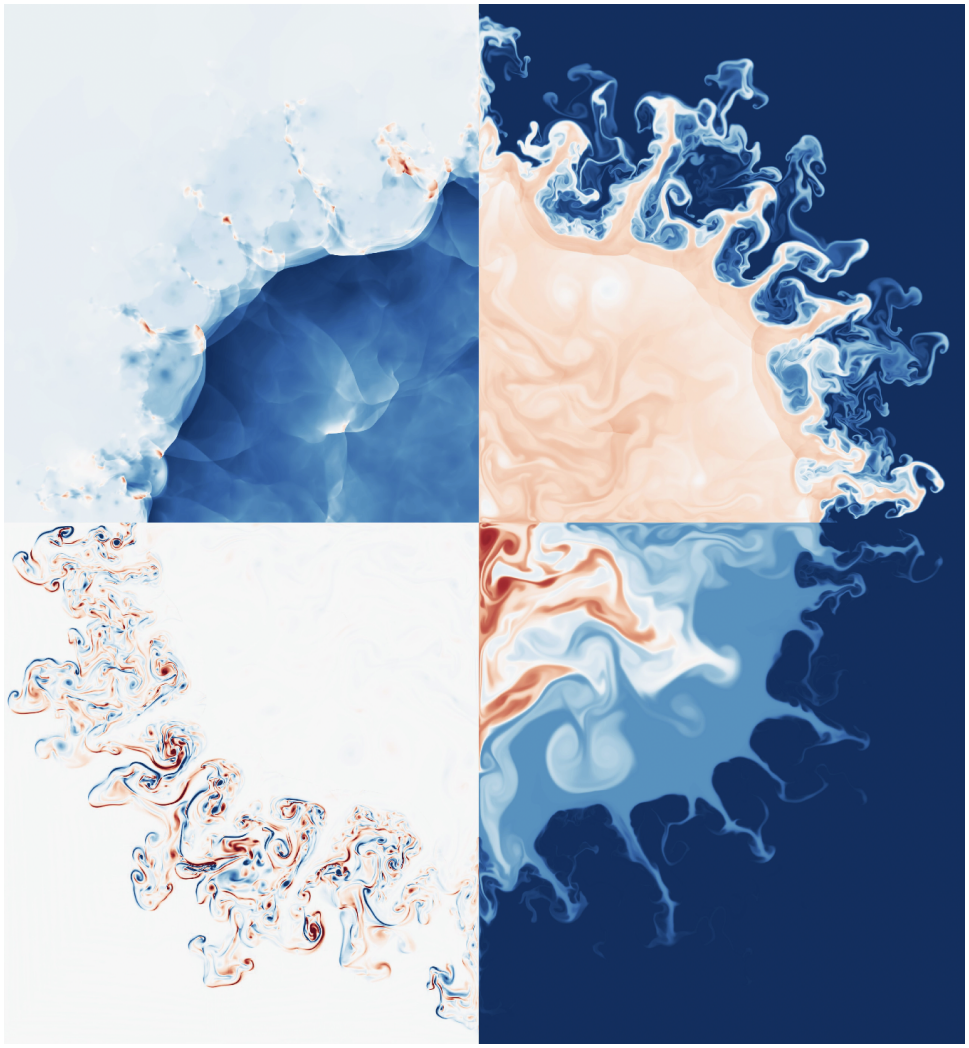
by

William J. White

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mechanical Engineering and Scientific Computing)
in the University of Michigan
2025

Doctoral Committee:

Professor Eric Johnsen, Chair
Doctor Joshua Dolence, Los Alamos National Laboratory
Professor Carolyn Kuranz
Professor Venkat Raman
Assistant Professor Aaron Towne



William J. White
wjwhit@umich.edu
ORCID iD: 0000-0003-2216-4500

© William J. White 2025

DEDICATION

To my parents, who taught me that life is what you make of it.

ACKNOWLEDGEMENTS

I would like to begin by thanking my advisor, Professor Eric Johnsen, for his many years of encouragement, patience, thoughtful advice, and professional mentorship, which have helped me develop as a scientist, an engineer, and an individual. Through sending me to conferences, rigorously examining my work, and discussing the nuances of academia, he has helped me develop from the undergraduate student I was when we met into a researcher; for this, I am truly thankful. I also extend my gratitude to my committee members, Dr. Joshua Dolence, Professor Carolyn Kuranz, Professor Venkat Raman, and Professor Aaron Towne, for their guidance and support both throughout my studies at Michigan and while serving on my committee.

My time in the Scientific Computing and Flow Physics Laboratory has been shaped by interactions with my wonderful colleagues. In particular, Shahab, thank you for introducing me to numerical methods, teaching me about bubble dynamics, and being patient as I learned the ropes. Michael, thank you for the many exciting conversations over simulation capabilities, instabilities, and pickleball; it has been a pleasure to study flow physics together. To Ziyang, a special thanks for the many hours of insightful conversations, debugging sessions, and whiteboard derivations. Our interactions greatly deepened my knowledge of numerical methods and interface models, and gave me a deeper appreciation for the beauty of an extensive mathematical proof.

To my friends at Michigan, I thank you for all of the great memories in Ann Arbor, especially Forrest, thank you for being an incredible friend, roommate, marathon running partner, and fellow football fanatic; and Joe, thank you for the many excellent games of pickleball, culminating in our shared intramural championship title.

Thank you to my parents, Joe and Sarah, for their unwavering love and support; none of what I have in life would be possible without you. To my sister Emily, thank you for being a kind, supportive, and understanding sibling, carving a path through many of life's nuances that I could then follow. I thank all of my grandparents for building the families within which I have been able to choose my own path. To my father-in-law, Chris, my mother-in-law, Lisa, and my brother-in-law, Jack, thank you for accepting me into your family; your support throughout the years, both personally and professionally, has been a gift.

Lastly, and most importantly, I would like to thank my wife, Kate, for her unwavering support throughout my time at Michigan. It is the greatest blessing to share my life with such a kind, patient, and compassionate woman, and I am eternally grateful for her.

This work is funded in part by the U.S. Department of Energy NNSA Center of Excellence under cooperative agreement number DE-NA0004146, Lawrence Livermore National Laboratory under subcontract B642551, Los Alamos National Laboratories under the project, "Algorithm/-Software/Hardware Co-design for High Energy Density applications," the Office of Naval Research through grant number N00014-22-1-2828 under Dr. Timothy Bentley, DOE grant DE-SC0023341 (under Dr. Kramer Akli), and the DOE Office of Nuclear Energy's Nuclear Energy University Programs (grant number DE-NE0008747).

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGEMENTS iii

LIST OF FIGURES viii

LIST OF TABLES xv

LIST OF APPENDICES xvi

ABSTRACT xvii

CHAPTER

1 Introduction 1

1.1 Representative Applications 2
1.2 Key Physical Phenomena 4
1.2.1 The Richtmyer–Meshkov Instability 5
1.2.2 Vortex-Core Instabilities 7
1.2.3 Cavitation Bubble Dynamics 7
1.3 Numerical Methods and Physical Models 8
1.3.1 Sharp, Diffuse, and Phase-Field Interface Models 10
1.3.2 High-Order Discretizations 11
1.3.3 Adaptive Mesh Refinement and High-Performance Computing 13
1.4 Dissertation Overview 14

I Methods and Models for Compressible Interfacial Flows 17

2 COMPAS: A GPU-Accelerated, Adaptive-Mesh Finite-Volume Framework for Consistent and Conservative Phase-Field Modeling of Compressible, N-Material, Interfacial Flows 17

2.1 Abstract 17
2.2 Introduction 18
2.3 Code Overview 21
2.3.1 Installation and Testing 23
2.3.2 Code Structure 23

2.3.3 Running Cases . . . 25
2.4 Consistent and Conservative Phase-Field Method . . . 26
2.4.1 Physical Model . . . 27
2.4.2 Numerical Approach . . . 29
2.5 Examples . . . 31
2.5.1 Isentropic Vortex . . . 31
2.5.2 Spherical Bubble Collapse . . . 33
2.5.3 Three Material Richtmyer-Meshkov Instability . . . 36
2.5.4 Triple-Point Problem . . . 40
2.5.5 Shock-Driven Multi-Material Implosion . . . 41
2.5.6 Taylor-Green Vortex Scaling . . . 42
2.6 Conclusions . . . 44
2.6.1 Choosing a Code . . . 45

3 A High-Order Discontinuous Galerkin Method for Compressible Interfacial Flows with Consistent and Conservative Phase Fields . . . 49

3.1 Abstract . . . 49
3.2 Introduction . . . 50
3.3 Equations of Motion . . . 53
3.3.1 Multiphase Model . . . 53
3.3.2 Consistent and Conservative Phase Field . . . 55
3.4 Non-Conformal Discontinuous Galerkin and Subcell Finite Volume Discretization 56
3.4.1 Runge-Kutta Discontinuous Galerkin Discretization of the Hyperbolic Step 57
3.4.2 Runge-Kutta Discontinuous Galerkin Scheme for the Phase-Field Step . . 67
3.4.3 Summary . . . 72
3.5 Verification and Validation . . . 73
3.5.1 Advection of an Isentropic Vortex . . . 73
3.5.2 Advection of an Air-Water Interface . . . 76
3.5.3 Multiphase Titarev-Toro Shock Tube . . . 78
3.5.4 Two-Dimensional Richtmyer-Meshkov Instability . . . 80
3.5.5 Underwater Explosion Problem . . . 83
3.6 Conclusions . . . 86

II Numerical Simulations of Compressible Interfacial Flow Physics 88

4 Emergence of Three-Dimensional Structures from Vortex-Pair Instabilities in Shocked Interfacial Flows . . . 88

4.1 Abstract . . . 88
4.2 Introduction . . . 89
4.3 Setup and Methods . . . 90
4.4 Results . . . 92
4.5 Conclusion . . . 98

5 Vortex-Core Instability in Shocked Converging Flows . . . 99

5.1 Abstract . . . 99

5.2 Introduction . . . 99
5.3 Setup and Methods . . . 101
5.4 Results . . . 103
5.5 Conclusions . . . 109

6 A Wave-Based Framework for Understanding Cavitation Damage Potential in Blast-Induced Traumatic Brain Injury . . . 111

6.1 Abstract . . . 111
6.2 Introduction . . . 112
6.3 Theoretical Framework . . . 114
6.4 Models and Methods . . . 118
6.4.1 Physical Models . . . 118
6.4.2 Numerical Methods . . . 120
6.5 Results . . . 122
6.5.1 One-Dimensional Models . . . 123
6.5.2 Two-Dimensional Model . . . 129
6.5.3 Discussion . . . 132
6.6 Conclusions . . . 134

7 Conclusion . . . 136

7.1 Summary of Key Findings . . . 136
7.2 Discussion and Limitations . . . 137
7.3 Future Work . . . 142
7.3.1 Accurate and Bound-Preserving Refinement and Coarsening Procedures for *h*-Adaptive Discontinuous Galerkin Schemes . . . 142
7.3.2 Linear Stability Theory of Converging Vortex Ring Pairs . . . 145
7.3.3 Competition Between Instabilities in Shocked Gaps . . . 146

APPENDICES . . . 147

BIBLIOGRAPHY . . . 179

LIST OF FIGURES

FIGURE

1.1 Supernova 1987A, as imaged by the James Webb Space Telescope, with its characteristic equatorial ring of bright spots. Credit: NASA, ESA, CSA, M. Matsuura (Cardiff University), R. Arendt (NASA's Goddard Spaceflight Center & University of Maryland, Baltimore County), C. Fransson (Stockholm University), J. Larsson (KTH Royal Institute of Technology), A. Pagan (STScI). 2

1.2 Prototypical design for a double-shell ICF capsule (left), and simulation results with different assembly join profiles (right), from, "Los Alamos National Laboratory Double Shell Program Target Development," Schmidt et al., *Fusion Science and Technology*, © 2023, reprinted by permission of Informa UK Limited, trading as Taylor & Francis Group, https://www.tandfonline.com [216]. 3

1.3 Schematic of the evolution of a direct-drive ICF experiment showing RTI growth due to asymmetric imprinting from the laser drive. Reproduced with permission from [56]. 5

1.4 Experiments of the single-mode, two-material Richtmyer-Meshkov instability at several different times. Reproduced with permission from [52]. 6

1.5 A pair of Crow-unstable vortex cores pinching and forming a vortex ring, as visualized by the Q-criterion at several times. The cores repeat periodically in space. 7

1.6 Acoustic wave solutions at three sequential times (red, orange, blue) from a fifth-order WENO finite volume scheme (top) and a fifth-order $p = 4$ DG scheme (bottom), with equal numbers of unknowns. 13

1.7 A schematic of patch-based adaptive mesh refinement (left) and a hydrodynamic triple-point simulation with AMR patches and density shown (right). 14

2.1 File tree for `COMPAS`, including an example case which has been compiled for two dimensions with MPI parallelization. 24

2.2 Density (left) and volume fraction (right) L_∞ errors for the isentropic vortex test case. Second-order accuracy is indicated by the dashed line in the density plot. 32

2.3 Initial grid for mesh levels two through the maximum level of six (zero-indexed) for the spherical collapse problem alongside the initial bubble radius R_0 and the pressure field. Reproduced with permission from [104]. 34

2.4 Time evolution of the radius for the inertial collapse of a spherical air bubble in water, as compared to the Keller-Miksis solution [122]. Driving pressure ratios of 10 (top left), 20 (top right), 50 (bottom left), and 50 with viscous terms (bottom right). Adapted with permission from [104]. 35

2.5 Time evolution of the radius for the inertial collapse of a spherical air bubble in water, as compared to the Keller Miksis solution [122] with and without the Phase Field active for the inviscid $P_R = 50$ case. 36
2.6 Time evolution of the interface positions (left) and perturbation amplitude (right) for the three-material Richtmyer-Meshkov instability with a flat (left) and perturbed interface (right). Experimental results are from [141]. 38
2.7 Interface thickness over time with and without Phase Fields for the three-material Richtmyer-Meshkov instability. 38
2.8 L_∞ errors in the volume fraction summation (left) and auxiliary material (right) over time for the three-material Richtmyer-Meshkov instability. 39
2.9 Density and phase indicator contours for the triple-point problem at $t = 5$ without interface regularization (left), with the conservative diffuse interface (CDI) method (middle), and with the conservative Allen-Cahn (CAC) Phase-Field mechanism (right). The AMR blocks are overlaid on the density for the result without Phase Fields (Euler). 40
2.10 Initial conditions for the shock-driven multi-material implosion problem. The density field, interfaces, and AMR grid are shown on the x-y plane with labels. Reproduced with permission from [104]. 42
2.11 Numerical Schlieren, interfaces, pressure, and density for the multi-material implosion at four times. Reproduced with permission from [104]. 43
2.12 Strong scaling results for the Taylor-Green vortex with 1024^3 cells on OLCF *Frontier*. 44

3.1 Example of square and triangular DG cells Ω_L with $p = 2$ solutions U_L bordering sets of subcells $\{\Omega_R^s\}_{s=0}^{K-1}$ with cell-average solutions $\{\bar{U}_R^s\}_{s=0}^{K-1}$. Quadrature points are denoted by • for points that are duplicated by two adjacent subcells and × for those that are not. Integration along the segment $\bar{\mathcal{I}}_0$ is evaluated for Ω_L and Ω_R^0 according to a unique quadrature rule such that the scheme is conservative. 65
3.2 Reconstructions of $\partial_x Q$ in two $p = 1$ DG cells, L and R. The unique interfacial reconstruction $\partial_x Q_{\mathcal{I}}$ is shown in orange between the two elements. Subcell face values of $\partial_x Q$ not used for the interface reconstruction are shown with • while the face values used for the interface reconstruction are shown with ▪. The interface reconstruction matches nodally with both the left and right cells, but not necessarily at quadrature points. Face gradients are defined on a staggered grid. 71
3.3 L_2 (left column) and L_∞ (right column) error norms of the mixture density ρ (top row) and the volume fraction α_1 (bottom row) for the isentropic vortex advection problem. Grey guidelines correspond to the expected order of accuracy for each p and are labeled as such. 75
3.4 Top left: time evolution of the maximum error in velocity (solid), pressure (dashed), and temperature (dash-dotted); top right: volume fraction conservation error; bottom left: interface thickness; and bottom right: number of flagged cells for limiting or subcells for the interface advection problem. Colors represent each scheme: DG+HRL (red), DG+SC (orange), DG+SC+PF (blue), and DG+SC+PF (dark blue, 256×256), bottom right only. 77
3.5 Density and α_1 fields at $t = 5$ for the multiphase-adapted Titarev and Toro shock-tube problem for the DG+SC, DG+SC+PF, and fifth-order finite volume schemes with 400 DOF. The reference simulation (2560 DOFs) is shown in gray in all cases. 79

3.6 Two-dimensional Richtmyer-Meshkov instability for the DG+HRL, DG+SC, and DG+SC+PF schemes for $\lambda_0/h \in [8, 128]$. Left: Time evolution of the perturbation amplitude for different resolutions. Right: Contours of α_1 for each resolution and scheme at $t = 5.2$ ms. Odd-numbered contours are flipped horizontally to show the difference in vertical positions of the bubble or spike between resolutions. 82
3.7 Time evolution of the interface thickness for the two-dimensional Richtmyer-Meshkov instability problem for $\lambda_0/h = 128$. Left: comparison between the DG+HRL, DG+SC, and DG+SC+PF schemes. Right: dependence on resolution. 83
3.8 Contours of Mach number and α_1 for the underwater explosion problem for two mesh resolutions with the DG+SC and DG+SC+PF schemes at $t = 0.019$. 85

4.1 Two-dimensional schematics of the x-y (top) and y-z (bottom) planes. The $y - z$ plane is shown at the initial time, just after the shock impinges upon the interface, and after the cores are developed (late-time RMI). 91
4.2 Contours of the material indicator (left and far left) and Q-criterion (right) at times $t \in [0, 5, 30, 40, 50, 60]$ for a case with At = 0.61 and $\beta = 0.914$. The heavy fluid is below the interface, and the shock comes from below. A slice is taken from the x-z view of the material indicator to reveal the interior structure of the mixed region (blue is the light fluid, yellow is the heavy fluid). 93
4.3 Contours of the material indicator (left and far left) and Q-criterion (right) at times $t \in [0, 5, 30, 50, 70, 90]$ for a case with At = 0.61 and $\beta = 3.2$. The heavy fluid is below the interface, and the shock comes from below. A slice is taken from the x-z view of the material indicator to reveal the interior structure of the mixed region (blue is the light fluid, yellow is the heavy fluid). 94
4.4 Non-dimensional growth rates calculated from simulations vs. non-dimensional wavenumber. Green squares, blue crosses, and yellow diamonds are for Atwood numbers 0.34, 0.61, and 0.78, respectively, and the black line is Eq. (4.4). The inset shows the normalized core perturbation amplitude over time for At = 0.78 and $\beta = 0.914$. 97

5.1 Schematic of the initial condition showing the initial interface, initial equatorial perturbation, and the shock. 102
5.2 Contours of the material indicator (left) and Q-criterion (right) at times $t \in [0, 0.2, 0.7, 1.0]$ for a case with $k_\theta = 8$. Slices are taken out of the view of the material indicator to see the interior structure of the mixed region (blue is the heavy fluid, yellow is the light fluid). 104
5.3 Contours of the material indicator (left) and Q-criterion (right) at times $t \in [0, 0.2, 0.7, 1.0]$ for a case with $k_\theta = 12$. Slices are taken out of the view of the material indicator to see the interior structure of the mixed region (blue is the heavy fluid, yellow is the light fluid). 106
5.4 Perturbation growth rates a_f versus angular wavenumber k_θ. The inset plot shows the normalized perturbation amplitude over time overlaid with the corresponding exponential for the $k_\theta = 4$ case. 107
5.5 Amplitude of the RMI perturbation vs. time at the crest (top) and trough (bottom) of the $k_\theta = 4$ case. 108

6.1 Schematic representation of a tension wave inside brain tissue interacting with a microbubble, leading to growth that may directly damage neurons. 113
6.2 Simple (left) and compound (right) Friedlander blast profiles, with properties similar to those in [269] and [268]. 117
6.3 One-, two- and three-dimensional models. The three-layer model demonstrates the effects of reflection and transmission, while the five-layer model shows the changes induced by the skull, as well as the resultant superposition. The two-dimensional model demonstrates geometric effects including focusing, while the three-dimensional model accounts for more realistic geometries. 119
6.4 Experiments from [218] for the interaction of a Mach 2.4 shock with a column of water. Reprinted from S. Sembian, M. Liverts, N. Tillmark, and N. Apazidis. Plane shock wave interaction with a cylindrical water column. *Physics of Fluids*, 28(5), 2016, with the permission of AIP Publishing. 122
6.5 Numerical Schlieren contours for the interaction of a Mach 2.4 shock with a column of water at times 12.5 μs, 16 μs, 22.5 μs, and 29.5 μs, corresponding to (b)-(e) in Figure 6.4. 122
6.6 Minimum pressure vs. number of AMR levels. 123
6.7 Three representative times for the three-dimensional model described in §6.4.1. In the first frame, the blast has just impinged upon the skull, and the transmitted wave is propagating through towards the brain. The leading transmitted wave is propagating through the brain in the middle frame, and in the final frame, reflection, superposition, and focusing lead to a region of tension in the posterior of the brain tissue. 124
6.8 Demonstrative *x-t* wave diagrams for the three-layer one-dimensional model with constructive (left) and destructive (right) superposition. The two-peak blast propagates in low-impedance air until it reaches the high-impedance brain surrogate at point (a), which transmits a wave that travels to point (b) where it is reflected. The second blast peak intersects with the reflected blast at point (c) and impinges upon the solid at point (d), where it transmits a wave that superposes with the wave in the brain surrogate and either increases (left) or decreases (right) the intracranial wave's magnitude. 124
6.9 Pressure contours in *x-t* space for a two-peak blast interacting with a brain surrogate. Left: Contour for blast peak lag time $\Delta t = 2$ ms. Right: Contour for blast peak lag time $\Delta t = 1.88$ ms. The dashed lines represent the boundaries of the brain surrogate. . 126
6.10 Left: A demonstrative $x - t$ wave diagram for the five-layer one-dimensional model. The blast propagates in low-impedance air until it reaches the high-impedance skull surrogate at point (a), which transmits a wave that travels to point (b), where it is partially reflected and partially transmitted into the brain surrogate. The intra-cranial wave propagates to point (c) where it is partially transmitted and reflected. At point (d), energy exits the head at the anterior skull-air interface and the wave inverts. At point (e), the wave front in the brain superposes with waves inside the skull and reflects, and at point (f), a new wave is transmitted into the brain surrogate. Right: Pressure contours in *x-t* space for a simple blast interacting with a brain surrogate. The dashed lines represent the boundaries of the brain and skull surrogates. 128
6.11 Left column: Pressure (top) and numerical Schlieren (bottom) contours for the two-dimensional blast-cylinder model at different times. Right: Pressure contours in *x-t* space along the cylinder's centerline. 131

6.12 Pressure over time at the position of peak tension in each model (top) and center-line pressure at the time of peak tension for each model (bottom), where the three-layer model with constructive superposition of the second blast peak is denoted "3LC", the three-layer model with destructive superposition is "3LD", the five-layer model is "5L", the two-dimensional model is "2D", and the three-dimensional model is "3D". . 133

7.1 Coarse solutions scaled around the coarse cell-average value (left) and around the refined cell averages (right). 144
7.2 Contours of the material tracer (top) and Q-criterion (bottom) for varying magnetic field strengths μ_z, increasing from left (no field) to right. 145

A.1 Primitive variables $\mathbf{W}$ for the double rarefaction problem at $t = 0.3$. Left: Solution from the proposed well-balanced scheme. Right: Solution from the described path-conservative scheme [154, 200]. Results are generated with 64 $p = 3$ elements. 148

B.1 Schematic of the Rayleigh collapse of a single bubble near two perpendicular rigid walls. 153
B.2 Contours of (numerical) Schlieren, pressure normalized by $\rho_l a_l \sqrt{\Delta p/\rho_l}$, and velocity magnitude normalized by $\sqrt{\Delta p/\rho_l}$ for $(\delta_x, \delta_z) = (1.5, 3.0)$. Contours are of times $t = 0.163$, $t = 1.155$, $t = 1.199$, and $t = 1.294$, normalized by the Rayleigh collapse time $0.915 R_0 \sqrt{\rho_l/\Delta p}$. 156
B.3 Contours of (numerical) Schlieren, pressure normalized by $\rho_l a_l \sqrt{\Delta p/\rho_l}$, and velocity magnitude normalized by $\sqrt{\Delta p/\rho_l}$ for $(\delta_x, \delta_z) = (2.5, 3.0)$. Contours are of times $t = 0.163$, $t = 1.102$, $t = 1.198$, and $t = 1.261$, normalized by the Rayleigh collapse time $0.915 R_0 \sqrt{\rho_l/\Delta p}$. 157
B.4 Outlines for (a) $(\delta_x, \delta_z) = (1.5, 3.0)$ and (b) $(\delta_x, \delta_z) = (2.5, 3.0)$, and the measurement of λ. 160
B.5 Non-sphericity ξ vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty. 160
B.6 Minimum volumes vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty. 161
B.7 Collapse times scaled by the Rayleigh collapse time vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$; gray circles: other single-wall cases. Corner cases are filled and single-wall cases are empty. 162

B.8 Jet angle θ_j as a function of (a) the ratio of the initial stand-off distances, δ_x/δ_z; red diamonds: direct simulations; green circles: experimental data from [34]; blue triangles: experimental data from [233]; black line ($\arctan(\delta_x/\delta_z)$): Eq. (B.6); gray dashed line ($\arccos(\mathrm{Im}(\vec{u}_j)/|\vec{u}_j|)$): Eq. (B.8) and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$; gray circles: other single-wall cases. Corner cases are filled and single-wall cases are empty. 163
B.9 Bubble migration angle θ_j as a function of (a) the ratio of the initial stand-off distances, δ_x/δ_z; red diamonds: direct simulations; green circles: experimental data from [34]; black line ($\arctan(\delta_x/\delta_z)$): Eq. (B.6); gray dashed line ($\arccos(\mathrm{Im}(\vec{u}_j)/|\vec{u}_j|)$): Eq. (B.8) and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$; gray circles: other single-wall cases. Corner cases are filled and single-wall cases are empty. 165
B.10 (a) $\min(\delta_{xc}, \delta_{zc})$ vs. $\min(\delta_x, \delta_z)$; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty. (b) Migration from initial stand-off distances to collapse location vs. initial stand-off distance; red circles: Δ_{xm} vs. δ_x; blue squares: Δ_{zm} vs. δ_z. The black dashed curve is a direct fit of Δ_{xm} and Δ_{zm} to $a\zeta^{3/5} + b$ with $a \approx 1.9$ and $b \approx -0.06$. 165
B.11 Maximum jet speed vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty. 167
B.12 Time history of wall pressure at the location of the maximum wall pressure (red solid line), at the corner (green dashed-dotted line), $(x, y, z) = (\delta_x, 0, 0)$ (gray dash-dot-dot), and $(x, y, z) = (0, 0, \delta_z)$ (blue dashed line) for (a) $(\delta_x, \delta_z) = (1.5, 3.0)$ and (b) $(\delta_x, \delta_z) = (2.5, 3.0)$. 168
B.13 Location of maximum wall pressure δ_{zp}/δ_z; gray region: $p_r \geq 1$. 170
B.14 Maximum wall pressure vs. (a) $\min(\delta_x, \delta_z)$ and ξ (inset), and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty. 171
B.15 Maximum wall temperature (UHMWPE) vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty. 172
B.16 Contours of (numerical) Schlieren, pressure normalized by $\rho_l a_l \sqrt{\Delta p/\rho_l}$, and velocity magnitude normalized by the characteristic interfacial speed $\sqrt{\Delta p/\rho_l}$ for $(\delta_x, \delta_z) = (0.75, 1.25)$ at $t = 1.387$ and $t = 1.607$. 173
B.17 Contours of (numerical) Schlieren, pressure normalized by $\rho_l a_l \sqrt{\Delta p/\rho_l}$, and velocity magnitude normalized by the characteristic interfacial speed $\sqrt{\Delta p/\rho_l}$ for $(\delta_x, \delta_z) = (0.25, 0.75)$ at $t = 1.319$ and $t = 1.493$. 174

B.18 (a) Collapse time, (b) jet speed, (c) jet angle, and (d) maximum wall pressure vs. δ_z; blue diamonds: $\delta_x = 0.25$; green squares: $\delta_x = 0.5$; red deltas: $\delta_x = 0.75$. Corner cases are filled and single-wall cases are empty. 176

LIST OF TABLES

TABLE

2.1 Comparison of several open-source and laboratory multiphase simulation codes. . . . 46

6.1 Typical densities, sound speeds, and acoustic impedances of head layers [106]. 116
6.2 Pressure-amplitude reflection (left) and transmission (right) coefficients, R_{ij} and T_{ij}, respectively, or a wave incident from medium i (row) into medium j (column). 116

B.1 Relevant gas and liquid properties. 154

LIST OF APPENDICES

A Path-Conservative Scheme for the Divergence-Form Volume Fraction Equation . . . 147

B Pressure Fields Produced by Single-Bubble Collapse Near a Corner 150

Abstract 150
Introduction 151
Problem Description and Physical Model 152
Detached Bubbles 155
Attached Bubbles 172
Conclusions 177

ABSTRACT

Compressible interfacial flows involve interactions between fluids of different phases or materials. In such flows, variations in density and pressure significantly influence the dynamics and evolution of the interface. These flows are central to applications including astrophysics, high-speed droplet dynamics, inertial confinement fusion, and cavitation bubble dynamics. Steep gradients, smooth features, varying time and length scales, and dramatic topology changes make studying these flows challenging. Thus, numerical solutions are crucial for understanding the underlying flow physics. However, predictive simulations are often constrained by numerical schemes that either smear interfaces, violate conservation laws, or require complex mesh management. Few software libraries address these concerns within a GPU-accelerated, adaptive mesh framework. This dissertation tackles these deficiencies by advancing numerical methodology, delivering a scalable open-source code, and applying that capability to interrogate shock-induced instabilities and cavitation phenomena.

An open-source simulation framework has been developed to address the need for a GPU-portable, adaptive-mesh framework that implements consistent, conservative interface regularization techniques. This package implements the general, consistent, and conservative Phase-Field method for N-material flows alongside the massively parallel, GPU-accelerated AMReX package for patch-based adaptive mesh refinement. Through several points of verification, validation, and performance testing, the software's accuracy and capability to resolve flows of interest are demonstrated.

A fully conservative, bound-preserving Phase-Field formulation is derived and embedded in a high-order discontinuous Galerkin method featuring non-conforming sub-cell finite-volume shock-capturing, building on this platform. The scheme achieves high-order accuracy for smooth problems. It also maintains consistently thin interfaces throughout long-time simulations and demonstrates robustness across density and pressure jumps of many orders of magnitude. Together, these features provide a high-order alternative to traditional second-order-accurate implementations of diffuse-interface methods.

The developed numerical framework is subsequently applied to two classes of physical problems. Three-dimensional adaptive-mesh simulations demonstrate that the planar Richtmyer-Meshkov instability can seed Crow-type vortex-line instabilities, yielding three-dimensional structure and

potentially altered advective mixing rates. Measured perturbation growth rates match theoretical trends, supporting this hypothesized pathway to structure formation. Extensions to cylindrical geometries provide a quantitative demonstration of exponential perturbation growth and its deleterious impact on implosion symmetry, which is key to double-shell inertial confinement fusion experiments. A wave-based model hierarchy, complemented by numerical simulations completed with the developed framework, shows the pathway from blast exposure to tensions that may cause damaging cavitation in blast-induced traumatic brain injury, emphasizing the need for refined injury-risk metrics. An appendix quantifies the pressures produced during single-bubble inertial collapse near rigid corners, yielding empirical relationships for predicting damage.

CHAPTER 1

Introduction

Compressible interfacial flows arise in many applications, including astrophysics [129, 224, 259], high-speed droplet dynamics [37, 191], inertial confinement fusion (ICF) experiments [56, 144, 216], cavitation bubble dynamics [25, 28, 30, 189, 233, 260, 264], and biomedicine [7, 23, 55, 114, 192, 273]. The study of each of these areas is challenging due to the presence of material interfaces, shock waves, and high-speed dynamics, making it difficult for laboratory diagnostics to capture sufficient spatio-temporal resolution, while analytical theory is limited to simplified systems or geometries. As a consequence, numerical simulations are often useful for understanding the underlying multiscale physics of these phenomena.

To address these complex flow phenomena, high-resolution numerical simulations must rely on tools capable of accurately representing interfacial physics, shock waves, and vorticity dynamics. Achieving this while efficiently utilizing heterogeneous high-performance computing (HPC) systems often requires techniques like adaptive mesh refinement (AMR) [26, 27, 79, 190, 203, 282]. High-order accurate methods can take even greater advantage of HPC systems due to their computational intensity, whereas low-order methods are often bandwidth-limited. Hybrid techniques can be designed to efficiently resolve both sharp and smooth features, such as interfaces and vortices [9, 49, 66, 205, 266, 280]. While these approaches offer enhanced resolution, special consideration is still required for handling material interfaces. Approaches include either complex geometric reconstructions or interface regularization techniques, the latter of which can sacrifice sharpness [86, 87, 93, 107, 152, 176, 223]. Numerical simulations are commonly used to study phenomena such as the Rayleigh-Taylor [235] and Richtmyer-Meshkov [162, 201] instabilities, which are characterized by the growth of interfacial perturbations driven by buoyancy or acceleration following the passage of a shock. These and other instabilities are commonly found in astrophysical systems [119], laboratory fluids [52], atmospheric flows [247], and high-energy-density physics (HEDP) applications like inertial confinement fusion (ICF) [56, 216]. Although the Richtmyer-Meshkov instability (RMI) has been documented in many contexts, questions remain about how secondary mechanisms such as vortex instabilities affect structure formation and material mixing in interfacial flows [57, 256, 259]. Cavitation bubbles, which can produce shock waves [28, 115, 228],



Figure 1.1: Supernova 1987A, as imaged by the James Webb Space Telescope, with its characteristic equatorial ring of bright spots. Credit: NASA, ESA, CSA, M. Matsuura (Cardiff University), R. Arendt (NASA's Goddard Spaceflight Center & University of Maryland, Baltimore County), C. Fransson (Stockholm University), J. Larsson (KTH Royal Institute of Technology), A. Pagan (STScI).

destroy tissue [51, 114], and potentially cause brain injury [72, 82, 159, 181], are often characterized by numerical simulations, but the conditions that lead to their inception are not always clear in some systems.

The objectives of this dissertation are, thus, to advance the methods, models, and software capabilities used to represent compressible interfacial flows and to understand their physics, with applications to interfacial instabilities, vortex-core instabilities, and cavitation bubble dynamics. This introduction provides background on the physics, numerical methods, models, and applications that underpin the remaining chapters of this dissertation.

1.1 Representative Applications

In young supernova remnants, dense ejecta accelerate through lighter layers, triggering RMI and Rayleigh-Taylor instability (RTI) growth that develops to filamentary structures [129]. Additionally, it has been suggested that the hot spots observed around supernova 1987A are due to cylindrical vortex-ring instabilities, including the Crow instability (CI) [259]. Figure 1.1 shows an image from the James Webb Space Telescope of SN1987a with the hotspots clearly illuminated. The hypothesized vortex rings originate from mass stripping in a former binary star merger that produced

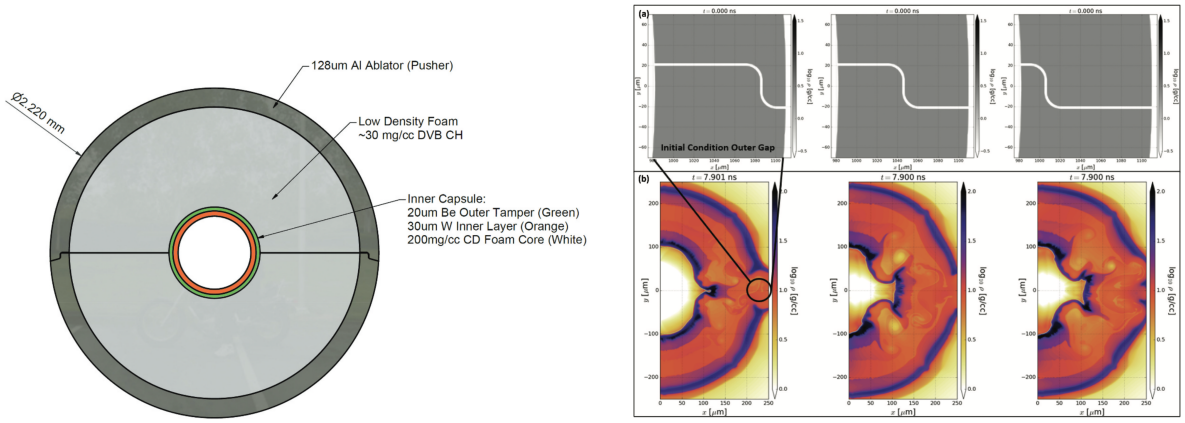


Figure 1.2: Prototypical design for a double-shell ICF capsule (left), and simulation results with different assembly join profiles (right), from, "Los Alamos National Laboratory Double Shell Program Target Development," Schmidt et al., *Fusion Science and Technology*, © 2023, reprinted by permission of Informa UK Limited, trading as Taylor & Francis Group, https://www.tandfonline.com [216].

a circumstellar ring; following the merger, the solar wind deposited vorticity on the ring, leading to a pair of unstable vortex rings [259]. Recent studies have investigated the excitation of the CI from a shocked cylinder, suggesting that the CI may be present in compressible interfacial flows, including those in the HEDP regime [256], and supporting the hypothesis that an accelerated cylinder can develop into Crow-unstable cores.

In double-shell ICF experiments [168, 210], multiple shocks reflect between the successive layers of high- and low-density materials that compose the shells, generating significant RMI and RTI structure [84]. Of key importance in double-shell experiments is the assembly joint, which is known to be a significant source of asymmetry [84, 210, 216]. Various joint profiles have been explored to reduce the deleterious effects, but an optimal configuration remains elusive. Figure 1.2 shows a sample design for a double-shell capsule alongside simulation results from [216] demonstrating the impact of different joint profiles. The instabilities due to the joint reduce the hot-spot symmetry and therefore suggest stringent guidelines on manufacturing techniques, joint profiles, and pulse shapes. Compared with single-shell ICF experiments, which lack an assembly joint, double-shell capsules can develop equatorial features that may be subject to instability mechanisms beyond the classically investigated RMI and RTI. Assessing the potential for such instabilities is thus of interest in characterizing the performance of experimental platforms.

Cavitation, the nucleation, growth, and collapse of gas bubbles, occurs in a variety of applications to include fluid machinery [18, 69], fuel systems [45], microfluidics [196, 284], the Spallation Neutron Source [160], blast-induced traumatic brain injury (bTBI) [181], and shockwave lithotripsy [51, 55, 114]. In soft biological tissues such as the brain, the growth of bubbles can

lead directly to cell death [70]. Because bTBI is currently the most common source of brain injury among service members globally [72], and mounting evidence implicates intracranial cavitation as a primary damage pathway [70, 82, 130, 159, 181], a detailed understanding of bubble dynamics and onset criteria is essential.

These applications span a wide range of regimes, as characterized by the Reynolds, Mach, and Atwood numbers. The Reynolds number $\mathrm{Re} = \rho|\boldsymbol{u}|L/\mu$ identifies the importance of inertial versus viscous effects, where L is the characteristic length scale and μ is the fluid viscosity, and typically varies from $O(10^4)$ [30] in bubble dynamics to $O(10^{14})$ or greater in supernovas [271], indicating that viscous forces are negligible. As such, the physics of interest rely primarily on the Mach and Atwood numbers. The Mach number $\mathrm{Ma} = |\boldsymbol{u}|/c$ describes compressibility, and spans from $O(10^{-2})$ in biomedical applications to $O(10^2)$ in astrophysical and HEDP systems [148]. The Atwood number at an interface between fluids of densities ρ_1 and ρ_2, $\mathrm{At} = (\rho_2 - \rho_1)/(\rho_2 + \rho_1)$ measures density gradient effects, and is typically lower for gas-gas problems like many RMI experiments [52, 141] ($\mathrm{At} \sim 0.5$) and higher ($\mathrm{At} \sim 1$) for metal-gas interfaces like those in double-shell ICF experiments [216]. The wide range of scales across application areas necessitates methods that capture the physics in each regime.

1.2 Key Physical Phenomena

Discontinuities in fluid density or material subjected to acceleration or shear experience a suite of classical instabilities driven by vorticity, as described by the vorticity-transport equation,

$$\frac{\partial \boldsymbol{\omega}}{\partial t} + (\boldsymbol{u} \cdot \nabla)\boldsymbol{\omega} = (\boldsymbol{\omega} \cdot \nabla)\boldsymbol{u} - \boldsymbol{\omega}(\nabla \cdot \boldsymbol{u}) + \frac{\nabla\rho \times \nabla P}{\rho^2}, \tag{1.1}$$

where $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$ is the vorticity. Rayleigh-Taylor (RT) growth is driven by the buoyancy of a light fluid under a heavy one [197, 235]; perturbations along the interface between the two fluids lead to baroclinic torque on the interface that causes the perturbation to grow, as seen in the last term of Eq. 1.1. Classical linear analysis predicts that the RTI grows exponentially in time, potentially leading to significant advective mixing. The Kelvin–Helmholtz instability (KHI) develops at shear layers, where perturbations to the layer – which may coincide with an interface – lead to the generation of vorticity and subsequent exponential growth in the linear regime [89]. These two interfacial instabilities are common in HEDP and astrophysical settings [56, 129, 144, 224], and considerable effort has been dedicated to preventing their occurrence in ICF experiments to maximize implosion symmetry. Figure 1.3 shows a schematic representation of a direct-drive ICF experiment with the RTI. The RTI adds significant asymmetries to converging ICF capsules and gives rise to substantial mixing such that heavy material enters the hot spot and cools it, or fuel

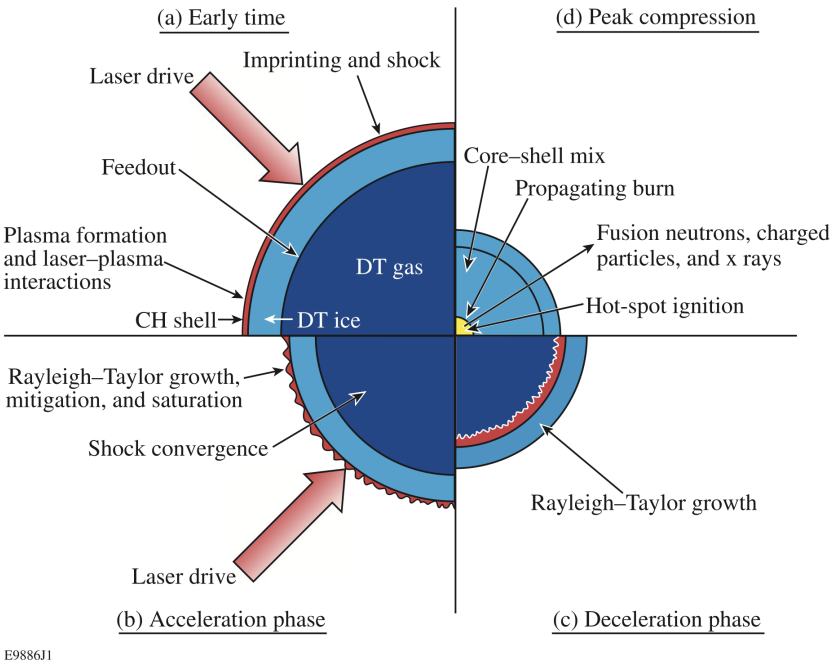


Figure 1.3: Schematic of the evolution of a direct-drive ICF experiment showing RTI growth due to asymmetric imprinting from the laser drive. Reproduced with permission from [56].

escapes the hot spot, both of which reduce the potential energy gain [56]. Interfacial instabilities can also lead to the ejection of vortical structures such as vortex rings, which transport mass away from the interface, including in supernovae [209, 257]

1.2.1 The Richtmyer–Meshkov Instability

The RMI is similar to the RTI in that it is driven by a misalignment between the pressure and density gradients, thus leading to the deposition of vorticity via baroclinic torque [197, 235]. In the RMI, however, the pressure gradient is induced by an impulsive acceleration, typically in the form of a shock. This impulsive source leads to linear perturbation growth instead of the exponential growth observed in the RTI. Figure 1.4 shows canonical experimental results demonstrating the RMI between air and SF_6 [52]. As the shock impinges on the interface, the sinusoidal perturbation compresses, but then grows due to the deposited vorticity. Early in time, the amplitude can be predicted from linear theory as,

$$\frac{da}{dt} = kA\Delta V a_0, \tag{1.2}$$

where a is the perturbation amplitude, k is the perturbation wavenumber, ΔV is the velocity imparted on the interface by the shock, and a_0 is the initial perturbation amplitude. Later, the growth saturates, and the vorticity sheet at the interface rolls into counter-rotating vortices. The

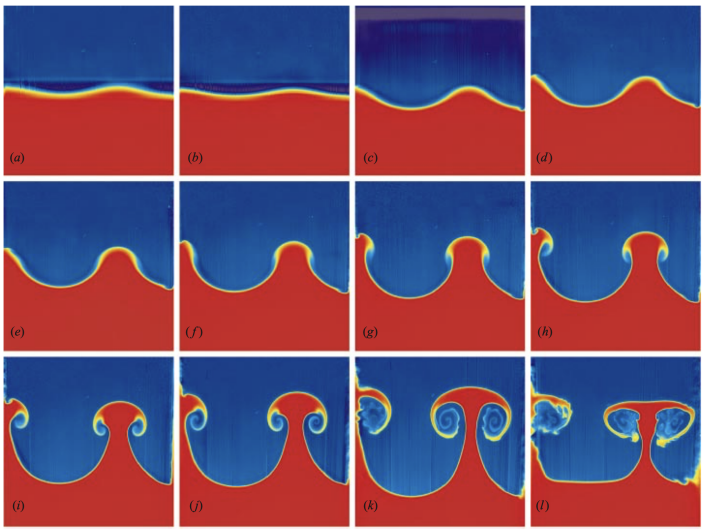


Figure 1.4: Experiments of the single-mode, two-material Richtmyer-Meshkov instability at several different times. Reproduced with permission from [52].

interpenetrating fluids form bubbles and spikes whose shapes depend on the shock strength, the perturbation profile, and the material properties. Figure 1.4 shows an example of this behavior. Multi-modal interfaces can exhibit additional behaviors, such as the ejection of vortex pairs [257]. Studies of more complex interfaces, such as those likely to be present in experiments, demonstrate that RMI can significantly alter mixing rates depending on the amplitude and distribution of perturbations [236]. Curved interfaces also experience RMI, with the local curvature determining the degree of departure from planar theory. Substantial effort has been devoted to characterizing this modification due to its relevance across a variety of systems [21, 146, 149, 272, 276].

Significant effort has also gone into characterizing the behavior of re-shocked RMI, where a secondary shock is introduced to a, usually late-time, RMI structure [131, 155, 166]. As the secondary shock interacts with the bubble-spike structures, the vorticity is significantly enhanced, accelerating the transition to small-scale structures and thus furthering advective mixing. In converging systems, the rebound of the incident shock intrinsically leads to re-shock and thus significantly altered interface dynamics [140, 252]. These effects are especially pertinent to single- and double-shell ICF targets, where multiple driven and reflected shocks interact with various types of interfaces [56, 144]. Accurate prediction of the resulting effect on mixing is thus needed to predict experimental performance, so simulation tools that efficiently resolve shocks, interfaces, and vorticity dynamics are key.

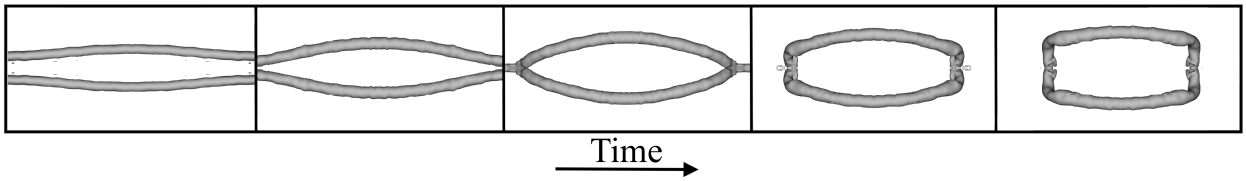


Figure 1.5: A pair of Crow-unstable vortex cores pinching and forming a vortex ring, as visualized by the Q-criterion at several times. The cores repeat periodically in space.

1.2.2 Vortex-Core Instabilities

Vortex cores and rings emerge from a variety of sources, including in the trailing vortices behind airplanes [57], in supernova remnants [259], and in the late-time RMI [52, 236]. Adjacent and isolated vortex cores and rings interact and evolve under the influence of a variety of instabilities, including the elliptical [123] and Widnall [267] instabilities. The elliptical instability is a mechanism by which planar flows with elliptical streamlines generate three-dimensional structure; resonance between the strain field and inertial waves amplifies small perturbations, leading to instability. This process is present in a variety of systems, including planetary interiors [135]. The Widnall instability is a short-wavelength mechanism in vortex rings where helical disturbances grow and trigger vortex breakdown. Together, these instabilities have the potential to break down coherent vortex structures into smaller-scale features with implications on mixing and turbulent transition [123].

The Crow instability (CI) is a mutual-induction instability of counter-rotating vortex pairs [57], as opposed to the local resonance of other core instabilities. In the CI, sinusoidal perturbations to nominally parallel vortex lines grow because small velocity disturbances on one core amplify the perturbation on the other core. As this amplification progresses, the perturbations grow exponentially (in the linear regime) and can eventually lead to vortex interactions and the formation of vortex rings. Originally identified in aircraft trailing vortices, the CI sets a length-scale for wake breakdown and therefore guides aircraft design [57]. The CI is relatively unique in that it is generally excited at longer wavelengths than at shorter ones, in direct contrast to the RMI. Additionally, the Crow mechanism can act on scales much larger than the core radius and is relatively insensitive to small-scale core structure. Figure 1.5 shows a pair of Crow-unstable vortex cores. The initial sinusoidal perturbation in the two cores grows over time until vortex reconnection forms a vortex ring. Theoretical results for this mechanism in cylindrical geometries have been explored as well [255], showing good agreement with experimental results [161].

1.2.3 Cavitation Bubble Dynamics

When a blast wave impinges on the head, a compressive pressure wave enters the skull and is transmitted into the brain. Skull flexure, impedance mismatches at the cerebrospinal fluid

(CSF)/brain interface, and geometric focusing of waves can lead to local tensile stresses that fall below the cavitation threshold of CSF and interstitial fluid [82, 181]. Experimental and numerical results have suggested bubble nucleation and growth can occur under such tensile spikes [159], while other studies have demonstrated direct cell death with sufficient bubble growth due to mechanical stretching of tissue [70].

Approximating the behavior of spherical bubbles can be done via the Keller-Miksis equation [122], given by,

$$\left(1-\frac{\dot{R}}{c}\right)R\ddot{R}+\frac{3}{2}\dot{R}^2\left(1-\frac{\dot{R}}{3c}\right)=\left(1+\frac{\dot{R}}{c}\right)\frac{1}{\rho_l}\left[P_B(R,t)-P_\infty\right]+\frac{R}{\rho_l c}\frac{dP_B(R,t)}{dt}, \tag{1.3}$$

where R is the bubble radius, $\dot{R}$ is the bubble wall velocity, $\ddot{R}$ is the acceleration, ρ_l is the liquid density, c is the speed of sound in the liquid, P_B is the bubble pressure, and $P_\infty(t)$ is the exterior driving pressure. This model, which can be adapted to account for effects such as viscosity and surface tension, is useful for predicting how much a bubble will grow under an external pressure field. In bTBI, the intracranial waves induced by the blast can generate sufficient tension to drive R beyond biological thresholds; once the hoop stretch exceeds neuronal limits, direct injury can occur [70]. Thus, understanding the mechanisms underlying these conditions is useful for refining safe exposure limits.

When cavitation occurs in proximity to a rigid boundary, the asymmetry of the surrounding pressure field leads to the formation of a high-speed re-entrant jet that pierces the bubble and impacts the bubble surface [115, 132, 139, 143, 264]. The resulting jet and the shock waves emitted during the final stages of collapse can cause significant surface damage, and in biological contexts, potentially perforate cell layers. While this dissertation concentrates on the pathway from blast waves to cavitation in bTBI, an appendix reports complementary simulations of single-bubble collapse near a corner [264]. Those results illustrate how geometry and stand-off distance modulate jet direction and peak wall pressure, thereby improving predictions of damage.

1.3 Numerical Methods and Physical Models

Accurate simulations of compressible, multi-material flows hinge on the representation of the interface and the discretization of the governing equations. To simulate practical systems, methods must stably capture shocks and conservatively track, capture, or resolve interfaces, while achieving sufficient accuracy to model smooth dynamics such as turbulence. Interface treatments fall into "sharp" and "diffuse" paradigms, each with distinct algorithmic nuances and physical assumptions. Discretizations vary from first-order (robust but dissipative) to spectral (high accuracy but fragile at discontinuities), with many general frameworks landing somewhere between the two. When

applied to modern heterogeneous HPC systems alongside AMR, sufficient throughput can be achieved for predictive three-dimensional simulations if accurate and efficient methods and models are employed.

The fundamental continuum model of viscous, compressible fluid flow is given by the Navier-Stokes equations,

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \nabla\cdot(\rho \boldsymbol{u}) &= 0 && \text{(mass conservation)}, \\ \frac{\partial \rho \boldsymbol{u}}{\partial t} + \nabla\cdot(\rho \boldsymbol{u}\otimes\boldsymbol{u} + P\mathbb{I}) &= \nabla\cdot\boldsymbol{\tau} && \text{(momentum balance)}, \\ \frac{\partial \rho E}{\partial t} + \nabla\cdot(\boldsymbol{u}(\rho E + P)) &= \nabla\cdot(\boldsymbol{u}\cdot\boldsymbol{\tau} + \boldsymbol{q}), && \text{(energy balance)}, \end{aligned} \tag{1.4}$$

where ρ, $\boldsymbol{u}$, P, $\mathbb{I}$, E, $\boldsymbol{\tau}$, and $\boldsymbol{q}$ are the density, velocity, pressure, identity tensor, total energy, viscous stress tensor, and heat flux, respectively, and we have ignored body forces like gravity. The equations of motion for interfacial flows are the Navier-Stokes equations supplemented by an indicator function, such as a level set [219] or a fluid volume fraction [19], with appropriate jump conditions. Only reduction-consistent models for multi-material flows recover the Navier-Stokes equations in pure materials, which is a crucial property [94]. The system is closed by an equation of state, $e = e(\rho, P)$ and $T = T(\rho, P)$, where the form of these expressions depends on the fluid model, and from which the sound speed c is derived. Such a model can be used to capture the dynamics of gases and liquids. Radiative and electromagnetic effects are readily included, with various different formulations [12, 29, 42, 77, 83, 163].

A variety of numerical methods can be applied to approximate solutions to the governing equations and study the flows of interest; finite-volume methods are a robust choice for hyperbolic conservation laws [68, 81, 250], and are given generally in one dimension by,

$$\frac{\partial \boldsymbol{q}}{\partial t} + \frac{\partial \boldsymbol{F}(\boldsymbol{q})}{\partial x} = 0 \quad \rightarrow \quad \frac{\partial \bar{\boldsymbol{q}}_i}{\partial t} + \frac{\hat{\boldsymbol{F}}_{i+1/2} - \hat{\boldsymbol{F}}_{i-1/2}}{\Delta x_i} = 0, \tag{1.5}$$

where $\boldsymbol{q}$ is a vector of conserved variables, $\boldsymbol{F}$ is the advective flux, $\Delta x_i = x_{i+1/2} - x_{i-1/2}$ is the cell size, $x_{i\pm1/2}$ are the locations of the cell faces, $\bar{\boldsymbol{q}}$ is the cell-average value of $\boldsymbol{q}$, and $\hat{\boldsymbol{F}}$ is an upwind flux for $\boldsymbol{F}$, typically computed with a Riemann solver [243]. This type of scheme is readily applied to nonlinear systems, and parabolic contributions are often added via finite-difference methods applied to cell averages. Temporal integration is typically done explicitly via a Runge-Kutta type method-of-lines scheme [85]. More sophisticated schemes, such as space-time ADER [279] and spectral elements [207], can be applied to these problems but often struggle with robustness.

1.3.1 Sharp, Diffuse, and Phase-Field Interface Models

The treatment of interfaces is paramount to the behavior and accuracy of multi-material flow simulations. Sharp interface models like the ghost-fluid method [20] or front-tracking methods [245] keep interfaces sharp, but can struggle with topological changes, dynamically created or destroyed interfaces, mesh conformity, and strict conservation [20, 245, 246]. Level-set methods [5, 80, 176, 219] evolve a signed distance function $\phi(\mathbf{x}, t)$ via an advection equation,

$$\frac{\partial \phi}{\partial t} + \boldsymbol{u} \cdot \nabla \phi = 0, \tag{1.6}$$

where $\phi(\mathbf{x}, t) = 0$ defines the interface. These methods simplify the calculation of normal vectors and curvatures for effects such as surface tension, but often struggle to preserve conservation due to the diffusion of the distance function.

Diffuse interface methods [13, 121], on the other hand, spread the interface over a few cells via a smooth indicator function, allowing for simpler treatment of dynamic topological changes at the cost of artificial mixing, altered wave dynamics, and potentially pressure and temperature issues [116]. One such model is the five-equations system [13], which is a simplification of the Baer-Nunziato seven-equations model [19] under the assumption of equilibrium velocities and pressures for all materials in a cell. The volume fraction of a given material p in a cell, α_p, is transported in this model via,

$$\frac{\partial \alpha_p}{\partial t} + \nabla \cdot (\alpha_p \boldsymbol{u}) = \alpha_p \nabla \cdot \boldsymbol{u}. \tag{1.7}$$

This model is not strictly hyperbolic due to the nonconservative term on the right-hand side. Further, like contact discontinuities, the wave that transports α_p is a wave of a linearly degenerate characteristic field, such that there is no steepening mechanism to act against numerical diffusion as there is for shock waves. As a result, any numerical diffusion leads to artificial mixing at material interfaces. Methods like THINC [223] attempt to reconstruct interfaces with a specified profile, such as a hyperbolic tangent, to keep them from spreading over more than a few cells, but can struggle with consistency. The volume-of-fluid (VOF) method tracks a smooth indicator function like level-set and diffuse interface methods, but uses geometric reconstructions of the indicator to approximate a sharp interface in each cell [92]. While this interface representation is attractive, conservation, consistency, and N-material extensions have all proven complex [16, 17, 145].

Having origins in techniques for materials science [14, 39], Phase-Field models bridge the gap between sharp and diffuse interface models by incorporating additional fluxes into the governing equations that drive the interface to a known profile, typically over one time step [96, 99, 100]. These models have been applied to incompressible flows to include solid-liquid-gas phase-change dynamics [101] and contact line physics [102]. More recently, formulations have been presented for

compressible multiphase flows [93, 107, 109, 110, 124]. One such formulation is the conservative diffuse interface method [109], where the volume fraction transport equation is given for a two-phase flow by,

$$\frac{\partial \alpha_1}{\partial t} + \nabla\cdot(\alpha_1 \boldsymbol{u}) = \alpha_1 \nabla\cdot \boldsymbol{u} + \nabla\cdot\left[\Gamma_I\left\{\epsilon_I \nabla\alpha_1 - \alpha_1(1-\alpha_1)\boldsymbol{n}_1\right\}\right], \tag{1.8}$$

where Γ_I is an artificial velocity scale for the regularization, ϵ_I controls the interface thickness, and $\boldsymbol{n}_1 = \nabla\alpha_1/|\nabla\alpha_1|$ is the interface normal [109]. Another formulation is the consistent and conservative Phase-Field method, which guarantees several key properties related to accuracy, consistency, conservation, and boundedness [93], and has been applied to N-material flows [94]. Phase-Field methods, however, do require assumptions about the interface's behavior that vary across models, and some models impose strict time-step constraints or require implicit time steps.

Regardless of the chosen interface model, satisfying consistency and conservation is essential. Lack of conservation can directly impact wave propagation [81, 137], leading to incorrect shock speeds and interface motion. Schemes that are not reduction consistent not only fail to recover the Navier-Stokes equations in pure material regions, but also can lead to voids and overfull cells, which violate key assumptions of the governing equations [93, 94]. Additionally, the choice of interface model directly affects numerical stability [93, 107, 109] and introduces additional considerations for implementations within AMR and HPC frameworks.

1.3.2 High-Order Discretizations

Increasing the order of accuracy of the spatial discretization of the governing equations trades memory (in the form of cell count) for computational effort. Increasing the accuracy of the numerical scheme not only improves the overall solution accuracy but also reduces numerical diffusion at the interface in diffuse-interface methods. One of the most popular schemes to improve upon Godunov's first-order method [81] is the MUSCL scheme [250], where linear reconstructions of the solution provide second-order accuracy in smooth regions (although higher-order extensions exist [251]) while remaining robust by satisfying the total-variation diminishing property with limiters [250]. Slope limiters are applied in the reconstruction to avoid spurious oscillations, with the reconstructed states for the face at $x_{i+1/2}$ given by,

$$q^L_{i+\frac{1}{2}} = q_i + \tfrac{1}{2}\,\phi(r_i)\,(q_{i+1} - q_i), \qquad q^R_{i+\frac{1}{2}} = q_{i+1} - \tfrac{1}{2}\,\phi(r_{i+1})\,(q_{i+2} - q_{i+1}), \tag{1.9}$$

where the ratio of successive gradients is,

$$r_i = \frac{q_i - q_{i-1}}{q_{i+1} - q_i}, \tag{1.10}$$

and ϕ is a slope limiter, which has common forms,

$$\phi_{\text{minmod}}(r) = \max\left(0, \min(1, r)\right), \qquad \phi_{\text{vanLeer}}(r) = \frac{r + |r|}{1 + |r|}. \tag{1.11}$$

The upwind flux is then computed as,

$$\hat{F}_{i+1/2} = \mathcal{R}\left(q^L_{i+\frac{1}{2}}, q^R_{i+\frac{1}{2}}\right), \tag{1.12}$$

where $\mathcal{R}$ is a Riemann solver.

High-order accuracy (greater than second) can be achieved by a variety of methods: Compact finite differences achieve high-order accuracy using implicit relations between derivative values at neighbor nodes, providing spectral-like resolution on uniform grids [136]. The finite-volume scheme, Eq. 1.5, can be extended to high-order accuracy via reconstructions like ENO or WENO, which adaptively select the smoothest local stencil to avoid spurious oscillations near discontinuities [111, 113, 147]. These methods, however, are challenging to apply to complex geometries and require filtering or limiting to remain stable at discontinuities.

Discontinuous Galerkin (DG) methods [199], alternatively, can achieve arbitrarily high-order accuracy in smooth regions without complex reconstructions or implicit steps, and can be readily applied to unstructured and non-conformal meshes [49]. Further, DG schemes can achieve super-convergence, with orders of accuracy of up to $O(h^{2p+1})$ in the cell-average and at Radau points for hyperbolic conservation laws in some cases [9, 40, 205], and even greater for other types of equations [117]. The formulation for DG methods begins with the expansion,

$$u(x \in \Omega_e, t) \approx u^e(x, t) = \sum_{j=1}^{p+1} a_j^e(t) v_j(x), \tag{1.13}$$

where Ω_e is the domain of a cell e, u_e is the finite-dimensional expansion of u in cell e, $a_j^e(t)$ are temporally varying expansion coefficients, and $v_j(x)$ are basis functions. The weak formulation then follows after multiplying the governing equations by the basis functions and integrating over each cell. The resultant update scheme is computed via numerical quadrature. Compared to finite volume schemes of the same order of accuracy, DG schemes tend to be less dissipative. Figure 1.6 shows this property by computing the evolution of an acoustic wave with a fifth-order accurate WENO scheme and a fifth-order accurate ($p = 4$) DG scheme, each with the same number of unknowns. In this example, an acoustic wave travels from left to right in an ideal gas; the exact solution is that the wave remains unaltered as it propagates to the right. The DG method produces a much more accurate solution than the fifth-order WENO scheme with the same amount of data.

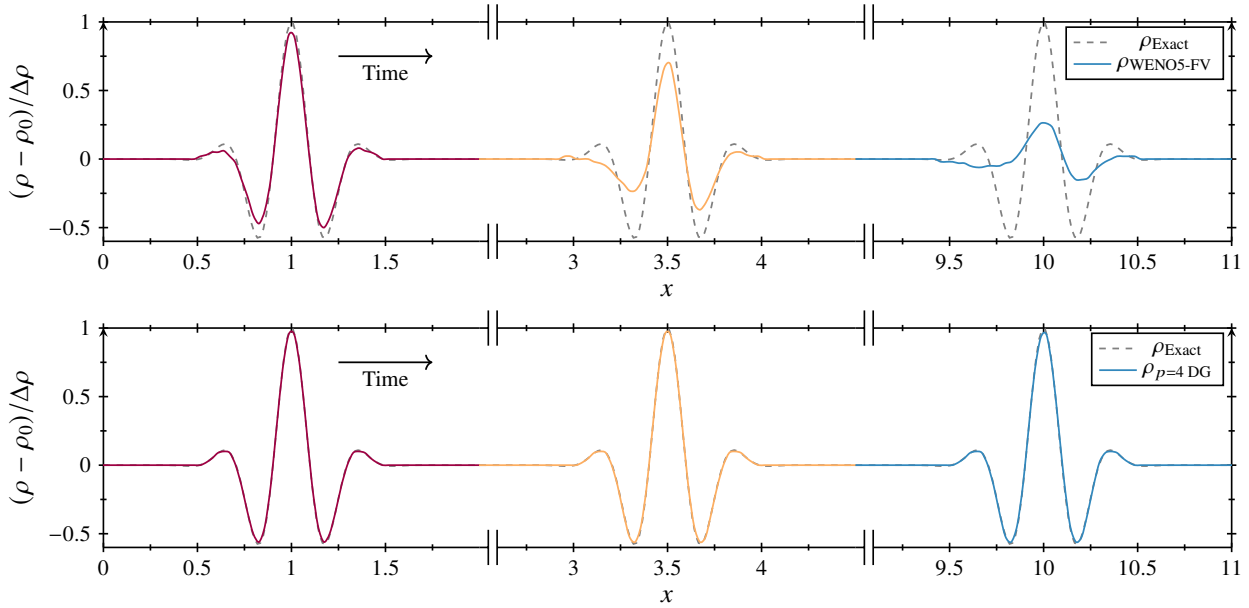


Figure 1.6: Acoustic wave solutions at three sequential times (red, orange, blue) from a fifth-order WENO finite volume scheme (top) and a fifth-order $p = 4$ DG scheme (bottom), with equal numbers of unknowns.

This property makes DG methods attractive for large-scale calculations, where memory is often the limiting factor, as greater accuracy can be achieved by performing more calculations on the same compact stencil of data than a comparable finite-volume scheme. Stability, however, is not trivial for DG methods, and limiting schemes [49, 126], artificial viscosity [185], and subcell capturing methods [66, 112, 206, 207] are all applied to maintain robustness.

1.3.3 Adaptive Mesh Refinement and High-Performance Computing

While increasing the order of accuracy of a scheme generally improves solution quality, high mesh densities are still required to accurately capture discontinuous features such as shocks and interfaces, as well as smooth features such as vortices [118]. Uniformly increasing the mesh density is sufficient for some basic problems, but this technique becomes intractable for many applications. As such, AMR has become an essential tool in many numerical settings. AMR concentrates cells around discontinuities or other key features, such as vortices, often reducing total cell counts by orders of magnitude compared to simulations with the same cell size but uniform grids. For steady problems, mesh refinement can be performed *a posteriori* using a variety of sophisticated error metrics [253, 270], whereas for dynamic problems, simple error measures, such as gradients, often suffice. This technique is not new [6, 26, 27], but some versions have suffered in the past from implementations that do not scale on massively parallel computing systems. One counterexample,

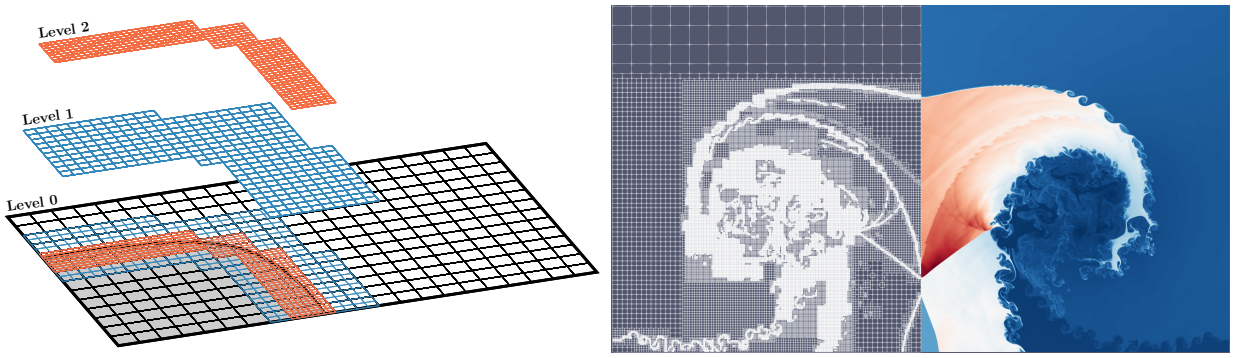


Figure 1.7: A schematic of patch-based adaptive mesh refinement (left) and a hydrodynamic triple-point simulation with AMR patches and density shown (right).

however, is the AMReX library [282], which utilizes patch-based mesh refinement to simplify solver implementations and improve kernel and data efficiency [41, 170, 220].

In patch-based refinement, logically rectangular boxes of cells are nested, with decreasing cell size, and conservative averaging and integration techniques are applied at coarse-to-fine boundaries. Figure 1.7 shows a schematic of one such mesh, as well as a solution to the hydrodynamic triple-point problem [178, 182, 283] computed with AMR. In the triple-point example, the mesh patches concentrate around large density gradients to resolve contact discontinuities and shocks with greater accuracy than the nearly constant background flow. The simplified grid structure in patch-based methods reduces overheads in connectivity information and enables efficient parallelization. AMReX, specifically, supports multiple levels of parallelism on the CPU, thread, and GPU levels, enabling efficient utilization of modern HPC architectures. Further, the portable CPU/GPU parallelism in AMReX makes derived software readily adapted to a variety of systems, increasing usability.

Running simulations efficiently on HPC systems requires methods that scale; patch-based mesh refinement has shown good results in this area, particularly when combined with efficient dynamic load balancing [282]. To maximize the efficiency of simulations of interfacial flows, interface treatments that allow for even load balancing should be used. Unlike interface-tracking methods, which require substantial computational effort focused on very small regions of the domain, diffuse interface methods are agnostic to the interface from a computational perspective and have been shown to achieve near-optimal scaling on HPC systems [36].

1.4 Dissertation Overview

This dissertation advances the state of the art across the areas of methods, software, and physics in the field of compressible interfacial flows. A comprehensive numerical toolkit for the simulation

of compressible interfacial flows is presented in Chapter 2. This software serves as the basis for the remaining chapters, including algorithmic advances for high-order-accurate multiphase simulations, as presented in Chapter 3. Chapters 4 and 5 apply the code to investigate the role of vortex pair instabilities in the generation of three-dimensional structure from Richtmyer-Meshkov flows. The wave dynamics that may lead to cavitation in the brain are described in Chapter 6, with theory supplemented by simulations. Finally, Chapter 7 provides a summary of this dissertation, discusses the limitations of the current work, highlights the research enabled by this work, and suggests future directions. The research contributions from the chapters and appendices of this dissertation include:

1. An open-source, GPU-enabled, AMR code for simulations of compressible interfacial flows with consistent and conservative Phase Fields. This simulation tool is a unique combination of methods, models, and software capabilities that has been used for studies across a range of applications, including high-speed droplet dynamics [37, 38]. astrophysics [259], and methods development [104, 266].

2. A high-order DG method for consistent and conservative Phase-Field modeling of compressible interfacial flows (W. J. White, Z. Huang, and E. Johnsen, A high-order discontinuous Galerkin method for compressible interfacial flows with consistent and conservative Phase Fields. *Journal of Computational Physics* 527:113830, 2025). With subcell finite volume shock and interface capturing, Phase-Field interface regularization, and boundedness, this scheme provides a robust and accurate method for interfacial flow simulations.

3. A novel pathway for the generation of three-dimensional structure from nominally two-dimensional RMI flows. Numerical simulations demonstrate that an extruded RMI flow with a transverse perturbation can develop large-scale vortex rings, and that the perturbation growth rates tend to those predicted by Crow instability theory as the Atwood number increases.

4. A demonstration of the instability of pairs of vortex rings in converging RMI systems, with relevance to double-shell ICF. Simulations of an imploding cylindrical gas-gas interface show clear exponential growth in pairs of counter-rotating vortex rings, as well as non-monotonic behavior of the measured growth rates with respect to wavenumber, suggesting vortex-pair instabilities.

5. A wave-based framework for understanding cavitation damage potential in blast-induced traumatic brain injury. Using a sequence of models ranging from one to three dimensions, along with accompanying simulations, the generation of cavitation-relevant tension is observed and explained.

6. The characterization of cavitation bubble dynamics in corner geometries and the resultant damage potential (W. White, S. A. Beig, and E. Johnsen, Pressure fields produced by single-bubble collapse in a corner, *Physical Review Fluids*, 8(2):023601, 2023). Quantification of wall pressures and temperatures resulting from bubbles at various positions near a corner enables improved damage predictions.

Part I

Methods and Models for Compressible Interfacial Flows

CHAPTER 2

COMPAS: A GPU-Accelerated, Adaptive-Mesh Finite-Volume Framework for Consistent and Conservative Phase-Field Modeling of Compressible, N-Material, Interfacial Flows

2.1 Abstract

This work presents COMPAS (COmpressible Multimaterial multiPhysics Adaptive Solver), a GPU-accelerated adaptive-mesh framework for compressible, interface-resolved simulations with shocks and N materials. The recently introduced consistent and conservative Phase-Field method is implemented in a robust, accurate, and efficient manner for high-performance computing (HPC) systems. The code is written with the adaptive mesh refinement (AMR) library AMReX and leverages its multi-level heterogeneous parallelization, grid and data management, and parallel I/O operations. The structured patch-based AMR enabled by AMReX allows simple kernel definitions and index-based solver design, with appropriate boundary operations to ensure

consistency and conservation across AMR levels. This software implements several recent methods for compressible N-material interfacial flows and Phase Fields, including a consistent limiting procedure to ensure consistency of reduction in the reconstruction step, and a stable, consistent time-advancement strategy for the Phase Field on AMR grids. Modular, object-oriented software design enables straightforward development and testing and makes the code ready for future multi-physics extensions. Several Phase-Field mechanisms are implemented, including the conservative Allen-Cahn, Cahn-Hilliard, and (accurate) conservative diffuse-interface models, all of which are incorporated into the consistent and conservative Phase-Field framework. The ability of `COMPAS` to accurately resolve a variety of flows is demonstrated through several examples, and the performance is quantified. Overall, `COMPAS` serves as an accurate, robust, and efficient computational framework for studying compressible interfacial flows with Phase-Field interface regularization and adaptive mesh refinement on HPC systems.

Program summary

Program Title: `COMPAS`

Programming language: C++

External libraries/routines: AMReX (https://github.com/AMReX-Codes/amrex)

Nature of problem: Solving the compressible, N-material Navier-Stokes equations for interfacial flows. Numerically thin interfaces are maintained via the recently introduced consistent and conservative Phase-Field framework, which includes different interface regularization mechanisms, including the conservative Allen-Cahn equation and the (accurate) conservative diffuse interface formulation.

Solution method: A finite-volume approach is used to discretize the governing equations in space, with a second-order MUSCL reconstruction and various limiters. Explicit time marching of various orders advances the advection, diffusion, and source terms of the system, while the Phase-Field terms are applied via a fractional step approach, for which novel implicit and explicit formulations are available. Patch-based adaptive mesh refinement and heterogeneous multi-level parallelism are enabled by the AMReX library.

2.2 Introduction

Compressible interfacial flows are ubiquitous in science and engineering systems, including inertial confinement fusion experiments [56, 144], cavitation [114, 189], astrophysics [15, 172, 224, 259], and high-speed droplet dynamics [191]. Each of these systems is characterized by transient nonlinear phenomena, including shocks, dynamic interface motion, the onset and growth of hydrodynamic instabilities, and potentially multi-physics effects such as radiation or electromagnetic effects. These

phenomena make experimental investigations challenging and also impose significant numerical difficulties due to steep gradients, multiscale physics, and modeling complexities [3, 245]. As such, numerical methods that robustly and accurately represent the physics of compressible interfacial flows are useful tools for research into otherwise intractable systems.

The treatment of material interfaces between phases (gases, liquids, solids) has been the primary challenge in extending the catalogue of numerical techniques for single-phase flow to multi-material systems. While some interfaces are macroscopically thick and studied with interface-resolved simulations [174], immiscible interfaces present several unique challenges. Many strategies have been proposed for immiscible interfaces with notable strengths and weaknesses. Early methods like particle-in-cell [86] and marker-and-cell [87] found success but struggled with diffusion or complex topological changes, limiting their applicability to problems that exhibit highly nonlinear effects. Interface-capturing schemes have become popular for modeling multiphase flows, with notable variants including the volume-of-fluid (VOF) [92] and level-set (LS) [175] methods. VOF methods are often coupled with interface reconstruction methods, such as piecewise linear interface calculations [202, 278]. Level-set methods instead evolve a signed distance function via an advection equation [5]. While each of these techniques has found success, their application to compressible, N-material interfacial flows remains complicated due to the need for geometric reconstructions, difficulties in ensuring conservation, and numerical issues when interfaces interact with strong shocks or when several materials interact simultaneously.

Diffuse interface methods (DIM), in contrast, rely on a smooth representation of the interface alongside thermodynamically consistent mixture relations to capture both shocks and interfaces simultaneously [3, 13, 19] without distinct numerical treatment. Flow models within the Baer-Nunziato hierarchy provide multi-material evolution equations that are rigorously conservative, with varying degrees of interface modeling [19]. In the context of adaptive mesh refinement (AMR) and high-performance computing (HPC) systems, DIMs are particularly attractive because they avoid explicit interface treatment that can lead to load-balancing issues, accommodate dramatic topological changes, and readily accommodate additional physics, all within an Eulerian framework.

Despite these advantages, DIMs are hampered by significant limitations. The use of upwind-biased advection and standard finite-volume discretizations introduces numerical diffusion across material interfaces, leading to spurious material mixing, the possible generation of fictitious phases [94], and possibly errors in the pressure and temperature [2, 116]. Increasing mesh resolution reduces the thickness of interfaces in calculations but does not eliminate the issue, and long-time and large-scale calculations make such a mitigation strategy impractical. Further, ensuring that the flow variables remain bounded (e.g., positive density) is not trivial, and many schemes fail this property in reconstructions.

Adaptive mesh refinement (AMR) addresses these challenges by enabling significant increases

in mesh resolution at shocks and interfaces through dynamic grid refinement. AMR is not trivial, however, and involves trade-offs between the complexities of mesh connectivity, non-conformal integration, or periodic remeshing for tree-based and unstructured meshes. Patch-based refinement, as implemented in frameworks such as AMReX [282], Chombo [6], and others [26, 27], trades some extraneous storage and synchronization for logically rectangular, hierarchically refined meshes, which readily support local refinement and multi-level parallelism. Such a technique is particularly useful for compressible interfacial flows where the key dynamics span several cells, making hyper-local refinement unnecessary. AMReX natively supports multi-level heterogeneous CPU/GPU parallelism, parallel input/output operations, and comprehensive compilation routines for HPC systems, making it ideal for high-resolution three-dimensional simulations.

Despite the high mesh resolution enabled by AMR and modern HPC systems, long-time, large-scale simulations remain challenging for DIMs, motivating the development of techniques to address these challenges. The THINC method [274] seeks to sharpen a smooth interface profile by applying an algebraic reconstruction of the interface within just a few computational cells. Phase-field models, introduced in materials science [124], provide physical regularization of the interface width while enforcing conservation and enabling natural coupling to thermodynamics via an auxiliary field governed by partial differential equations such as the Allen-Cahn or Cahn-Hilliard equations. The accurate conservative diffuse-interface (ACDI) scheme [110] implements these ideas for two-material configurations, enabling conservation and better interfacial normal calculation. With these methods, interfaces are kept thin and uniform throughout simulations. Each of these approaches yields improvements in interface sharpness, conservation, or extensibility, but their adaptation to compressible, consistent, conservative, N-material flow solvers with AMR remains an ongoing area of development.

There are numerous open-source platforms for simulating compressible multi-material flows, with a wide range of numerical methods and models within each reflecting the broad interest in this flow regime. In MFC [36] and ECOGEN [215], the Baer-Nunziato family of multiphase models is implemented and includes multi-physics extensions. FLASH [79] is widely used for designing fluid and plasma experiments, while Basilisk [190] targets general-purpose interface-tracking within a unified infrastructure for a variety of flows, but with a focus on lower Mach numbers and two-phase flows.

Against this background, several needs can be identified: a) reduction-consistent, conservative interface handling in compressible, general N-material flows; b) systematic integration of Phase-Field interface regularization with robust shock- and interface-capturing and adaptive mesh refinement; c) scalability and adaptability for heterogeneous HPC architectures; and d) modular, open-source frameworks that can accommodate simple method extension and multi-physics coupling. This work addresses these challenges by introducing an open-source simulation frame-

work for compressible, N-material interfacial flows, built upon AMReX. The software centers on the consistent and conservative Phase-Field approach [93, 94, 100], which incorporates several interface-regularization models, such as ACDI [107] and conservative Allen–Cahn [48], and is developed for general N-material systems. The resulting platform enables users to combine advanced finite-volume and finite-difference solvers with flexible Phase-Field modeling, patch-based AMR, and portability to various HPC systems. The object-oriented and modular design readily supports the extension to other models in the Baer-Nunziato family [19], as well as to additional interface-regularization mechanisms. Separating physics and numerics wherever possible and abstracting the governing equations into kernels facilitates the addition of multi-physics effects (such as radiation hydrodynamics, magnetohydrodynamics, and hyper-elasticity) in the future without the need for significant code modification.

Verification, validation, and performance examples are provided through select test cases, including: isentropic vortex advection (measuring the order of accuracy and verifying reduction consistency); spherical bubble collapse (comparing Euler and Navier-Stokes models with theory); multi-material shock-driven interfacial instabilities (verifying consistency of reduction for more than two materials and validating against experiments); and a Taylor-Green vortex (demonstrating GPU scaling). These test cases show that the proposed method, model, and implementation accurately solve problems of interest in a computationally efficient manner.

Relative to widely used simulation platforms such as MFC [36], FLASH [79], Basilisk [190], ECOGEN [215], CASTRO [15], and MHIT36 [203], the present solver offers access to the consistent and conservative Phase-Field framework in a modular, adaptive-mesh, GPU-accelerated package. Through the general Phase Field implementation, the code incorporates several interface regularization models within a single, consistent, and bound-preserving numerical framework, enabling direct comparison and serving as a practical foundation for ongoing developments. These features make `COMPAS` unique and position it as a tool with a low barrier to entry for HPC-scale simulations of compressible interfacial flows.

The rest of the study is organized as follows. Section 2.3 describes the implementation of the code, running jobs, and pre- and post-processing data. Section 2.4 describes the consistent and conservative Phase-Field method [93, 94, 101] and its application in `COMPAS`, as well as the fluid model. Verification, validation, and scaling results are shown alongside example cases in Section 2.5, with conclusions provided in Section 2.6.

2.3 Code Overview

The distinguishing feature of `COMPAS` compared to other open-source simulation tools is its implementation of the general, consistent, and conservative Phase-Field method of [93] for compressible

N-material flows [94] into a GPU-enabled AMR framework [282]. Phase-Field interface regularization maintains thin interfacial regions of uniform thickness, and the present framework ensures this is done consistently and conservatively. This general framework incorporates different interface regularization techniques, including the popularly used conservative Allen-Cahn [14], Cahn-Hilliard [39], conservative diffuse interface [109], and accurate conservative diffuse interface [107] mechanisms, with an accurate and robust numerical implementation. Each of these mechanisms treats mixture regions differently; the conservative Allen-Cahn model, for example, conservatively diffuses small-scale structures, while the conservative diffuse interface mechanism breaks structures into small persistent droplets. With several of these models implemented, users can freely choose the mechanism that best suits their modeling needs. This characteristic, in combination with the AMReX framework, thus serves as both a highly capable tool for large-scale simulations of compressible interfacial flows and a useful tool for comparing interface modeling techniques.

In general, `COMPAS` solves hyperbolic-parabolic partial differential equations with source terms using finite-volume and finite-difference methods (although high-order extensions exist [266]), with specific emphasis on the Baer-Nunziato class of multi-component flow models. The core model is the two-/N-material compressible five-equations model [13, 121], which is the focus of this open-source version. A general convex equation of state that captures the ideal gas and stiffened gas models closes the system. A second-order spatial discretization underpins the main solvers, with the LLF, HLL, and HLLC Riemann solvers available [243]. Various limiters are also implemented, including minmod, Superbee, Van Leer, and MC [232, 249, 250]. Crucially, the consistent limiter and reconstructions of [93] ensure absolute reduction consistency for N-material flows. Additionally, care is taken to ensure bounds are preserved in the reconstructions [103]. Full state, periodic, zero-gradient, slip-wall, no-slip-wall, and reflecting boundary conditions are available for user selection. A generalized inhomogeneous Robin boundary condition is provided for user-defined boundary conditions. With and without Phase Fields, the models and methods in `COMPAS` are strictly conservative by definition, including with AMR. Explicit Runge-Kutta time stepping [85] is used for the main hyperbolic-parabolic system; calculation of the bound-preserving Phase-Field fluxes may employ either explicit or semi-implicit techniques, depending on the model [103].

The AMReX library [282] is a fundamental component of `COMPAS`, and provides a robust, portable, and effective framework for patch-based adaptive mesh refinement. AMReX is used for data structures such as the `MultiFab` object, for portable parallelization through a variety of operators and functions like the GPU-enabled `ParallelFor` operator, for parallel input/output through read/write functions like `WriteMultiLevelPlotfile`, for adaptive meshing via functions like `regrid`, and for more intricate operations like conservative flux registration at

coarse-fine mesh boundaries via `Reflux`. In short, AMReX splits the user-specified problem domain into a set of logically indexed boxes, on top of which additional, more refined boxes (i.e., patches) are placed where refinement is needed. Iterators loop over these boxes and call the solver kernels that compute the update for each cell. The heavy use of kernels separates much of the present implementation from the source code, making modification, addition, and debugging simpler. The lightweight and modular structure of `COMPAS` also makes it a potent testbed for additional models and methods, with current efforts developing extensions for the Baer-Nunziato family, magnetohydrodynamics, radiation hydrodynamics, and hyper-elastic flows.

2.3.1 Installation and Testing

The `COMPAS` repository is publicly available on GitHub. Its installation is designed for both ease and portability across hardware systems. The required dependencies are limited to those of AMReX itself, namely C++17 or later, `MPI` for the base parallelism (`OpenMPI` and `MPICH` are supported), `GNU Make` for compiling, and one of `CUDA`, `HIP`, or `SYCL` if GPU parallelism is desired. Additional parallelism is available through `openMP`. The integration with AMReX only requires cloning it into the `lib` directory, or trivially changing the path to the local installation.

Testing the code can be done via example cases, such as uniform advection or a shock-induced bubble collapse, or via convergence tests, which the code can automatically run and generate convergence plots.

2.3.2 Code Structure

`COMPAS` is based on the AMReX class `AmrCore` and is designed to be lightweight and flexible, using derived classes where possible to reduce code duplication across modules. Figure 2.1 shows the file tree for the code and an example case.

The `include` directory contains C++ header files with kernel implementations for finite-volume and finite-difference operators, physics-specific functions, Riemann solvers, and the class headers for the major physical and numerical modules, as well as helper functions. Kernels are generally written to operate directly on AMReX logical patches and thus do not require mesh connectivity data. For many operator kernels, particularly those that use stencil data, there is a router kernel that takes in the (i, j, k) cell index and full AMReX `Array4` objects and extracts the appropriate cells before passing those data points to the computation kernel, which does the actual calculation. Such an abstraction layer improves code cleanliness and flexibility and makes error tracking simpler without much added cost, since the router kernel takes its arguments by reference almost exclusively. The AMReX `ParallelFor` loop acts on these kernels by iterating over the

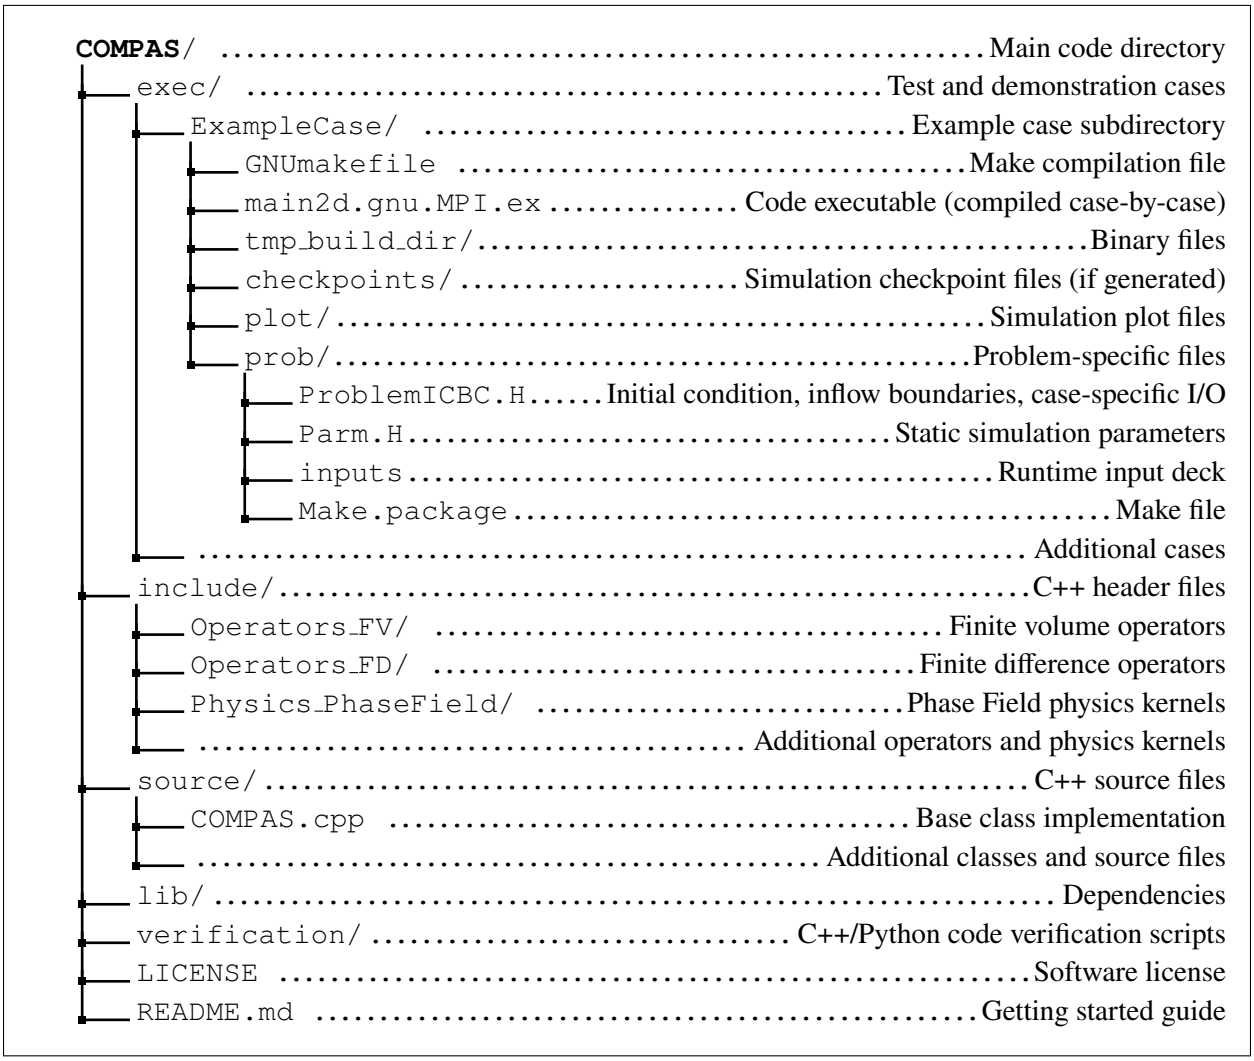


Figure 2.1: File tree for COMPAS, including an example case which has been compiled for two dimensions with MPI parallelization.

boxes/patches within a given level of the solution, as stored in a `MultiFab` object, and calls the appropriate kernels. An example of such is,

```
amrex::ParallelFor(bx,
[=] AMREX_GPU_DEVICE (int i, int j, int k){
    FVM_Conservative2FluxHyperbolic_K(i, j, k, ... );
});
```

where `FVM_Conservative2FluxHyperbolic_K` computes the hyperbolic flux along the d-axis at face $(i - \delta_{x,d}/2, j - \delta_{y,d}/2, k - \delta_{z,d}/2)$ from the conservative variables. This function is executed on GPUs when available and falls back to the CPU equivalent when they are not.

The `source` directory contains all C++ source code for the main loops, time integration techniques, AMR management routines, and driver logic. The AMReX library itself resides in the `lib` directory by default, but referencing an external AMReX installation to accommodate various workflows or system deployments is trivial. The `verification` directory holds Python and C++ scripts for automatically running convergence tests. License and instructions are found in the `LICENSE` and `README.md` files, respectively. Simulation setups are managed in the `exec` directory, which contains problem-specific subdirectories, as will be discussed below.

2.3.3 Running Cases

Simulations in `COMPAS` are run from the `exec` folder, where a run directory can be made or selected for each case. As shown in Figure 2.1 for an example case, within each directory is a set of files that are specific to that problem; all required information for each simulation exists in this case directory. The initial condition is specified through a kernel in `ProblemICBC.H` that fills the state variables at each (x, y, z) position in the domain. Likewise, user-specified full-state boundary conditions are defined by a similar kernel in the same file. Problem-specific input/output operations are also defined in this file, if desired; an example application of this functionality is to track integral quantities at every timestep and write them to a file during the simulation, rather than relying on post-processing. Functions for taking lineouts and slices are also provided. The user output function has access to the full simulation object, enabling complex analyses at runtime (e.g., tracking material interface positions over time) and reducing the need for post-processing.

Default equation of state constants, solver options, and Phase-Field parameters are set from the `Parm.H` file, although AMReX-enabled parsing functionality allows all of these properties to be set at runtime from the input deck. The `input` file also contains runtime parameters, such as the simulation domain size, the number of cells in each direction, AMR criteria, and domain decomposition options.

Compiling is done case-by-case with `make`, and the `GNUmakefile` defines broad properties

of the simulation, including the dimensionality, parallelization strategy (options for `MPI`, `OpenMP`, and GPU parallelization), and the physical model. While compiling the code case-by-case requires additional compilation time if many distinct cases are to be run, this strategy eliminates the need for the user to modify any source code while still having broad control over initial conditions, boundary conditions, I/O, and solver selection. Further, because of the minimal dependencies of `COMPAS`, compile times are usually < 1 minute when using `make -j<T>` for `T` threads.

Running the code once it has been compiled is simply executed by issuing a standard `mpirun`-type command (if running in parallel) from the problem directory, providing the number of cores and the input deck, e.g., `mpirun -n 4 ./main2d.gnu.MPI.ex ./prob/inputs`, which launches a two-dimensional executable with `MPI` parallelization that was generated with a `gnu` compiler on four `MPI` ranks. Additional flags are often required on heterogeneous HPC systems to specify how compute resources are used on a node. Similar commands are used when starting via job schedulers such as `slurm`.

Once the executable is launched, the code reads the input deck, AMReX creates the grids and data structures, and the initial conditions are populated. An initial time step is calculated, after which the time-stepping loop begins. At each time step, AMReX manages dynamic mesh adaptivity, ghost-cell filling, temporal subcycling coordination, flux reconciliation across AMR levels, and parallel I/O. Plot and checkpoint restart files can be written at step or time intervals, or at specific times; checkpoint files are used to restart a simulation after it has completed or failed. The plot files are written to the native AMReX file format, which can be read directly by both Paraview [10] and VisIt [47]. Both of these tools have GUIs and Python interfaces that make post-processing straightforward; see the AMReX documentation [282] for more information.

2.4 Consistent and Conservative Phase-Field Method

The consistent and conservative Phase-Field method [93, 94, 100] provides a general framework for coupling different Phase-Field mechanisms with compressible multiphase flows, via a physically consistent model and robust numerical approach. This coupling enables quantitative control over interface thickness and uniformity, resulting in an accurate representation of material interfaces even when complex wave-interface interactions occur in compressible multiphase flows. In this section, both the model and numerical approach are summarized.

2.4.1 Physical Model

The governing equation for compressible N-material flows with Phase-Field mechanism [100] is

$$\frac{\partial}{\partial t}\underbrace{\begin{bmatrix}\{(\alpha_p\rho_p)\}_{p=1}^N\\(\rho\mathbf{u})\\(\rho E)\\\{\alpha_p\}_{p=1}^N\end{bmatrix}}_{\mathbf{U}}+\nabla\cdot\underbrace{\begin{bmatrix}\{(\alpha_p\rho_p)\mathbf{u}\}_{p=1}^N\\(\rho\mathbf{u})\otimes\mathbf{u}+P\mathbf{I}\\((\rho E)+P)\mathbf{u}\\\{\alpha_p\mathbf{u}\}_{p=1}^N\end{bmatrix}}_{\mathbf{F}^{HB}}=\underbrace{\begin{bmatrix}\{0\}_{p=1}^N\\\mathbf{0}\\0\\\{\alpha_p(1+K_p)\}_{p=1}^N\end{bmatrix}}_{\mathbf{H}^{HB}}\nabla\cdot\mathbf{u}+\nabla\cdot\underbrace{\begin{bmatrix}\{\mathbf{0}\}_{p=1}^N\\\boldsymbol{\tau}\\\mathbf{u}\cdot\boldsymbol{\tau}+\mathbf{q}\\\{\mathbf{0}\}_{p=1}^N\end{bmatrix}}_{\mathbf{F}^{DF}}+\nabla\cdot\underbrace{\begin{bmatrix}\{\mathbf{J}_p\rho_p\}_{p=1}^N\\\sum_{p=1}^N\mathbf{J}_p\rho_p\otimes\mathbf{u}\\\sum_{p=1}^N\mathbf{J}_p\rho_p E_p\\\{\mathbf{J}_p\}_{p=1}^N\end{bmatrix}}_{\mathbf{F}^{PF}}, \tag{2.1}$$

where $\mathbf{U}$ is the vector of state variables, $\mathbf{F}^{HB}$ the hyperbolic flux vector, $\mathbf{H}^{HB}\nabla\cdot\mathbf{u}$ the hyperbolic source term vector, $\mathbf{F}^{DF}$ the diffusion flux vector, and $\mathbf{F}^{PF}$ the Phase-Field flux vector. ρ, $\mathbf{u}$, E, P, and α denote the density, velocity, specific total energy, pressure, and volume fraction, respectively, with the subscript p referring to the phase index. The viscous stress tensor is $\boldsymbol{\tau}$, and $\mathbf{q}$ is the heat flux. $\mathbf{J}$ is the flux resulting from any Phase-Field mechanism. The system in Eq. (2.1) is completed by supplementing the equations of state $e_p = \mathrm{e}_p(\rho_p, P)$ and $T_p = \mathrm{T}_p(\rho_p, P)$ for individual phases, where e ($= E - \frac{1}{2}\mathbf{u}\cdot\mathbf{u}$) is the specific internal energy and T is the temperature. The mixture properties are defined by $\mathcal{Y} = \sum_{p=1}^N \alpha_p\mathcal{Y}_p$ for $\mathcal{Y} = \{\rho, \rho E, \rho e\}$, from which the mixture pressure and temperature can be computed. To be consistent with $\sum_{p=1}^N \alpha_p = 1$ and the second law of thermodynamics, the following requirements must be met: $\sum_{p=1}^N \alpha_p K_p = 0$ and $\sum_{p=1}^N \nabla\cdot\mathbf{J}_p = 0$, where K depends on the speed of sound, as explained below.

We use a simple four-parameter convex equation of state that captures stiffened and ideal gases, with internal energy,

$$\rho_p e_p = A_p P + B_p + \rho_p D_p, \tag{2.2}$$

and temperature,

$$T_p = \frac{1}{\rho_p C_p}\left(\rho_p e_p - \frac{B_p}{1+A_p} - \rho_p D_p\right) \tag{2.3}$$

parameterized by four constant coefficients, A_p, B_p, C_p, and D_p, which are either experimentally fit or reduce to common expressions, e.g., $A_p = 1/(\gamma_p - 1)$ and $B_p = \gamma_p P_p^\infty/(\gamma_p - 1)$ for a stiffened gas, where γ_p is the ratio of specific heats and P_p^∞ is the stiffening pressure. The sound speed for

material p is derived from Eq. (2.2),

$$\rho_p c_p^2 = \frac{(1 + A_p)P + B_p}{A_p}, \tag{2.4}$$

and has the mixture relationship,

$$\rho c^2 = \frac{\sum_{p=1}^{N} \alpha_p A_p \rho_p c_p^2 (1 + K_p)}{\sum_{p=1}^{N} \alpha_p A_p}. \tag{2.5}$$

The choice of K_p leads to different multiphase models, with $\{K_p = 0\}_{p=1}^{N}$ resulting in the model of Allaire et al. [13] and $\{K_p = \frac{\rho c^2}{\rho_p c_p^2} - 1\}_{p=1}^{N}$ reducing to the model of Kapila et al. [121]. The advantage of such an EOS formulation is that it does not require the computation of any phasic properties unweighted by the volume fraction, i.e., operations like $\rho_p = (\alpha_p \rho_p)/\alpha_p$. This property is important for multi-material flows; $\rho_{p'}$ $(p' \neq p)$ is not well behaved in "pure" regions of material p, leading to the necessity for *ad hoc* changes to avoid division by zero, which introduces inconsistency between the mass and volume fraction and generates singular densities [93, 100].

For an arbitrary number of phases ($N \geq 1$), Eq. (2.1) satisfies:

- mass conservation for individual phases;
- momentum and energy balance for the multiphase mixture;
- second law of thermodynamics;
- Galilean invariance;
- kinematic, mechanical, and thermal equilibrium at material interfaces;
- $\sum_{p=1}^{N} \alpha_p = 1$ (no local voids or overfilling);
- consistency of reduction (no fictitious phases).

The related analysis is detailed in [100].

The modular structure of `COMPAS` allows for the straightforward addition of new systems, with additional models in the Baer-Nunziato hierarchy, magnetohydrodynamics, radiation hydrodynamics, and hyperelasticity being actively implemented by various users. Further, the class structure used makes the inclusion of Phase Fields straightforward in any of these systems.

2.4.2 Numerical Approach

A finite-volume-based numerical approach [93, 94, 100] is developed to robustly solve Eq. (2.1) on adaptive-mesh-refinement grids. The MUSCL reconstruction (with different slope limiters) [138] is primarily used, although higher-order reconstructions based on WENO [111] are available. The approach starts with the *hyperbolic* step that evolves

$$\begin{aligned}&\frac{\partial \overline{\mathbf{U}}}{\partial t} + \nabla \cdot \hat{\mathbf{F}}^{HB}\left(\mathbf{U}^L, \mathbf{U}^R\right) \\ &= \hat{\mathbf{H}}^{HB}\left(\overline{\mathbf{U}}\right) \nabla \cdot \hat{\mathbf{u}}^{HB}\left(\mathbf{U}^L, \mathbf{U}^R\right) + \nabla \cdot \hat{\mathbf{F}}^{DF}\left(\overline{\mathbf{U}}, \nabla \overline{\mathbf{U}}\right),\end{aligned} \tag{2.6}$$

where $\overline{\mathbf{U}}$ is the cell-averaged state vector, $\mathbf{U}^{L,R}$ is the reconstructed data at cell faces, and $\hat{\mathcal{F}}$ represents a numerical approximation to $\mathcal{F}$. The LLF, HLL, and HLLC numerical fluxes [242] are available for $\hat{\mathbf{F}}^{HB}$. To prevent any production of fictitious phases, local voids, or overfilling, $\hat{\mathbf{u}}^{HB} = \hat{\mathbf{F}}^{HB,\alpha_p}|_{\alpha_p=1}$ is used [93, 100]. $\hat{\mathbf{H}}^{HB}$ is an evaluation of $\mathbf{H}^{HB}$ directly with the cell-averaged data, and $\hat{\mathbf{F}}^{DF}$ uses central differences for its approximation.

The *Phase-Field* step is then performed to include the Phase-Field contribution with the multiphase reduction-consistent formulation [93, 94, 100]. Given the Phase-Field mechanism (i.e., an admissible formulation of $\{[\nabla \cdot \mathbf{J}_p]\}_{p=1}^N$ and its discretization $\{[\nabla \cdot \hat{\mathbf{J}}_p]\}_{p=1}^N$), we first update temporary volume fractions $\{\tilde{\alpha}_p\}_{p=1}^N$ with a stable time advancement via

$$\frac{\partial \tilde{\alpha}_p}{\partial \tau^{PF}} = [\nabla \cdot \hat{\mathbf{J}}_p], \tag{2.7}$$

from $\tau^{PF} = 0$ to $\tau^{PF} = \Delta t$, where Δt is the time step in the hyperbolic step, and the initial condition is the volume fractions after the hyperbolic step. As a result, the effect of the Phase-Field mechanism on the volume fractions over Δt is obtained via

$$\widehat{[\nabla \cdot \mathbf{J}_p]} = \frac{\tilde{\alpha}_p|_{\tau^{PF}=\Delta t} - \tilde{\alpha}_p|_{\tau^{PF}=0}}{\Delta t}, \tag{2.8}$$

which is then mapped to auxiliary variables $\{Q_p\}_{p=1}^N$ via

$$\nabla \cdot \left(\alpha_p(1-\alpha_p)\nabla Q_p\right) = \widehat{[\nabla \cdot \mathbf{J}_p]}. \tag{2.9}$$

Thereafter, the vector of conservative variables is updated consistently via

$$\frac{\partial \overline{\mathbf{U}}}{\partial \tau} = \nabla \cdot \hat{\mathbf{F}}^{PF}\left(\mathbf{U}^L, \mathbf{U}^R; \{\nabla Q_p\}_{p=1}^N\right), \tag{2.10}$$

from $\tau = 0$ to $\tau = \Delta t$, where the initial condition is the output of the hyperbolic step. Different options of $\hat{\mathbf{F}}^{PF}$ are available, including the Upwind-Downwind (UD) flux and HLL flux with the base (mixed) and efficient formulations [93, 94, 100]. It is clear that different Phase-Field mechanisms can be incorporated into our approach; only the right-hand side of Eq. (2.7) needs to be changed. The conservative Allen-Cahn model [48, 96] is used by default, while, with minor modifications, other Phase-Field models have been explored [103], including the Cahn-Hilliard [39] model, the (accurate) conservative diffuse-interface (Phase-Field) [109, 165] model, and their variants.

The finite-volume reconstruction is performed based on the primitive variables $\mathbf{V} = [\{(\alpha_p\rho_p\}_{p=1}^N, \mathbf{u}, P, \{\alpha_p\}_{p=1}^N]^T$ and is consistent between $(\alpha_p\rho_p)$ and α_p, so that kinematic, mechanic, and thermal equilibria at isolated material interfaces are maintained for any equations of state [94]. During the reconstruction, the *consistent limiter* [93] is implemented to guarantee $\sum_{p=1}^N \alpha_p = 1$ without producing fictitious phases. Regardless of which Phase-Field mechanism is selected, our approach preserves the mass positivity and volume fraction boundedness with appropriate CFL numbers for Eqs. (2.6) and (2.10), respectively, depending on the reconstruction order. The robustness of the approach can be further improved by considering the energy bound [94]. In addition to increasing the reconstruction order, the discontinuous Galerkin method has also been implemented to improve accuracy [266].

When AMR is used, Eqs. (2.6) and (2.10) are integrated over time with subcycling [104], and consistent synchronization [104] is enforced to preserve the consistency between $\hat{\mathbf{u}}^{HB}$ and $\hat{\mathbf{F}}^{HB}$ (for eliminating fictitious phases, local voids, or overfilling) as well as mass, momentum, and energy conservation. Eqs. (2.7) and (2.8) are critical for enabling a thin interface thickness on the order of the finest AMR grid size without producing fictitious interface oscillations [104].

Options for time integration are the first- to fourth-order Runge-Kutta (RK) schemes, with the third-order TVD RK scheme [85] most frequently used:

$$\begin{aligned}
\mathbf{U}^{(1)} &= \mathbf{U}^n + \Delta t\, \mathbf{R}(\mathbf{U}^n), \\
\mathbf{U}^{(2)} &= \tfrac{3}{4}\mathbf{U}^n + \tfrac{1}{4}\left(\mathbf{U}^{(1)} + \Delta t\, \mathbf{R}\left(\mathbf{U}^{(1)}\right)\right), \\
\mathbf{U}^{n+1} &= \tfrac{1}{3}\mathbf{U}^n + \tfrac{2}{3}\left(\mathbf{U}^{(2)} + \Delta t\, \mathbf{R}\left(\mathbf{U}^{(2)}\right)\right).
\end{aligned} \tag{2.11}$$

In AMR calculations, the entire RK step is taken on a single level before computing the step on the next level, as opposed to advancing a given stage on every level before moving to the next RK stage [104].

Our approach is modularized; the hyperbolic and Phase-Field steps are independent modules. As a result, the performance of the hyperbolic step is unaffected whether the Phase-Field step

is activated or not, which is desirable for users who do not wish to include the Phase-Field mechanism. Furthermore, the same time step (or CFL number) can be used even after including the Phase-Field mechanism, because the Phase-Field step takes into account the stability requirement associated with the Phase-Field contribution automatically via Eqs. (2.7) and (2.10). Specifically, if $\Delta t > \Delta\tau^{PF}$ for Eq. (2.7) or $\Delta t > \Delta\tau$ for Eq. (2.10), then multiple time advancements are performed in a corresponding fashion.

2.5 Examples

To verify and validate the proposed model and numerical approach with a series of test cases in two and three dimensions, and demonstrate the ability of `COMPAS` to simulate flows of interest, including gas-liquid flows and interfacial instabilities. These examples highlight the importance of using the Phase Field and the discussed limiting procedure when resolving compressible flows with multiple materials and interfaces to avoid numerical effects that lead to voids/overfilling, fictitious phases, and continuously diffusing interfaces.

2.5.1 Isentropic Vortex

To verify that the proposed solution method achieves the expected order of accuracy and that the Phase Field numerically satisfies reduction consistency, we consider the advection of a two-dimensional, inviscid isentropic vortex [277]. As this flow is smooth and involves no shocks or interfaces, the solution should converge at a second-order rate.

The flow is initialized as in [226] according to,

$$\begin{aligned} \rho_1 &= \left(1 + \delta T\right)^{A_1}, \\ u &= M_\infty \cos\theta + \delta u, \\ v &= M_\infty \sin\theta + \delta v, \\ P &= \frac{A_1}{A_1 + 1}\left(1 + \delta T\right)^{A_1+1}, \end{aligned} \tag{2.12}$$

where the perturbations are given by

$$\begin{aligned} \delta u &= -\frac{y}{R}\,\omega, \\ \delta v &= \frac{x}{R}\,\omega, \\ \delta T &= -\frac{1}{2A_1}\,\omega^2, \end{aligned} \tag{2.13}$$

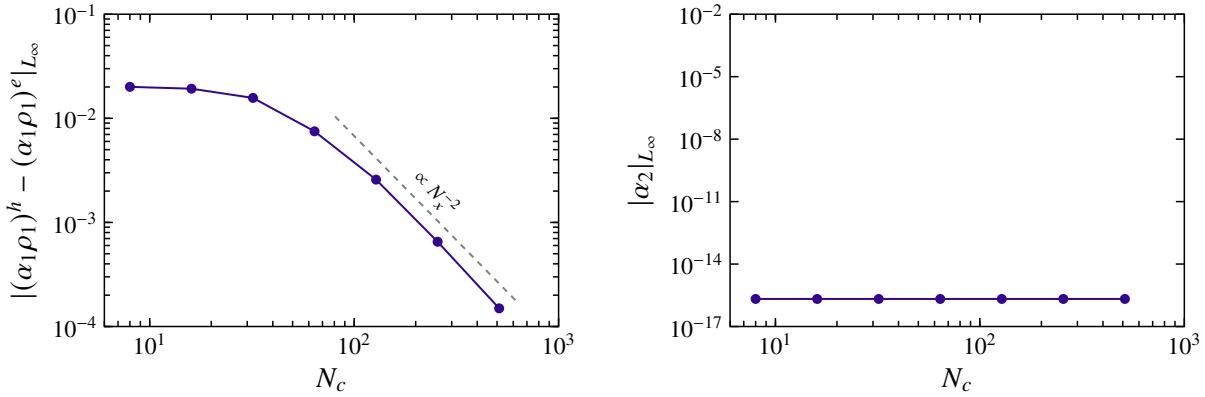


Figure 2.2: Density (left) and volume fraction (right) L_∞ errors for the isentropic vortex test case. Second-order accuracy is indicated by the dashed line in the density plot.

and the vortex profile is given by

$$f(x, y) = -\frac{1}{2\sigma^2}\left[\left(\frac{x}{R}\right)^2 + \left(\frac{y}{R}\right)^2\right], \tag{2.14}$$

where $\omega = \beta e^f$ is the perturbation function, $\beta = 1/5$ is the maximum perturbation amplitude, $\sigma = 1$ and $R = 1/2$ determine the vortex width, $A_1 = 5/2$, $M_\infty = 2\sqrt{2}$ is the translational Mach number, and $\theta = \pi/4$ is the translation angle. The vortex is initially located at $(0, 0)$ in a periodic domain $\{x, y\} \in [-2, 2]^2$, which is sufficiently large to avoid vortex-coupling effects from the periodic boundaries as the vortex translates over two cycles to its original position. While this flow involves only one material, and no interfaces are expected to be generated, a second material is evolved in addition to the principal fluid. Any value exceeding round-off in the auxiliary phase volume fraction indicates that either the model or method is not reduction-consistent, and that the auxiliary phase is being produced spuriously. The auxiliary material is given the same density as the primary fluid, and the volume fractions are given by $\alpha_1 = 1 - \alpha_2 = 1 - \epsilon$, where $\epsilon = 10^{-16}$.

Grids of $N_c \times N_c$ cells are used, with $N_c \in [8, 512]$. The HLLC Riemann solver [243] computes face fluxes based on a second-order MUSCL reconstruction with the MC limiter [249, 250].

The L_∞ norm of the primary phase density is given by,

$$|(\alpha_1\rho_1) - (\alpha_1\rho_1)_{\text{exact}}|_{L_\infty} = \max_{\Omega} |(\alpha_1\rho_1) - (\alpha_1\rho_1)_{\text{exact}}|. \tag{2.15}$$

To test reduction-consistency, we also track the auxiliary material via $|\alpha_2|_{L_\infty} = \max_\Omega |\alpha_2|$, which indicates inconsistency if it does not remain at round-off. Figure 2.2 shows the density error norm and the auxiliary volume fraction norm versus N_c on a log-log scale. Second-order convergence is observed as N_c increases, confirming the expected accuracy. The volume fraction error remains

at round-off level across all resolutions, verifying the consistency of reduction property in pure materials, even with the Phase Field active. If the calculation of the Phase-Field mechanism were not theoretically and numerically reduction-consistent, then the volume fraction error would grow and alter the dynamics of the primary fluid. This test case demonstrates that the present method and implementation are numerically reduction consistent in pure materials, i.e., the single-phase dynamics are recovered. Also, the Phase Field does not impact the nominal accuracy of the underlying scheme in pure material regions. This property will be demonstrated in subsequent test problems involving dynamic multiphase and multi-material interfaces interacting with shocks.

2.5.2 Spherical Bubble Collapse

The inertial collapse of a gas bubble in a liquid is a canonical compressible, multiphase phenomenon relevant to cavitation, biomedical ultrasound, and laser-generated bubbles [189]. An accurate numerical description, therefore, demands (i) correct acoustic coupling between gas and liquid, (ii) proper treatment of large density and pressure jumps, and (iii) the ability to follow an interface whose radius changes by more than an order of magnitude and whose volume, hence, changes by more than three orders of magnitude. The present test case demonstrates that the compressible Euler/Navier–Stokes solver coupled to the Phase-Field (PF) model recovers the classical Keller–Miksis (KM) theory [122] and verifies that the PF framework is numerically stable with adaptive mesh refinement.

A spherical air bubble of initial radius $R_0 = 0.5\,\mathrm{mm}$ is centered at the origin of a cubic domain $\{x, y, z\} \in [0, 80R_0]^3$. Symmetry conditions are applied at the origin, so only an eighth of the bubble is simulated, with outflow elsewhere. The fluid is initially at rest everywhere, with the pressure outside the bubble given by,

$$P(r) = P_L + \frac{R_0}{r}(P_B - P_L), \tag{2.16}$$

such that the far-field liquid pressure $P_L = \mathcal{P}_R P_B$ decays to the bubble pressure $P_B = 10^4\,\mathrm{Pa}$ at the bubble surface. The bubble is an ideal gas with $\gamma_B = 1.4$, $\rho_B = 1\,\mathrm{kg/m^3}$, and $P_B^\infty = 0\,\mathrm{Pa}$, and the liquid is a stiffened gas with $\gamma_L = 2.35$, $\rho_L = 1000\,\mathrm{kg/m^3}$, and $P_L^\infty = 10^9\,\mathrm{Pa}$. Pressure ratios of $\mathcal{P}_R \in \{10, 20, 50\}$ are run without viscosity, while the case with $\mathcal{P}_R = 50$ is run with $\mu_L = 100\mu_w$ to emphasize the damping effect, where μ_w is the viscosity of water. The expected behavior is a cycle in which the external high pressure drives the bubble spherically inward, followed by a rebound once the bubble pressure is sufficiently high.

The grid has $32 \times 32 \times 32$ cells with six levels of refinement, dynamically set based on the presence of a bubble, the distance from the origin, or a large pressure gradient. Adaptive meshing is useful for problems such as this, as a fixed grid with the same cell and domain size would require

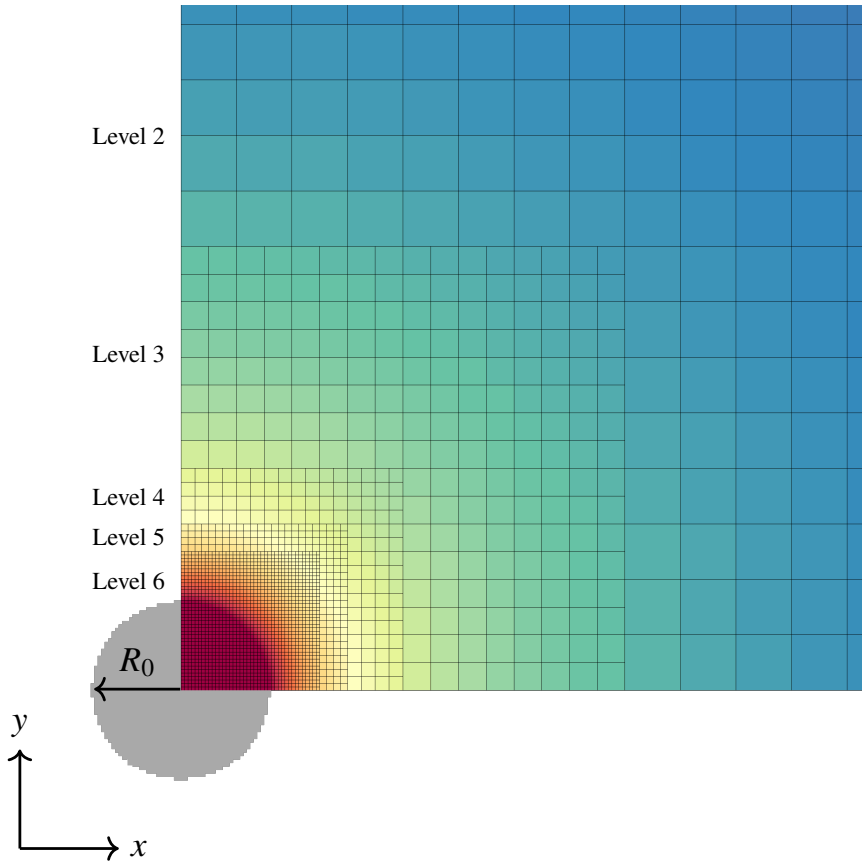


Figure 2.3: Initial grid for mesh levels two through the maximum level of six (zero-indexed) for the spherical collapse problem alongside the initial bubble radius R_0 and the pressure field. Reproduced with permission from [104].

~ 8.5 billion cells instead of the $\sim 125,000$ used here. Flexible adaptive meshing indicators allow users to automatically refine based on gradients or thresholds (as with the pressure and volume fraction in this case), or on a user-defined spatio-temporal indicator function.

Figure 2.3 shows the initial AMR grid, bubble radius, and pressure profile. Figure 2.4 shows the time history of the bubble radius for each case. Time is non-dimensionalized by the classical collapse time $t_c = 0.915 R_0 \sqrt{\rho_L / P_L}$, and radius by the initial radius. Each case is simulated until $t/t_c = 4$ to capture multiple collapses and rebounds, which represent significant volumetric changes in the bubble and interface motion over many cells. In all four cases, the radius trajectory matches well with the Keller-Miksis solution for both collapses and rebounds despite the modest resolution of ~ 25 cells per initial bubble radius, indicating that the proposed method and model capture bubble dynamics well. The bubble shape at the end of the first rebound and at the second collapse is also visualized for each case; the interface remains consistently thin throughout the dramatic volumetric changes, and the bubble shape remains spherical, as shown in [104]. Compared to

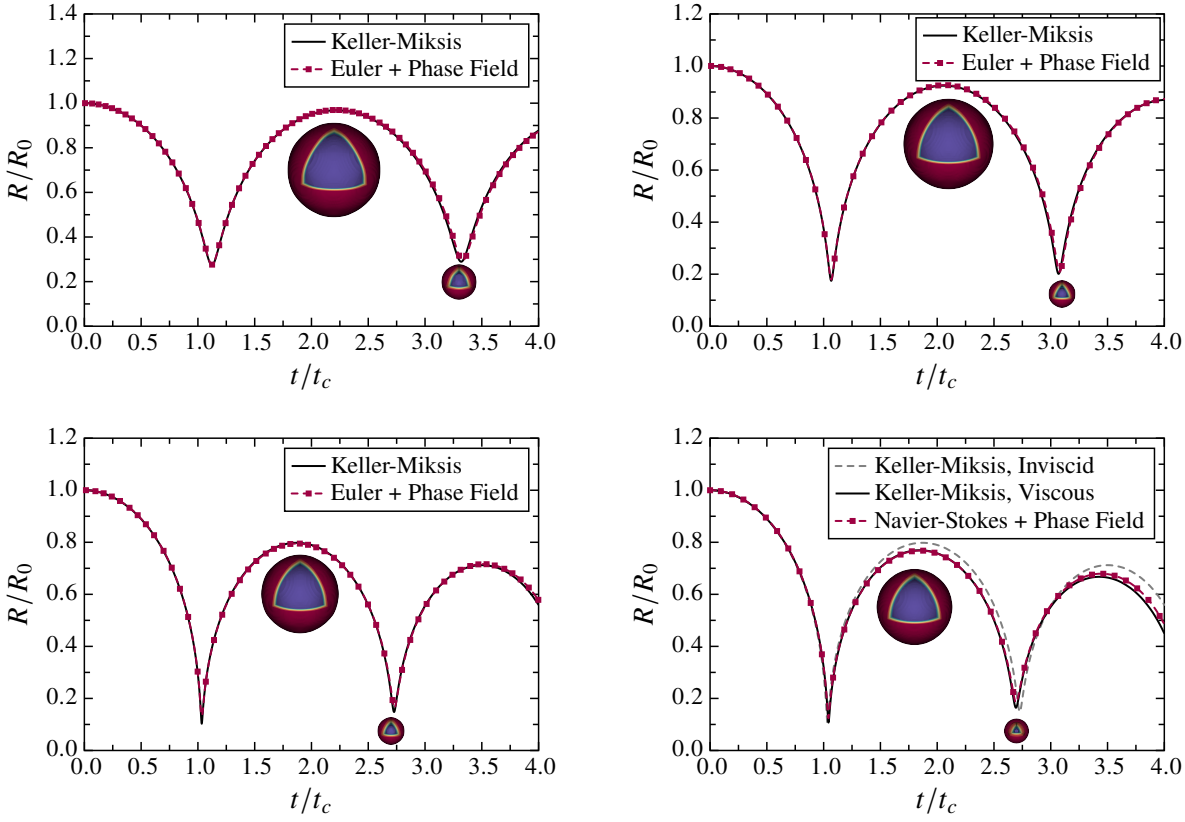


Figure 2.4: Time evolution of the radius for the inertial collapse of a spherical air bubble in water, as compared to the Keller-Miksis solution [122]. Driving pressure ratios of 10 (top left), 20 (top right), 50 (bottom left), and 50 with viscous terms (bottom right). Adapted with permission from [104].

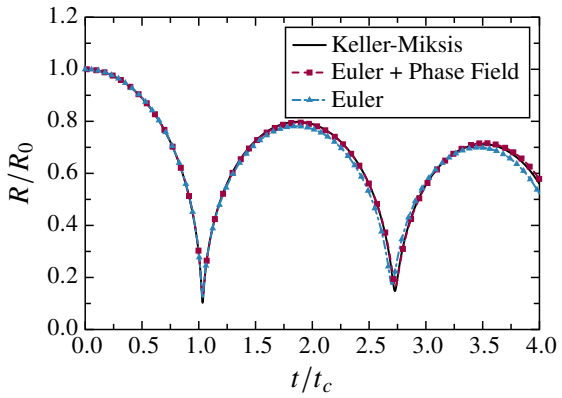


Figure 2.5: Time evolution of the radius for the inertial collapse of a spherical air bubble in water, as compared to the Keller Miksis solution [122] with and without the Phase Field active for the inviscid $P_R = 50$ case.

simulations without Phase Fields active, the Phase-Field approach more closely matches Keller-Miksis, particularly during the second collapse and rebound. Figure 2.5 shows the radius versus time with and without Phase Fields active for the inviscid $P_R = 50$ case. These results demonstrate the suitability of the proposed method for bubble collapse problems and, more generally, for gas-liquid flows with large gradients in density, pressure, and compressibility. The inclusion of viscous effects also shows that the Phase Field does not negatively affect Navier-Stokes calculations.

2.5.3 Three Material Richtmyer-Meshkov Instability

Next, we simulate the interaction of a shock wave with two successive material interfaces, based on the experimental work of [141], to demonstrate the ability of the proposed model and method to accurately capture the shock, interface, and vorticity dynamics that govern the motion of the interface and instability growth while maintaining a thin interface. Additionally, this case is used as another point of verification of the consistent limiter, ensuring the consistency of reduction property without introducing any voids/overfilling. Finally, this case is used for validation by comparing the interface positions and instability amplitudes for flat and perturbed interfaces with [141].

A rectangular domain with $x \in [-20, 130]$ mm, and $y \in [0, 30]$ mm captures a half wavelength, with symmetry boundary conditions on the upper and lower boundaries and outflow elsewhere. Material 1 is SF_6, material 3 is air, and material 2 is a mixture of air and SF_6 with 0.67 mass-fraction SF_6 placed between materials 1 and 3. The equation of state parameters for each material are $\{\gamma_p\}_{p=1}^3 = \{1.1074, 1.2040, 1.400\}$ and the densities are $\{\rho_p\}_{p=1}^3 = \{6.14, 4.48, 1.79\}$. The interface between materials p and q is denoted $I_{p,q}$, and has mean position $x_{I_{p,q}}$. The initial distance between interfaces $I_{1,2}$ and $I_{2,3}$ is 10 mm. The shock is initialized with Mach number 1.24. Initially,

flat and sinusoidally perturbed interfaces with wavelength $\lambda = 60\,\text{mm}$ are considered.

A base grid of 80×16 cells is used, with five levels of refinement, yielding an effective resolution of 2560×512. Refinement is based on the mixture density gradient, capturing both shocks and material interfaces. The interfaces positions are tracked over time by interpolating α_2 where $\alpha_2^L \leq 0.5 \leq \alpha_2^R$ and $\alpha_2^L \geq 0.5 \geq \alpha_2^R$. The perturbation amplitude for interface $a_{p,q}(t)$ is computed as with the interface positions, but sampled at the crest and trough of the perturbation. The interface thickness is computed as a domain-averaged number of cells across the interface, given by,

$$(N_I)_p = \frac{1}{\Delta x}\frac{\int_\Omega \mathcal{H}(\alpha_p)d\Omega}{\int_\Omega |\nabla\alpha_p|d\Omega}, \tag{2.17}$$

where the indicator is given by,

$$\mathcal{H}(\alpha_p) = \begin{cases} 1, & (\alpha_p)_{\min} \leq \alpha_p \leq (\alpha_p)_{\max} \\ 0, & \text{else.} \end{cases} \tag{2.18}$$

We choose $(\alpha_p)_{\min} = 0.05$ and $(\alpha_p)_{\max} = 0.95$; results are largely insensitive to this choice.

In the flat case, there is no instability growth, and the interfaces uniformly translate under the shock and rarefaction dynamics. Specifically, the shock impinges on the first interface, where a rarefaction propagates away from the interface and a shock travels to the second interface, generating another shock and rarefaction. The wave reverberation between the two interfaces leads to nearly identical translational velocities for both interfaces after the initial transience. Figure 2.6 (left) shows the change in interface positions over time ($\Delta x_{I_{p,q}} = x_{I_{p,q}}(t=0) - x_{I_{p,q}}(t)$) from simulations overlaid with the experimental results of [141]. The first interface is accelerated by the shock, after which a transmitted shock accelerates the second interface. This behavior is well captured by the simulations, and good agreement with experiments is observed.

With a sinusoidal perturbation amplitude of 2 mm along each interface, the flow is unstable to the Richtmyer-Meshkov instability. Figure 2.6 (right) shows the perturbation amplitudes over time from simulations and experiments. As the shock impinges upon the first interface, it compresses the interface and deposits baroclinic vorticity. The transmitted shock propagates to the second interface and does the same. The rarefaction reflected from the second interface returns to the first interface and nearly halts the perturbation growth. The second interface, however, undergoes a phase inversion and continues to grow over time. Good agreement is observed between the present simulations and experiments, indicating that the shock-interface interactions, vorticity dynamics, and interfacial instabilities in this case are captured well, including when the Phase Field is active.

Figure 2.7 shows N_I for material 2 over time with and without the Phase Field. Without the Phase Field, the interface thickness grows precipitously throughout the simulation, potentially

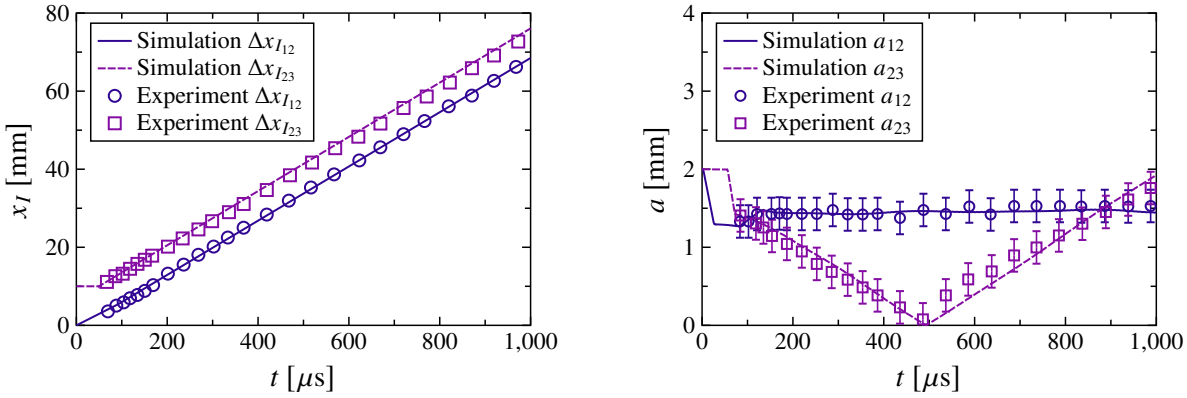


Figure 2.6: Time evolution of the interface positions (left) and perturbation amplitude (right) for the three-material Richtmyer-Meshkov instability with a flat (left) and perturbed interface (right). Experimental results are from [141].

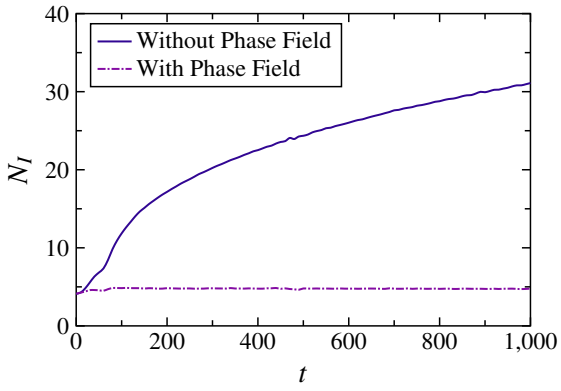


Figure 2.7: Interface thickness over time with and without Phase Fields for the three-material Richtmyer-Meshkov instability.

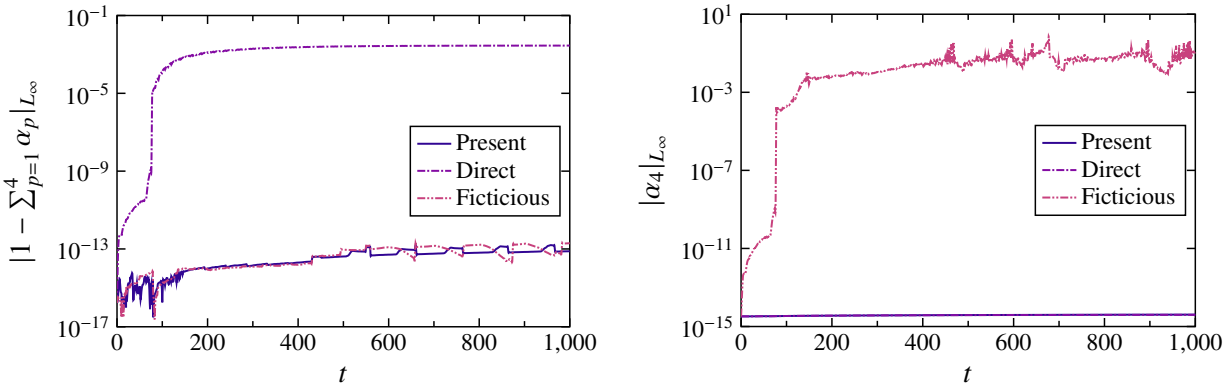


Figure 2.8: L_∞ errors in the volume fraction summation (left) and auxiliary material (right) over time for the three-material Richtmyer-Meshkov instability.

deleteriously impacting wave dynamics at interfaces due to variable mixture sound speed in the diffuse interface region, potentially delaying wave propagation [212]. With the Phase Field, however, the interface remains a consistent, uniform thickness over time despite the migration of interfaces and the growth of instability.

We repeat the simulation above with a fourth auxiliary material that is not initially present in the flow with properties of the third material, along with two algorithmic changes to demonstrate the importance of reduction consistency in N-material problems, as provided by the proposed method [93]. Firstly, we consider the "Direct" method, which does not include the consistent limiter and is thus expected to produce voids/overfilling. Secondly, we consider the "Fictitious" method, in which $\{\alpha_p\}_{p=1}^{3}$ are reconstructed using the standard approach and α_4 is computed as $1-\sum_{p=1}^{3}\alpha_p$; this approach is thus expected to produce no summation errors but spurious amounts of the auxiliary material α_4 instead. Two levels of refinement are used to provide sufficient numerical diffusion to avoid the "Direct" and "Fictitious" approaches producing unrealizable states. Figure 2.8 (left) shows the summation error for all three methods. The present method and the "Fictitious" method both prevent the production of voids or overfilled cells, but the "Direct" method exhibits summation errors that grow over ten orders of magnitude. Such errors violate the constraints on the original system and the Phase Field model, and can lead to out-of-bounds states computed by the equation of state. Figure 2.8 (right) shows the L_∞ norm of the initially absent material α_4. In this case, both the present and "Direct" methods prevent the production of the fictitious material, while the "Fictitious" method produces significant quantities of α_4, despite no initial presence in the flow or physics that would lead to its production. These discrepancies between the "Direct" and "Fictitious" methods are problematic and indicate that consistency of reduction is violated; one cannot guarantee that a pure material remains pure. Only the present method avoids both types of

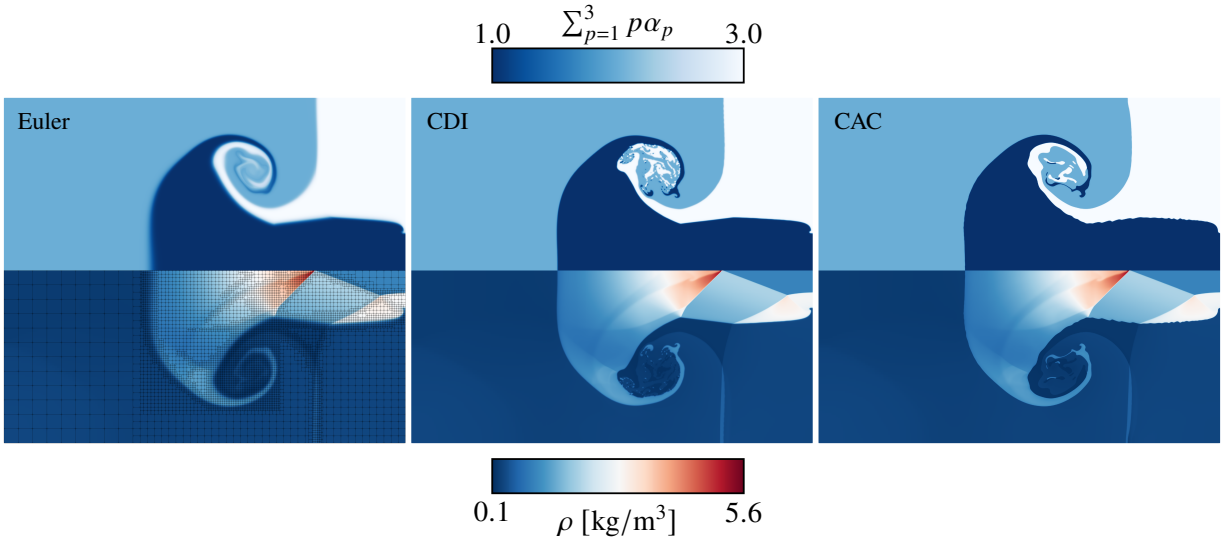


Figure 2.9: Density and phase indicator contours for the triple-point problem at $t = 5$ without interface regularization (left), with the conservative diffuse interface (CDI) method (middle), and with the conservative Allen-Cahn (CAC) Phase-Field mechanism (right). The AMR blocks are overlaid on the density for the result without Phase Fields (Euler).

errors [93], emphasizing its usefulness for accurately simulating N-material systems.

2.5.4 Triple-Point Problem

To demonstrate the effect of the choice of different Phase Field mechanisms, we simulate the triple-point problem [178, 182, 283] as in [182], with no Phase Fields active (Euler), with the conservative Allen-Cahn mechanism (CAC), and with the conservative diffuse interface mechanism (CDI).

In this problem, a high-pressure gas is placed next to two low-pressure gases with different densities and equation-of-state properties. A shock, contact, and rarefaction form at the pressure discontinuity, and the interfaces roll up at the triple point. The vortex that forms leads to thin filaments of each material, making this problem an illustrative example of the behavior of different interface models. Three gases are simulated with equation of state parameters $\{\gamma_p\}_{p=1}^3 = \{1.4, 1.5, 1.5\}$ and densities $\{\rho_p\}_{p=1}^3 = \{1.0, 1.0, 0.125\}$. The initial pressure is set according to $P = \sum_{p=1}^3 \alpha_p P_p$ where $\{P_p\}_{p=1}^3 = \{0.1, 1.0, 0.1\}$, and the volume fractions are set via the inversion of a level set.

The domain $\Omega = [0, 7] \times [0, 3]$ is discretized with 224×96 cells with three levels of refinement based on the density gradient and the gradient of the phase indicator $\sum_{p=1}^3 p\alpha_p$. Reflecting boundary conditions are applied on all boundaries.

Figure 2.9 shows the density and phase indicator contours for the Euler, CDI, and CAC simulations at $t = 5$. The AMR tracks the shocks, contacts, and interfaces in the flow well. The solutions from all three methods qualitatively show similar results for the shock and contact locations. The

behavior of the interfaces, however, differs greatly across the three methods: without any interface modeling, the filament formed at the triple point is continuously diffused as the vortex grows, leading to significant numerical mixing among the three materials. The CDI mechanism causes the filament to break into small droplets rather than allowing it to diffuse once the interface cannot be resolved. The CAC mechanism produces some large "droplets" from the filament's breakup, but small droplets are not persistent and instead conservatively diffuse once under-resolved. These results elucidate the significant differences in flow features produced by different interface modeling choices, particularly when interfacial structures are too small to be resolved on the mesh. The most useful or accurate model is likely problem-dependent. `COMPAS` includes several popular Phase Field models to facilitate direct comparisons, enabling users to make an informed choice when simulating a problem.

2.5.5 Shock-Driven Multi-Material Implosion

We simulate a three-dimensional implosion with materials matching those in [104]. This problem is a gas analogue that mimics the implosion of an inertial confinement capsule and is used to highlight the capability of `COMPAS` to capture three-dimensional, $(N > 2)$-material interfacial instabilities.

A spherically converging shock impinges upon an interface between air and SF_6, which transmits a shock to a perturbed interface between SF_6 and helium. The transmitted shock deposits baroclinic vorticity on the perturbed interface, leading to its roll-up due to the RMI. The shock transmitted into the helium converges and ultimately rebounds at the core, re-shocking the RMI and leading to high-aspect-ratio bubbles and spikes. A sphere of helium is surrounded by a concentric layer of SF_6, itself surrounded by air. The interface between the air and SF_6 is spherical, while the SF_6-helium interface is perturbed and given by,

$$\frac{r}{r_{\text{He}}} = 1 - a_0 \cos\left[\frac{2\pi}{\lambda}\arccos\left(\frac{x}{r}\right)\right]\cos\left[\frac{2\pi}{\lambda}\arccos\left(\frac{y}{r}\right)\right]\cos\left[\frac{2\pi}{\lambda}\arccos\left(\frac{z}{r}\right)\right], \tag{2.19}$$

where $a_0 = 0.05$, $r_{\text{He}} = 0.85$, and $\lambda = \pi/4$. A Mach 1.2 shock is initialized in the air at a radius of $r_{\text{shock}} = 1.05$, and the distance between the SF_6-helium and air-SF_6 interfaces is $\Delta r = 0.15$. The specific heats ratio and density of the gases are $\{\gamma_p\}_{p=1}^3 = \{1.276, 1.093, 1.667\}$, and the densities are $\{\rho_p\}_{p=1}^3 = \{1.351, 5.494, 0.214\}$. Figure 2.10 shows the initial condition schematically, as well as the initial density field and AMR grid.

Four total levels of adaptive grids are used, with refinement based on the gradient of the phase indicator $\sum_{p=1}^3 p\alpha_p$. Symmetry boundaries are used to simulate one-eighth of the domain, with outflow conditions elsewhere.

Figure 2.11 shows contours of numerical Schlieren, air-SF_6 and SF_6-helium interfaces, pressure

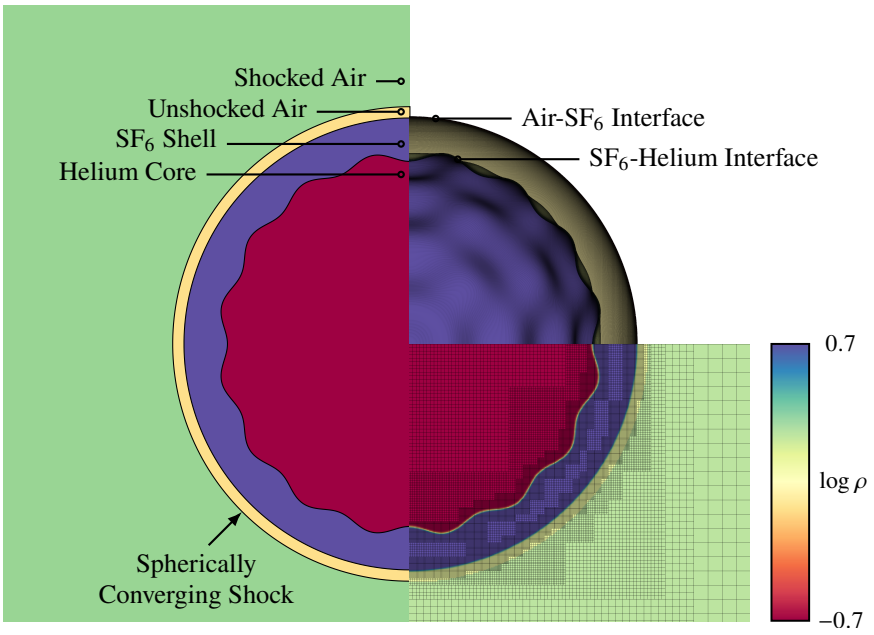


Figure 2.10: Initial conditions for the shock-driven multi-material implosion problem. The density field, interfaces, and AMR grid are shown on the x-y plane with labels. Reproduced with permission from [104].

field, and density overlaid with the grid at four time instants. The shock is transmitted into the SF_6 at the air-SF_6 interface and propagates to the helium. As expected, the perturbation in the SF_6-helium interface rolls up due to the RMI, such that the SF_6 penetrates deeply into the core of the capsule. This case demonstrates that the N-material Phase-Field model is applicable to three-dimensional flows with shocks and can maintain a consistently thin interface throughout complex interfacial instability dynamics between more than two gases. Additionally, this case shows that the AMR effectively tracks the complex three-dimensional interfaces present in the flow.

2.5.6 Taylor-Green Vortex Scaling

Finally, we consider a single-phase Taylor-Green vortex to measure the parallel performance of `COMPAS` for both CPU-only and CPU-GPU calculations.

The domain is discretized with a 1024^3-cell grid, and no mesh refinement or interface regularization is used. Simulations are run on the OLCF *Frontier* machine with both CPU-only and CPU-GPU results presented. Each *Frontier* node contains 64 CPU cores (56 in practice) via one AMD "Optimized 3rd Gen EPYC" CPU (with two hardware threads per physical core) and four AMD MI250X accelerators, each with two Graphics Compute Dies, for essentially eight GPUs per node. The CPU-only calculations use the Message Passing Interface (MPI) and OpenMP libraries to parallelize loops over boxes and cells, while the CPU-GPU calculations also use HIP.

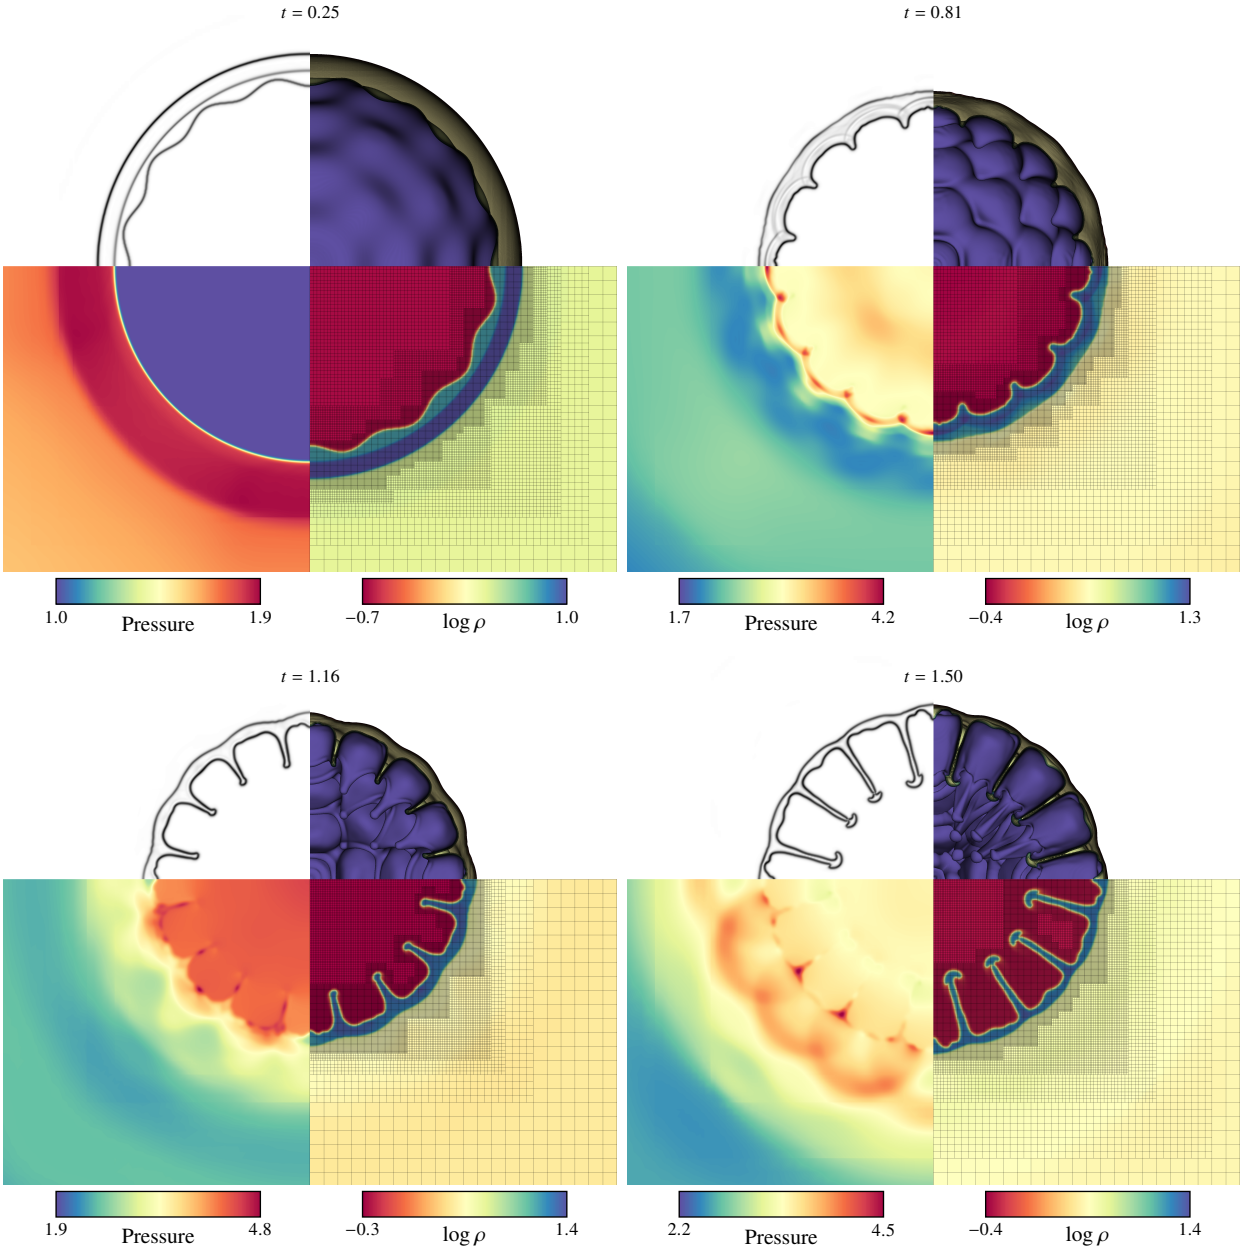


Figure 2.11: Numerical Schlieren, interfaces, pressure, and density for the multi-material implosion at four times. Reproduced with permission from [104].

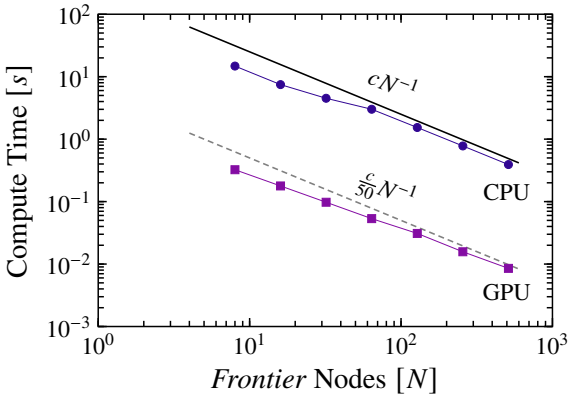


Figure 2.12: Strong scaling results for the Taylor-Green vortex with 1024^3 cells on OLCF *Frontier*.

Figure 2.12 shows strong-scaling results on up to 512 nodes (28, 672 CPUs, 4096 GPUs). Both the CPU and GPU scaling results show at worst 60% efficiency between 8 and 512 nodes (with better performance once the nodes are saturated), demonstrating that our software layer works well with the AMReX parallelization paradigm, and that efficient, large-scale, GPU-enabled calculations are possible with COMPAS.

2.6 Conclusions

We present COMPAS, an open-source, GPU-accelerated, adaptive-mesh-refinement framework enabling robust, accurate, and scalable simulations of compressible, N-material, interfacial flows with consistent, conservative Phase Fields. Through integration with the AMReX-based code base, COMPAS provides a unique framework for addressing the challenges of excessive and non-uniform numerical diffusion at material interfaces, conservation, and reduction consistency in multiphase flow computations.

Through several verification and validation cases, we demonstrate that the proposed approach achieves design-order accuracy, is conservative and reduction consistent, and maintains consistently thin interfaces even in the presence of strong shocks, large density and pressure ratios, and dynamic flow instabilities. Across several tests, including isentropic vortex advection, spherical bubble collapse, multi-material Richtmyer-Meshkov instabilities, and shock-driven capsule implosions, COMPAS consistently produces accurate results. Simulations across flow regimes confirm quantitative agreement with theory, experimental data, and established benchmarks, highlighting the applicability of the proposed method and model.

The consistent and conservative Phase-Field framework implemented in COMPAS keeps interfaces uniform and thin and enables stable, accurate, and reduction-consistent simulations of N-

material systems with various Phase-Field models, including the conservative Allen-Cahn, Cahn-Hilliard, and (accurate) conservative diffuse-interface models. Voids, overfilling, and fictitious phases are all prevented to round-off levels by the consistent reconstruction and limiting procedure, and the convex equation of state subverts the calculation of undefined quantities in pure material regions. The use of patch-based AMR enables logical indexing in the solver kernels, eliminating the need for substantial connectivity information. The modular C++ implementation separates user code from the source, simplifies running cases and updating the code, and makes algorithmic and model changes straightforward. Implementation with AMReX provides state-of-the-art parallelization, AMR procedures, and portability across HPC systems with both CPU and GPU compute nodes.

Although the presented model, method, and implementation mark a novel capability for simulating compressible interfacial flows, certain areas remain open to further improvement. As with all Eulerian interface techniques, under-resolved features such as films and ligaments must be handled with care. The Phase Field tends to either break these features into small "droplets" or diffuse them into the nearest like material interface, depending on the Phase-Field mechanism. Such behavior is a modeling choice whose implications are not well understood and warrant further study. Additionally, coupling the Phase Field to detailed physics, such as phase change and molecular mixing, is an open area of development that is crucial across a number of fields.

2.6.1 Choosing a Code

As discussed in Section 2.2, there is a wide variety of tools available for simulations of compressible interfacial flows, with the features of each shaped by a combination of scientific requirements, computational resources, and software philosophies. For each of these codes, varying modeling priorities and historical choices have led to a range of distinct implementations that researchers must wade through. As such, we briefly discuss some of the most popular tools in the field, their unique features, and how a user should select one for their application. Table 2.1 summarizes several of these platforms, providing context for the range of approaches encountered in the field.

Clearly, there is significant variability in methods, models, software implementations, and intended applications. The choice of code is typically determined by the types of physical processes prioritized for simulation. Some platforms, like FLASH, are optimized for large-scale multiphysics problems, while others, such as Basilisk and OpenFOAM, offer extensive flexibility and user customization, catering to community-driven development and a broad spectrum of fluid mechanics applications. Specialized codes, including MFC and ECOGEN, focus on high-order accuracy or particular multiphase models, sometimes at the cost of generality in multiphysics coupling or adaptive meshing. Similarly, `COMPAS` sacrifices some specialization in models for the flexibility

Table 2.1: Comparison of several open-source and laboratory multiphase simulation codes.

Code	Type	Interface treatment	Physics scope	Parallelism	Strengths	Limitations	Typical use
FLASH [79]	Eulerian finite-volume with block-structured AMR	Diffuse interface	Compressible hydrodynamics, MHD, radiation, gravity, EOS	MPI	Very mature multiphysics framework; strong community	Limited interface treatments; steep learning curve	Astrophysics, high-energy-density physics (shocks, supernovae)
MFC [36]	Eulerian finite volume, high-order (WENO)	Diffuse interface	Compressible multiphase, bubbly flows, hypo- and hyper-elastic solids	MPI + GPU	High-order accuracy; efficient for compressible multiphase	No AMR, interface regularization	Cavitation, bubble and droplet dynamics, fluid-solid interactions
ECOGEN [215]	Eulerian finite volume with AMR	Diffuse interface	Compressible multiphase, cavitation, evaporation/condensation	MPI	Modular diffuse-interface with thermodynamic consistency; supports unstructured meshes	Less developed for incompressible or multi-physics	Cavitation, boiling, compressible multiphase flows
Basilisk [190]	Eulerian finite volume with tree-AMR	VOF sharp interfaces	Incompressible or weakly compressible two-phase flows, capillary phenomena	MPI	Sharp interfaces; efficient adaptive mesh	Less developed for compressible flows or multi-physics	Droplet dynamics, vorticity-driven flows, layered flows
OpenFOAM [263]	Eulerian finite volume with AMR	VOF, Phase Fields	Incompressible/ compressible CFD, RANS/LES, reacting flows, multiphase	MPI	Very broad solver library; flexible customization	Complex code base, less optimized	Industrial multiphase flows, sprays, reacting/thermal flows
COMPAS (In preparation)	Eulerian finite-volume with AMR (AMReX)	Diffuse interfaces with reduction-consistent Phase Fields	Compressible hydrodynamics, multi-physics extensions	MPI + GPU	Efficient, modular, and portable GPU+AMR code; several Phase-Field models	Less mature multi-physics support; model choice is problem-dependent	Shock-interface physics, interfacial instabilities, vortex dynamics, bubbles, and droplets

of its Phase Field implementation, AMR strategy, and parallelization.

Generally, the selection of a code should not begin with the search for maximal feature lists, but with a clear identification of the systems of interest and workflows that must be supported. Considerations of computational scale are important, as some codes can leverage modern heterogeneous computing systems. Attention is often given to open-source accessibility and community support, though proprietary or export-controlled codes may remain necessary for selected government or industrial projects involving sensitive physics. From a development standpoint, attempting to make a single code solve every possible system is a doomed endeavor; instead, users should either pick a code that directly incorporates the necessary physics or implement their specific needs. In `COMPAS`, users have adopted this philosophy and are actively using or developing physics such as radiation hydrodynamics and coupled hyper-elastic solids for specific use cases.

A review of Table 2.1 shows that each simulation code occupies a distinct regime defined by its interface-handling strategies, parallelization approach, and physical modeling choices. For example, codes such as Basilisk and OpenFOAM are widely recognized for their sharp interface capturing via volume-of-fluid (VOF) methods and flexible scripting frameworks, making them well-suited for incompressible flows and problems where sharp interfaces and modularity are crucial. Conversely, FLASH leverages diffuse-interface techniques within robust multiphysics platforms, supporting adaptive mesh refinement and modules for gravity, magnetohydrodynamics, and radiation, making it well-suited to a variety of HEDP systems. MFC and ECOGEN excel at simulating models in the Baer-Nunziato hierarchy, with exascale support and AMR, respectively. In contrast to interface-tracking approaches, Phase-Field models consistently and conservatively handle spurious interfacial diffusion, scale well on HPC systems, and excel in problems with shocks and dynamic topological changes. The various Phase-Field models implemented in `COMPAS` offer additional choices for a user but also enable further customization for a given application.

The `COMPAS` was originally designed to study cavitation, high-speed droplet dynamics, and interfacial instabilities, and as such was built on top of the AMReX platform to simulate multi-material flows with robust numerical methods. Inaccuracies arising from spurious interface diffusion motivated the development and implementation of the consistent and conservative Phase-Field method, and questions around the influence of various interface treatments led to the incorporation of other Phase-Field mechanisms. User applications motivated additional physics, an active area of development, but limitations surrounding more intricate phenomena, such as laser-plasma interactions and phase change, remain.

Ultimately, the best code for a user is one that i) accurately and efficiently simulates the physics of interest, ii) enables straightforward use and modification when needed, iii) is verified, validated, and documented, and iv) has an active user or developer community. With each of these satisfied, a code is ready to be used to make an impact. As discussed, `COMPAS` does not solve every problem

in the field, but it does represent a tool that meets all the above requirements and is thus a useful tool for scientific exploration of methods, models, and physics. Table 2.1 provides a practical snapshot of this landscape and may serve as a starting point for someone seeking to conduct a physics study or implement a model/method, but selection is best informed by specific problem demands and sustained attention to ongoing developments in the field.

CHAPTER 3

A High-Order Discontinuous Galerkin Method for Compressible Interfacial Flows with Consistent and Conservative Phase Fields

The material in this chapter is adapted with minor corrections from W. J. White, Z. Huang, and E. Johnsen, A high-order discontinuous Galerkin method for compressible interfacial flows with consistent and conservative Phase Fields. *Journal of Computational Physics* 527:113830, 2025.

3.1 Abstract

Excessive and non-uniform numerical diffusion poses challenges for accurate simulations of compressible interfacial flows with shocks using diffuse interface methods. Even with high-order accurate methods, material interfaces continually diffuse, thus making material regions ambiguous and deleteriously impacting wave propagation across interfaces. Simulations with low-order methods are particularly affected by this issue, requiring a significant number of degrees of freedom to resolve interfaces. At the same time, Phase-Field methods are able to maintain interfaces with constant and uniform thickness over time. However, while implemented into second-order accurate solvers, a general fully conservative, bounds-preserving, and high-order accurate treatment for high-speed flows is lacking. To address these issues, we adapt the general consistent and conservative Phase-Field model and the associated theoretical and numerical analysis to high-order accurate discontinuous Galerkin schemes with non-conformal subcell finite volume discontinuity capturing. In the proposed approach, discontinuities are detected *a priori* by physical sensors and captured by robust finite volume schemes on uniform subcell grids. By computing the Phase-Field mechanism using subcell data and reconstructing high-order representations in each cell, the calculation of the Phase-Field fluxes is robust, efficient, and can be directly incorporated into the high-order discontinuous Galerkin framework. Further, because of the linearity of the chosen projection and reconstruction operators, the properties of the general Phase-Field scheme are retained in the

high-order framework, including consistency of reduction, Galilean invariance, maintenance of appropriate interfacial conditions, and conservation. We verify the proposed scheme with several test cases that demonstrate its accuracy in smooth regions, the numerically exact reduction consistency of the solver, its ability to simultaneously represent smooth and sharp features accurately, the advantage of the discontinuous Galerkin discretization with non-conformal subcell finite volume over conventional finite volume schemes and high-order limiters, and the capability of the scheme to robustly simulate flows with large density, pressure, velocity, and material property gradients. Validation against experimental data is also included.

3.2 Introduction

A number of natural and engineering systems involve compressible flows of multiple substances or phases with dynamic interfaces, including astrophysical phenomena [1, 259], bubbly and cavitating flows [24, 114, 264], high-speed droplet dynamics [191], hydrodynamic instabilities [52, 256, 258], and inertial confinement fusion experiments [56]. The numerical representation of material interfaces in simulations of such problems is a key factor dictating the accuracy of solutions and the performance of schemes. For multiphase and multi-component flows under the assumption of immiscibility, sharp interfaces can be maintained in practice by explicitly tracking interfaces, either via the mesh or an indicator [16, 20, 71, 80, 92, 175, 219, 230, 245, 248], such that the interface has a zero thickness. However, conservation, grid entanglement, and load balancing can be problematic with these schemes. Diffuse-interface methods, on the other hand, capture interfaces over multiple cells in a manner similar to shocks [211]. These methods can be implemented in a conservative fashion, satisfy the second law of thermodynamics [2, 19, 121, 184], and lead to favorable load balancing as the interface is not treated differently from other regions of the domain. However, interfaces may thicken indefinitely without an additional model due to numerical diffusion and the lack of physical compression mechanisms. Additionally, this numerical diffusion can be non-uniform, depending on velocity gradients and the directionality of Riemann solvers and the mesh. This phenomenon is problematic for numerical studies, particularly those that run over many characteristic times, as the distribution of materials can become ambiguous, so that the application of interfacial processes such as surface tension and phase change and the propagation of waves through the interface may become inaccurate [211]. Approaches like interface compression [222] and sharpening (e.g., THINC) [223] have been proposed to counteract numerical diffusion of the interface, but can have issues with robustness and consistency.

An approach bridging the modeling gap between sharp and diffuse interfaces is the Phase-Field method, with which the interface is captured over grid points as in a diffuse interface framework, but whose width and uniformity are controlled throughout the domain and

over time [93]. This framework has been successfully applied to incompressible multiphase flows [95, 96, 96, 97, 98, 99, 100, 101, 102] and has recently been extended to compressible multiphase flows [93, 94]. In this framework, the physical model incorporates the Phase-Field mechanism in a consistent and conservative manner such that the exact solution has a finite thickness interface with a known equilibrium distribution, facilitated by conservative fluxes introducing both compression and diffusion where appropriate. Consistency implies Galilean invariance, maintenance of appropriate interfacial conditions (i.e., constant pressure, velocity, and temperature across isolated interfaces), consistency of reduction (i.e., the dynamics of a flow with $N-1$ materials are recovered in regions of a simulation of N materials where only $N-1$ materials are present), and recovery of the incompressible model of [100] in the incompressible limit [93]. Conservation refers to ensuring that the Phase Field does not affect mass, momentum, and energy conservation, and does not affect the total volume fraction. Additionally, this Phase-Field formulation also extends to N-phases/materials [94]. Other works have considered the Phase-Field mechanism for compressible two-phase flows, but are limited in their models and discretizations such that some of the aforementioned conditions are not satisfied [222, 238]. In [109], consistency, consistency of reduction, and conservation were satisfied simultaneously, however, the model relied on reference values of the material densities that make that model less general than that in [93], and did not include a term in the volume fraction transport equation that is known to be important for inertial bubble collapse problems [214]. A second-order central difference scheme, suitable for low-speed flows without shocks, was used in that work, although an LAD [54] approach could be applied. The model in [110] included the missing term in the volume fraction equation and avoided use of reference values in the regularization terms.

The theoretical and numerical analysis in [93] was applied to a second-order (MUSCL) finite volume solver [250]. A high-order finite-difference scheme for a Phase-Field model was introduced in [110]; however, discrete conservation was only possible for a second-order method. Additionally, a discontinuous Galerkin scheme was recently applied to a Phase-Field model [241], but it is not discretely conservative, nor well-suited to high-speed flows or inertial bubble collapse problems. In this work, we seek to extend the general Phase-Field framework of [93] to high-order accuracy via the discontinuous Galerkin method in a fully conservative manner for high-speed flows. The DG method, as introduced by Reed and Hill to model neutron transport [199], combines the discontinuous solution representation of finite volume schemes with the high-order accuracy of (continuous) finite element methods. Reconstructions of cell face data are not needed for high-order surface flux integrals as in finite volume schemes due to the local functional representation of the solution in each cell, and the discontinuous nature of the data on cell faces decouples the degrees of freedom in different elements so spatial derivatives can be computed using a compact stencil of neighbor data without the global step of continuous finite element schemes. Where the solution

is smooth, DG methods exhibit super-convergence in the cell-average error, which makes them capable of far outperforming finite volume schemes, i.e., requiring far fewer degrees of freedom to achieve a particular level of error [8, 9, 205]. Further, high-order DG schemes are readily applied on unstructured meshes without special treatments like those needed for high-order unstructured finite difference and finite volume schemes, making them well-suited for complex geometries. The high-order polynomial representation of the solution within a cell can pose issues for DG schemes with discontinuity capturing; however, if a cell contains a discontinuity, the Gibbs phenomenon may cause the solution to be physically inadmissible within the cell, potentially leading to failure. Several approaches have been developed to combat this issue [49, 66, 91, 105, 126, 158, 185, 207, 275]; however, all discontinuity-capturing schemes are at most first-order accurate at discontinuities [137]. Effective techniques for regularizing discontinuities are robust, keep high-order information, require few tuning parameters, and are highly local, but achieving all of these properties simultaneously is challenging.

The two objectives of the present work are interrelated: to incorporate the consistent and conservative Phase-Field formulation into the DG framework and to extend the Phase-Field formulation to high-order accuracy. An isoparametric subcell finite volume scheme similar to that in [112] is used to capture discontinuities robustly in regions identified by physical sensors, so that the solution is high-order accurate elsewhere. With this method and the present Phase-Field model, interfaces and shocks/contacts are consistently captured over a width typically less than one DG cell. The subcell solution representation also increases the efficiency and robustness of the Phase-Field calculation, with special attention to the evaluation of the Phase-Field flux vector in a DG framework. The proposed scheme represents significant improvements both in accuracy over finite volume schemes in smooth regions and DG schemes for multiphase flows without Phase-Fields at interfaces. As compared to previous works, the current Phase-Field method is uniquely able to satisfy key properties such as conservation, consistency, boundedness, and consistency of reduction in a high-order discontinuous Galerkin framework, and does so with a general Phase-Field framework that captures a variety of previously proposed Phase-Field models. We also take special care to ensure that the overall scheme is well-balanced, does not produce spurious oscillations at interfaces, and satisfies consistency of reduction theoretically and numerically. Each of these properties is important, as existing schemes fail one or multiple of these properties. The rest of the paper is organized as follows: In Section 3.3, we describe the physical model for compressible interfacial flows of interest, including a general convex equation of state and Phase Fields. We then detail the Runge-Kutta discontinuous Galerkin and subcell finite volume discretization with consistent and conservative Phase Fields in Section 3.4. The importance of the well-balanced property for the non-conservative volume fraction equation is highlighted and demonstrated, and novel high-order schemes for the Phase-Field discretization are presented. The resolving power and efficiency of

the proposed scheme are demonstrated through numerical tests in Section 3.5, with comparisons to various other schemes and experiments. Finally, conclusions are presented in Section 3.6.

3.3 Equations of Motion

3.3.1 Multiphase Model

To simulate the compressible interfacial flows of interest, we solve the five-equations multiphase model, which describes the evolution of two immiscible fluids,

$$\begin{aligned}
&\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}^{\mathrm{HB}} = \mathbf{S}^{\mathrm{HB}}, \\
&\mathbf{U} = \begin{bmatrix} (\alpha_1\rho_1) \\ (\alpha_2\rho_2) \\ \rho\mathbf{u} \\ \rho E \\ \alpha_1 \end{bmatrix}, \quad \mathbf{F}^{\mathrm{HB}} = \begin{bmatrix} \mathbf{u}(\alpha_1\rho_1) \\ \mathbf{u}(\alpha_2\rho_2) \\ \rho\mathbf{u}\otimes\mathbf{u} + P\mathbf{I} \\ \mathbf{u}(\rho E + P) \\ \mathbf{u}\alpha_1 \end{bmatrix}, \quad \mathbf{S}^{\mathrm{HB}} = \begin{bmatrix} 0 \\ 0 \\ \mathbf{0} \\ 0 \\ \alpha_1 \left(1 + K_1\right) \nabla \cdot \mathbf{u} \end{bmatrix},
\end{aligned} \tag{3.1}$$

where $\mathbf{U}$ is the vector of the state variables, $\mathbf{F}^{\mathrm{HB}}$ is the vector of the hyperbolic fluxes, $\mathbf{S}^{\mathrm{HB}}$ are the non-conservative source terms, $\alpha_m \in [0, 1]$ is the volume fraction of material m, ρ_m is the density of material m, $\mathbf{u}$ is the flow velocity, E is the total energy per unit mass, and P is the pressure. The moniker *non-conservative* is used for the source terms to follow convention. As in [93], the present work admits any formulation of K_m that obeys

$$\sum_{m=1}^{N} \alpha_m K_m = 0, \tag{3.2}$$

two examples of which are provided below, and $N = 2$. For consistency of reduction, and thus to avoid voids, fictitious phases, or overfilling, the volume fraction must obey the summation constraint [93]

$$\sum_{m=1}^{N} \alpha_m = 1, \tag{3.3}$$

which is consistent with Eq. (3.1) with the requirement in Eq. (3.2). The phasic masses are required to be positive, i.e., $\{\alpha_m\rho_m \geq 0\}_{m=1}^{N}$. With the mixture density $\rho = \sum_{m=1}^{N} \alpha_m\rho_m$, this system trivially recovers the single-fluid Euler equations in regions of pure materials. A convex phasic equation of

state closes the system, which we define generally as

$$\rho_m e_m = A_m P + B_m + \rho D_m = \rho_m C_m T + \frac{B_m}{1 + A_m} + \rho D_m, \quad 1 \leq m \leq N \tag{3.4}$$

where A_m, B_m, C_m, and D_m are analytic or experimentally fit parameters that describe material properties. Different choices of these coefficients correspond to ideal gas, stiffened gas, and Mie-Grüneisen equations of state [91, 177]. The mixture internal energy is defined

$$\begin{aligned} \rho e = \rho E - \frac{1}{2}\mathbf{u}\cdot\mathbf{u}\sum_{m=1}^{N}\alpha_m\rho_m &= \sum_{m=1}^{N}\left(\alpha_m A_m P + \alpha_m B_m + \alpha_m\rho_m D_m\right) \\ &= \sum_{m=1}^{N}\left(\alpha_m\rho_m C_m T + \alpha_m\frac{B_m}{1+A_m} + \alpha_m\rho_m D_m\right) \end{aligned} \tag{3.5}$$

with which the mixture pressure and temperature are defined,

$$P = \frac{\rho e - \sum_{m=1}^{N}\alpha_m B_m - \sum_{m=1}^{N}\alpha_m\rho_m D_m}{\sum_{m=1}^{N}\alpha_m A_m}, \quad T = \frac{\rho e - \sum_{m=1}^{N}\alpha_m\frac{B_m}{1+A_m} - \sum_{m=1}^{N}\alpha_m\rho_m D_m}{\sum_{m=1}^{N}\alpha_m\rho_m C_m}. \tag{3.6}$$

The sound speed for material m is derived from Eq. (3.4),

$$\rho_m c_m^2 = \frac{(1 + A_m)P + B_m}{A_m}, \quad 1 \leq m \leq N. \tag{3.7}$$

The mixture speed of sound is,

$$\rho c^2 = \frac{\sum_{m=1}^{N}\alpha_m A_m\rho_m c_m^2(1 + K_m)}{\sum_{m=1}^{N}\alpha_m A_m}, \tag{3.8}$$

where choosing $\{K_m = 0\}_{m=1}^{N}$ results in the model of Allaire [13] and $\{K_m = \frac{\rho c^2}{\rho_m c_m^2} - 1\}_{m=1}^{N}$ reduces to the model of Kapila [121]. The corresponding sound speeds are,

$$\rho c^2 = \begin{cases} \frac{\sum_{m=1}^{N}\alpha_m A_m\rho_m c_m^2}{\sum_{m=1}^{N}\alpha_m A_m}, & \{K_m = 0\}_{m=1}^{N} \\ \left(\sum_{m=1}^{N}\alpha_m/(\rho_m c_m^2)\right)^{-1}, & \{K_m = \frac{\rho c^2}{\rho_m c_m^2} - 1\}_{m=1}^{N}. \end{cases} \tag{3.9}$$

This equation of state does not require the computation of any phasic properties unweighted by the volume fraction, such as $\rho_m = (\alpha_m\rho_m)/\alpha_m$, which are not well-defined for materials in regions where $\alpha_m = 0$ for any m. More complex equations of state have been considered in other works [11, 183]. While we do not consider viscous stresses, heat transfer, or surface tension, their inclusion is straightforward.

3.3.2 Consistent and Conservative Phase Field

The system Eq. (3.1) does not describe the process of mass diffusion by virtue of its definition for immiscible materials, with the underlying assumption that material interfaces instantaneously reach their thermodynamical equilibrium state. Material interfaces should thus be thin, if not sharp. The critical issue with Eq. (3.1) is that there is no physical mechanism maintaining this assumption of instantaneous relaxation of material interfaces to equilibrium when experiencing numerical diffusion. The numerical diffusion inherent to discontinuity-capturing schemes causes interfaces to diffuse while velocity gradients in multiple dimensions lead to local compression or expansion of the interfaces, thus resulting in continuously smearing interfaces with non-uniform thicknesses. This effect is detrimental to the accuracy of multiphase flow simulations, particularly when the simulation runs over multiple characteristic times, as material regions become increasingly challenging to distinguish from each other. This trait of linearly degenerate fields is in contrast to, for example, shocks, whose nonlinearity forces their thickness to be on the order of a few grid points under a shock-capturing scheme. To combat these challenges with representing material interfaces, we adapt the consistent and conservative Phase-Field method of [93]. Phase-Field terms are added to the governing equations according to

$$\begin{aligned}
&\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}^{\mathrm{HB}} = \mathbf{S}^{\mathrm{HB}} + \nabla \cdot \mathbf{F}^{\mathrm{PF}}, \\
&\mathbf{F}^{\mathrm{PF}} = \begin{bmatrix} \mathbf{J}_1\rho_1 \\ -\mathbf{J}_1\rho_2 \\ (\mathbf{J}_1\rho_1 - \mathbf{J}_1\rho_2) \otimes \mathbf{u} \\ (\mathbf{J}_1\rho_1 - \mathbf{J}_1\rho_2)\,\frac{\mathbf{u}\cdot\mathbf{u}}{2} + \mathbf{J}_1\rho_1 e_1 - \mathbf{J}_1\rho_2 e_2 \\ \mathbf{J}_1 \end{bmatrix},
\end{aligned} \tag{3.10}$$

where $\mathbf{F}^{\mathrm{PF}}$ is the Phase-Field flux vector and $\mathbf{J}_m$ is the spatial vector of the Phase-Field mechanism for phase m, which is a function of the volume fractions and their derivatives. The new system Eq. (3.10) recovers the relaxation mechanism of material interfaces in a thermodynamically consistent manner to counteract numerical diffusion. This system ensures consistency of volume fraction conservation, consistency of mass conservation, consistency of mass and momentum transport, and kinetic energy balance [97, 98, 100]. It also satisfies the second law of thermodynamics and is Galilean invariant. Importantly, this system does not introduce spurious oscillations at isolated interfaces in a uniform flow and ensures consistency of reduction, i.e., that when a substance is not present in the flow Eq. (3.10) reduces analytically and numerically to the same model with $N-1$ materials, which in the case of $N = 2$ corresponds to the single-fluid Euler equations. The form of $\mathbf{J}_m$ is chosen such that $\sum_{m=1}^{N} \mathbf{J}_m = 0$, which satisfies the second law of thermodynamics and preserves velocity, pressure,

and temperature equilibrium at isolated material interfaces [93]. The original model Eq. (3.1) is recovered with $\{\mathbf{J}_m = \mathbf{0}\}_{m=1}^{N}$, i.e., all material interfaces relax to their thermodynamic equilibrium state. This model, Eq. (3.10), is a simplification of the more complete model introduced in [93], and the interface regularization model in [109] is captured by this model under some simplifications.

3.4 Non-Conformal Discontinuous Galerkin and Subcell Finite Volume Discretization

Our goal is to extend the Phase-Field formulation of [93] for the two-phase Euler equations (Eq. (3.1)) to high-order accuracy with the Runge-Kutta discontinuous Galerkin (RKDG) method [49, 199]. The key difficulties lie in solving the volume fraction equation with its non-conservative term, robust and accurate discontinuity capturing, which we detail in Section 3.4.1, and computation of the Phase-Field flux vector for the reduction-consistent formulation [93] with DG, which we describe in Section 3.4.2. We first recall the RKDG approach for hyperbolic equations,

$$\frac{\partial \mathbf{U}(\mathbf{x},t)}{\partial t} + \nabla\cdot\mathbf{F}(\mathbf{U}(\mathbf{x},t)) = 0, \tag{3.11}$$

with

$$\mathbf{U}(\mathbf{x},t=0) = \mathbf{U}_0(\mathbf{x}), \quad \mathbf{U}(\mathbf{x}\in\partial\Omega,t) = \mathbf{U}_B(\mathbf{x},t), \quad \mathbf{x}\in\Omega\subset\mathbb{R}^d, \quad t\in\mathbb{R}^+, \tag{3.12}$$

where $\mathbf{F}$ is the flux vector, d is the spatial dimension, Ω is the physical domain, $\mathbf{U}_0$ is the initial condition, $\mathbf{U}_B$ is the boundary condition, and $\partial\Omega$ is the domain boundary. Equation (3.10) can fully be written in this form, except for the volume fraction equation.

We divide the domain into a set $\mathcal{E}_h$ of N_e non-overlapping elements, with $\Omega = \bigcup_{e=1}^{N_e}\Omega_e$. Uniform Cartesian elements are used in the present work such that the element edge lengths $\Delta x = \Delta y = \Delta z = h$ are equal; the analysis herein does not preclude other element shapes or unstructured grids unless otherwise noted. Within each element, the solution is approximated as a tensor-product series expansion of order p, with $\mathcal{K} = (p+1)^d$ total basis functions,

$$\mathbf{U}(\mathbf{x}\in\Omega_e,t) \approx \mathbf{U}_e(\mathbf{x},t) = \sum_{i=0}^{\mathcal{K}-1}\mathbf{a}_i^e(t)v_i(\mathbf{x}), \tag{3.13}$$

where $\mathbf{a}_i^e(t)$ are the time-varying expansion coefficients of the approximate solution in element Ω_e. The finite element space for this approximation is defined,

$$V_h^p = \left\{v\in L_2(\Omega) : v|_{\Omega_e}\in V^p(\Omega_e) \quad \forall\Omega_e\in\mathcal{E}_h\right\}, \tag{3.14}$$

where V^p are polynomial functions of order up to p. The present study uses orthonormalized Legendre polynomials for both the test and basis functions such that $\bar{\mathbf{U}}_e = \mathbf{a}_0^e 2^{-d/2}$, where $\overline{(\cdot)}$ is the cell-average value. Given the inner product $\langle a, b \rangle_e = \int_{\Omega_e} ab \, d\Omega$, the initial condition is projected onto the selected function space as $\langle v_i, U_e(t=0) \rangle_e - \langle v_i, \mathbf{U}_0 \rangle_e$. A similar treatment is applied to the boundary condition to achieve weak-enforcement, which is more accurate than strong enforcement for DG schemes [53]. To write Eq. (3.11) in its weak form, we multiply it by the test functions v_i, integrate over each element, and apply the divergence theorem,

$$\int_{\Omega_e} v_i \frac{\partial \mathbf{U}_e}{\partial t} d\Omega + \int_{\partial \Omega_e} v_i \hat{\mathbf{F}} \cdot \mathbf{n} dS - \int_{\Omega_e} \mathbf{F} \cdot \nabla v_i d\Omega = 0, \quad 0 \leq i < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h, \tag{3.15}$$

where $\widehat{(\cdot)}$ is a unique flux evaluation at element interfaces. Inserting the discrete approximation Eq. (3.13) and rearranging yields the semi-discrete weak form,

$$\mathbf{M}_{ij}^e \frac{\partial \mathbf{a}_j^e}{\partial t} = \frac{\partial \mathbf{a}_j^e}{\partial t} \int_{\Omega_e} v_i v_j d\Omega = \int_{\Omega_e} \mathbf{F} \cdot \nabla v_i d\Omega - \int_{\partial \Omega_e} v_i \hat{\mathbf{F}} \cdot \mathbf{n} dS, \quad 0 \leq \{i, j\} < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h. \tag{3.16}$$

where $\mathbf{M}^e$ is the mass matrix, which for the selected basis and Cartesian grids is the identity matrix multiplied by the volume $|\Omega_e|$. The weak form is now a system of ordinary differential equations, which can be integrated in time with explicit Runge-Kutta schemes, e.g., using total variation diminishing schemes [85]. Integrals are computed using Legendre-Gauss-Lobatto quadrature of order at least $p + 2$.

To solve Eq. (3.10) with the RKDG method, we follow a fractional-step approach [93, 94], where the hyperbolic part of Eq. (3.10) is first advanced (hyperbolic step), and the Phase-Field part of Eq. (3.10) is advanced thereafter (Phase-Field step). The non-conservative term in the volume fraction equation (see Eq. (3.1)) requires special attention when formulating a well-balanced weak form. Further, the proposed scheme is high-order accurate and stable in smooth regions, but the presence of discontinuities in the flow may lead to inadmissible solutions in the absence of limiting. Finally, the computation of the numerical Phase-Field flux vector in the Phase-Field step in the RKDG framework requires care to ensure that the resulting scheme is accurate and not stiff. In the following sections, we describe the discretization of Eq. (3.10) and address each of these issues.

3.4.1 Runge-Kutta Discontinuous Galerkin Discretization of the Hyperbolic Step

The hyperbolic step follows directly from the above RKDG description, and the numerical flux vector $\hat{\mathbf{F}} = \hat{\mathbf{F}}^{\mathrm{HB}}$ is computed with any consistent Riemann solver. However, the volume fraction equation requires special attention to prevent bound violation and production of fictitious phases. It

is also important to prevent spurious oscillations at discontinuities to keep the solution admissible. In the present work, we consider the Lax-Friedrichs, HLL, and HLLC Riemann solvers [242, 243], with discussion of their application in Section 3.4.1.2.

3.4.1.1 Weak-Form Discretization of the Volume Fraction Equation

To cast the volume fraction equation into the RKDG framework, we begin by writing the weak form,

$$\int_{\Omega_e} v_i \frac{\partial \alpha_{1,e}}{\partial t} d\Omega + \int_{\partial\Omega_e} v_i \mathbf{n}\cdot \widehat{(\alpha_1 \mathbf{u})} dS - \int_{\Omega_e} \alpha_1 \mathbf{u}\cdot \nabla v_i d\Omega \\ - \int_{\Omega_e} v_i \alpha_1 (1 + K_1) \nabla\cdot \mathbf{u} d\Omega = 0, \quad 0 \leq i < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h. \tag{3.17}$$

The first three terms in this weak form can be evaluated directly as in Eq. (3.16). The fourth term, however, does not allow for application of the divergence theorem to offload derivatives solely onto the test functions, which is problematic as derivatives of $\mathbf{u}$ do not exist in the selected finite-element space. Different schemes for terms of this type have been proposed [43, 44, 60], including in the context of discontinuous Galerkin methods [91, 200]. However, as noted in [4], consistent approximate solutions to non-conservative hyperbolic systems cannot generally be found using path-conservative type schemes. Similarly, as noted in [179, 180], for schemes that solve non-conservative equations to achieve the well-balanced property [179], the evaluation of the derivatives in the non-conservative terms must match that from the Riemann solver, where well-balanced in this context manifests in preservation of features in α_m under constant-velocity advection in a uniform pressure field. Specifically, the numerical discretization of $\nabla \cdot \mathbf{u}$ in the non-conservative term must be equal to the discretization of $\nabla \cdot (\mathbf{u}\alpha_1)$ for $\alpha_1 = 1$ in the flux term [93] to prevent generating Phase 2 in Phase 1. To this end, we rewrite the weak form of the volume fraction equation as:

$$\int_{\Omega_e} v_i \frac{\partial \alpha_{1,e}}{\partial t} d\Omega + \int_{\partial\Omega_e} v_i \mathbf{n}\cdot \widehat{(\alpha_1 \mathbf{u})} dS - \int_{\Omega_e} \alpha_1 \mathbf{u}\cdot \nabla v_i d\Omega \\ - \frac{1}{|\Omega_e|} \int_{\Omega_e} \alpha_1 \left(1 + K_1\right) d\Omega \int_{\Omega_e} v_i \nabla\cdot \mathbf{u} d\Omega = 0, \quad 0 \leq i < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h, \tag{3.18}$$

where the approximation $\alpha_1(1+K_1) \approx \frac{1}{|\Omega_e|}\int_{\Omega_e} \alpha_1(1+K_1) d\Omega = \bar{\alpha}_1 + \overline{\alpha_1 K_1}$ is made, which is different from [280] in that the additional approximation $\overline{\alpha_1 K_1} \approx \bar{\alpha}_1 \bar{K}_1$ made therein is not assumed. We demonstrate that this scheme still bounds the cell-average of α_1 under an appropriate time step constraint, and in the case of $\{K_m = 0\}_{m=1}^N$, these two methods are equivalent. Approximating this expression introduces an $O(\Delta x)$ error, but it is crucial as it allows the application of the divergence

theorem to offload all derivatives to the test functions so a weak form can be well defined in the finite element space. Further, because our modeling approach maintains thin interfaces, the overall order of accuracy is not impacted by the limitation that discontinuity capturing schemes are always at most first-order at discontinuities, in contrast to models that allow for arbitrarily smooth mixing regions. Applying the divergence theorem to the fourth term, we arrive upon the desired weak form, containing derivatives of only the test functions,

$$\begin{aligned}\int_{\Omega_e} v_i \frac{\partial \alpha_{1,e}}{\partial t} d\Omega + \int_{\partial\Omega_e} v_i \mathbf{n}\cdot \left[\widehat{(\alpha_1 \mathbf{u})} - \left(\bar{\alpha}_1 + \overline{\alpha_1 K_1}\right)\hat{\mathbf{u}}\right] dS \\ - \int_{\Omega_e} \left[\alpha_1 - \left(\bar{\alpha}_1 + \overline{\alpha_1 K_1}\right)\right] \mathbf{u}\cdot\nabla v_i d\Omega = 0, \quad 0 \leq i < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h,\end{aligned} \tag{3.19}$$

which is well-balanced as the numerical approximation to $\nabla \cdot \mathbf{u}$ is consistent with the solution to the Riemann problem for $\widehat{(\alpha_1 \mathbf{u})}$ by definition. Consistency of the cell average with finite volume schemes can be demonstrated by examining the zeroth mode of the solution ($v_0 = 2^{-d/2}$) and exploiting the orthogonality of the selected basis,

$$\begin{aligned}\int_{\Omega_e} v_0 \frac{\partial \alpha_{1,e}}{\partial t} d\Omega + \int_{\partial\Omega_e} v_0 \mathbf{n}\cdot \left[\widehat{(\alpha_1 \mathbf{u})} - \left(\bar{\alpha}_1 + \overline{\alpha_1 K_1}\right)\hat{\mathbf{u}}\right] dS \\ - \int_{\Omega_e} \left[\alpha_1 - \left(\bar{\alpha}_1 + \overline{\alpha_1 K_1}\right)\right] \mathbf{u}\cdot\nabla v_0 d\Omega = 0, \quad 0 \leq i < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h,\end{aligned} \tag{3.20}$$

$$\implies \quad \frac{\partial \bar{\alpha}_{1,e}}{\partial t} + \frac{1}{|\Omega_e|}\int_{\partial\Omega_e} \mathbf{n}\cdot \left[\widehat{(\alpha_1 \mathbf{u})} - \left(\bar{\alpha}_1 + \overline{\alpha_1 K_1}\right)\hat{\mathbf{u}}\right] dS = 0, \quad \forall \Omega_e \in \mathcal{E}_h, \tag{3.21}$$

which by setting $\bar{\alpha}_1 + \overline{\alpha_1 K_1} = \bar{\alpha}_1(1 + K_1(\bar{\mathbf{U}}))$ exactly matches various finite volume schemes under an appropriate polynomial reconstruction for the term in divergence form [113, 212, 214, 215].

We also require the scheme to be exact away from interfaces for all modes to mitigate the error from the above approximation, which we demonstrate by setting $\alpha_{1,e}(x) = \bar{\alpha}_1 = 1$ and manipulating to obtain,

$$\int_{\Omega_e} v_i \frac{\partial \alpha_{1,e}}{\partial t} d\Omega - \bar{\alpha}_1 \bar{K}_1 \left[\int_{\partial\Omega_e} v_i \mathbf{n}\cdot \widehat{\mathbf{u}} dS - \int_{\Omega_e} \mathbf{u}\cdot\nabla v_i d\Omega\right] = 0, \quad 0 \leq i < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h. \tag{3.22}$$

If $\bar{\alpha}_1 = 0$, the spatial integrals vanish and the scheme is trivially exact. To satisfy Eq. (3.2) we must have $\alpha_1 K_1 = -\sum_{m=2}^{N} \alpha_m K_m$. If $\alpha_1 = 1$, we have $\{\alpha_m = 0\}_{m=2}^{N}$ by Eq. (3.3) and K_1 must be equal to zero. As such, the spatial terms in Eq. (3.22) vanish and the scheme is exact.

The above approximation for the non-conservative term leads to $\bar{\alpha}_{1,e} \in [0, 1]$ under the CFL

condition,

$$\Delta t \leq \text{CFL}\, h \min\left[\frac{\hat{\omega}_1}{2},\ \frac{1}{2p+1}\right] \min\left[\frac{1}{||(|u|+c)||_\infty},\ \frac{1}{||\Delta u(1+K_1)||_\infty},\ \frac{1}{||\Delta u(1+K_2)||_\infty}\right]. \quad (3.23)$$

Theorem 1. *Consider a Legendre-Gauss-Lobatto quadrature rule on* $[-1/2, 1/2]$ *with* N_Q *nodes* ξ_q *and weights* $\hat{\omega}_q$ *such that* $\sum_{q=1}^{N_Q} \hat{\omega}_q = 1$. *Sufficient conditions for* $\bar{\alpha}_1^{n+1} \in [0, 1]$ *as evolved according to Eq.* (3.19) *and Eq.* (3.39) *are,*

i) $\alpha_1^n(x_q) \in [0, 1], \quad 1 \leq q \leq N_q$,

ii) $\bar{\alpha}_1^{n+1} \geq 0$,

iii) $\bar{\alpha}_2^{n+1} = 1 - \bar{\alpha}_1^{n+1} \geq 0$, *and*

iv) $\Delta t \leq \text{CFL}\, h \min\left[\frac{\hat{\omega}_1}{2},\ \frac{1}{2p+1}\right] \min\left[\frac{1}{||(|u|+c)||_\infty},\ \frac{1}{||\Delta u(1+K_1)||_\infty},\ \frac{1}{||\Delta u(1+K_2)||_\infty}\right]$.

Proof. To demonstrate the high-order DG scheme bounds the cell-average of the volume fraction, we follow [195] and use a convex combination of first-order schemes to find the sufficient time-step constraint. The $p = 0$ mode of Eq. (3.19) can be written,

$$\frac{\partial \bar{\alpha}_{1,e}}{\partial t} + \frac{1}{|\Omega_e|} \int_{\partial\Omega_e} \mathbf{n} \cdot \left[\widehat{(\alpha_1 \mathbf{u})} - \left(\bar{\alpha}_1 + \overline{\alpha_1 K_1}\right) \hat{\mathbf{u}}\right] dS = 0, \quad \forall \Omega_e \in \mathcal{E}_h. \quad (3.24)$$

While any convex Runge-Kutta scheme may be used, consider a Forward Euler step without loss of generality, which in one dimension leads to,

$$\begin{aligned} \bar{\alpha}_{1,k}^{n+1} = \bar{\alpha}_{1,k}^n &- \lambda \left[\hat{F}_{\alpha_1}(\mathbf{U}_{k+1/2}^{n,L}, \mathbf{U}_{k+1/2}^{n,R}) - \hat{F}_{\alpha_1}(\mathbf{U}_{k-1/2}^{n,L}, \mathbf{U}_{k-1/2}^{n,R})\right] \\ &+ \lambda \left(\bar{\alpha}_{1,k}^n + \left(\overline{\alpha_1 K_1}\right)_k^n\right) \left[\hat{u}(\mathbf{U}_{k+1/2}^{n,L}, \mathbf{U}_{k+1/2}^{n,R}) - \hat{u}(\mathbf{U}_{k-1/2}^{n,L}, \mathbf{U}_{k-1/2}^{n,R})\right], \end{aligned} \quad (3.25)$$

where k is the element index, $k \pm 1/2$ is the face location, L and R are the left and right states at the cell interface, and $\lambda = \Delta t/h$. The face values are evaluated as the polynomial solution at that point, i.e., $\mathbf{U}_{k\pm1/2} = \mathbf{U}_k(\xi = \pm 1/2)$ for a reference element defined on $[-1/2, 1/2]$. Any monotonic numerical flux $\hat{\mathbf{F}}$ and its corresponding velocity $\hat{u}$ are acceptable. Without loss of generality, we assume a Lax-Friedrichs flux,

$$\hat{\mathbf{F}}(\mathbf{U}_a, \mathbf{U}_b) = \frac{1}{2}\left[\mathbf{F}(\mathbf{U}_a) + \mathbf{F}(\mathbf{U}_b) - \sigma(\mathbf{U}_b - \mathbf{U}_a)\right], \quad \hat{u} = \frac{1}{2}(u_a + u_b), \quad (3.26)$$

where σ is the local maximum signal speed and $\mathbf{F}$ is the analytic conservative flux. After using

$$\bar{\phi} = \sum_{q=1}^{N_Q} \hat{\omega}_q \phi(\xi_q), \quad \sum_{m=1}^{N} \alpha_{m,k}^n(\xi_q) = 1, \quad \sum_{m=1}^{N} \alpha_{m,k}^n(\xi_q) K_m(\mathbf{U}_k^n(\xi_q)) = 0, \tag{3.27}$$

the implied evolution equation for $\bar{\alpha}_2$ from Eq. (3.25) is

$$\begin{aligned} \bar{\alpha}_{2,k}^{n+1} = 1 - \bar{\alpha}_{1,k}^{n+1} = \bar{\alpha}_{2,k}^n - \lambda \left[\hat{F}_{\alpha_2}(\mathbf{U}_{k+1/2}^{n,L}, \mathbf{U}_{k+1/2}^{n,R}) - \hat{F}_{\alpha_2}(\mathbf{U}_{k-1/2}^{n,L}, \mathbf{U}_{k-1/2}^{n,R}) \right] \\ + \lambda \left(\bar{\alpha}_{2,k}^n + \left(\overline{\alpha_2 K_2} \right)_k^n \right) \left[\hat{u}(\mathbf{U}_{k+1/2}^{n,L}, \mathbf{U}_{k+1/2}^{n,R}) - \hat{u}(\mathbf{U}_{k-1/2}^{n,L}, \mathbf{U}_{k-1/2}^{n,R}) \right]. \end{aligned} \tag{3.28}$$

The last two requirements in Eq. (3.27) follow from Eq. (3.3) and Eq. (3.2), which must be satisfied at every quadrature point. To obtain $\bar{\alpha}_{1,k}^{n+1} \in [0, 1]$, we only need to show that both $\bar{\alpha}_{1,k}^{n+1}$ and $\bar{\alpha}_{2,k}^{n+1}$ in Eq. (3.25) and Eq. (3.28) are non-negative. Since these equations have the same form, we only analyze Eq. (3.25). After some algebraic operations, Eq. (3.26) becomes

$$\begin{aligned} \bar{\alpha}_{1,k}^{n+1} = \sum_{q=2}^{N_Q - 1} \hat{\omega}_q \alpha_{1,k}^n(\xi_q)(1 + \lambda \Delta u (1 + K_1(\mathbf{U}_k^n(\xi_q)))) + \hat{\omega}_1 H_{\lambda/\hat{\omega}_1}^*(\mathbf{U}_{k-1/2}^{n,L}, \mathbf{U}_{k-1/2}^{n,R}, \mathbf{U}_{k+1/2}^{n,L}) \\ + \hat{\omega}_{N_Q} H_{\lambda/\hat{\omega}_{N_Q}}^*(\mathbf{U}_{k-1/2}^{n,R}, \mathbf{U}_{k+1/2}^{n,L}, \mathbf{U}_{k+1/2}^{n,R}), \end{aligned} \tag{3.29}$$

where the first-order scheme $H_{\lambda/\hat{\omega}_q}^*$ is,

$$\begin{aligned} H_{\lambda/\omega_q}^*(\mathbf{U}_a, \mathbf{U}_b, \mathbf{U}_c) = \alpha_{1,b}^n - \frac{\lambda}{\hat{\omega}_q} \left[\hat{F}(\mathbf{U}_b, \mathbf{U}_c) - \hat{F}(\mathbf{U}_a, \mathbf{U}_b) \right] \\ + \frac{\lambda}{\hat{\omega}_q} \left[\hat{u}(\mathbf{U}_b, \mathbf{U}_c) - \hat{u}(\mathbf{U}_a, \mathbf{U}_b) \right] \alpha_{1,b}^n \left(1 + K_1(\mathbf{U}_b^n) \right). \end{aligned} \tag{3.30}$$

As the quadrature rule is a convex combination, given the first sufficient condition satisfied via the scaling limiters, e.g., in [265, 280], simply ensuring that the quantities multiplied by the quadrature weights are positive is sufficient to make $\bar{\alpha}_{1,k}^{n+1}$ positive. The first sufficient condition is satisfied when necessary in practice via a simple limiter that scales the state vector around its cell-average, which is only possible in a fully-conservative manner if the cell-average is bounded. The first-order scheme can be rewritten as

$$\begin{aligned} H_{\lambda/\omega_q}^*(\mathbf{U}_a, \mathbf{U}_b, \mathbf{U}_c) = \frac{\lambda \sigma_{ab}}{2\hat{\omega}_q} \left[1 + \frac{u_a^n}{\sigma_{ab}} \right] \alpha_{1,a}^n + \frac{\lambda \sigma_{ab}}{2\hat{\omega}_q} \left[1 - \frac{u_c^n}{\sigma_{ab}} \right] \alpha_{1,c}^n + \\ \left[1 - \frac{\lambda \sigma_{ab}}{\hat{\omega}_q} \left(1 + \frac{(1 + K_1(\mathbf{U}_b^n))(u_a^n - u_c^n)}{2\sigma_{ab}} \right) \right] \alpha_{1,b}^n, \end{aligned} \tag{3.31}$$

which leads to the following condition for positivity of the first-order scheme:

$$\lambda \leq \frac{\hat{\omega}_q}{2\max\left(|u_a^n|, |u_c^n|, |(1+K_1(\mathbf{U}_b^n))(u_a^n - u_c^n)|\right)}. \tag{3.32}$$

The final condition comes from the truncated volume integral in Eq. (3.29). As the sign of Δu and K_1 are undetermined we write,

$$\lambda \leq \frac{1}{|\Delta u(1+K_1(\mathbf{U}_k^n(\xi_q)))|}, \tag{3.33}$$

with the difference between the face velocities given explicitly as,

$$\Delta u = \frac{1}{2}\left[u_{k+1/2}^{n,L} + u_{k+1/2}^{n,R} - u_{k-1/2}^{n,L} - u_{k-1/2}^{n,R}\right]. \tag{3.34}$$

The above requirements for λ guarantees $\bar{\alpha}_{1,k}^{n+1} \geq 0$, and we can obtain similar requirements for $\bar{\alpha}_{2,k}^{n+1} \geq 0$. A CFL condition for this scheme is thus given by Eq. (3.23) with norms computed at quadrature points; the CFL number depends on the Riemann solver and the order of the DG polynomials. Accordingly, $\bar{\alpha}_{1,k}^{n+1}$ as computed by Eq. (3.19) is bounded on $[0, 1]$ under the stated time-step constraint. □

An alternative approach would be to write the volume fraction equation as,

$$\frac{\partial \alpha_1}{\partial t} + \mathbf{u}\cdot\nabla\alpha_1 - \alpha_1 K_1 \nabla\cdot\mathbf{u} = 0, \tag{3.35}$$

which can be solved via the scheme of [91] or [46]. However, schemes for this form of the volume fraction equation can become inconsistent when the pressure varies strongly across a non-constant volume fraction [113]. Further, we demonstrate in Appendix A how schemes of this type fail to achieve consistency of reduction and cause non-physical states. Thus, to achieve a scheme that is well-balanced, exact away from interfaces, consistent across non-uniform pressure and volume fraction, and reduction consistent, the volume fraction equation is solved in the form presented in Eq. (3.19).

3.4.1.2 Subcell Discontinuity Capturing

Ensuring that the solution at quadrature points in each cell is in an admissible state before evaluating the discrete residual (right-hand side of Eq. (3.16)) is required for boundedness [280], but it can be challenging to satisfy this constraint near discontinuities due to the Gibbs phenomenon. Several techniques have been developed for discontinuous Galerkin schemes to minimize the Gibbs phenomenon when attempting to capture discontinuities, including solution limiters [49, 126, 275],

artificial viscosity [185], discontinuity tracking [90, 127, 128], and subcell schemes [66, 105, 112, 157, 186, 279]. Limiting schemes have historically been favored for explicit discontinuous Galerkin calculations for their straightforward implementations and efficiency. However, their resolving power for discontinuities has often been found to underperform finite volume schemes of the same order and DOF due to their non-locality, and their robustness and accuracy for high p is not clear. Artificial viscosity schemes can achieve greater locality in their discontinuity capturing, but typically rely on problem-specific parameters and can be more computationally expensive than limiters. Discontinuity tracking via adaptive meshing or by cutting cells can achieve high-order accuracy even in the presence of discontinuities, but comes with the cost of re-meshing or defining quadrature rules for irregular cut cells. Subcell schemes blend the highly local discontinuity capturing and robustness of finite volume schemes with the high-order accuracy of DG. Cells containing smooth data are evolved with DG, and cells containing discontinuities are evolved with a finite volume scheme defined on a subcell grid, typically with $(p+1)^d$ or $(2p+1)^d$ cells [66, 112]. The small cell size of the subcell grid allows for very local discontinuity capturing. Subcell schemes that blend the DG and finite volume residuals, such as in [157, 158, 206, 207], can significantly reduce numerical dissipation and improve the wavenumber resolution. However, they require computing two separate residuals, which increases computational expense. Additionally, while subcell schemes based on the Legendre-Gauss-Lobatto nodes have the advantage of not requiring a projection/reconstruction step, the resultant subcell grid is not uniform. Such subcell schemes thus underperform a finite volume scheme with the same number of DOF on a uniform grid if the blending primarily weighs the finite volume residual. Blending schemes based on uniform subcells exist, but require additional steps for conservation [158].

To achieve a reasonable combination of efficiency and accuracy, we use a subcell finite volume scheme similar to that of [112] without the boundary variation diminishing reconstruction and of [158] with no blending and with arbitrarily high-order quadrature. Excluding the blending procedure is done to avoid calculating both the DG and subcell residuals, and the boundary variation diminishing reconstruction is replaced by bound-preserving reconstructions [93]. We choose a uniform isoparametric subcell grid, i.e., the number of subcells is equal to the number of modes $\mathcal{K}$, to leverage its consistency with a finite volume scheme with the same DOF and the exact invertibility of the projection and reconstruction operators, as this choice ensures that they are square matrices. A cell to be evolved with subcells is defined $\Omega_e = \bigcup_{s=0}^{K-1} \Omega_e^s$, and the solution in each subcell is computed by integrating in each subcell and dividing by its volume,

$$\bar{\mathbf{U}}_e^s = \frac{1}{|\Omega_e^s|} \int_{\Omega_e^s} \mathbf{U}_e d\Omega = \frac{1}{|\Omega_e^s|} \sum_{i=0}^{\mathcal{K}-1} \int_{\Omega_e^s} \mathbf{a}_i^e v_i d\Omega, \quad 0 \leq s < \mathcal{K}, \tag{3.36}$$

where $\bar{\mathbf{U}}_e^s$ is the cell-average data in Ω_e^s. The projection matrix is then defined,

$$\frac{1}{|\Omega_e^s|}\int_{\Omega_e^s} v_i d\Omega = \mathbf{P}_{si}^e \in \mathbb{R}^{\mathcal{K}\times\mathcal{K}}, \tag{3.37}$$

which projects the DG polynomial to the cell-averaged values in individual subcells. As stated, the reconstruction matrix, which recovers the DG polynomial with the sub-cell data, is the inverse of the projection matrix, i.e., $\mathbf{R}^e = (\mathbf{P}^e)^{-1}$, as the projection matrix is square and non-singular. The condition numbers of the projection and reconstruction matrices appear to increase with $\exp(p)$. On an unstructured grid, these matrices would be different for every unique element due to the non-constant transformation Jacobian in the integral. However, with Cartesian grids, they are the same for all elements. The projection and reconstruction operations are, by construction, conservative. Additionally, because the projection and reconstruction operators are inverses of each other within machine precision, high-order information is not destroyed during either operation.

On the subcell grid, the finite volume scheme for the conservative equations in Eq. (3.1) is given in Ω_e by,

$$\frac{\partial \bar{\mathbf{U}}_e^s}{\partial t} + \frac{1}{|\Omega_e^s|}\int_{\partial\Omega_e^s} \hat{\mathbf{F}}\cdot\mathbf{n} dS = 0, \quad 0 \le s < \mathcal{K}. \tag{3.38}$$

We reconstruct the primitive variables $\mathbf{W}$ for the flux computation on the subcell grid using specific schemes to increase the accuracy while maintaining robustness and conservation, as reconstructing the conservative variables may not be conservative and may introduce spurious oscillations. For finite volume reconstructions in subcells, we need data from neighboring DG cells. Whenever requested, we perform the projection in the neighboring DG cells to provide the necessary data, even though we do not evolve that data over time. The reconstruction order is restricted to ensure the resultant scheme is compact on the DG grid. For $p \le 1$, we use a second-order MUSCL reconstruction [250]; for $p > 1$, we use a quartic WENO reconstruction [111] with midpoint-rule integration. Other schemes are available and can readily be implemented, but for $p \le 1$, MUSCL is robust, efficient, and uses the largest stencil that remains compact for the subcells; for $p > 1$, quartic WENO with mid-point integration provides less dissipation than MUSCL via a larger, but still compact stencil, and is still robust for many flows. The scheme for the volume fraction equation follows from Eq. (3.20),

$$\frac{\partial \bar{\alpha}_{1,e}^s}{\partial t} + \frac{1}{|\Omega_e^s|}\int_{\partial\Omega_e^s} \mathbf{n}\cdot\left(\widehat{(\alpha_1\mathbf{u})} - \bar{\alpha}_1\left[1 + K_1\left(\overline{\mathbf{U}}_e^s\right)\right]\hat{\mathbf{u}}\right) dS = 0, \quad 0 \le s < \mathcal{K}. \tag{3.39}$$

We note that the non-conservative term is evaluated using the subcell cell-averages as opposed to the DG cell-average. The subcell cell-averages are guaranteed to be bounded on $[0, 1]$ under the conditions in Section 3.4.1.1 with $\hat{\omega}_1 = 1$ and $p = 0$. The interface fluxes and the face velocity $\hat{\mathbf{u}}$ are

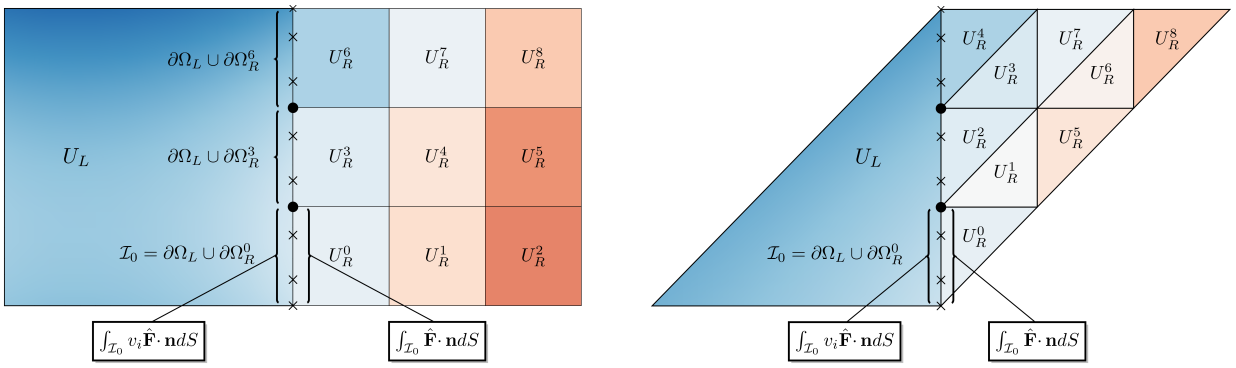


Figure 3.1: Example of square and triangular DG cells Ω_L with $p = 2$ solutions U_L bordering sets of subcells $\{\Omega_R^s\}_{s=0}^{K-1}$ with cell-average solutions $\{\bar{U}_R^s\}_{s=0}^{K-1}$. Quadrature points are denoted by $\bullet$ for points that are duplicated by two adjacent subcells and $\times$ for those that are not. Integration along the segment $\mathcal{I}_0$ is evaluated for Ω_L and Ω_R^0 according to a unique quadrature rule such that the scheme is conservative.

computed using Riemann solvers. Figure 3.1 shows the grid that results from the projection step, which is non-conformal; as such, on the interfaces between a DG cell and a set of subcells, sub-quadrature is necessary. The DG surface integral is evaluated by interpolating the DG solution to a set of quadrature points on each subcell face, integrating against the basis functions via quadrature, and summing the modal contribution from each face. The subcell finite volume surface integral is evaluated using the same quadrature points and quadrature rule, but the integration does not include the basis functions. The scheme is exactly conservative as the flux is uniquely defined at cell interfaces, and the same quadrature is used for the finite volume residual contribution as for the discontinuous Galerkin residual contribution. After a time step is complete, there may be a mixture of DG cells and subcells. In practice, the state of each DG cell is tracked between time steps or stages with a single integer, so the additional memory contribution is negligible, and as the projection is isoparametric, there is no extra memory required to store the solution in a projected cell. The finite volume reconstruction is performed according to [93] such that the volume fraction and phasic mass reconstructions are consistent and the scheme does not produce oscillations at interfaces, a key property for ensuring *numerical* consistency of reduction.

Crucially, the cells in which subcells are activated need to be determined. Schemes such as that in [66] detect cells *a posteriori* based on various factors (e.g., physical and numerical admissibility), then recompute the solution in flagged cells with subcells to achieve the desired admissibility properties, with a flux correction step for conservation. We instead select the cells to evolve with subcells *a priori* based on physical sensors. To detect contact waves, we use the characteristics-based sensor described in [91] for the Euler equations. Between two cell-averages

(L and R), the characteristic variables are given by,

$$\Delta\boldsymbol{\zeta} = \boldsymbol{\mathcal{R}}^{-1}\Delta\mathbf{U}, \quad \Delta\mathbf{U} = \mathbf{U}_R - \mathbf{U}_L = \sum_{i=1}^{d} \widetilde{\Delta\boldsymbol{\zeta}^{(i)}}\widetilde{\boldsymbol{\mathcal{R}}}^{(i)}, \tag{3.40}$$

where $\boldsymbol{\zeta}$ is the vector of characteristic variables, $\boldsymbol{\mathcal{R}}$ is the matrix of right eigenvectors, and $\tilde{\cdot}$ is an average. For a jump in the mixture density of $\Delta\rho = \sum_{m=1}^{N} \Delta(\alpha_m\rho_m)$, the strength of the contact is given by $\widetilde{\Delta\zeta^{(2)}} = (\Delta\rho\tilde{c}^2 - \Delta P)/\tilde{c}^2$. The sensor value Ξ is then given by the normalization,

$$\xi = \frac{|\widetilde{\Delta\zeta^{(2)}}|}{\rho_L + \rho_R}, \quad \Xi = \frac{2\xi}{(1+\xi)^2}. \tag{3.41}$$

To detect material interfaces, we measure the jump in the volume fraction to get the sensor value $Z = \max_m |\alpha_{m,R} - \alpha_{m,L}|$. As in [91], shocks are detected by the sensor of [149], which tests whether the Lax entropy condition is satisfied, $\mathbf{u}_L - c_L > \tilde{\mathbf{u}} - \tilde{c} > \mathbf{u}_R - c_R$. If this condition is satisfied for any component of the velocity, the pressure jump is normalized and used to get the sensor value Φ,

$$\phi_P = \frac{|P_R - P_L|}{P_L + P_R}, \quad \Phi = \frac{2\phi_P}{(1+\phi_P)^2}. \tag{3.42}$$

While these three sensors (Ξ, Z, and Φ) may be problem dependent, we find that they are quite insensitive to changes within an order of magnitude and recommend a starting point for all three of 0.01. We also check each quadrature point for admissibility, i.e., $\{\alpha_m\rho_m \geq 0\}_{m=1}^{N}$, $\{\alpha_m \in [0,1]\}_{m=1}^{N}$, and $P \geq P_{\min}$, where $P_{\min}$ is the pressure that satisfies $c^2 > 0$ and $\{c_m^2 > 0\}_{m=1}^{N}$. The sensors are computed before the hyperbolic residual is evaluated to determine where to apply subcells.

When evaluating the residual for the non-conservative term in Eq. (3.18), the DG cell-average is used, i.e., the approximation is made on the length scale of h. When an interface is evolved via subcells according to Eq. (3.39), however, the approximation is made on the scale of $h/(p+1)$, i.e., the non-conservative term is resolved at the sub-grid level. This sub-grid resolution increases the locality of the approximation and allows for sub-grid features in α_1 to be evolved accordingly, instead of evolving α_1 based on the DG cell-average. The interface sensor is inactive where α_m is constant or very close to constant, which is near pure material regions such that the non-conservative term is zero or very close to zero. As such, the non-conservative term introduces errors based on the DG cell-average only in cells away from interfaces where the non-conservative term is close to zero, and generates errors based on the subcell cell-averages at interfaces, where errors in the cell-averages are a factor of $p+1$ smaller than the DG cell-average. This method also ensures that the evolution of the interface in a subcell matches that of an equivalent finite volume scheme.

We either choose a single Riemann solver for the whole domain, including subcells, or we

choose the Riemann solver in a solution-adaptive fashion based on the described sensors. While the HLLC Riemann solver represents contacts more accurately than the HLL solver, it does not necessarily preserve bounds and may favor the carbuncle phenomenon [74, 221], particularly for high p. To minimize the impact of both of these issues in simulations with strong shocks, we can apply the HLL Riemann solver at shocks (when Φ is greater than the selected threshold) and where the solution is found to be inadmissible, and the HLLC Riemann solver elsewhere. At an interface between two cells, if either or both cells call for the application of the HLL solver, then HLL is used; otherwise, HLLC is applied at that interface. All subcells within a projected DG cell use the same Riemann solver. We find this scheme to be effective at avoiding the carbuncle instability while retaining the favorable properties of the HLLC solver at interfaces and contacts.

3.4.2 Runge-Kutta Discontinuous Galerkin Scheme for the Phase-Field Step

While DG can directly be used to compute the contribution of the Phase Field, accuracy, efficiency, and robustness are questionable in the computation of the Phase-Field mechanism. As such, we describe a straightforward and accurate scheme for the Phase-Field step, which follows largely the same procedure as the hyperbolic step, with the difference being in the calculation of the numerical Phase-Field flux vector. Additionally, due to the linearity of the proposed scheme, the consistency and conservation properties [93] of the Phase Field are retained. Splitting the Phase-Field step from the hyperbolic step is numerically convenient as both the hyperbolic and Phase-Field fluxes can be integrated explicitly in time while the calculation of $\mathbf{J}_1$ can be done using the data from the hyperbolic step, and can thus be thought of as a fractional step approach. That being said, like in [93], splitting is not required by the model. Both the hyperbolic and Phase-Field fluxes are evolved over a time Δt in each time step, unlike in a conservative level-set scheme that evolves the re-initialization terms to a steady state, and thus splits the model into distinct components. Further, it is possible and straightforward to implement the present approach in a non-split manner, by calculating the Phase-Field flux at the beginning of the time step. The system to be solved is Eq. (3.10) with the hyperbolic terms ignored. Following the *reduction-consistent formulation* [93], the Phase-Field step solves the conservation equation

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \left(-\tilde{\mathbf{F}}^{\mathrm{PF}}\right) = \mathbf{0}, \quad \tilde{\mathbf{F}}^{\mathrm{PF}} = \begin{bmatrix} \left(\tilde{\mathbf{J}}_1 \rho_1\right) \\ -\left(\tilde{\mathbf{J}}_1 \rho_2\right) \\ \left(\tilde{\mathbf{J}}_1 \rho_1\right) - \left(\tilde{\mathbf{J}}_1 \rho_2\right)) \otimes \mathbf{u} \\ \left(\tilde{\mathbf{J}}_1 \rho_1\right) - \left(\tilde{\mathbf{J}}_1 \rho_2\right)) \frac{\mathbf{u}\cdot\mathbf{u}}{2} + \left(\tilde{\mathbf{J}}_1 \rho_1 e_1\right) - \left(\tilde{\mathbf{J}}_1 \rho_2 e_2\right) \\ \tilde{\mathbf{J}}_1 \end{bmatrix}, \tag{3.43}$$

regardless of the formulation of $\mathbf{J}_1$ or $\nabla\cdot\mathbf{J}_1$, where

$$\begin{aligned}
\tilde{\mathbf{J}}_1 &= \alpha_1(1-\alpha_1)\nabla Q,\\
\left(\tilde{\mathbf{J}}_1\rho_1\right) &= \left(\alpha_1\rho_1\right)\left(1-\alpha_1\right)\nabla Q,\\
\left(\tilde{\mathbf{J}}_1\rho_2\right) &= \left(\alpha_2\rho_2\right)\alpha_1\nabla Q,\\
\left(\tilde{\mathbf{J}}_1\rho_1 e_1\right) &= A_1\tilde{\mathbf{J}}_1 P + B_1\tilde{\mathbf{J}}_1 + D_1\left(\tilde{\mathbf{J}}_1\rho_1\right),\\
\left(\tilde{\mathbf{J}}_1\rho_2 e_2\right) &= A_2\tilde{\mathbf{J}}_1 P + B_2\tilde{\mathbf{J}}_1 + D_2\left(\tilde{\mathbf{J}}_1\rho_2\right).
\end{aligned} \tag{3.44}$$

Here, $\tilde{\mathbf{J}}_1$ is a modified Phase-Field mechanism that satisfies the properties described in [93], and Q is a scalar that depends on the form of $\mathbf{J}_1$. To solve this system, we write a weak form of Eq. (3.43) and discretize it in a manner similar to that in Section 3.4, leading to,

$$\mathbf{M}^e_{ij}\frac{\partial \mathbf{a}^e_j}{\partial t} = \int_{\partial\Omega_e} v_i\hat{\tilde{\mathbf{F}}}^{\mathrm{PF}}\cdot\mathbf{n}dS - \int_{\Omega_e}\tilde{\mathbf{F}}^{\mathrm{PF}}\cdot\nabla v_i d\Omega, \quad 0\le\{i,j\}<\mathcal{K}, \quad \forall\Omega_e\in\mathcal{E}_h. \tag{3.45}$$

where $\hat{\tilde{\mathbf{F}}}^{\mathrm{PF}}$ represents the numerical Phase-Field flux vector defined as,

$$\hat{\tilde{\mathbf{F}}}^{\mathrm{PF}} = \begin{bmatrix}
\widehat{\left(\tilde{\mathbf{J}}_1\rho_1\right)}\\
-\widehat{\left(\tilde{\mathbf{J}}_1\rho_2\right)}\\
\left[\widehat{\left(\tilde{\mathbf{J}}_1\rho_1\right)} - \widehat{\left(\tilde{\mathbf{J}}_1\rho_2\right)}\right]\otimes\mathbf{u}^{(-)}\\
\left[\widehat{\left(\tilde{\mathbf{J}}_1\rho_1\right)} - \widehat{\left(\tilde{\mathbf{J}}_1\rho_2\right)}\right]\frac{\mathbf{u}^{(-)}\cdot\mathbf{u}^{(-)}}{2} + \widehat{\left(\tilde{\mathbf{J}}_1\rho_1 e_1\right)} - \widehat{\left(\tilde{\mathbf{J}}_1\rho_2 e_2\right)}\\
\hat{\tilde{\mathbf{J}}}_1
\end{bmatrix}. \tag{3.46}$$

Considering the upwind-downwind fluxes in [93], we have, at element faces, the volumetric flux

$$\hat{\tilde{\mathbf{J}}}_1 = \alpha_1^{(-)}(1-\alpha_1)^{(+)}\nabla Q. \tag{3.47}$$

where $f^{(+)}$ is the upwind face value of the f polynomial based on the direction of ∇Q, $f^{(-)}$ is the corresponding downwind face value of the f polynomial. The interfacial mass fluxes are given by,

$$\widehat{\left(\tilde{\mathbf{J}}_1\rho_1\right)} = \left(\alpha_1\rho_1\right)^{(-)}\left(1-\alpha_1\right)^{(+)}\nabla Q, \quad \widehat{\left(\tilde{\mathbf{J}}_1\rho_2\right)} = \left(\alpha_2\rho_2\right)^{(+)}\alpha_1^{(-)}\nabla Q. \tag{3.48}$$

The interfacial internal energy fluxes are given by,

$$\widehat{\left(\tilde{\mathbf{J}}_1\rho_1 e_1\right)} = A_1\hat{\tilde{\mathbf{J}}}_1 P^{(-)} + B_1\hat{\tilde{\mathbf{J}}}_1 + D_1\widehat{\left(\tilde{\mathbf{J}}_1\rho_1\right)}, \quad \widehat{\left(\tilde{\mathbf{J}}_1\rho_2 e_2\right)} = A_2\hat{\tilde{\mathbf{J}}}_1 P^{(+)} + B_2\hat{\tilde{\mathbf{J}}}_1 + D_2\widehat{\left(\tilde{\mathbf{J}}_1\rho_2\right)}. \tag{3.49}$$

To compute the volume integral, we evaluate the fluxes at quadrature points $\boldsymbol{\xi}_q$ inside elements. If a cell contains a discontinuity, the same subcell scheme as in Section 3.4.1.2 is used to evolve Eq. (3.45). The subcell weak form for the Phase-Field step is simply

$$\frac{\partial \bar{\mathbf{U}}_e^s}{\partial t} = \frac{1}{|\Omega_e^s|} \int_{\partial \Omega_e} \hat{\tilde{\mathbf{F}}}^{\mathrm{PF}} \cdot \mathbf{n} dS, \quad 0 \leq s < \mathcal{K}, \tag{3.50}$$

where the face values of $\mathbf{U}$ are reconstructed in the same manner as in Section 3.4.1.2.

The final element of the Phase-Field step is the calculation of Q and ∇Q based on the formulation of $\mathbf{J}_1$. A key component of the present formulation, and thus that in [93], is that a variety of Phase-Field mechanisms can be used while retaining key properties such as consistency, conservation, and consistency of reduction. Thus without loss of generality, we follow [93] and consider the conservative Allen-Cahn formulation described in [31],

$$\frac{\partial \phi}{\partial t} = M\nabla^2\phi - \frac{M}{\eta^2} g'(\phi) + M|\nabla\phi|\kappa + W(\phi) \frac{\int_\Omega \left(\frac{M}{\eta^2} g'(\phi) - M|\nabla\phi|\kappa \right) d\Omega}{\int_\Omega W(\phi) d\Omega} \tag{3.51}$$

with

$$W(\phi) = \sqrt{g(\phi)}, \quad g(\phi) = \phi^2(1-\phi)^2 \tag{3.52}$$

where ϕ is the order parameter, $M \geq 0$ is the mobility, $\eta \geq 0$ prescribes the interface thickness, $g(\phi)$ is the double-well potential, $g' = \partial_\phi g$, and $\kappa = -\nabla \cdot (\nabla\phi/|\nabla\phi|)$ is the curvature, computed with finite differences. This formulation has a regularization term $M\nabla^2\phi$ that smooths ϕ while the potential drives ϕ to either zero or one; the Lagrange multiplier ensures conservation. This model has been successfully applied to both compressible and incompressible flows [93, 96]. After manipulation, Eq. (3.51) is used to directly evaluate $\nabla \cdot \mathbf{J}_1$ as

$$\nabla \cdot \mathbf{J}_1 = M\nabla^2\phi - \frac{M}{\eta^2} g'(\phi) + M|\nabla\phi|\kappa + W(\phi) \frac{\int_\Omega \left(\frac{M}{\eta^2} g'(\phi) - M|\nabla\phi|\kappa \right) d\Omega}{\int_\Omega W(\phi) d\Omega}. \tag{3.53}$$

To ensure $\phi \in [0,1]$, [96] demonstrated $M\Delta t/\eta^2 \leq 1/4$, so we follow [93] and set $M = \eta \max||\mathbf{u}||_2/2$, such that $\max||\mathbf{u}||_2 \Delta t/\eta \leq 1/2$, which gives a less restrictive time step than that of the hyperbolic step. By setting $\phi = \alpha_1$, $\nabla \cdot \mathbf{J}_1$ can be directly evaluated using the data from the end of the hyperbolic step. While $\nabla \cdot \mathbf{J}_1$ can technically be directly evaluated using the DG solution approximation, such a scheme would have poor accuracy for the Laplacian term, particularly for non-smooth data. High-order schemes for solving elliptic equations with DG exist; however, they are often quite expensive and may be stiff. To balance accuracy with efficiency, we evaluate $\nabla \cdot \mathbf{J}_1$ using the subcell data of α_1 and second-order finite differences for the Laplacian term. As

the Phase-Field flux is non-zero only near interfaces due to consistency of reduction, where the discontinuity-capturing scheme is at most first-order accurate [137], this second-order approximation does not reduce the overall accuracy of the scheme. With $\nabla\cdot\mathbf{J}_1$ evaluated, Q and ∇Q can be calculated. For the specific formulation of $\nabla\cdot\mathbf{J}_1$ used here, the auxiliary variable is computed on the subcell grid of cell-averages as

$$\nabla\cdot\left(\bar{\alpha}_1\overline{(1-\alpha_1)}\nabla Q\right) = \nabla\cdot\mathbf{J}_1. \tag{3.54}$$

In this case, $\nabla\cdot\mathbf{J}_1$ is not exactly $\nabla\cdot\tilde{\mathbf{J}}_1$. Using $\nabla\cdot\tilde{\mathbf{J}}_1$ instead of $\nabla\cdot\mathbf{J}_1$ makes Eq. (3.54) a linear Poisson equation that is straightforward to approximate, but again the overall accuracy of the scheme is not impacted [93]. The left-hand side of Eq. (3.54) is discretized with linear interpolation and second-order finite differences; further discussion of the schemes for finding ∇Q is provided in Section 3.4.2.1. The gradient ∇Q is discretized in x as and $\partial_x Q|_{i+1/2} = (Q_{i+1} - Q_i)/(x_{i+1} - x_i)$ and similarly for each other direction. To avoid rank deficiency when solving for Q we follow [96] and add a small number $\delta = 10^{-6}$ to $\bar{\alpha}_1\overline{(1-\alpha_1)}$. We find that this value of δ works well, but results are not particularly sensitive to order of magnitude changes in δ. The linear solve is generally a lower cost than the hyperbolic step and the integration of the Phase-Field fluxes. Anecdotally, it typically takes between four and eight iterations to reduce the residual by ten orders of magnitude, and constitutes around 25% of the time step. This cost could be improved with tuning to the linear solver, but that is beyond the scope of this work.

To evaluate the interior and surface fluxes in Eq. (3.45), a polynomial representation of ∇Q is needed in each cell with unique values along the cell interface. However, there are more than $(p+1)^d$ values of ∇Q in each cell. Additionally, the face values of ∇Q are already uniquely defined, so any reconstruction should retain this property. In dimensions greater than one, direct volumetric interpolation of ∇Q in two DG cells does not, in general, result in matching face distributions, except at the interpolation nodes in the case of Lagrange interpolation. To overcome this issue, we relax the constraint that reconstructions in adjacent cells must be identical on shared surfaces and use Lagrange interpolation in each DG cell for the volumetric representation of ∇Q, directly using the values of the interpolation at quadrature points for use in the volume integral in Eq. (3.45). On faces, however, we reconcile the non-unique value of ∇Q by using an interpolation of dimension $d-1$ to evaluate the surface integrals. These face reconstructions match exactly with the volume reconstructions of the neighboring cells at $(p+1)^{d-1}$ interpolation nodes, but may differ at quadrature points. However, as the face fluxes are uniquely defined and surface integrals use the same quadrature rules across faces, the scheme is exactly conservative. This reconstruction is illustrated in Fig. 3.2. In DG cells that contain subcells, the interior face values of ∇Q are used directly; as such, conformal subcell fluxes exactly match those of an equivalent

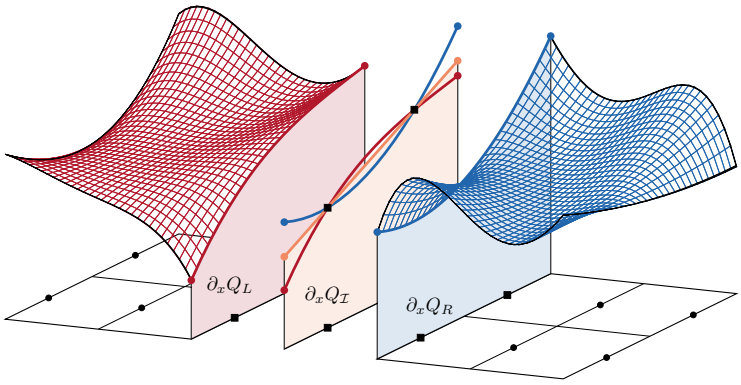


Figure 3.2: Reconstructions of $\partial_x Q$ in two $p = 1$ DG cells, L and R. The unique interfacial reconstruction $\partial_x Q_{\mathcal{I}}$ is shown in orange between the two elements. Subcell face values of $\partial_x Q$ not used for the interface reconstruction are shown with • while the face values used for the interface reconstruction are shown with ▪. The interface reconstruction matches nodally with both the left and right cells, but not necessarily at quadrature points. Face gradients are defined on a staggered grid.

finite volume scheme. On non-conformal faces, ∇Q is reconstructed and evaluated at appropriate quadrature points as in the conformal DG case, such that the scheme is exactly conservative. To integrate the Phase-Field step in time from $\tau = 0$ to $\tau = \Delta t$, we use the same Runge-Kutta scheme as in the hyperbolic step. However, the time step $\Delta\tau = \min(\Delta t, \Delta t^{\text{PF}})$ may be smaller than that in the hyperbolic step, where Δt^{PF} is determined for bound-preservation. As such, multiple Runge-Kutta steps may be necessary. The proposed scheme for the Phase-Field step does not require a specific formulation of the Phase-Field mechanism. Thus, the present work provides a generalized, arbitrarily high-order accurate framework for a variety of Phase-Field models that retains the consistency and conservation properties of its low-order counterparts.

3.4.2.1 Computation of ∇Q

The computation ∇Q is done via second-order schemes for robustness, efficiency, and because the intent here is to model sharp (thin) interfaces in flows with shocks. By Godunov's theorem [81], a discontinuity-capturing scheme can be at most first-order accurate. As such, at shocks, interfaces, and contacts, the present scheme will be at most first-order accurate. As demonstrated in Section 3.5.1, however, high-order accuracy can be achieved away from interfaces, even with the Phase-Field mechanism active, which highlights the numerical exactness of the consistency of reduction property. All components of the presented model are high-order accurate except for the linear solve, which, again, by the consistency of reduction property, only impacts the solution near interfaces. If thick interfaces are desired, a high-order method could be applied for the linear solve, which would make all components of the presented scheme formally high-order, but this is beyond

the scope of this work due to the problems of interest and the limitations of discontinuity-capturing schemes.

3.4.3 Summary

Given the solution at t^n, the resulting maximum wavespeeds sampled at quadrature points, and thus the time step size, the algorithm evolves the solution to the next step/stage according to the following algorithm:

1. Evaluate the hyperbolic step from t^n to t^{n+1} as in Section 3.4 with a Runge-Kutta scheme. At each stage:

   (a) Cells with discontinuities and inadmissible quadrature points are detected by the sensors in Section 3.4.1.2.

   (b) Unmarked cells are evolved with DG according to Eq. (3.16) and Eq. (3.19), and marked cells are evolved according to the subcell finite volume scheme in Eq. (3.38) and Eq. (3.39).

2. Project α_1^{n+1} to subcells to compute $\nabla \cdot \mathbf{J}_1$ according to Eq. (3.53).

3. Compute ∇Q on the subcell grid according to Eq. (3.54).

4. Evaluate the Phase-Field step from $\tau = 0$ to $\tau = \Delta t$ with a Runge-Kutta scheme in each cell via Eq. (3.45), with $\Delta\tau = \min(\Delta t, \Delta t^{\text{PF}})$, where Δt^{PF} is determined for bound-preservation, subcycling if necessary, see [93]. At each stage:

   (a) Cells with discontinuities and inadmissible quadrature points are detected by the sensors in Section 3.4.1.2.

   (b) Unmarked cells are evolved according to Eq. (3.45) with a volumetric reconstruction of ∇Q, and marked cells are evolved according to Eq. (3.50). Non-conformal faces and conformal DG faces use a surface reconstruction of ∇Q, while conformal finite volume faces directly use the subcell face values of ∇Q.

The scheme for computing the Phase-Field fluxes can be made more efficient, albeit at the cost of accuracy away from interfaces. By projecting the entire flow field to subcells, the framework of [93] can be directly applied in a finite volume sense. The calculation of ∇Q is the same as that presented above; however, the Phase-Field flux vector is computed and applied directly to the subcell data. High-order DG solutions are then simply reconstructed from the updated subcell solutions using the previously described invertible reconstruction operator. This scheme is less accurate away

from interfaces where the Phase-Field update would be computed using the high-order accurate DG scheme, Eq. (3.45); however, away from the interface, the Phase-Field fluxes are small, so the impact of the lower accuracy should not be dramatic. This scheme would be less accurate than the DG scheme in cases where the interface is far from its equilibrium thickness and gradients in α_1 are small, where the DG scheme would evolve the interface more accurately than the finite volume scheme. The high-order DG scheme for the Phase-Field step could be applied everywhere by deactivating the interface sensor and calculating ∇Q with a DG approximation. However, the Phase Field imposes strict requirements on the bounds of the polynomial representations of α_1, which are much easier to satisfy within the finite volume framework of the subcells [93].

In summary, we present a Runge-Kutta discontinuous Galerkin scheme that is high-order accurate away from discontinuities, evolves the non-conservative volume fraction equation in a well-balanced, bounded, and reduction-consistent manner, captures discontinuities robustly on the subgrid scale via an embedded subcell finite volume scheme, and keeps material interfaces uniformly thick throughout simulations via a general Phase-Field mechanism.

3.5 Verification and Validation

To quantify the interface thickness in some of the following test problems, we compute the average number of DG cells across the interface in a manner similar to that of [93],

$$N_I = \frac{1}{h} \frac{\int_\Omega \mathcal{H}(\alpha_1)\, d\Omega}{\int_\Omega |\nabla \alpha_1|\, d\Omega}, \quad \mathcal{H}(\alpha_1) = \begin{cases} 1, & (\alpha_1)_{\min} \leqslant \alpha_1 \leqslant (\alpha_1)_{\max} \\ 0, & \text{else} \end{cases} \tag{3.55}$$

As in [93] we use $(\alpha_1)_{\min} = 0.05$ and $(\alpha_1)_{\max} = 0.95$ to identify bounds of the diffuse interface. This quantity represents a volume in three dimensions, a surface area in two dimensions, and a length in one dimension. Additionally, we set $\eta = h/(p+1)$ for all test cases. For conciseness, while still demonstrating the ability of the presented scheme to resolve flows with various approximation orders, for test cases 3.5.2, 3.5.3, 3.5.4, and 3.5.5 we show results for a single value of p, which is different for the different problems; other values can be used with behavior consistent with that value.

3.5.1 Advection of an Isentropic Vortex

To verify the high-order accuracy of the proposed scheme for smooth flows, we compute the advection of an isentropic vortex [117, 261]. By setting the volume fraction to unity, this test also demonstrates that the schemes for the volume fraction equation and the Phase-Field step produce exact solutions in regions with no material interfaces, verifying consistency of reduction, and that

fictitious phases are not generated. This property is in contrast to other proposed schemes for non-conservative equations that may fail consistency of reduction and lead to errors, as demonstrated in Appendix A. The initial conditions are given by

$$\begin{bmatrix} \rho_1 \\ \rho_2 \\ u \\ v \\ P \\ \alpha_1 \end{bmatrix} = \begin{bmatrix} (1+\Delta T)^{A_1} \\ (1+\Delta T)^{A_2} \\ M_\infty \cos\alpha + \Delta u \\ M_\infty \sin\alpha + \Delta v \\ \frac{A_1}{1+A_1}(1+\Delta T)^{A_1+1} \\ 1 \end{bmatrix} \tag{3.56}$$

with $\Delta T = -\frac{1}{2}\omega^2/A_1$, $\Delta u = -y\omega/R$, $\Delta v = x\omega/R$, $\omega = \beta e^f$, $f = -\frac{1}{2}(x^2+y^2)/(R\sigma)^2$, $R = \frac{1}{2}$, $\sigma = 1$, $\beta = \frac{1}{5}$, $M_\infty = 2\sqrt{2}$, $\alpha = \pi/4$, $A_1 = A_2 = 5/2$, and $B_1 = B_2 = D_1 = D_2 = 0$. The vortex core should remain unchanged, only translated, over time, such that the exact solution is the initial condition shifted by $tM_\infty(\cos\alpha, \sin\alpha)$. To achieve design orders of accuracy, the simulation domain should be sufficiently large that $\omega \approx 0$ at the boundaries so that the effect of image vortices is not felt. We simulate this flow on $[x, y] \in [-2, 2]^2$ with doubly-periodic boundary conditions, which is sufficient at the cell sizes considered ($\omega \sim O(10^{-15})$ at boundaries). The flow is evolved until $t = 4$. Subcells are not activated anywhere in the domain, as the problem has no sharp features. The L_2 error norm is evaluated for a scalar quantity ϕ as,

$$||\phi_h - \phi_{\text{exact}}||_{L_2} = \sqrt{\frac{1}{|\Omega|}\sum_{e=1}^{N_e}\int_{\Omega_e}(\phi_h - \phi_{\text{exact}})^2\, d\Omega}, \tag{3.57}$$

where ϕ_{exact} is the exact solution of ϕ. This norm measures the integral error of the entire solution, and the expected convergence rate for smooth problems is $O(h^{p+1})$ for DG. The L_∞ error norm is calculated as,

$$||\phi_h - \phi_{\text{exact}}||_{L_\infty} = \max_{\Omega_e \in \mathcal{E}_h}\left|\frac{1}{|\Omega_e|}\int_{\Omega_e}(\phi_h - \phi_{\text{exact}})\, d\Omega\right|. \tag{3.58}$$

This norm measures the cell-average error, which for hyperbolic PDEs is expected to converge at $O(h^{2p+1})$ for DG and $O(h^{p_{\text{FV}}+1})$ for finite volume, where p_{FV} is the reconstruction order. This norm elucidates the resolving power of DG schemes for smooth regions of flows. Norms are computed using a Legendre-Gauss-Lobatto quadrature rule of order $2p + 5$ to minimize the quadrature integration error for non-polynomial integrands.

Figure 3.3 shows the L_2 and L_∞ errors in the mixture density and the volume fraction for the proposed DG scheme with different p and a second-order MUSCL finite volume scheme [250], where $\text{DOF}_x = (p+1)N_e^{1/2}$ is the number of degrees of freedom along an axis for DG and the

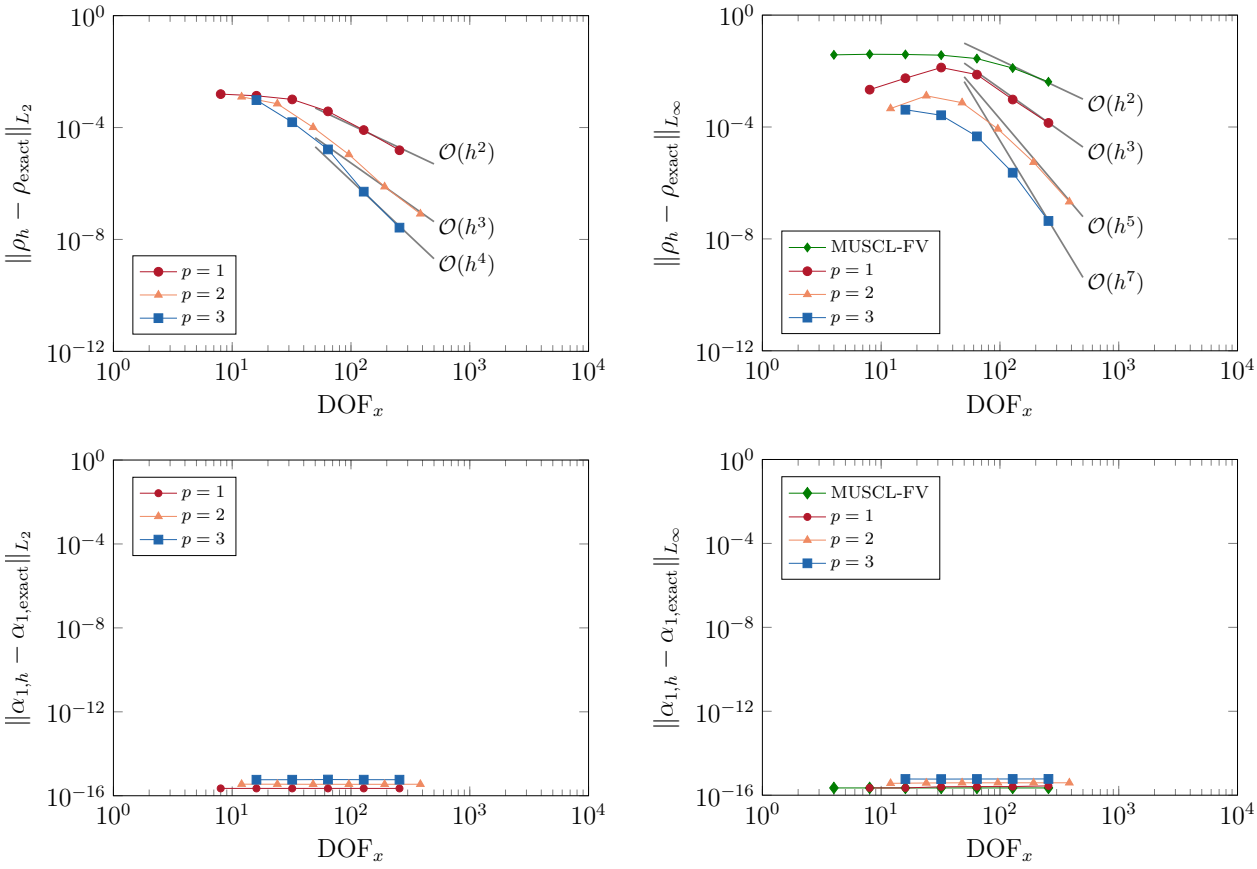


Figure 3.3: L_2 (left column) and L_∞ (right column) error norms of the mixture density ρ (top row) and the volume fraction α_1 (bottom row) for the isentropic vortex advection problem. Grey guidelines correspond to the expected order of accuracy for each p and are labeled as such.

number of cells along an axis for finite volume. We observe the expected convergence rates of $p+1$ (expected from the projection error) and $2p+1$ in the mixture density for all orders of p tested in the L_2 and L_∞ norms, respectively. The finite volume scheme converges quadratically, as expected. Additionally, the volume fraction error remains at round-off in both norms for all mesh resolutions and all p, as well as for the finite volume scheme. This test verifies that the DG scheme converges at the appropriate rate, the volume fraction equation is exact away from interfaces, that the projection of α_1 and reconstruction of ∇Q for computing the Phase-Field fluxes do not introduce errors, and that the Phase-Field solver itself does not introduce any measurable errors in pure material regions, even in the presence of velocity gradients.

3.5.2 Advection of an Air-Water Interface

To verify that the proposed scheme does not generate spurious oscillations across material interfaces and maintains a thin interface, we advect an interface between air and water ($A_1 = 5/2$, $B_1 = 0$, $D_1 = 0$, $A_2 = 5/17$, $B_2 = 132/17 \times 10^8$ Pa, $D_2 = 0$ in Eq. (3.4)) in a uniform velocity and pressure field, similar to [11]. In this case, the volume fraction is no longer constant, so the Phase-Field fluxes act to keep the interface at its equilibrium thickness. The initial conditions are,

$$(\rho_1, \rho_2, u, v, P, \alpha_1) = \left(1, 1000, 753.273, 753.273, 101325, \frac{1}{2}\left[1 - \tanh\left(\frac{I}{\eta_0\sqrt{2}}\right)\right]\right), \tag{3.59}$$

where $I = 2 - \sqrt{x^2 + y^2}$ and $\eta_0 = h(p+1)^{-1}/\sqrt{2}$. The problem is defined on $[x, y] \in [-2, 2]^2$ with doubly-periodic boundary conditions, and is evolved until $t = 4$. Three schemes are used to evolve the flow, all with $p = 3$: DG with a high-order limiter [275] (DG+HRL) and DG with the proposed subcell scheme and a quartic WENO reconstruction [111], without (DG+SC) and with (DG+SC+PF) Phase Fields. Each scheme is tested with 128×128 DOF; the Phase-Field discretization is also tested with 256×256 DOF. The runtimes of each scheme compared to the subcell scheme are are $t_{\text{DG+SC+PF}}/t_{\text{DG+SC}} = 1.92$ and $t_{\text{DG+HRL}}/t_{\text{DG+SC}} = 1.22$, demonstrating that the subcell scheme is more efficient than the limiting scheme, and the Phase-Field mechanism approximately doubles the cost.

Figure 3.4 shows the time evolution of the maximum error in velocity, pressure, and temperature, as well as the volume fraction conservation error, interface thickness, and number of flagged cells (for limiting or subcells). Since the velocity is expected to remain constant for all time in this problem, volume fraction must be conserved; this property is not true in general. The two subcell schemes produce no spurious errors in the velocity, pressure, or temperature beyond round-off level, which is not trivial [11, 156, 265]. This property is important for schemes that solve Eq. (3.1) as generating spurious velocities or pressure errors results in the activation of the non-conservative term in the volume fraction equation in a region of what should be uniform flow and thus acts as a source (or sink) for α_1; as a result, α_1 would not be conserved as would be expected for a solenoidal velocity field such as a uniform flow [265]. Avoiding spurious oscillations is also important to avoid prematurely triggering fluid instabilities. The high-order limiting scheme, in contrast, generates errors above machine precision in each of the velocity, pressure, and temperature, which is expected without special treatment of the conservative variables during limiting [91]. We verify that the proposed subcell schemes are conservative to machine round-off, where the conservation error is defined,

$$\epsilon_{\alpha_1} = \left| \sum_{e=1}^{N_e} \left(\int_{\Omega_e} \alpha_1 d\Omega - \int_{\Omega_e} \alpha_1(t=0) d\Omega \right) \right|. \tag{3.60}$$

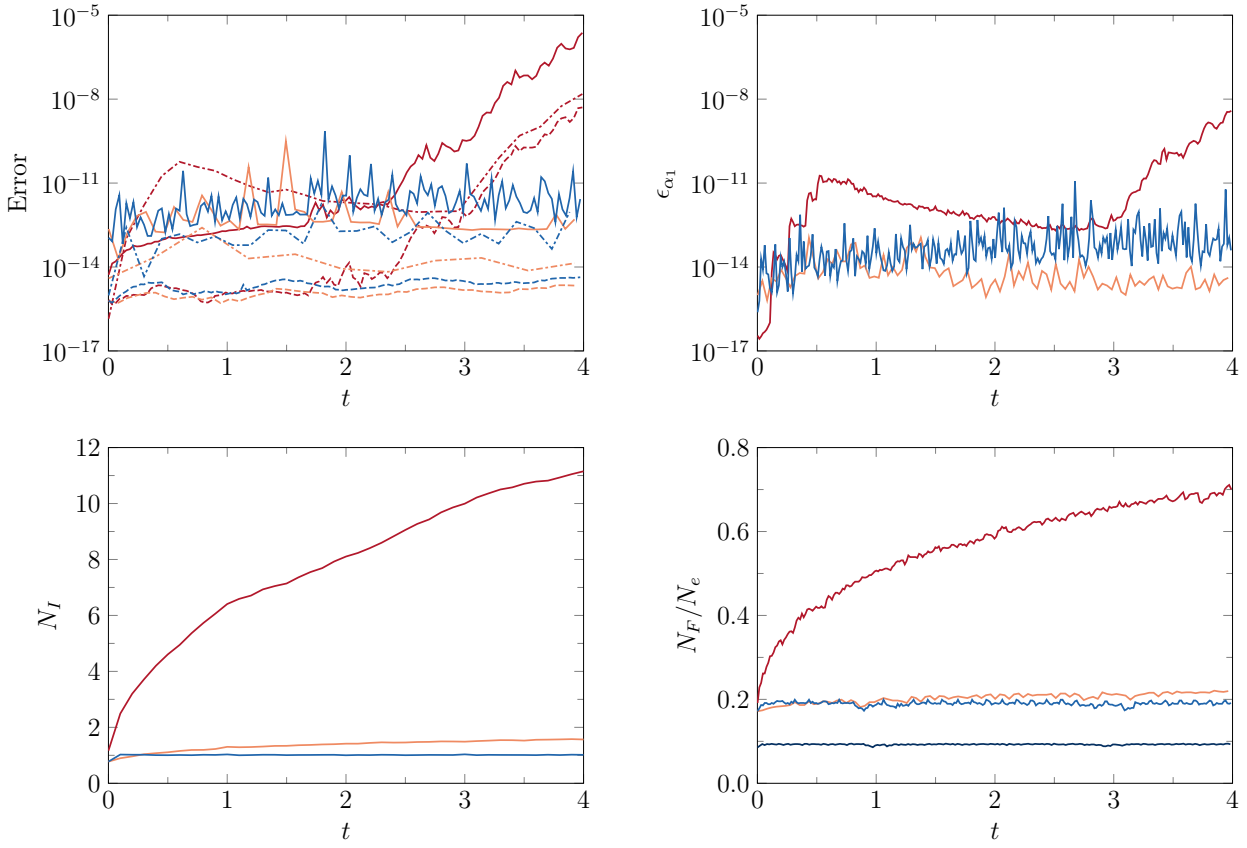


Figure 3.4: Top left: time evolution of the maximum error in velocity (solid), pressure (dashed), and temperature (dash-dotted); top right: volume fraction conservation error; bottom left: interface thickness; and bottom right: number of flagged cells for limiting or subcells for the interface advection problem. Colors represent each scheme: DG+HRL (red), DG+SC (orange), DG+SC+PF (blue), and DG+SC+PF (dark blue, 256×256), bottom right only.

Clearly, the subcell and the Phase-Field schemes are both non-oscillatory at interfaces and are conservative in α_1 in this special case, whereas the high-order limiting scheme is neither.

To see how well each scheme is able to resolve the interface over the course of the simulation, we examine the time evolution of the interface thickness. As expected, the scheme with the Phase-Field mechanism is able to maintain an essentially constant interface thickness of approximately one DG cell throughout the duration of the simulation, with only very slight deviation from the equilibrium. With the other two schemes (DG+HRL and DG+SC), the interface continually diffuses over time. The high-order limiting scheme produces the thickest interface with approximately eleven cells across the thickness by the end of the simulation, with continual growth. The subcell scheme produces a substantially thinner interface; however, it too grows continually over time. As the interface thickens, the number of cells that trigger the interface sensor potentially grows as well, which we quantify as the ratio of the number of flagged cells at the end of a time step to the number of total cells. As expected, the number of cells flagged by each scheme largely matches the trends in the interface thickness, with the limiting scheme flagging an increasing number of cells and the Phase-Field scheme flagging an approximately constant number, which is also the fewest. With twice the DOF in each direction (i.e., for the 256×256 case), the proportion of cells flagged decreases, demonstrating the resolution capabilities of the proposed scheme (DG + SC + PF), as the fraction of cells evolved with the high-order DG scheme (rather than the subcell finite volume scheme) is greater.

3.5.3 Multiphase Titarev-Toro Shock Tube

To test the ability of our proposed scheme to accurately resolve interfacial flows with both smooth and sharp features, we simulate a shock in one substance impinging upon a high-frequency density perturbation in a second substance, in a modification of the Titarev and Toro problem [237]. Previous work considered a multiphase extension of the Shu-Osher shock tube to demonstrate capturing of sharp and smooth features [91], but the Titarev-Toro shock tube contains more delicate flow structures that are more challenging to capture. A Mach 1.1 shock is initialized in air ($A_1 = 5/2$, $B_1 = 0$, $D_1 = 0$) and impinges upon a stationary Mie-Grüneisen material ($A_2 = 5/3$, $B_2 = 4/15$, $D_2 = 1/10$). The initial conditions are given by,

$$(\rho_1, \rho_2, u, P, \alpha_1) = \begin{cases} (1.515695, 0, 0.523346, 1.805, 1), & x < 1/2 \\ (0, 1 + \frac{1}{10}\sin 20x, 0, 1, 0), & x \geq 1/2. \end{cases} \tag{3.61}$$

The problem is defined on $x \in [0, 10]$, with outflow boundary conditions on both ends. The flow is evolved until time $t = 5$. This problem is particularly challenging due to the presence of high-frequency features interacting with a material interface and a shock wave. We test the

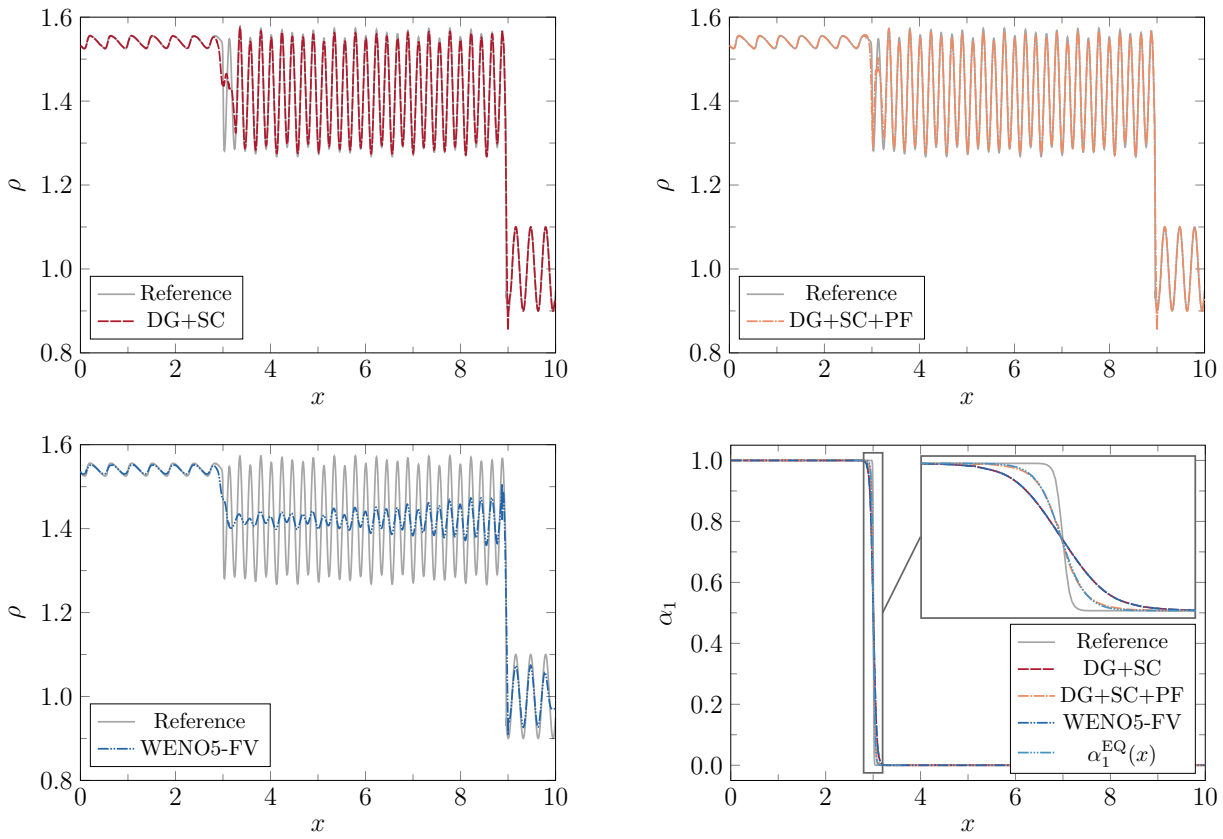


Figure 3.5: Density and α_1 fields at $t = 5$ for the multiphase-adapted Titarev and Toro shock-tube problem for the DG+SC, DG+SC+PF, and fifth-order finite volume schemes with 400 DOF. The reference simulation (2560 DOFs) is shown in gray in all cases.

subcell schemes described in Section 3.5.2 with $p = 4$, as well as a finite volume scheme with a fifth-order-accurate quartic WENO reconstruction (WENO5-FV). The HLL Riemann solver is used with all schemes for this case. Each solution is computed with 400 DOF, and a reference solution is computed with 2560 DOF. The runtimes of each scheme compared to the subcell scheme are $t_{\text{DG+SC+PF}}/t_{\text{DG+SC}} = 2.25$ and $t_{\text{WENO5-FV}}/t_{\text{DG+SC}} = 0.98$.

Figure 3.5 shows the flow density and a close-up view of the volume fraction at $t = 5$ for each solution. The shock interacts with the material interface and propagates in the positive x-direction, while the perturbation travels in the opposite direction. Downstream of the interface, the results from the two DG schemes are essentially identical, as expected due to consistency of reduction. Likewise, the two DG schemes produce very similar results upstream of the interface and only differ noticeably near the interface. The fifth-order WENO solution exhibits significantly more numerical diffusion than the DG schemes upstream and downstream of the interface.

Near the interface, the DG+SC scheme and the finite volume scheme produce essentially identical interfaces, as is expected due to the reconstruction method on the subcells scheme matching that

of the finite volume scheme. The DG+SC+PF scheme produces the sharpest interface. Observing the density near the interface, the three schemes produce noticeably different results, with the DG+SC+PF scheme producing the most accurate results, followed by the DG+SC scheme, then the finite volume scheme. Clearly, the Phase Field not only maintains a sharper interface than the standard DG and finite volume schemes, but it also reduces the error near the interface. Errors in the volume fraction field manifest in the other state variables via the mixture pressure; thus, by reducing the volume fraction error, the overall error is reduced in mixture regions. Another advantage of the Phase-Field scheme for DG is that by keeping the interface sharper, the number of cells required to regularize the interface is minimized, in contrast to schemes where the interface is allowed to thicken over time, which results in a corresponding increase in the use of the discontinuity-capturing scheme, as noted in Section 3.5.2. Importantly, the location of the interface is well predicted after including the Phase-Field mechanism, and agrees well with the reference solution. This test demonstrates that the Phase-Field mechanism not only preserves delicate flow features that interact with the interface (e.g., waves) but also decreases the error in that region. Finally, the strong agreement between the DG+SC+PF volume fraction field and the analytic equilibrium volume fraction distribution α_1^{EQ} in Fig. 3.5 clearly shows that the Phase-Field system is being solved faithfully and accurately by the proposed DG scheme.

3.5.4 Two-Dimensional Richtmyer-Meshkov Instability

For validation, we simulate the single-mode Richtmyer-Meshkov instability experiments of [52], which were also studied computationally by [91]. The flow is initialized with a Mach 1.21 shock in air ($\rho_{\text{air}} = 1.351\,\text{kg/m}^3$, $A_{\text{air}} = 3.623$) propagating towards a sinusoidal perturbation along the interface with sulfur hexafluoride ($\rho_{\text{SF}_6} = 5.494\,\text{kg/m}^3$, $A_{\text{SF}_6} = 10.753$). The perturbation has initial amplitude $\mathcal{A}_0 = 0.1833\,\text{cm}$ and wavelength $\lambda_0 = 5.933\,\text{cm}$. The problem is defined on $(x, y) \in [0, \lambda_0/2] \times [-8\lambda_0, 4\lambda_0]$, with symmetry boundary conditions at $x = 0$ and $x = \lambda_0/2$, and outflow elsewhere. The flow is simulated until $t = 7\,\text{ms}$. The volume fraction is initialized in a thermodynamically consistent manner according to,

$$\alpha_1 = \frac{1}{2}\left[1 - \tanh\left(\frac{\mathcal{A}_0 \cos(2\pi x/\lambda_0) - y}{\eta_0\sqrt{2}}\right)\right], \tag{3.62}$$

where $\eta_0 = h/\sqrt{8}$, such that the interface initially comprises the same number of grid cells at each resolution. We simulate the described flow with a $p = 2$ basis at several resolutions $\lambda_0/h \in [8, 128]$ and measure the perturbation amplitude, as well as the interface thickness according to Eq. (3.55). The interface is identified by the contour line $\alpha_1 = 0.5$, however, we note that the interface location is not a point, contour line, or iso-surface in one, two, and three dimensions respectively in numerical

solutions of Eq. (3.1), but is instead a width, surface area, and volume; here the material properties are mixed over a non-zero span. The three schemes applied in Section 3.5.2 are used here, with $p = 2$. The runtimes of each scheme compared to the subcell scheme are $t_{\text{DG+SC+PF}}/t_{\text{DG+SC}} = 2.55$ and $t_{\text{DG+HRL}}/t_{\text{DG+SC}} = 1.15$.

The time evolution of the perturbation amplitude is shown in Fig. 3.6 for each scheme, as well as contours of α_1 at $t = 5.2\,\text{ms}$ for different resolutions. The shock impinges upon the interface and deposits baroclinic vorticity that causes the initial perturbation to grow linearly at first, then non-linearly as the characteristic bubble and spike develop. All three schemes converge towards the experimental amplitude data monotonically with increasing resolution. There is also good agreement between all three schemes and the experimental data at the highest resolution. The DG scheme with limiting (DG+HRL) is outperformed by the two other schemes, particularly at small λ_0/h. It is not until $\lambda_0/h = 64$ that the scheme with limiting can be considered close to the experimental data, and only with $\lambda_0/h = 128$ is the agreement acceptable. In contrast, the DG+SC scheme is able to reach better agreement at $\lambda_0/h = 32$ than the scheme with limiting does at $\lambda_0/h = 64$. Additionally, beyond $\lambda_0/h = 8$, the bulk shape of the bubble and spike is quite consistent across resolutions by the DG+SC scheme, whereas the scheme with limiting generates very different bubble and spike shapes with varying h. Both schemes without the Phase-Field mechanism exhibit non-uniform diffusion due to the non-uniform velocity gradients, particularly in the roll-up region. The DG+SC+PF scheme not only converges towards the experimental data the fastest, but it also generates the most consistent bubble and spike shapes across resolutions. This behavior is consistent with that observed in Section 3.5.3; the scheme with the Phase-Field mechanism is more accurate near the interface, which results in lower error. All three methods agree well with the experimental data at high resolutions. For more complex flows such as a multi-mode Richtmyer-Meshkov instability, accurately representing fine-scale features is important to predict even bulk quantities (e.g., mixing). This regime is where the scheme with the Phase-Field mechanism excels, as it not only resolves fine-scale features with superior accuracy but also facilitates consistent identification of the interface across all scales and resolutions. With the Phase-Field mechanism active, the roll-up may eventually break up at points where the ligament becomes thinner than the interface, which is likely due to the scale invariance of the Euler equations. In reality, mechanisms like viscous stresses suppress the roll-up growth; increasing the resolution would not cause the ligaments to indefinitely become increasingly thin.

The time evolution of the interface thickness for these three schemes is shown in Fig. 3.7. The measured thicknesses are consistent with expectations, with the Phase-Field results generating both the thinnest and most consistent interface thickness over time. The high-order scheme with limiting gives rise to the thickest and most inconsistent interface. The limiter is very active at the beginning of the simulation, until the interface is diffuse enough to stop triggering the sensors. After this point,

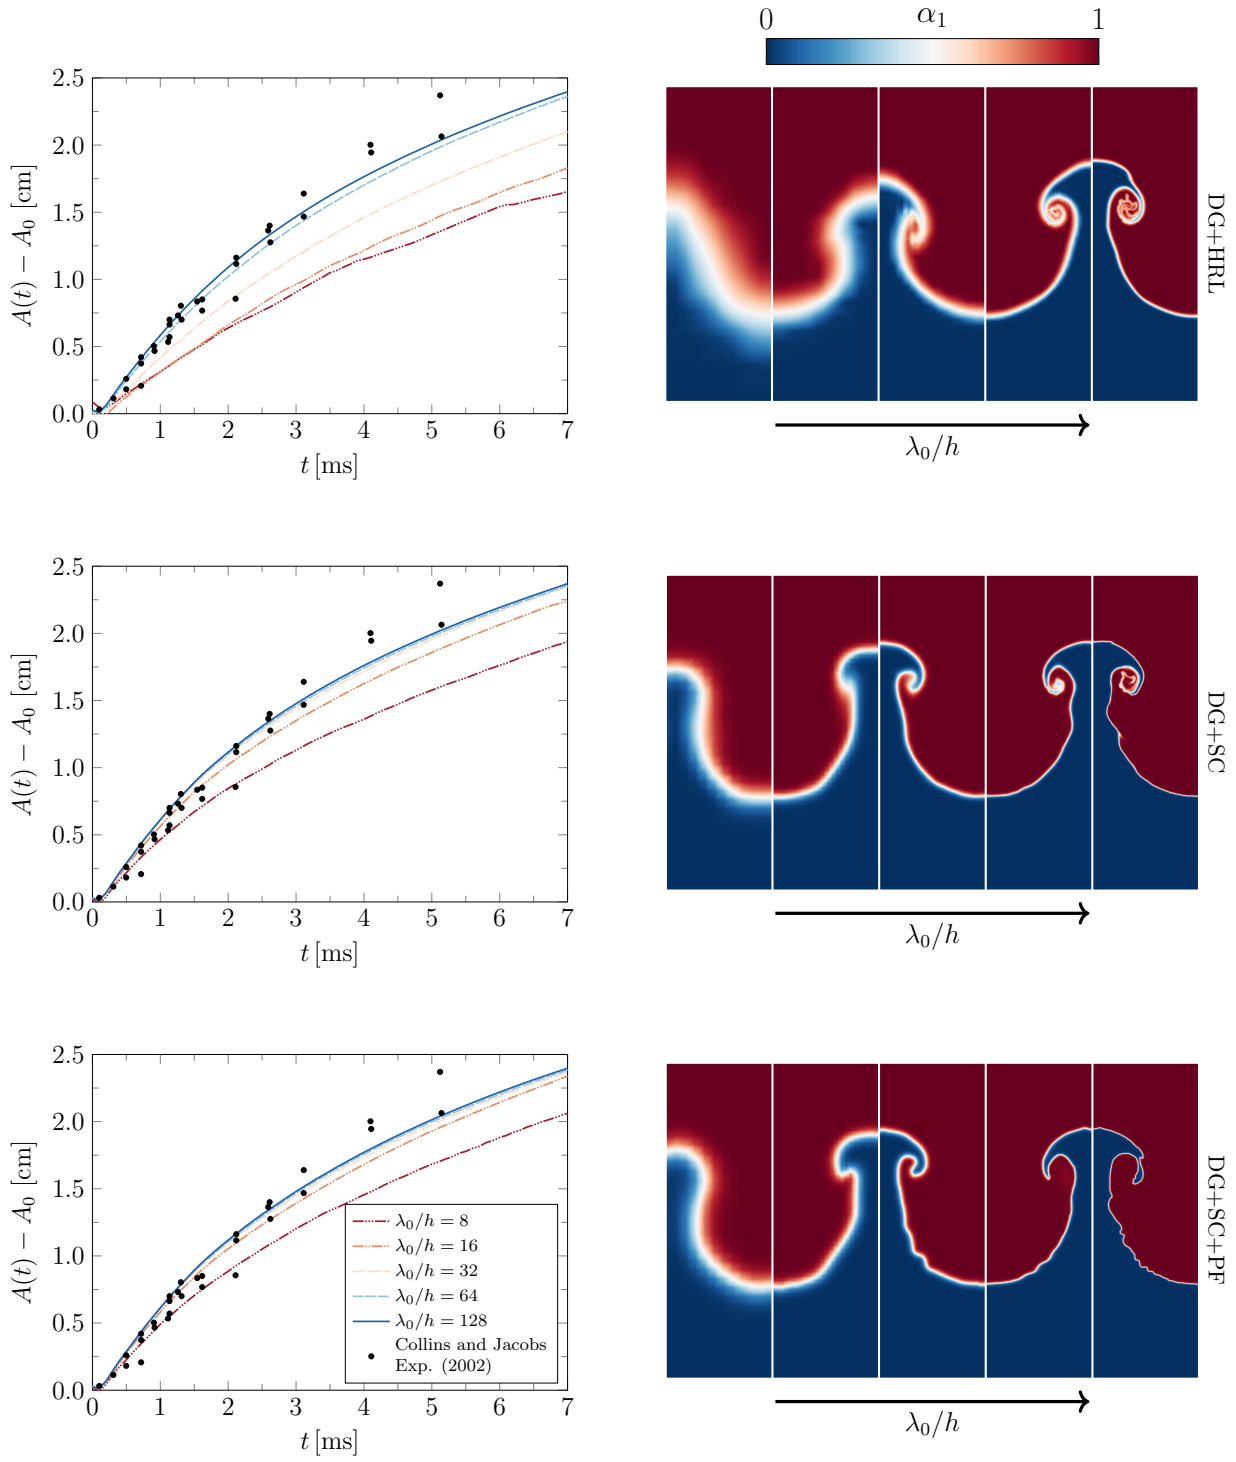


Figure 3.6: Two-dimensional Richtmyer-Meshkov instability for the DG+HRL, DG+SC, and DG+SC+PF schemes for $\lambda_0/h \in [8, 128]$. Left: Time evolution of the perturbation amplitude for different resolutions. Right: Contours of α_1 for each resolution and scheme at $t = 5.2$ ms. Odd-numbered contours are flipped horizontally to show the difference in vertical positions of the bubble or spike between resolutions.

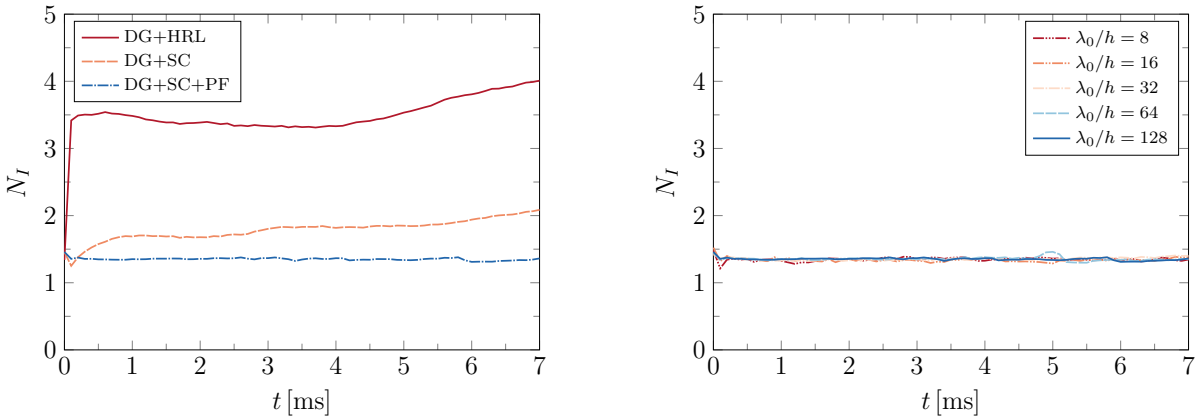


Figure 3.7: Time evolution of the interface thickness for the two-dimensional Richtmyer-Meshkov instability problem for $\lambda_0/h = 128$. Left: comparison between the DG+HRL, DG+SC, and DG+SC+PF schemes. Right: dependence on resolution.

there are between three and four DG cells in the interface, with the number growing substantially by the end of the simulation. The DG+SC scheme clearly generates a thinner interface than the DG+HRL scheme, with a maximum of approximately two DG cells in the interface at the end of the simulation. However, as with the DG+HRL scheme, this number continues to grow with time. The initial decrease in the interface thickness with this scheme is due to the impingement of the shock, which compresses the initially diffuse interface, after which point the thickness increases almost monotonically. Across all resolutions, the interface is consistently captured by just over one DG cell with the Phase Field. While the interface thickness is not identically constant in this case, the deviation from the equilibrium thickness is small and may be an artifact of transient flow features. Figure 3.7 also shows the time evolution of N_I for different resolutions with the DG+SC+PF scheme. As the resolution is increased, the interface thickness decreases proportionally, but the number of subcells across the interface remains constant. This result is important and expected, as in multi-scale simulations, the interface is always captured at the smallest allowable scale for a given resolution. While η could be chosen such that the interface is a constant thickness across resolutions (i.e. $\eta \neq \eta(h)$), the interface is captured over a consistent number of DOF for any h by making it directly proportional to the grid size. Choosing $\eta \ll h$ is allowed under the proposed scheme with an appropriate time step (tested in [93]); however, this value implies that the Phase-Field model is a sub-grid model whose impact has yet to be fully examined.

3.5.5 Underwater Explosion Problem

Finally, we consider the impulsive expansion of an air bubble in water near a free surface [20, 156, 222]. This case demonstrates the ability of the proposed scheme to capture large pressure, density,

and velocity gradients in a flow with disparate equation of state parameters, and still maintain a uniform interface. Additionally, this case is useful for studying grid-dependent effects with and without the Phase-Field mechanism. The initial conditions are given by

$$\begin{bmatrix} \rho_1 \\ \rho_2 \\ u \\ v \\ P \\ \alpha_1 \end{bmatrix} = \begin{bmatrix} \rho_a \mathcal{H}(y) + \rho_b (1 - \mathcal{H}(y)) \\ 1 \\ 0 \\ 0 \\ \left[\frac{1}{2} \tanh\left(\frac{I_1}{\eta_0 \sqrt{2}} \right) + \frac{1}{2} \right] (P_1 - P_2) + P_2 \\ \frac{1}{2} \tanh\left(\frac{\max(I_1, I_2)}{\eta \sqrt{2}} \right) + \frac{1}{2} \end{bmatrix} \tag{3.63}$$

with α_1 representing the air ($A_{\text{air}} = 5/2$) and α_2 representing water ($A_{\text{water}} = 5/17$, $B_{\text{water}} = 7764.71$), $\mathcal{H}$ is the Heaviside function, $\rho_a = 0.001225$, $\rho_b = 1.25$, $I_1 = 0.12 - \sqrt{x^2 + (y + 0.3)^2}$, $I_2 = y$, $P_1 = 10^4$, $P_2 = 1.01325$, and $\eta_0 = 1/320$, where all parameters have been non-dimensionalized. The flow is simulated on $(x, y) \in [-2, 2] \times [-3/2, 3/2]$ with wall boundary conditions on all sides, and is evolved until $t = 0.019$. The DG+SC and DG+SC+PF schemes are tested for $h \in \{1/10, 1/80\}$ and $p = 3$. The runtime of the Phase-Field scheme at the base resolution compared to the subcell scheme is $t_{\text{DG+SC+PF}}/t_{\text{DG+SC}} = 2.04$.

Contours of the Mach number and α_1 for two resolutions, $h = 1/10$ and $h = 1/80$, are shown with and without Phase Fields in Fig. 3.8. The high-pressure air rapidly expands, with an upward directional bias due to the initially compressive waves reflecting off the free surface as rarefactions and thus causing a low-pressure region above the bubble. The shock from the expansion is transmitted through the free surface, while subsequent reflections cause a series of waves in the air above the free surface. As the bubble grows, a thin film of water forms between the bubble and the air above, whose thickness decreases over time. The expected physical solution is that the film continually thins as the bubble expands, eventually breaking due to film instabilities. However, as effects like surface tension are not modeled here, the film should perpetually thin until its length scale is too fine to be resolved on the grid.

Between the higher-resolution cases ($h = 1/80$), the volume fraction and Mach contours are very similar, with notable differences only manifesting near the interface, thereby indicating that when structures can be resolved on the grid, the Phase-Field mechanism only serves to keep the interface uniformly thick and does not impact the material distribution otherwise. This behavior is expected, as the Phase-Field mechanism is guaranteed to only be active near the interface to satisfy consistency of reduction and has a strong effect when the interface is very far from equilibrium, which is only the case when a feature is not well-resolved on the grid. Given their agreement, these

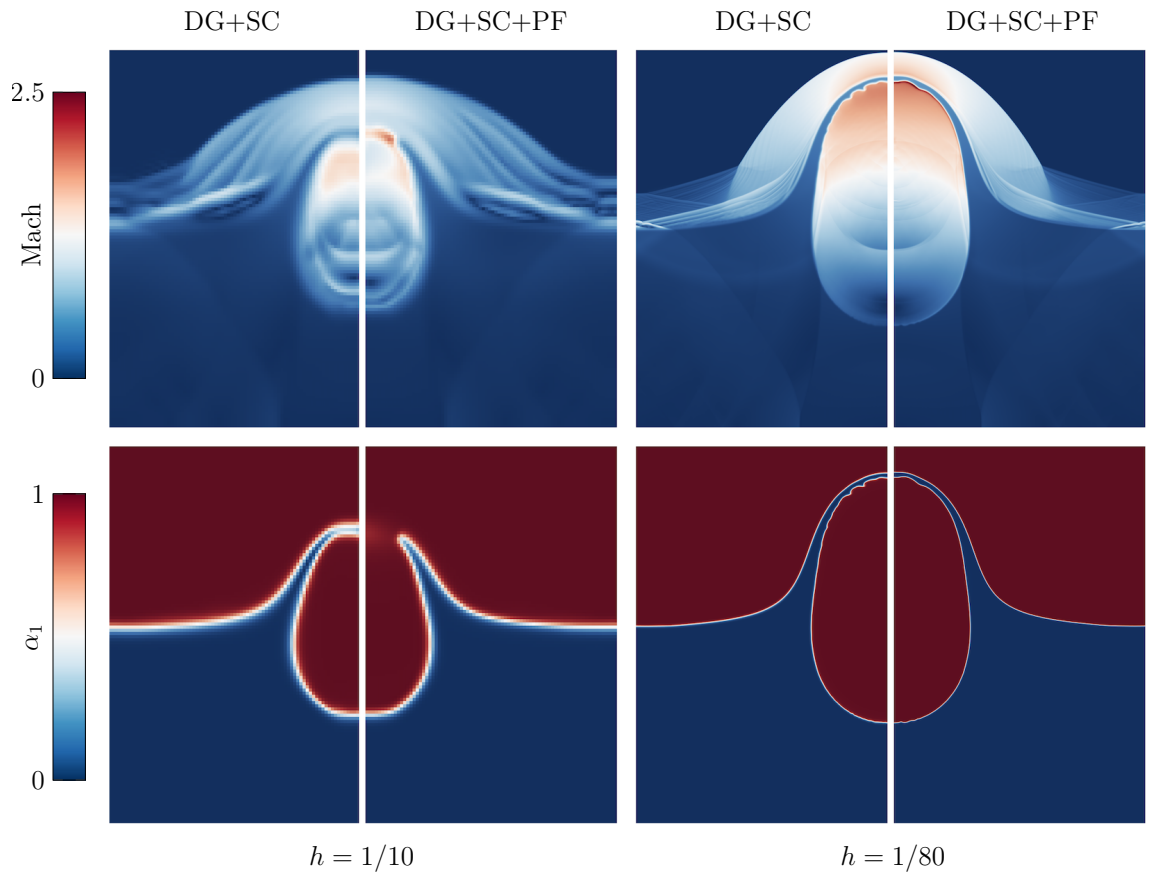


Figure 3.8: Contours of Mach number and α_1 for the underwater explosion problem for two mesh resolutions with the DG+SC and DG+SC+PF schemes at $t = 0.019$.

higher-resolution cases demonstrate that if a film is present at the top of the bubble, it is made of pure water, not a mixture.

The behavior of the Phase Field in the case of under-resolved features is examined with the lower resolution cases ($h = 1/10$), in which the volume fraction and Mach contours exhibit significant differences. Specifically, without the Phase-Field mechanism, the film at the top of the bubble is present and the bubble is not punctured, whereas with the Phase-Field mechanism active, the bubble is punctured and the upper film is not present. The composition of the top of the film without the Phase-Field mechanism, however, is not pure water as the more highly resolved cases would suggest, and is instead a mixture with a volume fraction close to half. In contrast, with the Phase-Field mechanism active, while the upper film is not present and the bubble is punctured, the remaining ligament is composed of pure water away from the interface. Clearly, the Phase-Field mechanism is acting to satisfy the model assumption of pure materials, at the cost of puncturing the bubble. This observation highlights the key difference in the modeling of under-resolved features between Eq. (3.1) and Eq. (3.10); Eq. (3.1) sacrifices immiscibility to retain under-resolved structures and assumes the modeling error in the mixture region, whereas Eq. (3.10) sacrifices under-resolved features to preserve immiscibility and maintain a uniform interface. Away from interfaces, the models are identical; near features that are well-resolved, the interface topology is not impacted by the Phase-Field mechanism beyond its thickness. The variation in the final bubble size between different resolutions has been previously observed [222], but its cause is beyond the scope of this work. The convergence of the two schemes towards a unique solution belies the consistency of Eq. (3.10) with Eq. (3.1) in the limit of a sharp interface.

3.6 Conclusions

In this work, we develop a high-order discontinuous Galerkin approach with a non-conformal subcell finite volume scheme to solve the general consistent and conservative Phase-Field model [93] for numerical simulations of compressible interfacial flows with shocks. This work both incorporates the Phase-Field formulation into the DG framework and extends the Phase-Field formulation to high-order accuracy.

The scheme consists of splitting the governing equations into two steps: a hyperbolic step that updates the solution without interface regularization, and a Phase-Field step driving the interface to its equilibrium distribution as dictated by the chosen Phase-Field model. Each step is discretized with an explicit Runge-Kutta method in time, and a discontinuous Galerkin scheme with non-conformal subcell finite volume discontinuity capturing in space. The first step solves the five-equations compressible multiphase model without Phase Fields in a high-order accurate, conservative, and well-balanced manner. We demonstrate the importance of the well-balanced

property for the volume fraction equation in ensuring *consistency of reduction* and show how other schemes may fail this condition. We also present a time-step constraint that provably bounds the volume fraction cell-average under the chosen approximation. The subcell discontinuity-capturing scheme is activated in cells triggered by a set of sensors, and invertible projection and reconstruction operators ensure conservation and accuracy when transferring the solution between the grid and subgrid. The second step drives the interface to its equilibrium distribution by calculating an auxiliary variable using the subcell volume fraction values and reconstructing high-order polynomials to compute the Phase-Field mechanism and integrate the associated fluxes. Special care is taken in computing the Phase-Field fluxes to ensure accuracy and conservation. Two additional methods for computing the Phase-Field fluxes are proposed, each with clear advantages and disadvantages. A suite of one- and two-dimensional multicomponent and multiphase problems is used to verify the analysis and code, and for validation.

A discontinuity sensor identifies regions where the subcell finite volume scheme is applied; other regions are treated with high-order DG. The proposed scheme is arbitrarily high-order accurate away from discontinuities, is well-balanced for the volume fraction equation, captures material interfaces on the sub-DG-grid scale, maintains material interfaces of uniform and constant thickness, satisfies *consistency of reduction*, and can model a variety of materials through a general equation of state. Consistent with the behavior reported in [93], including the Phase-Field mechanism gives rise to a smaller error in all flow variables near interfaces. The DG and non-conformal subcell finite volume schemes enable highly local, sub-DG-grid discontinuity capturing that dramatically outperforms high-order limiting [275] and does not generate spurious oscillations at material interfaces thanks to appropriate finite volume reconstructions [93].

The proposed model and scheme are able to capture both smooth and sharp features with fewer degrees of freedom than second-order schemes, and are well-suited for multi-material simulations with complex flow physics. Extensions to include additional physics such as viscous stresses, heat transfer, and surface tension are expected to require only small modifications. Phase-Field modeling of N-materials under the current framework is a logical next step that began in [94] and is ongoing work.

Part II

Numerical Simulations of Compressible Interfacial Flow Physics

CHAPTER 4

Emergence of Three-Dimensional Structures from Vortex-Pair Instabilities in Shocked Interfacial Flows

4.1 Abstract

Numerical simulations of the Richtmyer-Meshkov instability (RMI) with a transverse perturbation are conducted to investigate the transition from a nominally two-dimensional flow to three dimensions. A planar shock impinges upon a perturbed notch along the interface between two gases, leading to perturbed counter-rotating vortex lines. We show that the growth rates of perturbations in the transverse direction follow the behavior described by the Crow vortex-core instability rather than only the transverse RMI, with a distinctive feature being the formation of three-dimensional vortex rings. As such, classical vortex-core instabilities occur and may dominate the formation of three-dimensional structures, and possibly turbulence, in RMI flows.

4.2 Introduction

Interfacial mixing induced by shock waves plays a critical role in high-energy-density (HED) flows, inertial confinement fusion (ICF) experiments, and astrophysical systems. In ICF, shock-induced instabilities, seeded by defects, fill tubes, and assembly joints in single- and double-shell systems, degrade the capsule symmetry and reduce the potential fusion gain [56, 144, 216]. In astrophysical systems, similar mechanisms alter the mixing and ejection of materials in supernova remnants, influencing the distribution of elements in space [119, 209]. Studying these systems in gaseous experiments is not trivial due to the complex configurations that govern the relevant dynamics, so HED platforms with machined targets are often used at facilities such as the National Ignition Facility or the Omega EP laser [256].

Among interfacial instabilities, the Richtmyer-Meshkov instability (RMI) stands out due to the presence of shocks in these systems. The RMI occurs when a shock passes through a perturbed interface between fluids of different densities, where the misalignment of the density and pressure gradients leads to the deposition of baroclinic vorticity [162, 201]. This transient acceleration of the interface leads to initially linear growth of perturbations, with the late-time behavior characterized by interpenetrating bubbles and spikes of the light and heavy fluids, accompanied by rolled-up regions of high vorticity that manifest as vortex cores. The generation of these flow structures can alter mixing and potentially turbulent transition [52, 149, 150, 151, 208], and various studies have investigated different aspects of the two-dimensional RMI, including the dynamics of single-mode [193] and multi-mode [257] interfaces. The mechanisms by which these nominally two-dimensional flows become three-dimensional and eventually turbulent are still not completely understood, with geometric effects in three-dimensional systems, including reshock in converging geometries, further complicating the interfacial mixing dynamics and turbulent transition [61, 131, 153, 166]. The full spectrum of pathways between nominally two-dimensional RMI flows and three-dimensional structures remains poorly characterized, with recent work suggesting the presence of vortex-core instabilities in these compressible, multi-material flows [256]. Classical vortex instabilities such as the elliptical instability [123], the Widnall instability [267], and the Crow instability [57], deform perturbed vortex cores and break them down into secondary vortex structures, and are linked to the onset of turbulence [161]. These instabilities are typically studied in incompressible, single-fluid systems; their emergence in more complex compressible, multi-material flows like RMI is largely unexplored.

In this study, we hypothesize that long-wavelength vortex-pair instabilities, in particular the Crow instability (CI), may be triggered in nominally two-dimensional RMI flows with transverse perturbations, and may dominate the formation of three-dimensional structures. We test this hypothesis using three-dimensional simulations, with measurements of the resultant instability

dynamics quantifying their role in structure formation. We propose that transverse perturbations to a nominally two-dimensional RMI flow give rise to coherent perturbed vortex cores, whose disturbances grow exponentially due to vortex-core instabilities, leading to the formation of secondary vortex structures. Our simulations test this hypothesis across multiple Atwood numbers and transverse wavelengths, allowing for direct comparison with the classical CI growth rates. When non-dimensionalized, a characteristic non-monotonic dependence emerges between the growth rates and wavenumbers, with well-defined stability thresholds and magnitudes similar to theory. These results demonstrate that vortex-pair instabilities are a viable pathway to three-dimensional structure and mixing in RMI flows, a mechanism with direct relevance to ICF target performance and astrophysical mixing.

The remainder of this work proceeds as follows: Section 4.3 describes the setup of the system and the numerical methods. Results, including a qualitative description of the dynamics, the growth-rate measurements, and comparisons with theory, are given in Section 4.4. Section 4.5 concludes and discusses future directions.

4.3 Setup and Methods

Figure 4.1 shows the problem set-up. We initialize a Mach 1.21 shock in a heavy fluid propagating at normal incidence toward an interface with a light fluid. The initial perturbation along the interface consists of a rectangular notch of height $a_{0,z} = 0.15b$ and width $b = 0.4$; this geometry, readily manufacturable in HED experiments, enables control of the spacing between the neighboring resulting vortex cores. A single-mode perturbation of initial amplitude $a_{0,y} = 0.1b$ is added to the notch in the transverse direction. The initial interface between the two materials is thus described by,

$$\eta(x, y) = a_{0,z} \left\{ \mathcal{H}\left[y + a_{0,y} \cos(k_x x) + \frac{b}{2} \right] - \mathcal{H}\left[y - a_{0,y} \cos(k_x x) - \frac{b}{2} \right] \right\}, \tag{4.1}$$

where $a_{0,z}$ is the initial perturbation amplitude in the z direction, $\mathcal{H}$ is the Heaviside function, $k_x = 2\pi/\lambda_x$ is the transverse perturbation wavenumber, $a_{0,y}$ is the initial perturbation amplitude in the y direction, and x_0 is the x-y plane; there is no perturbation in x. With this configuration, we achieve clean, planar RMI growth and have control over the Mach number, Atwood number, transverse wavelength, and core separation.

Both fluids are ideal gases with equation of state parameters $\{\gamma_p\}_{p=1}^2 = \{1.093, 1.276\}$ and nominal unshocked densities $\{\rho_p\}_{p=1}^2 = \{5.494, 1.351\}$. The Atwood number is varied by varying the heavy gas density $\rho_1 \in \{2.747, 5.494, 10.988\}$, corresponding to At $\in \{0.34, 0.61, 0.78\}$. The largest of these Atwood numbers is similar to that in [256] (At ≈ 0.82), wherein a similar experimental setup was proposed for the Omega EP laser facility.

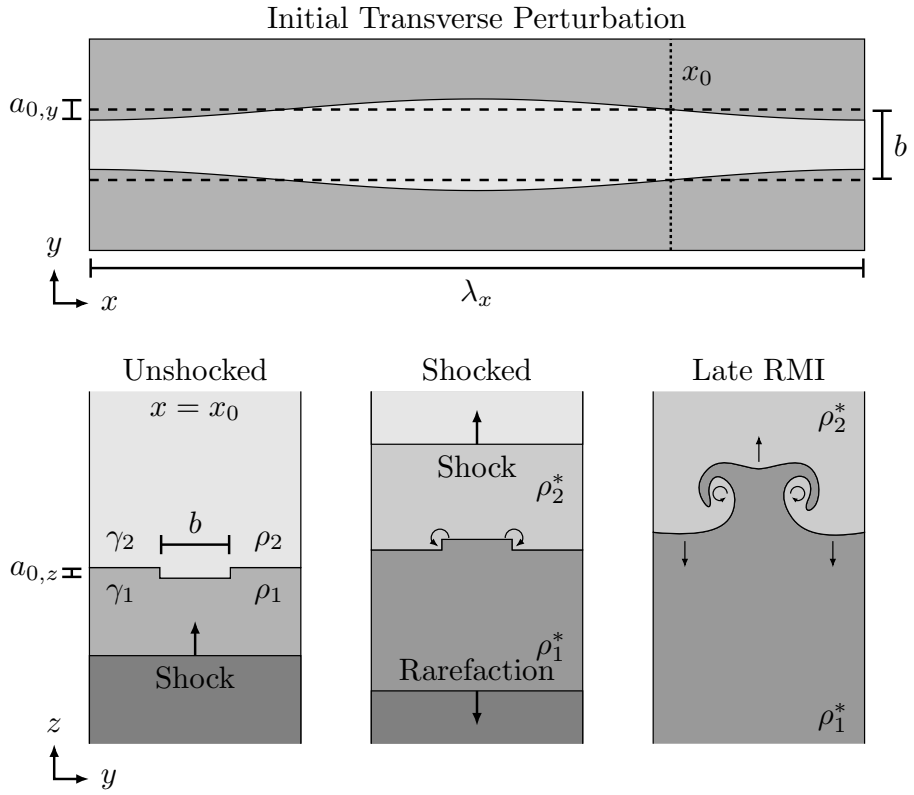


Figure 4.1: Two-dimensional schematics of the x-y (top) and y-z (bottom) planes. The $y-z$ plane is shown at the initial time, just after the shock impinges upon the interface, and after the cores are developed (late-time RMI).

This problem is simulated using the COMPAS code (Chapter 2), specifically the multi-material Euler solver. The spatial discretization is a second-order finite volume method with the MUSCL reconstruction [250] and the HLLC Riemann solver [243]. Temporal integration is performed using a third-order accurate total-variation-diminishing Runge-Kutta scheme [85]. Only one quarter of the domain, corresponding to one-half wavelength, is simulated via the use of symmetry boundaries. Periodic boundaries are applied in the x direction, and outflow is applied elsewhere. To efficiently resolve the evolution of the perturbation, adaptive mesh refinement based on the material indicator is applied via the AMReX library, which supports patch-based mesh refinement and straightforward parallelization.

To quantify the behavior of the flow, we track the positions of the vortex cores over time at the transverse locations of the crest and trough of the initial perturbation, x_n^C and x_n^T for core n, respectively. The cores are identified by positive values of the Q-criterion, which is commonly used to detect vortices and indicates regions where the flow rotation exceeds the strain.

4.4 Results

Simulations are performed for each Atwood number, with $\beta \in [0.4, 3.2]$ and $t \in [0, 100]$, where $\beta = k_x b$ is the non-dimensional wavenumber. Figure 4.2 shows the material indicator and Q-criterion for a case with At = 0.61 and $\beta = 0.914$ at six time snapshots. As expected, the shock traverses the interface, leading to an inversion of the notch and almost immediately generating a pair of vortex cores that share the notch's seeded sinusoidal perturbation. The RMI perturbation amplitude grows precipitously, with a spike straddled by counter-rotating cores. The transverse perturbation gives rise to a varying width of the notch, such that the RMI growth rate is largest at the trough of the initial perturbation and smallest at the crest. This transverse variation in growth rate would be present for any planar system with a transverse perturbation and may encourage the CI, which grows with an angle between the cores [57]. As the RMI continues to grow, the cores become increasingly disrupted until, around $t \approx 50$, vortex reconnection begins. As the cores completely reconnect and form rings, they pull the heavy fluid with them into clumps. The process of pinching and clumping is the key feature of this flow, as it leads to large-scale, independent, three-dimensional vortex structures unlike standard RMI bubbles and spikes.

To contrast the above results, which demonstrate instability, we show a case that is stable. Figure 4.3 again shows the material indicator and Q-criterion for a case with At = 0.61, this time with $\beta = 3.2$. The wavelength is 3.5 times shorter in this case, which would lead to a larger transverse RMI growth rate as the transverse distance between the crests and troughs is decreased, and the RMI is generally excited by shorter wavelengths [162, 201]. Additionally, any short-wavelength vortex instabilities should be more excited in this case than in the case in Figure 4.2 due

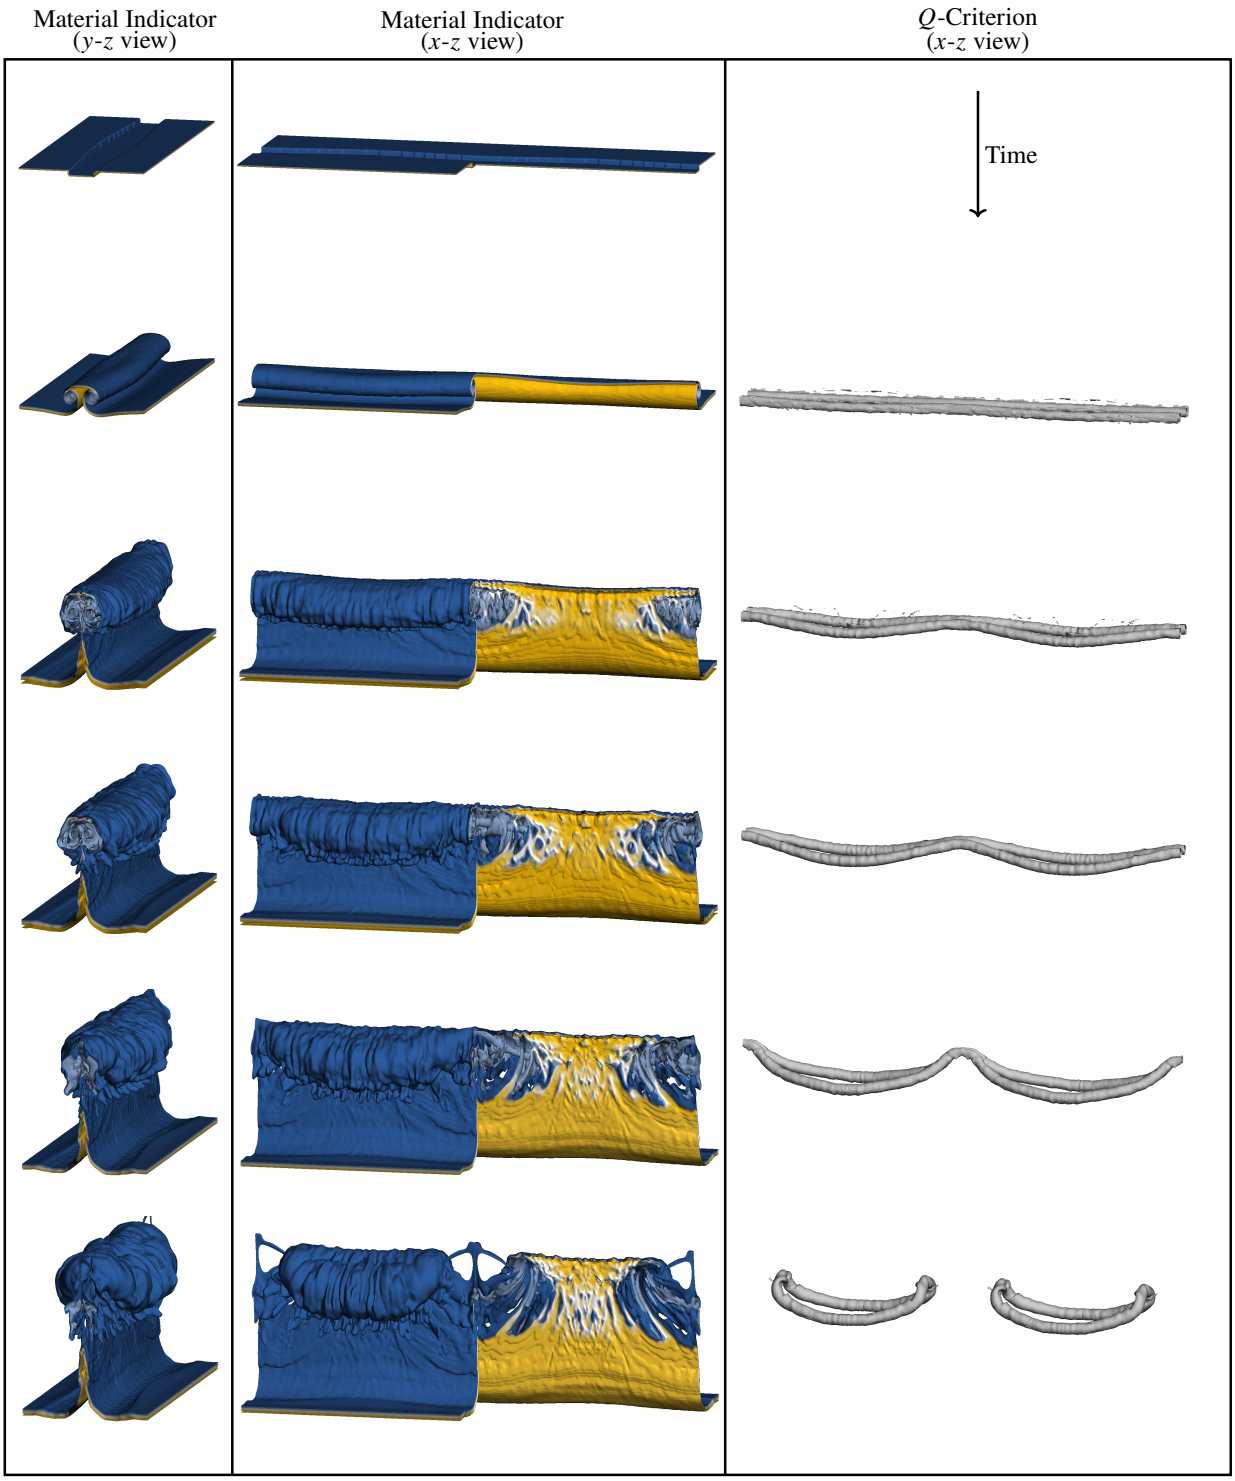


Figure 4.2: Contours of the material indicator (left and far left) and Q-criterion (right) at times $t \in [0, 5, 30, 40, 50, 60]$ for a case with At = 0.61 and $\beta = 0.914$. The heavy fluid is below the interface, and the shock comes from below. A slice is taken from the x-z view of the material indicator to reveal the interior structure of the mixed region (blue is the light fluid, yellow is the heavy fluid).

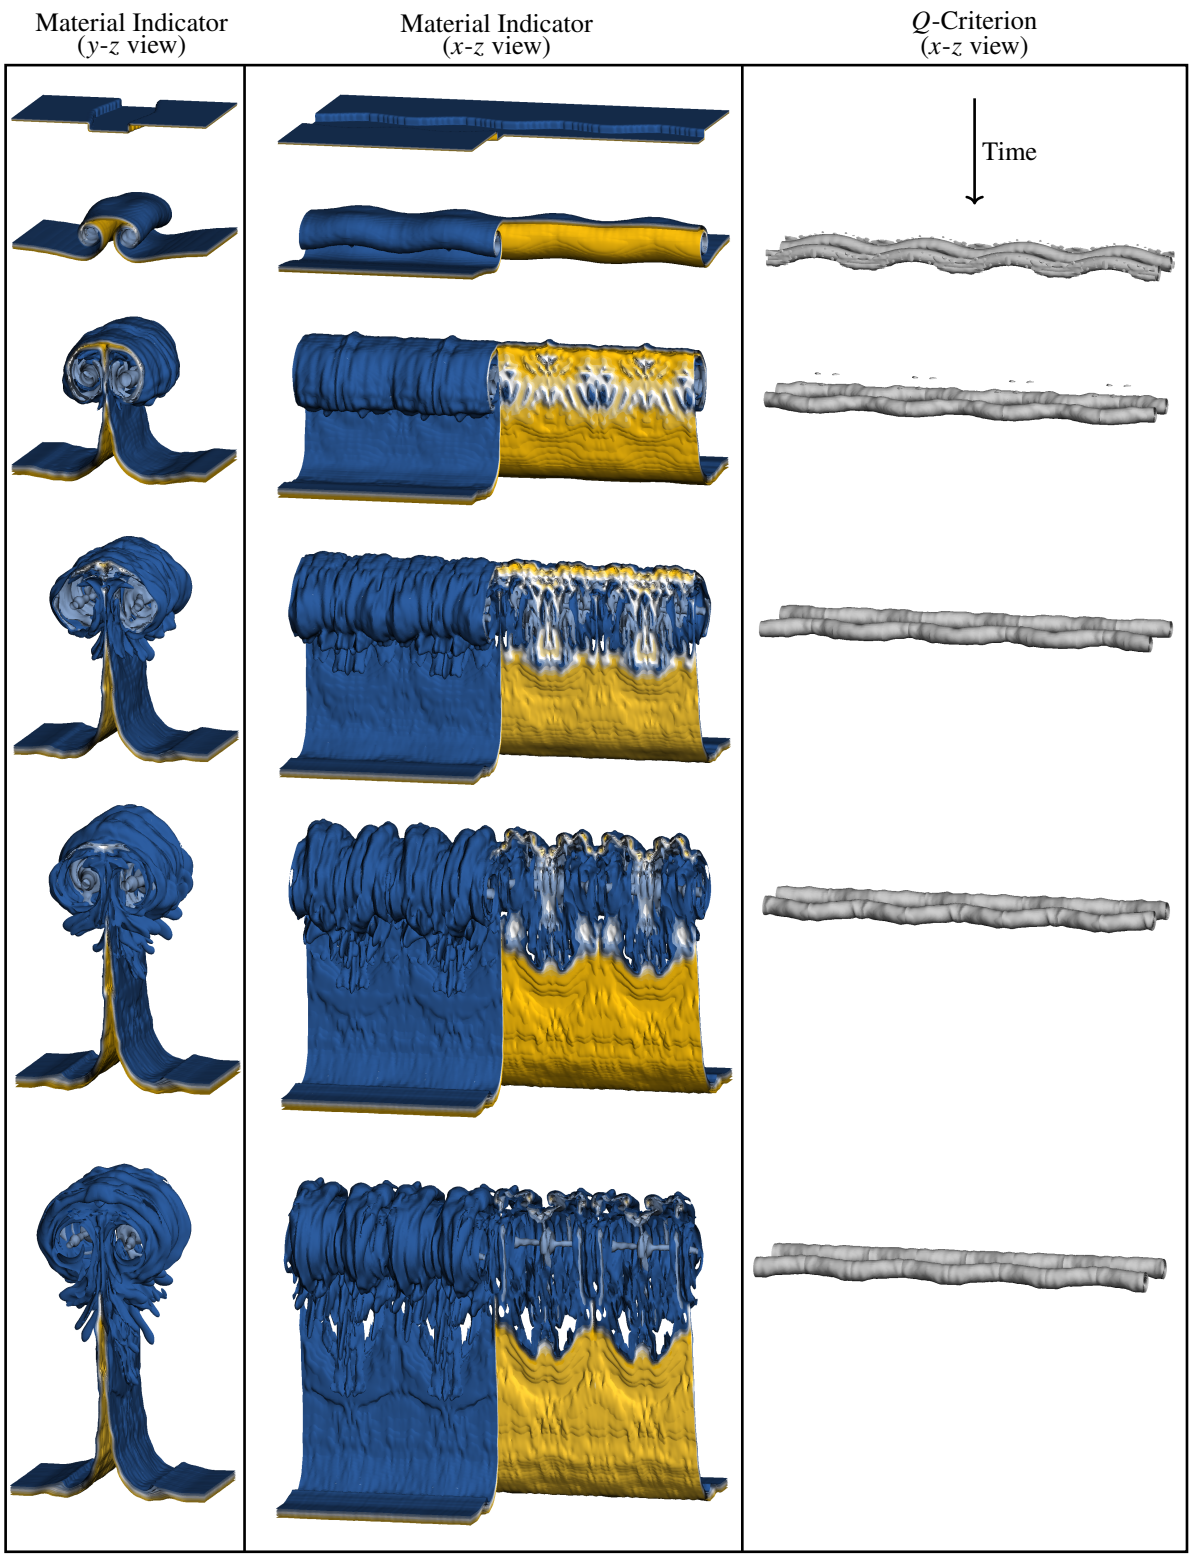


Figure 4.3: Contours of the material indicator (left and far left) and Q-criterion (right) at times $t \in [0, 5, 30, 50, 70, 90]$ for a case with At $= 0.61$ and $\beta = 3.2$. The heavy fluid is below the interface, and the shock comes from below. A slice is taken from the x-z view of the material indicator to reveal the interior structure of the mixed region (blue is the light fluid, yellow is the heavy fluid).

to the higher β here. In reality, the opposite is observed; despite the initial perturbation to the cores, reconnection is never observed. Further, the perturbation actually decays over time, simply rotating with the core. This behavior is evidence that neither transverse RMI nor short-wavelength vortex instabilities are responsible for the observed instability growth, since both mechanisms would be more excited in this case if they were active.

To bolster the case that vortex-pair instabilities are the dominant mechanism of three-dimensional structure formation, we compute the perturbation growth rate for the cores in each simulation. The vorticity and velocity are related in an incompressible and inviscid fluid via the Biot-Savart law,

$$\boldsymbol{U}_n = \sum_{m=1}^{2} \Gamma_m \int \frac{\boldsymbol{R}_{mn} \times d\boldsymbol{L}_m}{4\pi |\boldsymbol{R}_{mn}|^3} \tag{4.2}$$

where $n = 1, 2$ accounts for each vortex, $\boldsymbol{R}_{mn}$ is the separation between points on vortex m and n, $d\boldsymbol{L}_m$ is an infinitesimal length along core m, and Γ_m is the strength of core m [57]. Assuming a kinematic constraint that perturbations in the cores, $\boldsymbol{r}_n(x_n, t)$, move at the local flow velocity, taking a normal-mode ansatz of the form

$$\boldsymbol{r}_n(x_n, t) = \hat{\boldsymbol{r}}_n e^{at + ik_x x_n}, \tag{4.3}$$

and linearizing leads to an eigenvalue problem that can be solved to obtain the growth rates of the symmetric and antisymmetric modes. The symmetric modes are investigated here, which are predicted to have growth rates

$$\alpha_S = a\frac{2\pi b^2}{\Gamma_0} = \pm \left[(1 - \psi + \beta^2 \omega)(1 + \chi - \beta^2 \omega) \right]^{1/2}, \tag{4.4}$$

where α_S is the non-dimensionalized growth rate, Γ_0 is the core circulation, $\delta = 0.321 c k_x$ is the non-dimensional cutoff parameter [57], c is the core diameter, and

$$\psi = \int_0^\infty \frac{\cos\beta\xi + \beta\xi \sin\beta\xi}{(\xi^2 + 1)^{3/2}}\, d\xi, \tag{4.5}$$

$$\phi = \int_\delta^\infty \frac{\cos\xi + \xi \sin\xi - 1}{\xi^3}\, d\xi, \tag{4.6}$$

$$\chi = \int_0^\infty \frac{\cos\beta\xi}{(\xi^2 + 1)^{3/2}}\, d\xi. \tag{4.7}$$

The positive branch of α_S indicates instability. Further, this relationship demonstrates that the classical growth rate can be fully characterized by the core separation, core radius, and circulation. While the present flow is compressible, a multi-material flow clearly departs from the assumptions

that led to this theoretical result; compressibility effects are not expected to be large after shock passage. The vortex cores are close to circular after initial transience; as such, the core radius is approximated numerically as,

$$r_c = c/2 \approx \left[\frac{1}{\pi}\int I(\boldsymbol{x})d\Omega\right]^{\frac{1}{2}} \tag{4.8}$$

where $I(\boldsymbol{x})$ is an indicator function that is one in the region of positive Q-criterion that contains its maximum, and is zero elsewhere. The actual core separation is taken from simulations and is consistently approximately 90% of the mean notch width. Finally, the core circulation is calculated by integrating the vorticity field after shock passage. With these three quantities, the theoretical growth rates can be calculated, and the numerical growth rates can be non-dimensionalized. The amplitude of the perturbation over time $\zeta(t)$ is measured as the distance between the crest and trough of the vortex core, as projected onto a single z-y plane, i.e.,

$$\begin{aligned}\zeta(t;k_x) &= |\boldsymbol{r}_n(x_n^C,t) - \boldsymbol{r}_n(x_n^T,t)| \\ &= e^{at}\sqrt{\left(\hat{y}_n^C e^{ik_x x_n^C} - \hat{y}_n^T e^{ik_x x_n^T}\right)^2 + \left(\hat{z}_n^C e^{ik_x x_n^C} - \hat{z}_n^T e^{ik_x x_n^T}\right)^2} \\ &= \zeta_0 e^{at}\end{aligned} \tag{4.9}$$

Thus, if the CI is excited in the cores, then it is expected that $|\boldsymbol{r}_n(x_n^C,t) - \boldsymbol{r}_n(x_n^T,t)|$ computed numerically should be exponential. This hypothesis is tested by fitting an exponential curve across all wavelengths with the form,

$$\zeta_f(t;\beta_i) = \zeta_0 e^{a_f(\beta_i)t}, \tag{4.10}$$

where is ζ_0 is fixed and a_f is parametrized by β. This fit yields the dimensional exponential growth rates for each wavelength $a_f(\beta_i)$. Non-dimensionalizing using Eq. (4.4) allows direct comparison with theory. Figure 4.4 shows the non-dimensional growth rates from theory and from simulation vs. wavenumber for all three Atwood numbers, as well as the core perturbation amplitude over time for one case. Remarkable agreement is observed in both the trends and magnitudes of the growth rates as compared to theory for all three Atwood numbers.

Several features of the time series and growth rates suggest that the CI is the dominant mechanism driving structure formation when core instability is observed. Firstly, the symmetric long-wavelength sinusoidal perturbation growth, followed by vortex reconnection and ring formation, is characteristic of the CI and dissimilar to mechanisms such as the elliptical instability. Secondly, the observed exponential growth in the core perturbation is not attributable to the RMI, which grows at most linearly. This exponential behavior is observed long after baroclinic forcing from the shock ceases. Thirdly, the non-monotonic behavior of the growth rates with respect to the

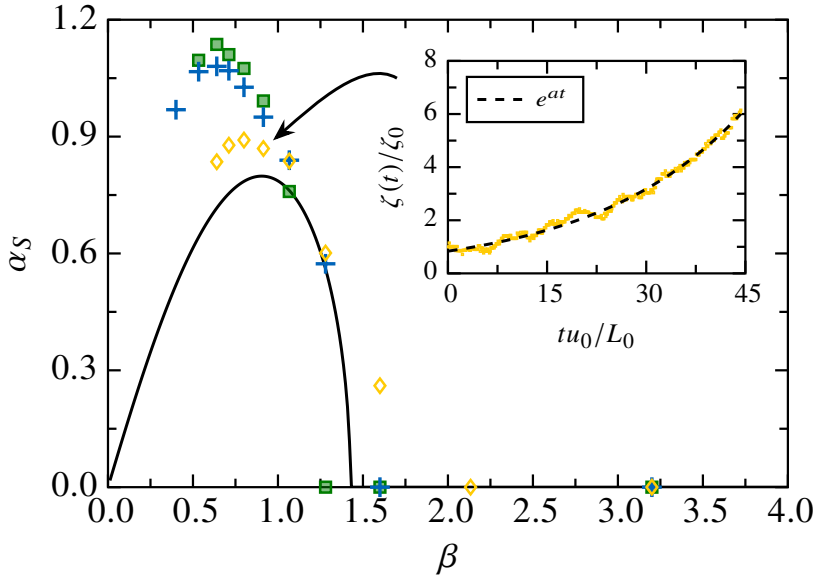


Figure 4.4: Non-dimensional growth rates calculated from simulations vs. non-dimensional wavenumber. Green squares, blue crosses, and yellow diamonds are for Atwood numbers 0.34, 0.61, and 0.78, respectively, and the black line is Eq. (4.4). The inset shows the normalized core perturbation amplitude over time for At = 0.78 and $\beta = 0.914$.

wavenumber and the presence of a cutoff wavelength are both characteristic of the CI. Further, the present results match well not only with the magnitude of the growth rates, but also with the value of the most-unstable wavenumber and the cutoff wavenumber, particularly for At = 0.78. Other instability modes like RMI or the elliptical instability are expected to be more excited as β is increased, but the opposite is observed. Taken together, these results directly support the CI as the dominant mechanism in the post-shock flow when instability is observed. As such, the emergence of three-dimensional structures from a nominally two-dimensional RMI flow may be dominated by vortex-pair instabilities, potentially altering interfacial mixing rates and the onset of turbulence.

With regard to the Atwood number dependence, the greatest degree of agreement with theory is observed in the At = 0.78 cases. The lower Atwood-number cases also show features that suggest the CI; however, the measured growth rates are larger, and the most unstable wavenumber is lower. As such, the results tend towards the theoretical result as the Atwood number increases. This behavior is initially counterintuitive as the incompressible and constant-density assumptions in the original analysis suggest that the closest agreement would be observed for $\mathrm{At} \to 0$. This suggestion, however, ignores that, for the shock to deposit circulation, the Atwood number must be non-zero. Instead, we hypothesize that in the large-Atwood limit, the vortex cores reside primarily in the heavy fluid, so that the light fluid has little effect on the kinematics of the heavy fluid, and the cores move as though they were in a single fluid. However, future work should investigate this further.

Additionally, quantifying the impact of numerical errors and grid effects, as well as the dynamics with smaller initial perturbations, would further strengthen the case for this instability pathway.

4.5 Conclusion

In this work, numerical simulations of a shocked, notched interface with a transverse perturbation demonstrate that vortex-pair instabilities are present in high-aspect-ratio RMI flows and can dominate the generation of three-dimensional structures. Further, we demonstrate agreement in the instability growth rates between classical Crow instability (CI) theory and simulations across three different Atwood numbers, with a tendency towards the classical result as the Atwood number increases.

As the shock passes through the interface, it deposits baroclinic vorticity, causing the perturbation to invert and roll up into a spike of heavy fluid, straddled by counter-rotating vortex cores. The cores share the transverse perturbation present in the notch, and in unstable cases, the sinusoidal perturbation grows exponentially until the cores are close enough to initiate reconnection. Coherent vortex rings form after reconnection, pulling mass from the interface into clumps and dramatically increasing the flow's three-dimensional nature. In cases where the cores are stable, reconnection does not occur; instead, the cores simply rotate and damp the initial perturbation from the notch. The cores are observed to be most unstable at moderately long wavelengths, with increased stability in the very long-wavelength limit and no instability at short wavelengths, both of which are consistent with CI theory. This result, along with the observed exponential perturbation growth and strong agreement between numerical results and theory in the non-dimensional growth rates, the most-unstable wavenumber, and the cutoff wavenumber, adds confidence that three-dimensional vortex-pair instabilities, the CI in particular, dominate the late-time behavior of this nominally two-dimensional RMI flow. Thus, incorporating this mechanism into mixing and turbulence models may improve predictions of ICF and astrophysical flows.

These results expose a new pathway for mixing, transition, and structure formation in shocked interfacial flows, with applications including inertial confinement fusion experiments and astrophysical systems. Future work should include a theoretical framework for predicting vortex-core instabilities in multi-material flows, further characterization of the effects of Atwood and Mach numbers, an assessment of the potential for this mechanism in systems such as ICF, and extensions to additional geometries, such as cylinders and spheres.

CHAPTER 5

Vortex-Core Instability in Shocked Converging Flows

5.1 Abstract

The vortex dynamics of the cylindrically converging Richtmyer-Meshkov instability are numerically investigated. A cylindrically converging shock impinges upon a sinusoidally perturbed, notched, heavy-to-light cylindrical interface, leading to equatorial counter-rotating vortex rings. Once baroclinic forcing ceases, the cores exhibit a symmetric long-wave bending whose amplitude grows exponentially long after the reshock. The measured growth rates exhibit a pronounced, non-monotonic dependence on angular wavenumber, with slower growth at both high and low wavenumbers, consistent with classical Crow-instability theory. The evolution produces secondary vortex rings through reconnection at late times. These signatures reveal a novel pathway for generating three-dimensional structure in compressible, multi-material, converging flows, with implications for astrophysics and inertial confinement fusion experiments.

5.2 Introduction

The Richtmyer-Meshkov instability (RMI) is present in a variety of astrophysical and high-energy-density (HED) systems, including supernovae [119, 209] and inertial confinement fusion (ICF) experiments [56, 144, 216]. The RMI occurs when a shock passes through a perturbed interface and deposits vorticity via baroclinic torque [162, 201]. The perturbation grows under this vorticity, leading to advective mixing of the two materials, typically observed in a bubble-spike formation [52] where the heavy and light fluids interpenetrate in opposite directions. In some cases, the RMI ejects vortex pairs and vortex rings, which can accelerate the development of the mixing layer [209, 236, 257]. In ICF experiments, the presence of the RMI reduces the implosion symmetry and thus truncates the potential gain [56, 144, 216]. While the RMI is typically studied as a two-dimensional mechanism, the pathway from nominally two-dimensional flows to three-dimensional structures is

of key importance in studies of multi-material mixing and turbulence; interfacial instabilities play a major role in this process.

Significant work has been conducted on the classical planar RMI [52, 193, 257], and converging geometries have also been the subject of many studies due to their relevance in systems like ICF [21, 65, 140, 146, 149, 272, 276]. Most of these studies are two-dimensional and consider perturbations of the form $\eta(\theta; k_\theta) = r_0 - a_0 f(\theta; k_\theta)$, where η is the interface radius, k_θ is an angular wavenumber, r_0 is the mean radius, a_0 is the initial perturbation amplitude, and $f(r; k_\theta)$ is the perturbation profile, which is typically sinusoidal. Such geometries give rise to unique effects, including the potential for halting growth due to the interplay between the RMI and other instabilities [65]. One particular aspect of cylindrical geometries is the convergence and rebound of the driving shock. The rebounding shock interacts with the RMI induced from the initial shock and deposits additional vorticity. Reshocked RMI has been investigated in planar and convergent geometries in order to assess growth rates, mixing widths, and potential turbulent transition [131, 149, 150, 151, 155, 166].

Vortex-core instabilities arise when small disturbances grow on otherwise coherent vortical structures such as vortex lines, rings, or pairs. In each case, the perturbation alters the local velocity field and feeds back on the vortex or vortices, causing the original cores to deform. As the instability enters its nonlinear stage, these deformations become complex and ultimately trigger either core breakdown (formation of small-scale turbulence) or discrete reconnection events in which segments of the original vortices join to form secondary vortex rings. The elliptical instability is induced by parametric resonance between fluid strain and rotational modes [123]. The Widnall instability manifests as the growth of azimuthal bending modes in vortex rings induced by self-rotation of waves in the vortex [267]. The Crow instability (CI), in particular, is a vortex-pair instability that manifests as exponential growth in sinusoidal perturbations to counter-rotating vortex lines, and is often observed behind airplanes [57]. A defining feature of the CI is that it is excited over a relatively narrow band of long wavelengths and has distinct stability cutoffs. This behavior distinguishes it from other mechanisms when measuring growth rates as a function of wavenumber. While these instabilities have typically been studied in single-fluid systems, recent works have demonstrated that they may be present in compressible multi-material flows [256], as well as in supernova remnants [259]. These works suggest that vortex-pair instabilities, such as the CI, manifest in compressible, multi-material systems, thereby exposing a variety of key mechanisms for structure formation in these flows.

This pathway has recently been studied in planar geometries (Chapter 4), where a shocked interface with a transverse perturbation was shown to develop coherent cores which pinch into vortex rings for certain wavenumbers. Further, the measured growth rates in that study showed non-monotonic behavior with respect to wavenumber, exhibited a distinct cutoff for large wavenum-

bers, and, when non-dimensionalized, largely matched classical Crow theory. This behavior was observed across different Atwood numbers with convergence towards the classical result at high density ratios, suggesting that vortex-pair instabilities can be directly excited from a nominally two-dimensional RMI. Similarly, the present work seeks to interrogate the presence of vortex-pair instabilities in a converging RMI and their effect on the generation of three-dimensional structures in a nominally two-dimensional flow. This problem couples the effects of RMI-driven compressible multi-material CI with the geometric complexities of converging RMI and the dynamic effects of reshock. This setup is relevant to systems such as double-shell ICF [216], where the hemispherical assembly joint may seed RMI that develops into vortex rings. The remainder of this work is as follows: Section 5.3 details the setup and numerical methods used to study this flow. Section 5.4 presents results and discussion, and Section 5.5 concludes.

5.3 Setup and Methods

To study the potential for vortex-pair instabilities in converging geometries, the setup from Chapter 4 is adapted to a cylinder, with a converging shock impinging upon a concentric cylinder. The cylinder is given a rectangular equatorial notch in order to seed the RMI and encourage the development of vortex cores. Instead of a continuous distribution of wavelengths like in the classical planar case, discrete mode numbers are prescribed by the angular wavenumber k_θ. A sinusoidal perturbation in the z-direction is given to this notch, seeding a single mode. The initial interface is then given by,

$$\eta(\theta, z; k_\theta) = r_0 + a_{0,r}\left\{\mathcal{H}\left[z + a_{0,z}\cos(k_\theta\theta) + \frac{b}{2}\right] - \mathcal{H}\left[z - a_{0,z}\cos(k_\theta\theta) - \frac{b}{2}\right]\right\}, \tag{5.1}$$

where $r_0 = 0.5$ is the unperturbed interface radius, $\mathcal{H}$ is the Heaviside step function, and $a_{0,z}$ is the perturbation in the z-direction. The radial perturbation amplitude $a_{0,r} = r_0/50$ seeds the initial RMI and the cores, while the mean core separation is approximately given by b. This setup is shown schematically in Figure 5.1. This setup produces a predictable RMI flow with direct control over the core separation (via b). The arcwise wavenumber $k = k_\theta/r_0$ and wavelength $\lambda = 2\pi r_0/k_\theta$ are useful in comparing to the planar case, as both the angular wavenumber and the arc length of the perturbation may be significant for stability when other variables are held constant.

The materials are both ideal gases with ratios of specific heat $\{\gamma_p\}_{p=1}^2 = \{1.093, 1.276\}$ and densities $\{\rho_p\}_{p=1}^2 = \{5.494, 1.351\}$, corresponding to the heavy and light fluids in Chapter 4. A converging Mach 1.21 shock is initialized via the Rankine-Hugoniot conditions just outside the interface. While this is a gas-gas system with an exaggerated gap width, similarities to the assembly joint in double-shell ICF are apparent [216]. As the shock converges, it develops into a blast, but the initial shock is initially positioned close to the interface, so its interaction is largely unaffected

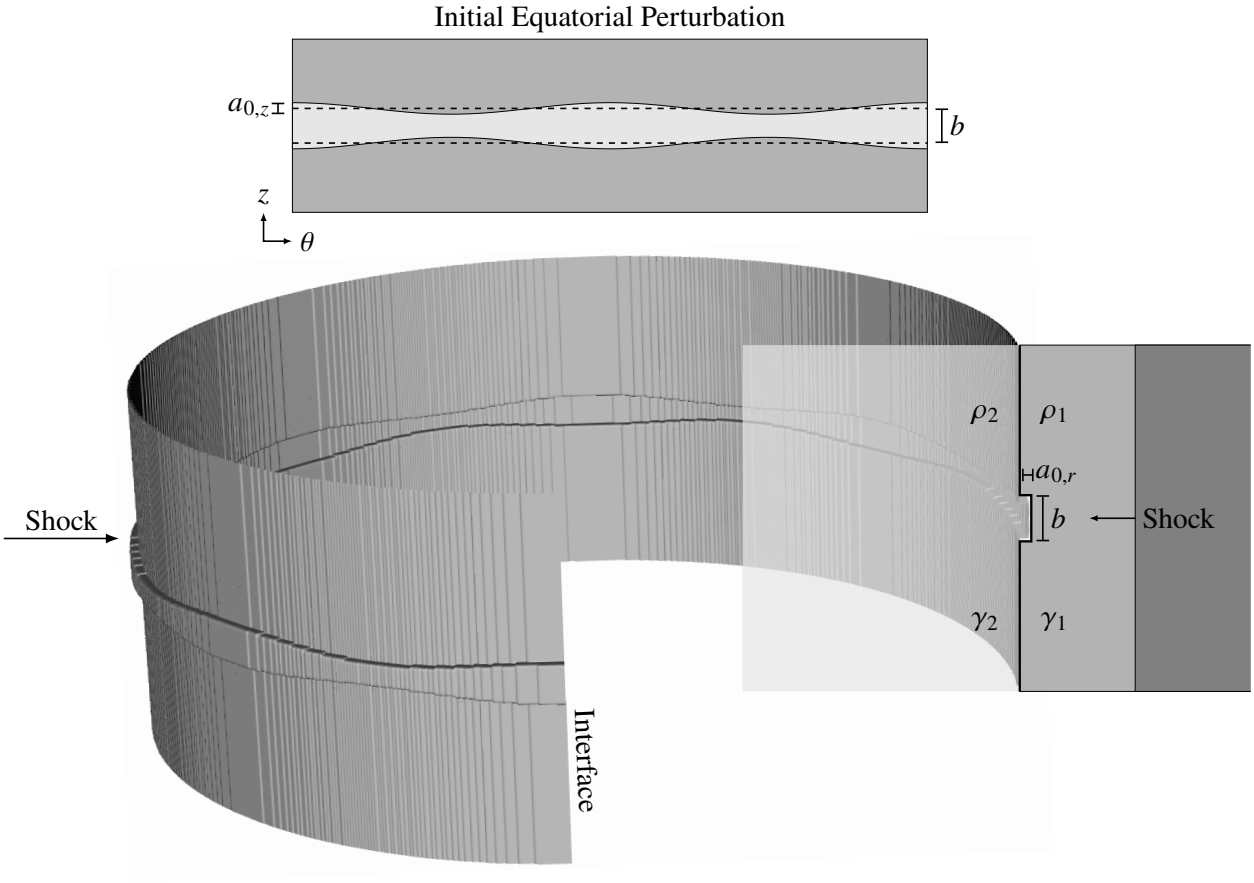


Figure 5.1: Schematic of the initial condition showing the initial interface, initial equatorial perturbation, and the shock.

by this geometric effect.

Numerical simulations are conducted using the multi-material Euler solver from the COMPAS code (Chapter 2). A second-order MUSCL reconstruction [250] is used alongside the HLLC Riemann solver [243] and a third-order total-variation diminishing Runge-Kutta scheme in time [85]. The AMReX library provides support for patch-based adaptive mesh refinement, which is done based on the gradient in the material indicator to refine the Cartesian grid around interfaces. Symmetry boundary conditions are applied to simulate only one eighth of the domain, with outflow conditions elsewhere. The use of outflow conditions in the z-direction ensures that there are no vortex-coupling effects from periodic boundaries.

The positions of the crests and troughs of the cores are tracked over time by computing the location of the maximum positive value of the Q-criterion at the crests and troughs of the perturbation, which measures where the magnitude of a flow's rotation rate exceeds its strain rate, indicating a vortex. When a crest or trough does not lie on an axis, its position is rotated so that both parts of the core lie in the same plane. The growth of the RMI bubbles and spikes is also tracked over time at the crest and trough angles by interpolating the interface position between cells.

5.4 Results

We simulate wavenumbers $k_\theta \in \{2, 4, 8, 12, 16, 24\}$ over $t \in [0, 1]$ to quantify the behavior of the cores across a range of wavenumbers. Each simulation begins with the converging shock impinging upon the interface, driving it radially inwards. Figure 5.2 shows the material indicator and Q-criterion for a case with $k_\theta = 8$ at four time snapshots. The notch, which initially has a greater radius than the unperturbed interface, is driven inwards faster than the surrounding flow, as the shock deposits vorticity that drives RMI growth. The step-like perturbation yields two perturbed vortex rings almost immediately, which are also driven inwards with the converging interface. The bubbles and spikes of the RMI grow at marginally different rates along the equator due to the variable perturbation width, seeding an angle between the perturbations which is favorable for the development of the CI [57]. The vortex rings converge without significant topological variation until the rebounding shock impinges upon the bubbles and spikes around $t = 0.3$, at which point significant additional circulation is deposited into each core, and the bubbles and spikes elongate significantly. Although the RMI is dramatically altered by the reshock, the core perturbation is almost unaffected. Additionally, away from the perturbation, the reshock essentially halts the interface's motion, while the spikes penetrate deeply into the cylinder's core. Following the reshock, there is no significant source of baroclinic vorticity generation in the flow, such that further perturbation growth is primarily vortex-driven. All simulated cases are observed to be unstable to varying degrees; the perturbations continued to grow until pinch-off.

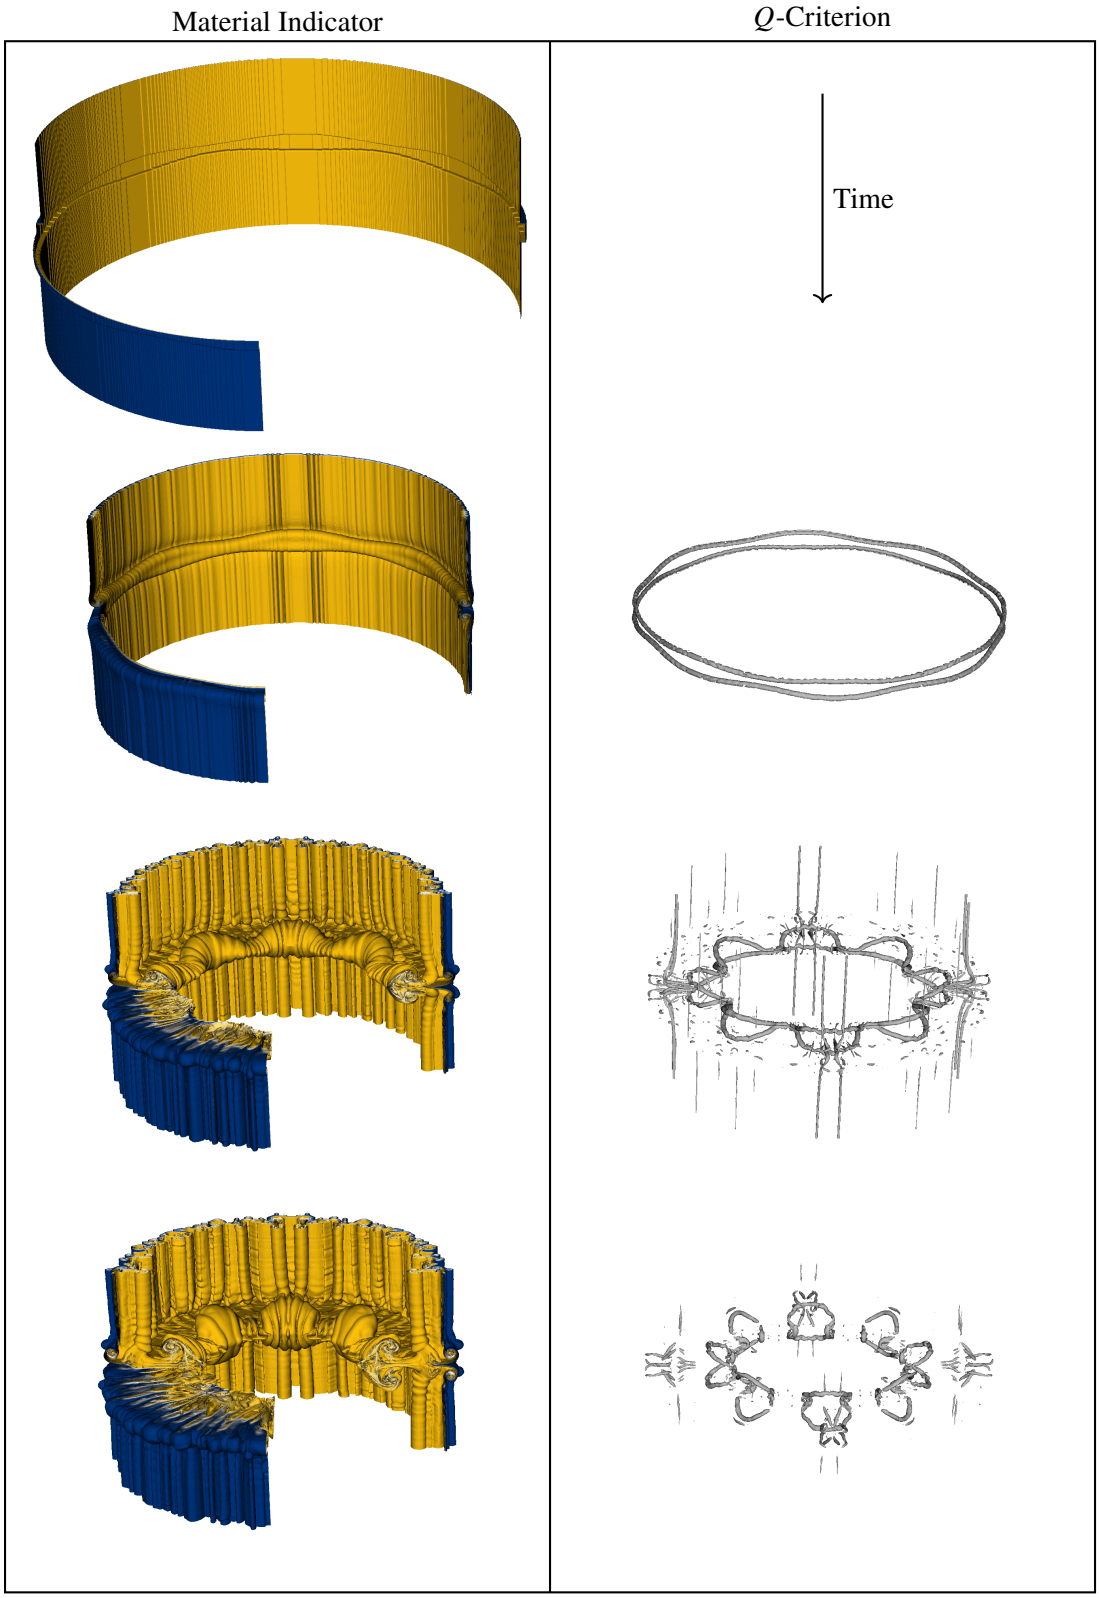


Figure 5.2: Contours of the material indicator (left) and Q-criterion (right) at times $t \in [0, 0.2, 0.7, 1.0]$ for a case with $k_\theta = 8$. Slices are taken out of the view of the material indicator to see the interior structure of the mixed region (blue is the heavy fluid, yellow is the light fluid).

This simulation demonstrates (i) the early-time RMI growth, (ii) the sinusoidal perturbation growth in the vortex cores, (iii) significant spike penetration following reshock, and (iv) vortex reconnection into secondary rings late in time. Additionally, the angle and growth behavior of the cores are reminiscent of that observed in previous studies [256, 259]. The observation of this phenomenon immediately suggests a pathway to three-dimensional structure formation in converging flows. Specifically, the interaction of the converging shock with the equatorial notch deposits vorticity that rolls up via the RMI into counter-rotating vortex rings. Perturbations to these vortex rings, which may be multi-modal, grow to the mutual effect of one ring on the other, leading to exponential growth and eventually reconnection into vortex rings.

This instability behavior is observed across wavenumbers. Figure 5.3 shows the evolution for a case with $k_\theta = 12$. In this case, the same pathway from the RMI to secondary vortex rings is observed at a higher frequency, as expected.

The amplitude of the perturbation $\zeta(t; k_\theta)$ to core n is measured from simulations as,

$$\zeta(t; k_\theta) = |\boldsymbol{s}_n(\theta_n^C, t) - \boldsymbol{s}_n(\theta_n^T, t)|, \tag{5.2}$$

where $\boldsymbol{s}_n$ is the position of core n, and θ_n^C and θ_n^T are the angular positions of the crest and trough of the sinusoidal perturbation, respectively. The classical CI theory uses an exponential ansatz for the normal-mode analysis; the same form is used here,

$$\begin{aligned}\zeta(t; k_\theta) &= |\boldsymbol{s}_n(\theta_n^C, t) - \boldsymbol{s}_n(\theta_n^T, t)| \\ &= e^{at}\sqrt{\left(\hat{r}_n^C e^{ik_\theta\theta_n^C} - \hat{r}_n^T e^{ik_\theta\theta_n^T}\right)^2 + \left(\hat{z}_n^C e^{ik_\theta\theta_n^C} - \hat{z}_n^T e^{ik_\theta\theta_n^T}\right)^2} \\ &= \zeta_0 e^{at}.\end{aligned} \tag{5.3}$$

Thus, we find the perturbation growth rates via a simultaneous exponential fit of the form,

$$\zeta_f(t; k_{\theta,i}) = \zeta_0 e^{a_f(k_{\theta,i})t}, \tag{5.4}$$

where $\zeta_f(t; k_{\theta,i})$ is the fit for a given wavenumber $k_{\theta,i}$, ζ_0 is fit to a single value across all unstable wavenumbers, and $a_f(k_{\theta,i})$ is the wavenumber-specific growth rate. Figure 5.4 shows the fit growth rates across all simulated wavenumbers. Remarkably, the measured growth rates exhibit a non-monotonic trend, with the largest growth rates occurring in the low-wavenumber regime. This behavior is characteristic of the Crow instability, which across a wide range of wavenumbers is the most excited by long-wavelength perturbations with reduced growth rates at short and very long wavelengths. While it is possible that the converging base flow alters the growth rate over time, the perturbation amplitude is observed to be very close to exponential (as shown in the inset of Fig. 5.4),

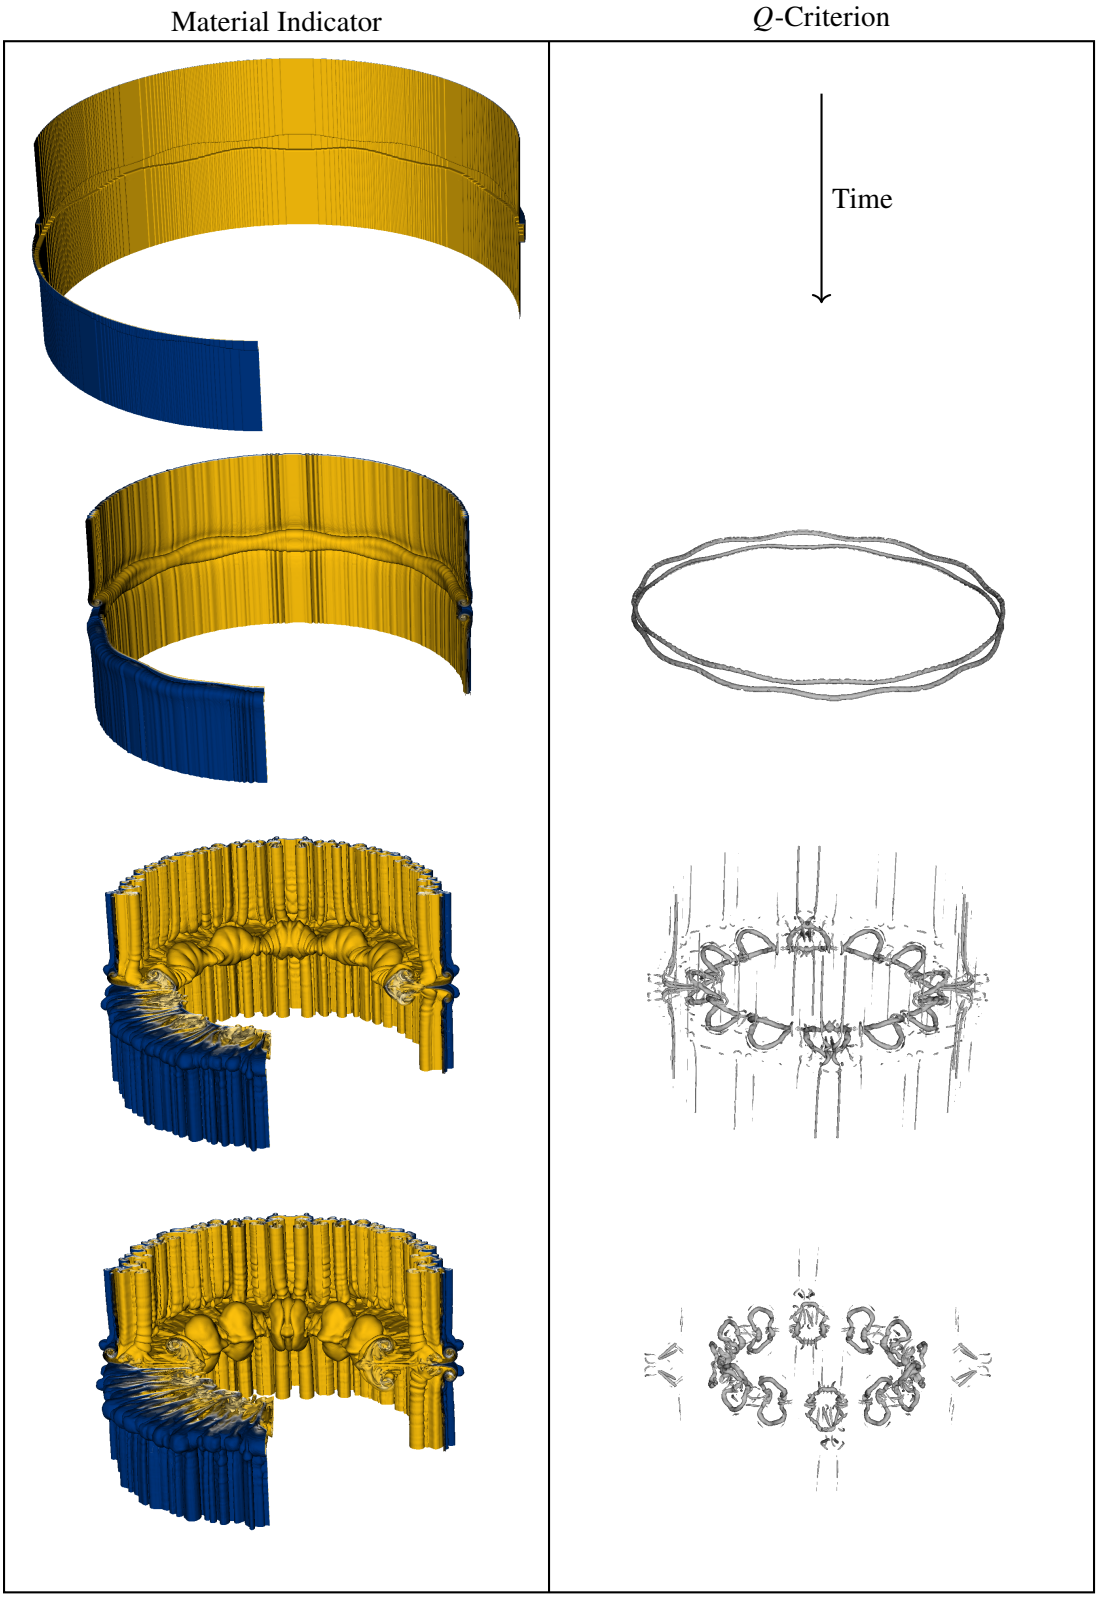


Figure 5.3: Contours of the material indicator (left) and Q-criterion (right) at times $t \in [0, 0.2, 0.7, 1.0]$ for a case with $k_\theta = 12$. Slices are taken out of the view of the material indicator to see the interior structure of the mixed region (blue is the heavy fluid, yellow is the light fluid).

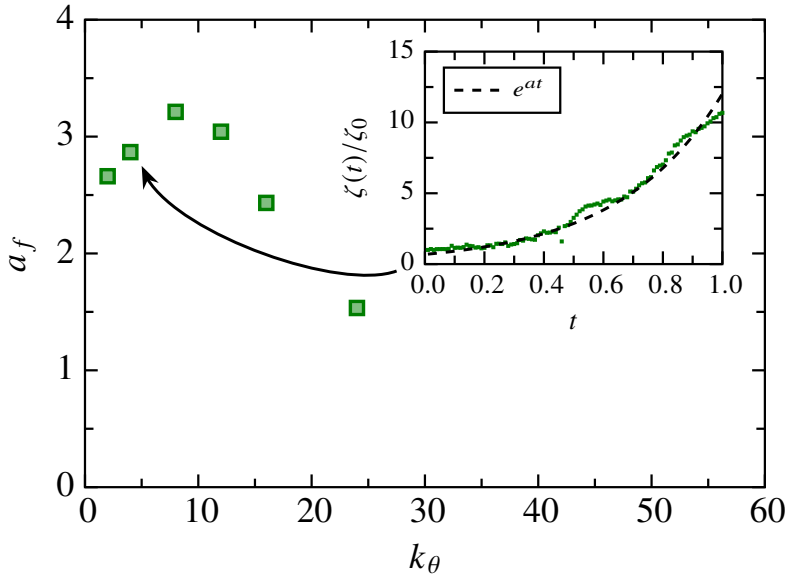


Figure 5.4: Perturbation growth rates a_f versus angular wavenumber k_θ. The inset plot shows the normalized perturbation amplitude over time overlaid with the corresponding exponential for the $k_\theta = 4$ case.

again consistent with the CI. Further, the reshock around $t \approx 0.3$ is almost indistinguishable in the perturbation amplitude over time, suggesting that the effect of the reshock on the narrow and wide regions of the spikes is not significant in the current system.

While these results do not conclusively point to a Crow-like mechanism for the observed pinch-off behavior, several independent features indicate that a long-wavelength instability drives the exponential growth observed in this cylindrical, reshocked configuration. Firstly, the saturated RMI stage leaves behind two coaxial, counter-rotating vortex rings whose centroids exhibit a sinusoidal, symmetric displacement, precisely the geometry predicted for the planar Crow instability. Clearly, the heavy seeding of this profile leads to such behavior, but the lack of other dominant modes emerging on the rings suggests that a single long-wavelength mode is most excited. Further, the measured perturbation amplitudes grow exponentially for several core-turnover times after the rebounding shock arrives at the interface, eliminating continued baroclinic forcing as the driving mechanism and implicating independent vortex-vortex interaction instead. Figure 5.5 shows the amplitude $A(t)$ of the equatorial RMI at the crest and trough angular positions. The perturbation growth, both pre- and post-reshock, is not exponential, supporting the notion that the observed growth in the core perturbation is due to vortex instabilities rather than transverse instabilities. The observed linear growth is unexpected from late-time RMI and is likely due to geometric effects, although future work should characterize this behavior. When plotted against the angular wavenumber, the measured growth rates exhibit a clear non-monotonic behavior. This trend is a

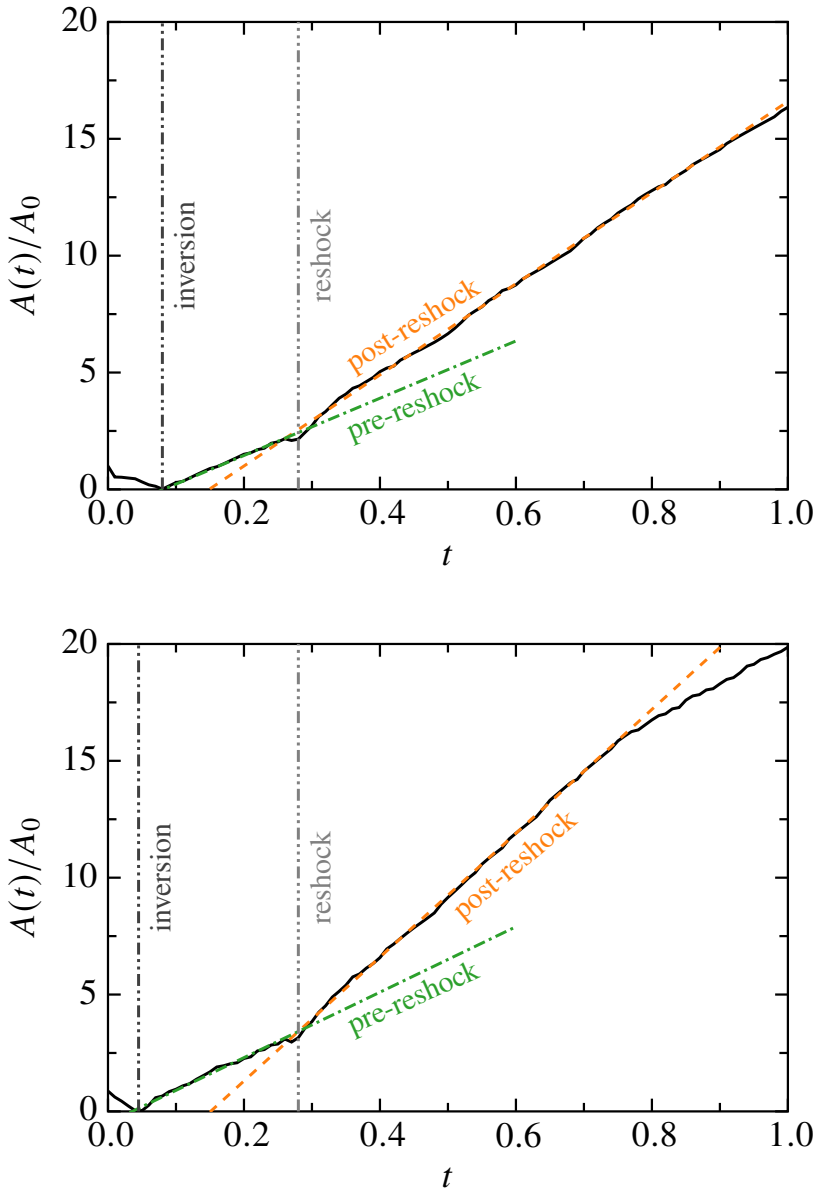


Figure 5.5: Amplitude of the RMI perturbation vs. time at the crest (top) and trough (bottom) of the $k_\theta = 4$ case.

hallmark of the CI; reduced growth rates at both the high and low limits of k_θ are not easily explained by other mechanisms, such as the elliptical instability or the Widnall instability, particularly given the bias towards lower values of k_θ. Finally, the development of secondary vortex rings via a clean pinch-off at late time is characteristic behavior for the classic CI, here in a cylindrical setting. Taken together, the observed core morphology, exponential temporal growth, and wavelength selectivity are compelling evidence that a Crow-like mechanism governs the core dynamics rather than residual RMI growth or other instability mechanisms. This is bolstered by the linear RMI growth both pre- and post-reshock, indicating that the exponential core perturbation growth is not due to a geometric effect on the RMI. Decisively identifying this behavior as the CI, rather than other potential mechanisms such as the Widnall instability or a geometric effect from the converging interface, could be done in three steps: first, a theoretical result on the emergence of the CI from the RMI in planar geometries should be developed. Then, the theory in [255] should be extended to converging vortex rings and combined with the results for the planar RMI-CI pathway. Finally, experiments should validate these theories, at which point numerical simulations can be used for further investigation.

5.5 Conclusions

In this work, numerical simulations of a shocked converging flow demonstrate the excitation of vortex-core instabilities from a cylindrical Richtmyer-Meshkov instability (RMI). These results establish a coherent picture of a nested instability pathway in which the deposition of baroclinic vorticity by a converging shock seeds an equatorial RMI, which develops into counter-rotating vortex cores that subsequently succumb to a Crow-type long-wavelength bending mode.

In this cylindrical configuration, the shock and reshock produce two coaxial vortex rings that exhibit symmetric sinusoidal perturbation growth that continues to grow exponentially long after baroclinic forcing has ceased, and ultimately pinch off into secondary vortex rings. The perturbation growth rates were observed to be non-monotonic with respect to the angular wavenumber, with the largest growth rates occurring near the low-wavenumber limit. These signatures of the core geometry, exponential time dependence, wavelength selectivity, and reconnection closely resemble those identified in the classical planar Crow instability, despite the compressible, multi-material, curvature, and reshock effects present here, and thus suggest a Crow-like mutual-induction mechanism, rather than residual RMI growth or other vortex instabilities, governs the late-time evolution of these shock-generated vortex cores.

This finding may have direct implications for double-shell inertial confinement fusion (ICF) experiments, where the assembly joint may seed equatorial RMI features that proceed in a manner similar to those studied here, leading to three-dimensional structure formation. The pinch-off

behavior we observe ejects coherent vortex rings that could transport shell material away from the interface, providing a long-wavelength mix mechanism absent from most current design codes. Assessing the potential of this instability mode under realistic experimental conditions should therefore become a priority for evaluating performance and for designing more symmetric experimental platforms. Ongoing work is addressing a number of aspects of this system, including determining the Atwood-number dependence, quantifying radius effects in terms of the reshock and arcwise wavelength, and synthesizing theoretical expressions for the expected instability growth rates.

CHAPTER 6

A Wave-Based Framework for Understanding Cavitation Damage Potential in Blast-Induced Traumatic Brain Injury

6.1 Abstract

The risk of blast-induced traumatic brain injury (bTBI) is a serious concern for armed service members globally, but much is not known about the causal mechanisms due to substantial complexity in the dynamics of a blast interacting with the head (fluid dynamics, solid mechanics, biomechanics, etc.). Further, historically proposed damage pathways tend to underpredict the observed injury in bTBI, suggesting that another mechanism may be responsible for some of the damage. In this work, we build on the hypothesis that cavitation may occur at micro and nanoscale nucleation sites within brain tissue due to tension induced by intracranial acoustic waves. To do this, a set of simplified models is used to demonstrate the fundamental physical mechanisms that, in concert, can result in cavitation-relevant conditions in the brain tissue, i.e., regions of high tension. One-dimensional models are used alongside a theoretical discussion to demonstrate the fundamental wave dynamics that serve as building blocks for more complex models. A two-dimensional model is used to describe focusing and divergence effects, and a three-dimensional model is used to demonstrate the occurrence of cavitation-relevant conditions in a realistic geometry. Relevant nondimensional parameters are calculated to justify the use of acoustic principles in the models, and numerical simulations quantify the generation of tension. Higher-order effects are discussed, including the consequences of compound blasts, the multi-layered nature of the head, and nonlinear wave dynamics. The mechanisms leading to cavitation-relevant conditions are shown in each model, demonstrating that acoustic effects alone may be responsible for cavitation damage in bTBI.

6.2 Introduction

The mechanisms that lead to blast-induced traumatic brain injury (bTBI) are numerous and complex, and current guidelines for safe exposure may not be adequate as bTBI currently represents the largest source of brain injury in armed service members globally [72]. Blast exposures that lead to compression and decompression of the brain tissue are a known pathway to synaptic damage that may progress to neurodegenerative disorders [194]. Studying the potential for damage from this pathway, however, is made difficult by situational complexities such as compound blast structures (potentially from multiple sources or in confined geometries), as well as diagnostic challenges in characterizing the blasts and their resultant loads on the head. Further, *in-vivo* characterization of the wave dynamics inside the brain and *in-vitro* characterization of the resultant damage each present unique challenges, such as inaccuracies in surrogate models and insufficient spatio-temporal diagnostic resolution, and a lack of accurate three-dimensional neuronal-damage platforms, respectively [22]. Thus, to protect those in situations where blasts may be present, more scrupulous damage thresholds must be determined by studying the dynamics that lead to the compression and decompression of the brain tissue.

Blast waves are generated by the rapid deposition of a large amount of energy in a small volume, typically by detonations, and are characterized by a shock front followed by a subsonic rarefaction with a tail of negative gauge pressure. When a blast interacts with an obstruction, portions of the wave's energy are distributed as reflected and transmitted waves, with magnitudes and phases determined by the obstruction's physical properties, particularly its impedance. When a blast in the air interacts with an exposed head, the change in impedance leads to a wave of greater pressure than the incident blast to be transmitted into the head and then into the brain, relatively unaltered by the intermediate layers of the head other than the skull itself, due to minimal changes in impedance. This, now close to acoustic, transmitted wave then propagates inside the skull cavity, compressing and decompressing the brain as it reflects and inverts phase. As mentioned, these compressions and decompressions are known to lead to neuronal injury [22]. With this potential injury mechanism in mind, however, the damaging strains observed in experiments are hypothesized to be more likely produced by phenomena such as skull flexure and cavitation in the cerebrospinal fluid [82] than the blast directly, which produces relatively small strains independently [159, 164]. The high-pressure intracranial waves induced by the blast may, however, be responsible for cavitation directly inside the brain tissue, and thus characterizing this injury pathway is critical for safety [130].

Cavitation inside the brain tissue has increasingly been considered a serious potential damage pathway in bTBI, with both computational and experimental studies investigating and demonstrating it as a serious risk [7, 120, 130, 262]. The mechanism by which this occurs may be directly attributable to the blast-induced waves; tension waves produced by intracranial reflections of the

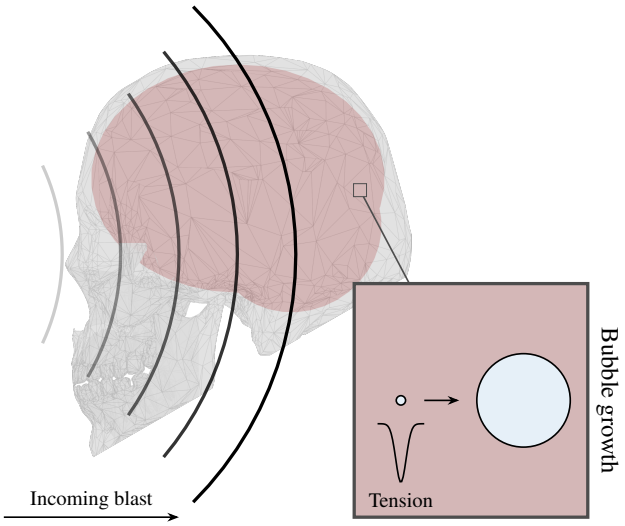


Figure 6.1: Schematic representation of a tension wave inside brain tissue interacting with a microbubble, leading to growth that may directly damage neurons.

transmitted blast interact with microbubbles in brain tissue, leading to growth. This growth, depending on the degree of cortical stretch, can lead to significant neuronal death, directly due to physical deformation rather than apoptosis or necrosis [70]. Figure 6.1 illustrates this phenomenon schematically. Accurately predicting the pressure fields resulting from the blast-head interaction is thus critical for characterizing the damage potential from cavitation. Current state-of-the-art methods (e.g., anatomically relevant finite element models) are, however, not well suited to predicting these transient, sharp, nonlinear phenomena (blasts and blast-interface interactions). Further, these models are often computationally very expensive and difficult to analyze. Therefore, in this work, we present a sequence of simplified models that elucidate simple phenomena that lead to tension and, thus, potentially to cavitation in bTBI. Previous work has investigated simple models such as one-dimensional multi-layer systems [62, 231] and more complex three-dimensional models [281], but has not connected these to the potential for cavitation in brain tissue. We show, through a minimalist sequence of analytic and numerical models, that even without skull flexure, rotation, or viscoelastic dissipation, wave interactions alone can generate tension within the brain of sufficient magnitude to induce damaging levels of cavitation.

Beginning with one-dimensional models, the wave dynamics that govern planar blast interaction in multi-layer systems with bTBI-relevant properties are illustrated. Further, the values of relevant non-dimensional quantities suggest that acoustics predominantly govern the pressure dynamics after the blast deposits its energy in the head, allowing simple relationships to predict tensions and build intuition about the physical mechanisms. Subsequent two- and three-dimensional models demonstrate geometric effects, such as focusing, that lead to distinctly multidimensional dynamics. This hierarchy of models is used to build an intuitive understanding of the physical mechanisms

underlying intracranial tension waves in bTBI and to provide a tool for interpreting experimental data through this lens. Time-resolved numerical simulations of each model supplement the theoretical and qualitative discussions throughout and demonstrate quantitatively that cavitation-relevant tensions can be generated purely through fluid mechanisms. Discussions of higher-order effects, such as anatomically correct geometries, elasticity, and biological considerations, are included. As such, this work provides a complete wave-based framework for understanding the potential for cavitation damage in bTBI. The remainder of this work is organized as follows: The theoretical models and numerical methods used herein are discussed in § 6.4. The hierarchy of models, along with numerical results, is presented in § 6.5. Discussion of additional considerations and cavitation potential is provided in § 6.5.3, and conclusions are provided in § 6.6.

6.3 Theoretical Framework

The process by which a blast wave deposits energy in the head is complex and nonlinear. As the blast wave impinges upon the head, a portion of the incident wave's energy is deposited into the head, some of which is expended in skull flexure and some is transmitted into the head as an acoustic wave. At this point, several mechanisms distribute energy throughout the head, including, but not limited to, elastic wave transport within the skull, viscous and thermal dissipation, and acoustic wave propagation. This process is complicated by the biological structure of the human head, with each layer inside the head representing a change in acoustic impedance and a complex geometric change. Despite these complexities, we are interested in modeling the waves inside the head as acoustic, primarily linear. To justify this simplification, consider first the Reynolds and Peclét numbers,

$$\mathrm{Re} = \frac{\rho u L}{\mu} \approx 10^3, \qquad \mathrm{Pe} = \frac{u L}{\kappa} \approx 10^4, \tag{6.1}$$

where the reference viscosity, density, and thermal conductivity are the approximate experimental values, the length scale is the radius of the head, and the velocity scale is the displacement speed of the tissue inside the head. It is immediately clear that, on the time scales here, viscous dissipation and heat transfer effects are not expected to be substantial, so an inviscid approach without heat transfer is suitable. Next, consider the aforementioned displacement velocity and the speed of sound in the brain tissue, which together give the local Mach number,

$$\mathrm{Ma} = \frac{u}{c} \approx 10^{-4}. \tag{6.2}$$

With the Mach number much less than unity, the waves inside the brain tissue that result in tissue displacement are very well approximated as acoustic, and the dynamics can be studied via multi-material fluid models with appropriate equations of state (e.g., stiffened gas). As such, effects such

as nonlinear steepening and focusing (in multiple dimensions) are negligible once the blast enters the comparatively high-impedance skull/brain system. Elastic effects, particularly shear waves, are an important consideration in the skull and brain. Due, however, to the short time scales over which intracranial pressure waves travel ($c_{\text{Brain}} \approx 1500\,\text{m/s}$), shear waves are not expected to be important as they travel orders of magnitude slower than the waves considered here [234]. This is true in both the skull and the brain, suggesting a purely fluid approximation of the dynamics is reasonable for studying wave propagation. Further, the use of a stiffened-gas equation of state for the skull and brain allows us to approximately match the bulk elasticity of each material without additional modeling. With these results, we posit that the use of acoustic relationships to study the wave dynamics of the intracranial waves is reasonable for the purposes of this work.

Using a fluid/acoustic approach for this problem is a departure from much of the theoretical and computational work on this subject [181], but, as shown above, it is a reasonable leading-order approximation for the dynamics of interest. Further, it removes various sources of complexity and their respective uncertainties, leading to an incomplete but tractable theoretical framework that can be used to interpret experimental results and build topical intuition.

The equations of linear acoustics yield three fundamental relationships that describe the one-dimensional interaction of a wave with an interface,

$$R = \frac{Z_2 - Z_1}{Z_1 + Z_2}, \qquad T = \frac{2Z_2}{Z_1 + Z_2}, \qquad 1 = R^2 + \frac{Z_1}{Z_2}T^2, \tag{6.3}$$

where R is the reflection coefficient, T is the transmission coefficient, and the third expression relates their magnitudes (a manifestation of conservation of energy). The acoustic impedance $Z_k = \rho_k c_k$ for material k is a measure of how easily a wave can propagate through it, where ρ_k is the material density and c_k is the sound speed. The behavior of a (linear) wave impinging upon an interface depends entirely on the change in impedance $\Delta Z = Z_2 - Z_1$ across the interface, as suggested by Eq. (6.3). If $Z_2/Z_1 \to \infty$, the solid acts as a rigid boundary and the incoming wave is entirely reflected without a phase inversion ($R > 0$). Conversely, if $Z_2/Z_1 \to 0$, then the interface acts as an open boundary, and the wave is reflected with the opposite phase of the incident wave ($R < 0$). In the case $Z_1 = Z_2$, an acoustic wave will be unaffected by the interface, even if there is a material change. When Z_2 and Z_1 are similar but not equal, the incident wave will be partially transmitted and partially reflected, with magnitudes given by the reflection and transmission coefficients.

These relationships reveal another important result from the given assumptions: when the change in impedance across an interface is small, the wave is largely unaltered. This is very useful, as many of the layers within the skull cavity have very similar acoustic impedances and, from a wave perspective, are not dissimilar to water. That leaves only the skull and brain to be considered for the purposes of this work. Table 6.1 shows the impedance values calculated from experimentally

Table 6.1: Typical densities, sound speeds, and acoustic impedances of head layers [106].

Layer	ρ (kg m^{-3})	c (m/s)	$Z = \rho c$ (kg m^{-2}s^{-1})
Air	1.23	340	4.18×10^2
Scalp / soft tissue	1109	1624	1.80×10^6
Skull bone	1908	2770	5.29×10^6
Cerebrospinal fluid (CSF)	1007	1504	1.51×10^6
Dura mater	1174	1500	1.76×10^6
Blood	1050	1578	1.66×10^6
Brain tissue (parenchyma)	1046	1546	1.62×10^6

measured densities and sound speeds. Clearly, only the skull has a substantially different impedance than any of the other layers considered here, although the dura mater is an outlier amongst the brain, blood, and CSF, with an impedance about 9% larger than the brain. Due to the thinness of the dura mater (~ 0.5 mm), however, this layer will only manifest as a negligible time delay ($\sim 0.1\ \mu$s) in the wave as it propagates through the interface. Otherwise, intracranial waves will not "see" the other layers, such that modeling only the brain and skull is reasonable. Table 6.2 shows the reflection and transmission coefficients that result from the interactions of interest. These values clearly

Table 6.2: Pressure-amplitude reflection (left) and transmission (right) coefficients, R_{ij} and T_{ij}, respectively, or a wave incident from medium i (row) into medium j (column).

R_{ij}	Air	Skull	Brain
Air	-	0.9998	0.9995
Skull	−0.9998	-	−0.5311
Brain	−0.9995	0.5311	-

T_{ij}	Air	Skull	Brain
Air	-	1.9998	1.9995
Skull	0.0002	-	0.4689
Brain	0.0005	1.5311	-

show the behavior of a wave in the system; at the air-skull and air-brain interfaces, the transmitted wave nearly doubles in magnitude. A wave impinging upon a brain-air, skull-air, or skull-brain interface will result in phase inversion, while all other interactions do not. These inversions are a mechanism for producing tension in the brain as the blast deposits a positive-phase wave into the head. Finally, at the brain-skull and skull-brain interfaces, the incident wave is significantly altered, and a substantial portion of its energy remains in the incident medium. These behaviors dominate the dynamics to first order and clearly can lead to tensions.

The blast itself is a strongly nonlinear wave that, as mentioned, is characterized by a shock front followed by a rarefaction. In contrast to a Sedov-type blast [217], the blast waves typically measured in the field are Friedlander-type blasts [78], whose rarefactions have a period of negative

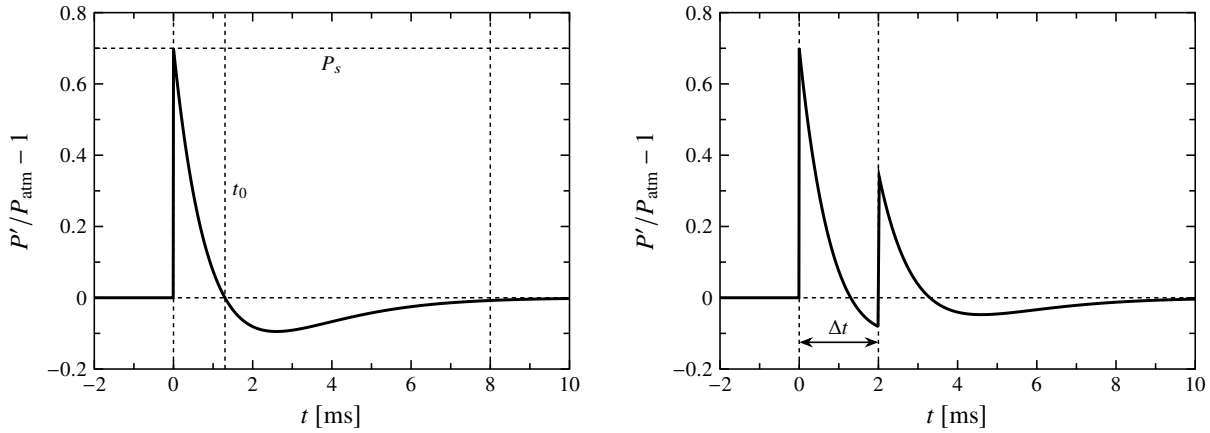


Figure 6.2: Simple (left) and compound (right) Friedlander blast profiles, with properties similar to those in [269] and [268].

gauge pressure [269]. A canonical Friedlander blast given by

$$P'(t) = P_s e^{-t/t_0}\left(1 - \frac{t}{t_0}\right), \tag{6.4}$$

where P' here is the gauge pressure, P_s is the peak pressure, and t_0 is the earliest time at which the blast gauge pressure is zero, corresponding to the beginning of the negative period. Figure 6.2 shows this profile, as well as a compound Friedlander blast. In terms of wave dynamics, this period of negative gauge pressure is not particularly influential because the speed of sound in the tissues of the head is substantially higher than the blast speed. Computing the magnitude of the wave transmitted into the head by the blast can be done via the solution to a Riemann problem; however, to first order, the above relationships for reflection and transmission capture the behavior of interest. Computationally modeling the blast in a fluid context is straightforward, and numerical solutions of a blast interacting with another fluid are well-behaved and validated. Many situations may lead to a compound blast profile, including multiple blasts, confined geometries, or weapons-specific discharges [268]. This is typically captured by fitting multiple Friedlander profiles additively, with the time delay between two shock fronts given by Δt, as shown in figure 6.2. If a secondary peak in a blast interacts constructively, then the magnitude of the waves inside the head may increase substantially due to superposition. Superposition is the phenomenon in which two interacting waves form a new wave whose amplitude is the sum of the amplitudes of the original waves, for linear waves. For example, if a positive-phase wave enters the brain from a secondary blast peak and interacts with a negative-phase (tension) wave inside the brain from the first blast peak, then the resultant wave will be smaller in magnitude than either of the incident waves. Conversely,

if two negative-phase or positive-phase waves interact, the tension or overpressure will increase, respectively, potentially leading to increased cavitation damage. Thus, considering compound blast profiles is important for understanding pathways to tension in the brain. In multiple dimensions, geometric effects such as focusing can make these analyses more complicated in theory, but the outlined principles still hold, particularly along axes of symmetry.

Connecting these results to cavitation is straightforward; if a tension wave of sufficient magnitude and duration (wavelength) interacts with a microbubble in the brain tissue, then growth will occur. Nanobubbles are suggested to naturally occur in the brain, although inhomogeneities inside the skull cavity, such as microvasculature, may act as preferential nucleation sites [7, 159]. Injury is expected if a bubble grows large enough to cause physical damage; this degree of growth is typically described via the *hoop* or *cortical* stretch [70]. Quantitatively, spherical bubble dynamics can be modeled with the Keller-Miksis equation [122]. This equation describes the evolution of a bubble under the influence of an external pressure field and can be solved semi-analytically by specifying the external pressure and integrating the resulting ordinary differential equation. From a pressure field, such as that from a simulation, the pressure at a given location can be extracted over time and used as input to the Keller-Miksis equation. The maximum bubble radius can then be compared to the original radius to compute the hoop stretch, which is a reliable metric for predicting neuronal damage [70]. In the context of the above theoretical results, it is clear that a positive-phase wave, as deposited in the head by blast, will reflect at the cranium and become a tension wave that, depending on its magnitude and wavelength, can directly lead to cavitation in the brain tissue.

6.4 Models and Methods

6.4.1 Physical Models

With the above simplifications, we propose a sequence of one-, two-, and three-dimensional models to investigate different aspects of the wave that may lead to cavitation-inducing tension waves. Figure 6.3 shows these models, as well as the boundary conditions for each that would be applied in a numerical context. The three-layer model is the simplest, assuming that the blast is planar and infinite in the transverse direction, while the head is modeled as a slab with water-like properties and thickness consistent with the front-to-back thickness of an adult male head. With this model, the ideas of reflection and transmission are easily observed without the added complexity of additional layers or geometric effects. Additionally, tension waves can be observed directly in this model after the blast reflects and inverts at the posterior boundary, thereby directly implying cavitation potential. The five-layer model is a direct extension of the three-layer model and includes a skull layer at the anterior and posterior boundaries of the brain. The fluid effects of the skull

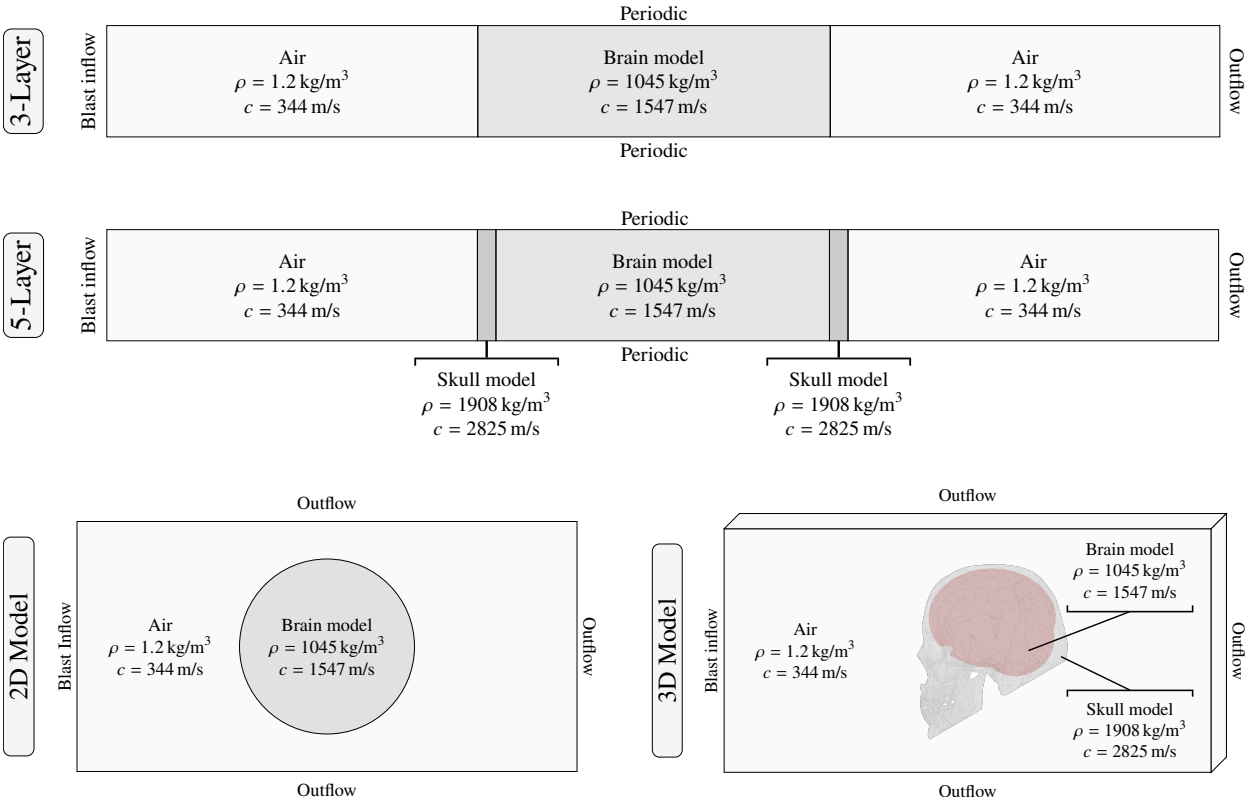


Figure 6.3: One-, two- and three-dimensional models. The three-layer model demonstrates the effects of reflection and transmission, while the five-layer model shows the changes induced by the skull, as well as the resultant superposition. The two-dimensional model demonstrates geometric effects including focusing, while the three-dimensional model accounts for more realistic geometries.

are directly observable in this model, with the expected behavior, as shown in Table 6.2, being that the skull will transmit some of the incident energy into the brain, while a substantial portion will be reflected back into the skull. The wave transmitted into the brain will then propagate to the posterior brain-skull interface, where, via a sequence of reflection and transmission events, the wave will invert in a similar manner to the three-layer model. Clearly, these models ignore the many additional layers of the head, but, as detailed above, only the skull will substantially change the wave propagation.

These one-dimensional models predict that cavitation will occur in the anterior of the brain, if it were to occur. This is in contrast to finite element and experimental results, which suggest that large tensions are expected to occur in the posterior of the brain [281]. This is easily explained by and corrected in the two-dimensional model, which recovers geometric effects via a cylindrical head surrogate. In short, the curvature of the cylinder changes the dynamics in that the transmitted wave has an inherent curvature, as does the reflecting surface (the posterior cranial wall), both of which lead to the generation of a tension wave in the posterior of the brain as opposed to the anterior, all still within the context of linear acoustics. One could include a two-dimensional model with a skull, but, as demonstrated by the three- and five-layer models, the qualitative dynamics are not substantially changed, and tensions of similar magnitude still emerge.

The three-dimensional model serves to elucidate the complexity of the real blast interaction, even in this simplified context, but using a simple skull-brain surrogate. Despite the added geometric complexity, this model still produces significant tension in the posterior of the brain via the same acoustic wave mechanisms as in the simpler models. Thus, this sequence of models demonstrates specific wave phenomena that lead to tension in brain tissue using only the principles of acoustics, thereby suggesting a wave-based pathway to cavitation.

6.4.2 Numerical Methods

In order to study the described dynamics in each model quantitatively, we use numerical simulations of the multi-species Euler equations, given by,

$$\begin{aligned}
\frac{\partial \alpha_k \rho_k}{\partial t} + \partial_{x_j}\left(\alpha_k \rho_k u_j\right) &= 0, \\
\frac{\partial \rho u_i}{\partial t} + \partial_{x_j}\left(\rho u_i u_j + \delta_{ij} P\right) &= 0, \\
\frac{\partial \rho E}{\partial t} + \partial_{x_j}\left(\rho E u_j + P u_j\right) &= 0, \\
\frac{\partial \alpha_1}{\partial t} + \partial_{x_i}\left(\alpha_1 u_j\right) - \alpha_1 \partial_{x_j} u_j &= 0,
\end{aligned} \tag{6.5}$$

where α_k and ρ_k are the volume fraction and density of material k, respectively, $\mathbf{u} = [u_1, u_2, u_3]$ is the Eulerian fluid velocity, $E = e + u_i u_i/2$ is the fluid total energy, P is the fluid pressure, and δ_{ij} is the Kronecker delta. The system is closed by a convex equation of state,

$$\rho_k e_k = A_k P + B_k = \rho_k C_k T + \frac{B_k}{1 + A_k} \tag{6.6}$$

where A_k, B_k, and C_k are experimentally fit properties for the air, skull, and brain surrogates, and choosing $B_k = 0$ recovers the ideal gas equation of state. This leads to the mixture relationships,

$$\rho e = \rho E - \frac{1}{2} u_i u_i \sum_{k=1}^{N} \alpha_k \rho_k = \sum_{k=1}^{N} (\alpha_k A_k P + \alpha_k B_k) = \sum_{k=1}^{N} \left(\alpha_k \rho_k C_k T + \alpha_k \frac{B_k}{1 + A_k} \right). \tag{6.7}$$

and sound speed

$$\rho c^2 = \frac{\sum_{k=1}^{N} \alpha_k A_k \rho_k c_k^2}{\sum_{k=1}^{N} \alpha_k A_k} \tag{6.8}$$

with $N = 2$ or $N = 3$ depending on the model. This equation of state allows us to match both the density and the sound speed in brain and skull surrogates with experimentally computed values. We discretize this system with a second-order MUSCL finite volume scheme [250], with the weak form given by,

$$\frac{\partial \bar{U}_j}{\partial t} + \frac{1}{|\Omega|} \int_{\partial\Omega} \hat{F}_{ij} n_i dS = 0 \tag{6.9}$$

where $\partial\Omega$ is the surface of any given cell Ω, $\bar{U}_j$ is the cell-average of the j-th component of the state vector $\mathbf{U}$, $\hat{F}_{ij}$ is the numerical flux of the j-th component, and n_i is the outward-pointing normal vector on the surface $\partial\Omega$. Adaptive mesh refinement (AMR) is enabled via the AMReX framework [282]. The `COMPAS` code is used for all simulations (Chapter 2). The weak form, Eq. (6.9), is integrated in time with a third-order total-variation-diminishing Runge-Kutta scheme. In all models, the blast is weakly enforced as a pressure Dirichlet boundary condition according to the Friedlander profile, Eq. (6.4). Outflow boundaries are used elsewhere except in the one-dimensional models, which are simulated in two dimensions with periodic transverse boundaries to produce exactly one-dimensional solutions.

To validate the method and model, we compare our numerical solutions with the experimental results of [218] for a Mach 2.4 shock. This case is useful for validation as it includes nearly all of the wave phenomena of interest, both linear and nonlinear, for fluids with relevant properties. Figure 6.4 shows experimental results from from [218]. We compare these directly with simulations. Figure 6.5 shows numerical Schlieren contours at four times with the present method. Strong agreement is observed, which gives us confidence to apply the described method to the above models. Figure 6.6 additionally shows the minimum pressure observed after the first reflection

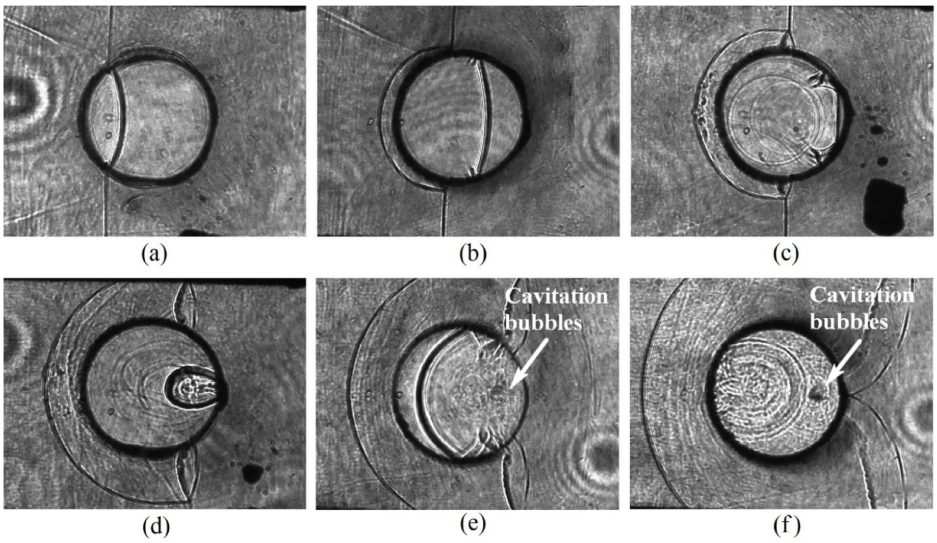


Figure 6.4: Experiments from [218] for the interaction of a Mach 2.4 shock with a column of water. Reprinted from S. Sembian, M. Liverts, N. Tillmark, and N. Apazidis. Plane shock wave interaction with a cylindrical water column. *Physics of Fluids*, 28(5), 2016, with the permission of AIP Publishing.

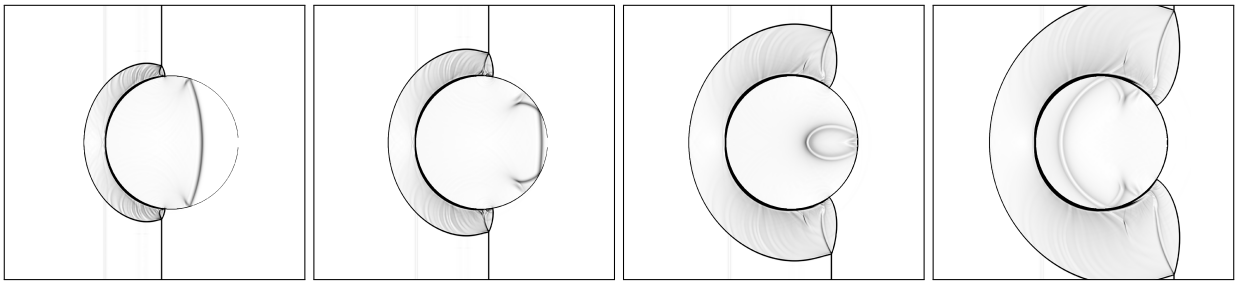


Figure 6.5: Numerical Schlieren contours for the interaction of a Mach 2.4 shock with a column of water at times 12.5 μs, 16 μs, 22.5 μs, and 29.5 μs, corresponding to (b)-(e) in Figure 6.4.

in the three-layer model. The minimum pressure quickly levels off with increasing AMR level, particularly after two levels, suggesting that three levels of refinement are sufficient to capture the pressures of interest.

6.5 Results

We begin with the three-dimensional model described in §6.4.1. In this model, the geometries of the skull and brain surrogates complicate the wave dynamics as compared to a planar model. Regardless, the acoustic principles outlined in §6.3 still hold, and tension is expected to be observed in the brain tissue. Figure 6.7 shows numerical results of the pressure field at three times. As

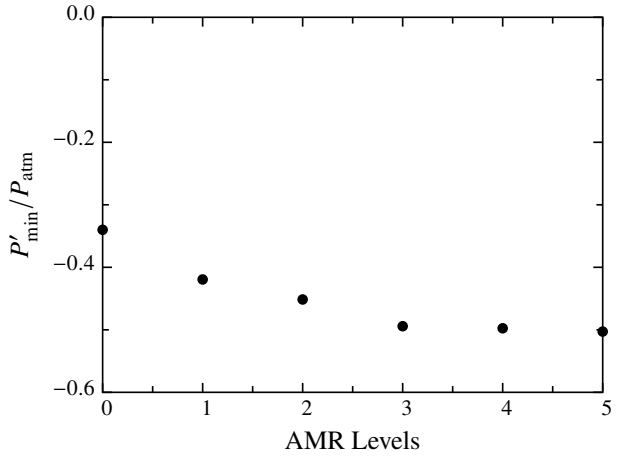


Figure 6.6: Minimum pressure vs. number of AMR levels.

expected, the wave deposited by the blast leads to tensions in the posterior of the brain surrogate. The wave dynamics underlying this behavior are, however, complex, and studying the details of the one- and two-dimensional models will further elucidate the mechanisms driving the observed tension.

6.5.1 One-Dimensional Models

6.5.1.1 Three-Layers

The simplest model we consider is the three-layer system shown in Figure 6.3, which models a planar brain surrogate subjected to a blast. When the blast impinges upon the brain surrogate, a positive-phase wave is reflected back with nearly the same magnitude as the incident wave, while a transmitted wave is induced, which is nearly twice the magnitude of the incident wave, as shown in Table 6.2. Based on the arguments above, this transmitted wave is expected to be a linear acoustic wave, and the qualitative behavior of the waves follows directly from §6.3. Figure 6.8 shows a representative x-t diagram for the wave front. The transmission of the blast into the brain surrogate occurs at point (a), where the propagation speed of the wave increases due to the change in impedance. This wave propagates to the posterior of the brain surrogate largely unaltered until point (b) where $Z_2/Z_1 \ll 1$ such that reflection is the dominant mechanism. The vast majority of the wave is reflected back into the brain surrogate while only a small portion is transmitted into the air. As seen in the reflection and transmission coefficients in Table 6.2, due to the sign of ΔZ, the new reflected wave is phase inverted from the incident wave, potentially leading to tension. If no other effects are considered, then this inverted wave would simply propagate from one side of the brain surrogate to the other, inverting phase and losing a small amount of energy each time

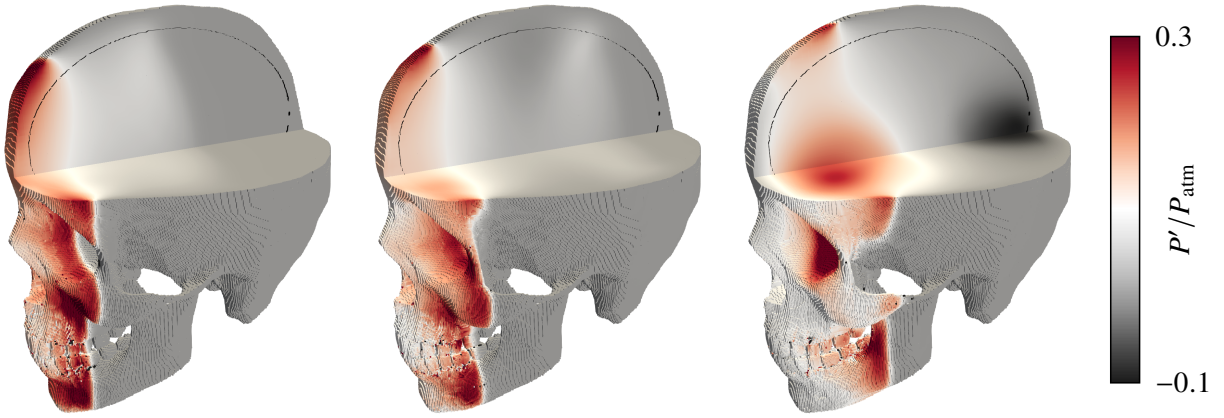


Figure 6.7: Three representative times for the three-dimensional model described in §6.4.1. In the first frame, the blast has just impinged upon the skull, and the transmitted wave is propagating through towards the brain. The leading transmitted wave is propagating through the brain in the middle frame, and in the final frame, reflection, superposition, and focusing lead to a region of tension in the posterior of the brain tissue.

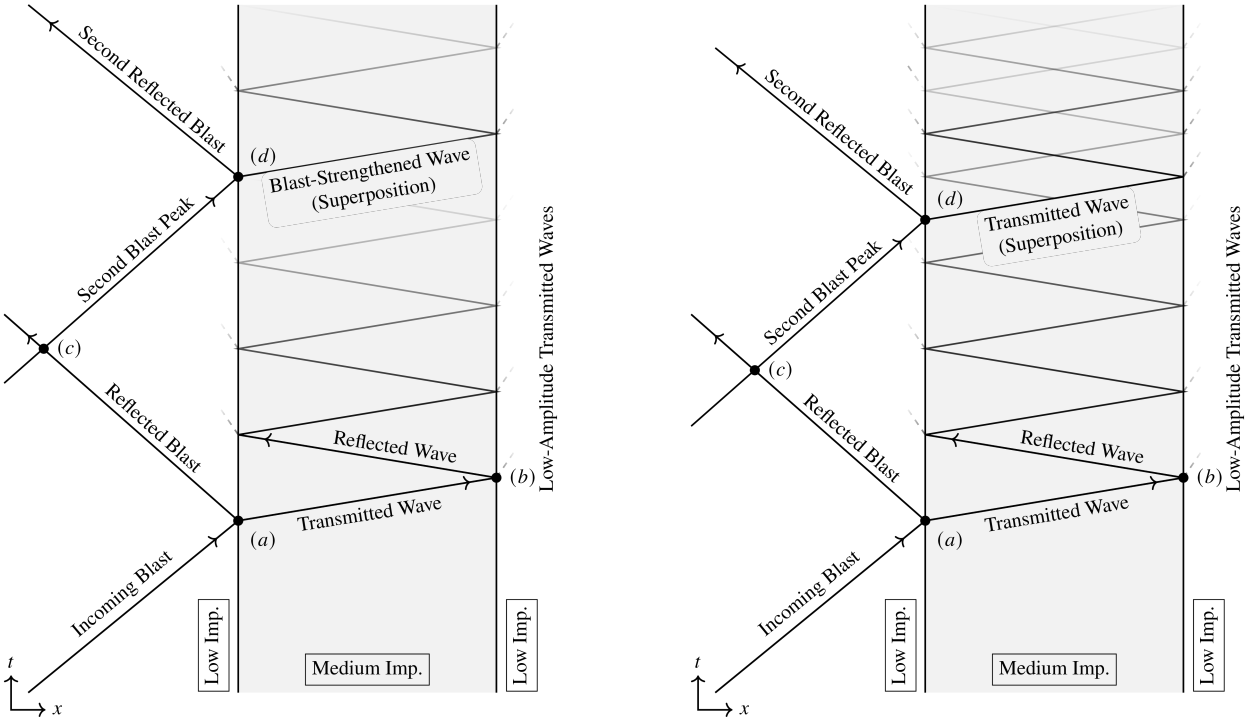


Figure 6.8: Demonstrative x-t wave diagrams for the three-layer one-dimensional model with constructive (left) and destructive (right) superposition. The two-peak blast propagates in low-impedance air until it reaches the high-impedance brain surrogate at point (a), which transmits a wave that travels to point (b) where it is reflected. The second blast peak intersects with the reflected blast at point (c) and impinges upon the solid at point (d), where it transmits a wave that superposes with the wave in the brain surrogate and either increases (left) or decreases (right) the intracranial wave's magnitude.

as low-amplitude transmitted waves. This phase-inversion process immediately suggests acoustic wave propagation as a mechanism for inducing cavitation.

Consider a compound blast profile, as discussed §6.3, which has two limiting behaviors; if the second peak reaches the air-brain interface at the same point as the interior wave is reflecting at the anterior interface then the wave deposited by the second blast peak will contribute to the magnitude of the interior wave and the resultant wave will have a larger magnitude, potentially leading to stronger tension waves. This behavior is shown qualitatively in Figure 6.8 (left). Conversely, if the second blast peak reaches the anterior interface while the interior wave is at the posterior interface, the positive-phase of the newly-deposited wave from the second blast peak will superpose with the tension wave such that the resultant wave is weaker, leading to lower tensions. This destructive behavior is shown qualitatively in Figure 6.8 (right). In reality, it is unlikely that secondary blast peaks will be exactly coincident or opposite the interior wave, and the resultant behavior will be a (linear) combination of the two effects. This demonstrates that the additive nature of the intracranial waves needs to be considered when developing blast threshold guidelines. The complex nature of many realistic blast profiles often yields multiple peaks [269] that deposit energy into the head and may interact constructively to produce conditions favorable for cavitation.

We simulate the above systems using the method described in §6.4.2 to observe the pressure field produced in each case. We first simulate a case like that in Fig. 6.8 (left), in which the second blast peak arrives at the anterior interface of the brain surrogate coincident with the interior wave front, i.e., $\Delta t = 2r_h/c_{\text{Brain}}$. Figure 6.9 (left) shows this numerically. The wave structure is largely the same as in Fig. 6.8 (left), and the described features are clear. The increase in pressure at point (a) is now clearly seen, as well as the reflected and transmitted waves. The reflection at point (b), and subsequent interior reflections are as expected, with a phase inversion at each. The magnitude of the interior wave is also clearly decaying over time, due to a combination of the decreasing exterior pressure on the proximal side of the solid and, to a lesser extent, the exit of low-amplitude acoustic waves. Point (c) is unremarkable, but point (d) clearly shows the pressure amplification that arises upon the impingement of the second blast onto the brain surrogate, which leads to higher pressures and tensions in the brain. The pressure and tension experienced after the second blast are greater than those caused by the initial blast, despite the first blast peak having a larger amplitude. This demonstrates how superposition can lead to more extreme conditions inside the brain than would be expected from two distinct blasts. As such, lower amplitude blasts with compound profiles should be considered potentially dangerous when developing exposure thresholds.

Consider $P(t)$ along $x/r_h = -0.9$ in Figure 6.9; the peak tension pulse has a duration of $\sim 20\,\mu$s, tension magnitude of $P'/P_{\text{atm}} \sim 0.75$, and a nearly Gaussian profile. Numerically solving the Keller-Miksis equation for an initially $10\,\mu$m bubble and the described Gaussian tension pulse leads to a predicted hoop stretch of around two, suggesting damage may be significant [70]. This

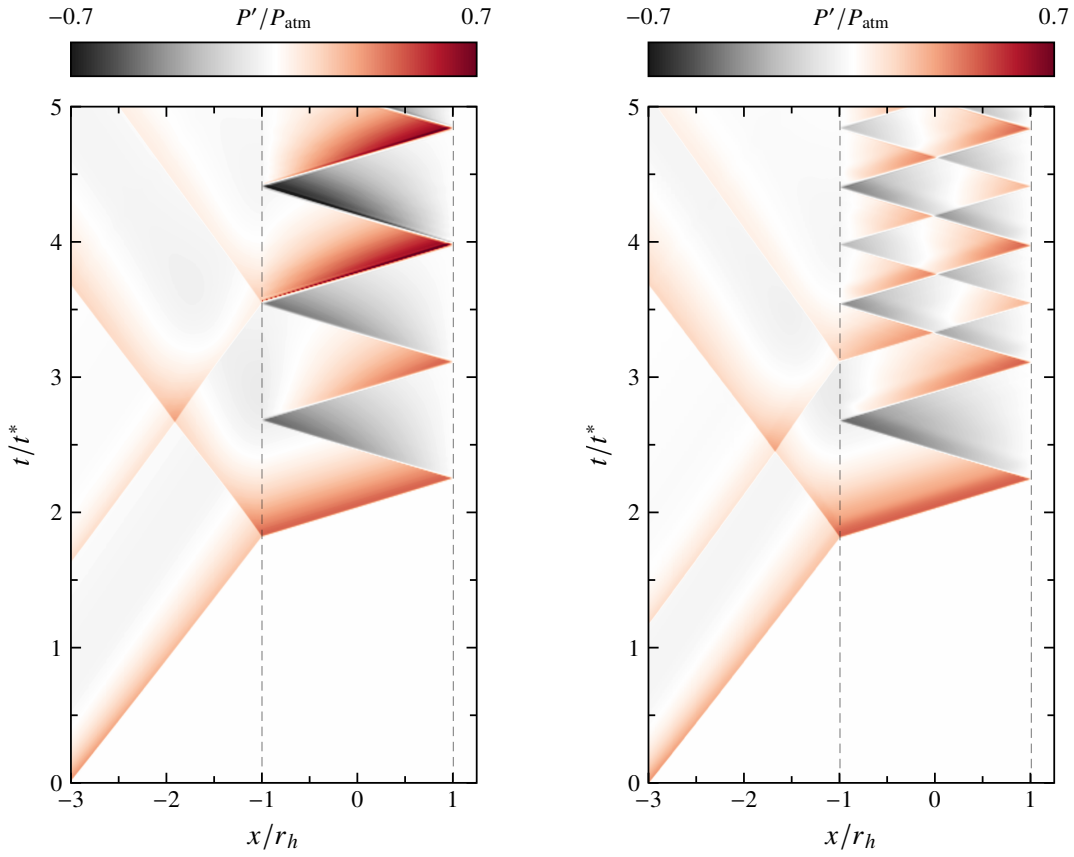


Figure 6.9: Pressure contours in x-t space for a two-peak blast interacting with a brain surrogate. Left: Contour for blast peak lag time $\Delta t = 2\,\text{ms}$. Right: Contour for blast peak lag time $\Delta t = 1.88\,\text{ms}$. The dashed lines represent the boundaries of the brain surrogate.

rough calculation indicates that cavitation should be studied in this regime.

Now consider a blast with $\Delta t = r_h/c_{\text{Brain}}$ such that the second blast peak arrives at the anterior air-brain interface while the interior wave is at the posterior brain-air interface, as shown in Figure 6.9 (right). At point (d), in this case, the second blast peak transmits a wave into the brain surrogate and reflects, as in the prior case. The transmitted wave does increase the pressure inside the solid, but not nearly as much as in the coincident case. Due to the linearity of the waves inside the solid, the interior wave structures from both blasts are clearly seen overlaid, but this leads to the important consequence that the lowest pressure region of the first wave is superposed with a high-pressure region of the second wave, essentially eliminating the tension observed in the coincident case. The interior waves then propagate back and forth inside the solid, losing a small amount of energy upon each reflection. This result contradicts the conclusion from the previous blast profile, but this discrepancy just highlights that the damage potential is sensitive to small changes in exposure conditions. Further, despite the observed reduction in tension, the degree remains significant and should be carefully considered. In this case, the measured tension is less severe than in the case above, with a peak tension of $P'/P_{\text{atm}} \sim 0.5$. Regardless, the Keller-Miksis equation again predicts growth, with a hoop stretch of about 1.5, suggesting damage may occur. As such, both one-dimensional models predict potential cavitation damage, at least for the parameters considered here.

6.5.1.2 Five-Layers

We now investigate the effect of the skull on the wave propagation via the five-layer model in Figure 6.3. Figure 6.10 (left) shows these dynamics qualitatively. The interaction of the incoming blast with the skull surrogate at point (a) is very similar to that with the brain surrogate in the previous section, as the change in impedance is similar. Once the wave is inside the skull, however, the dynamics change significantly. When the wave inside the skull encounters the skull-brain interface, point (b), it is inverted and reflected with just over half of its original magnitude, while a transmitted wave of just under half the incident magnitude propagates into the brain surrogate. The wave inside the brain surrogate propagates to the posterior brain-skull interface, where another reflection-transmission event occurs at point (c), with magnitudes given in Table 6.2. At this point, a reflected wave travels back into the brain without a phase inversion and about half the amplitude of the incoming wave, and a transmitted wave propagates in the skull with about 1.5 times the amplitude of the incoming wave. This transmitted wave propagates to the posterior skull-air interface, point (d), where a very low amplitude wave is transmitted, and the majority of the incoming wave is reflected back into the skull with a phase inversion. This reflected wave then reaches the skull-brain interface at point (f), where another wave is transmitted into the brain. This wave lags the front induced at point (c), and superposition leads to their amplitudes summing.

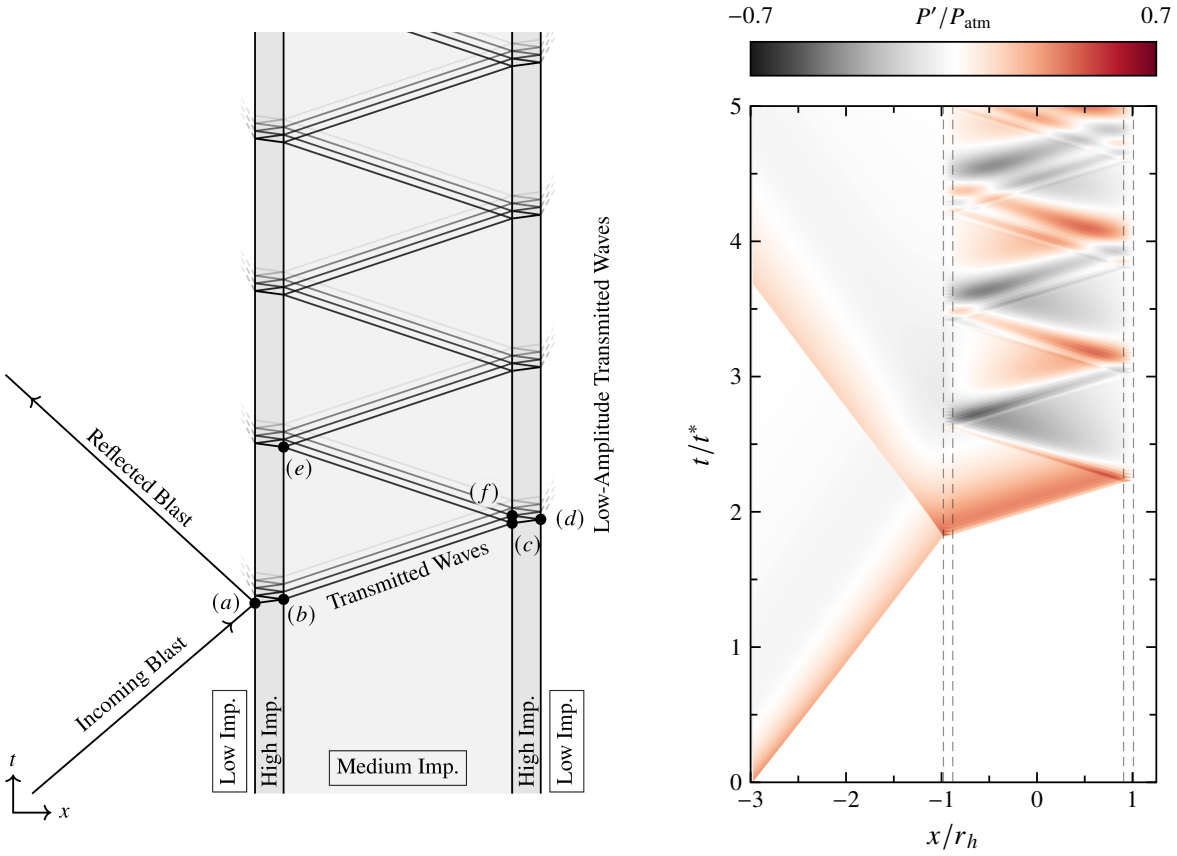


Figure 6.10: Left: A demonstrative $x-t$ wave diagram for the five-layer one-dimensional model. The blast propagates in low-impedance air until it reaches the high-impedance skull surrogate at point (a), which transmits a wave that travels to point (b), where it is partially reflected and partially transmitted into the brain surrogate. The intra-cranial wave propagates to point (c) where it is partially transmitted and reflected. At point (d), energy exits the head at the anterior skull-air interface and the wave inverts. At point (e), the wave front in the brain superposes with waves inside the skull and reflects, and at point (f), a new wave is transmitted into the brain surrogate. Right: Pressure contours in x-t space for a simple blast interacting with a brain surrogate. The dashed lines represent the boundaries of the brain and skull surrogates.

Subsequent interactions within the skull layers generate additional waves that act as a spreading mechanism, with each successive reflection depositing substantial energy between the skull and brain and decaying in magnitude. This sequence of waves, induced by the finite thickness of the skull, spreads the wave front such that, while the amount of energy deposited into the head is very similar to that in the three-layer model, the resultant tensions are expected to have a lower intensity, but still be substantial.

We simulate this system with the same blast conditions as in §6.5.1.1, but with a single blast peak, as shown in Figure 6.10 (right). Immediately, the overall wave structure appears similar between the five- and three-layer models, with alternating high-pressure and tension phases. As expected, however, the skull layers cause a sequence of reflections and transmissions that spread the wave front. This phenomenon is clearly observable in the numerical results as high-frequency reflections in both the anterior and posterior skull layers, as well as layered wave fronts. Another key distinction, though, is that, unlike the three-layer models, the dynamics are not qualitatively periodic due to the skull's effect. Over time, the tension and peak pressure decay as each wave front is split via reflection and transmission at the skull-brain interfaces. If the skull were substantially thicker, the front separation would be much more consequential, and the overall wave structure would significantly differ from that in the three-layer models. If the skull were substantially thinner, the reflections inside the skull would occur quickly, such that the resultant wave front inside the brain surrogate would be less affected. The measured tension and pulse width in this five-layer case, despite the more complex dynamics, are substantial and are expected to lead to microbubble growth.

This model expands upon the three-layer model via the addition of the skull layers and importantly demonstrates the following: the wave dynamics are qualitatively similar with and without the skull, with alternating compression and tension phases; the skull spreads the wave front by a sequence of fast reflections within the skull; the dynamics are not qualitatively periodic in time due to the spreading effect of of the skull, such that the first period is most consequential; and measured tensions are of sufficient duration and magnitude to be relevant for microbubble cavitation. This model again shows that the dynamics leading to cavitation-relevant conditions can be described solely through acoustic analyses.

6.5.2 Two-Dimensional Model

With the one-dimensional dynamics characterized via the three- and five-layer models, we move to the two-dimensional model shown in Figure 6.3 to investigate geometric effects. Based on the conclusion that the skull does not substantially affect the dynamics, particularly the first posterior and anterior reflections, we neglect it in the two-dimensional model.

Figure 6.11 (left) shows pressure contours of the wave interaction with the cylinder. As the blast reaches the anterior air-brain interface, it generates a wave in the brain surrogate, as in the planar model. Due, however, to the curvature of the cylinder, the resultant wave within the cylinder is also curved. Further, because the wave spreads as it propagates across the cylinder, its magnitude decays. As the cylinder's curvature inverts, the wave reflects from the now-converging surface, and a reflected wave propagates toward a focal point at the rear of the cylinder. This effect, combined with the reflection at the posterior brain-air interface, leads to a point of high tension in the posterior region of the brain surrogate. This is contrary to the location of peak tension observed in one-dimensional models but is consistent with both theoretical expectations and past work [281]. The location of the peak tension depends on the curvature of the head (geometric focal point), the curvature of the transmitted wave (which weakly depends on the blast strength), and the direction of the incoming blast (e.g., side blast versus back blast), but generally occurs on the side opposite to the incoming blast.

We simulate this system, as shown in Figure 6.11 (right). The cylinder is a circle of radius $r = r_h = c_H/2\pi$, where $c_H = 0.57\,\mathrm{m}$ is a typical head circumference for an adult male, and the blast is planar with the same properties as in the one-dimensional models. The numerical results match theoretical expectations and show the propagation of the transmitted wave across the cylinder, as well as the peak tension occurring in the rear of the cylinder. The pressure in x-t space clearly shows again the amplification of the blast as it transmits into the brain surrogate and the amplitude decay of the transmitted wave as it propagates to the posterior brain-air interface. The most distinctive feature of this model is the peak tension observed following the first reflection at the posterior interface, which results from the inversion of the curved wave front as it reflects along the cylinder's boundary. This behavior is well documented [281] and is a clear mechanism for the generation of tensions in the rear of the brain in this cylindrical model.

The side reflections, as well as the mismatch between the curvature of the cylinder and the curvature of the wave front over time, lead to a complex wave structure following the first posterior reflection, as seen in Figure 6.11. Thus, as with the five-layer model, it is implied that the first posterior reflection is likely to produce the greatest tension, as the superposition of compression and tension waves generally leads to lower peak amplitudes. That being said, cavitation-relevant tensions may still occur later in time, and at various positions, but likely with a smaller magnitude. Combining this geometrically induced transience with the transience due to the skull suggests that, indeed, the first posterior reflection is expected to be the most consequential.

In this case, the pulse duration for the peak tension is around $100\,\mu$s and has a similar magnitude to those in the previous cases, $P'/P_{\mathrm{atm}} \sim 0.5$, such that we would expect microbubbles to grow significantly. In water, the growth of cavitation bubbles at the location predicted by the above results is consistent with experiments [218]. This implies that the predicted tensions are indeed

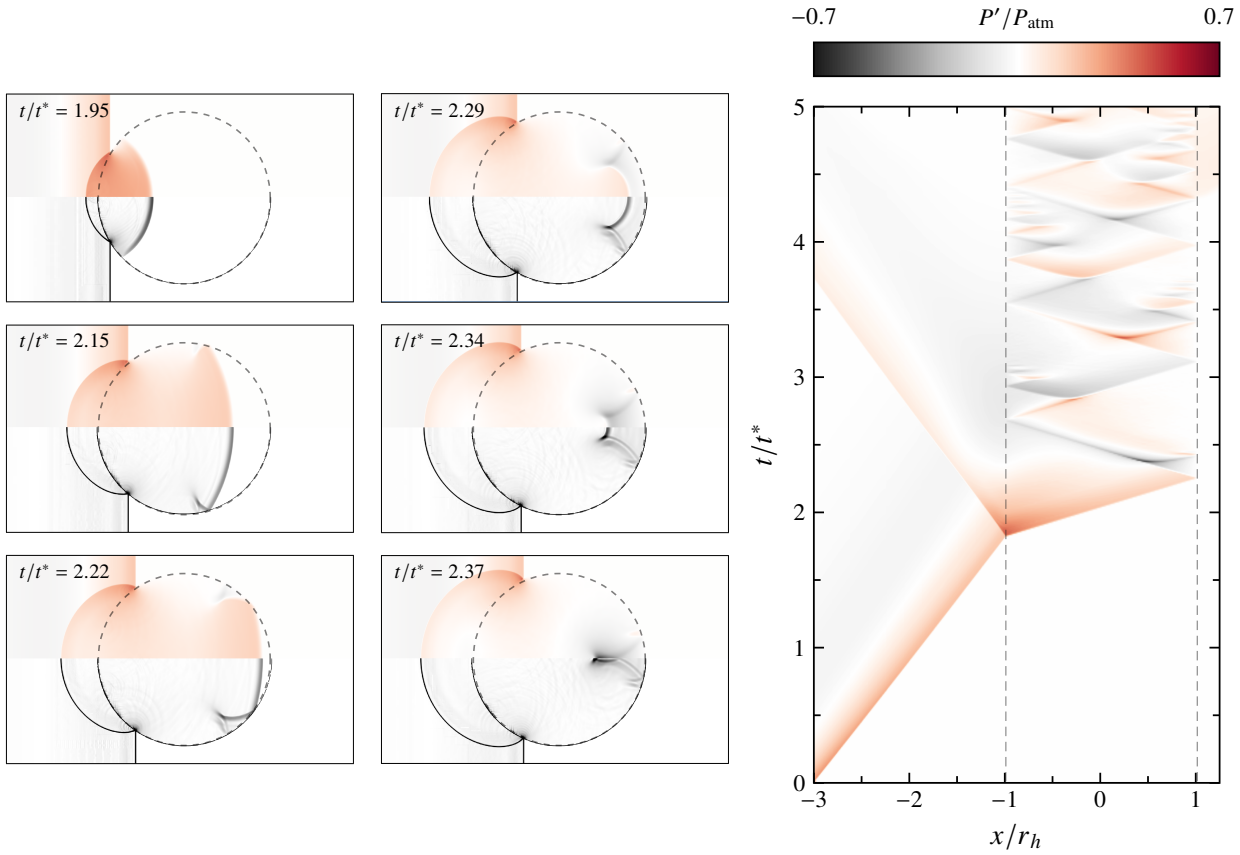


Figure 6.11: Left column: Pressure (top) and numerical Schlieren (bottom) contours for the two-dimensional blast-cylinder model at different times. Right: Pressure contours in x-t space along the cylinder's centerline.

physical and are a plausible mechanism for damage.

6.5.3 Discussion

The above sequence of models demonstrates that acoustic mechanisms can directly lead to cavitation-relevant conditions, i.e., tension pulses of sufficient magnitude and duration to induce significant radial bubble growth. In each model, acoustic theory accurately predicts the wave behavior observed in fully nonlinear numerical fluid simulations, with the fundamental wave behaviors of transmission, reflection, and superposition (which can manifest as linear focusing in multiple dimensions) capturing all key dynamics. Specifically, the three-layer model was useful for showing the deposition of the blast into the head, the effect of compound blast profiles, and the wave dynamics that lead to tension. Additionally, using the numerically-computed tension waves alongside the Keller-Miksis equation suggested that the tensions produced from a moderate blast ($P'/P_{\text{atm}} \sim 0.5$) can produce conditions that may lead to damaging growth of cavitation bubbles. The five-layer model added physical accuracy by including the skull and demonstrated that, because the skull's impedance is higher than the brain's, the blast-induced wave front becomes stratified by high-frequency reflections within the skull. Further, the sequence of waves released from the skull layers led to lower tensions over a short time, unlike in the three-layer model. Despite this, the overall wave dynamics were not substantially altered, and cavitation-relevant conditions were observed once again. Both one-dimensional models predicted that the peak tension would occur in the anterior region of the brain. Geometric focusing was accounted for in the two-dimensional model, while the skull's effect was ignored. This model importantly reconciles the inconsistency between experiments and simulations, which suggest that the peak tension (and potentially cavitation) will occur in the posterior region of the brain, rather than the one-dimensional results that suggest it will occur in the anterior region. This shift is purely a product of the geometry, with the curvature of the cylinder inverting as the wave front propagates from anterior to posterior and the edges of the front reflecting continuously. This behavior still leads to significant tensions, as predicted to be relevant to microbubble cavitation damage.

Combining the results from the one- and two-dimensional models makes the observed dynamics in the three-dimensional model clear, as shown in Figure 6.7. The blast deposits a wave into the skull surrogate, which propagates to the brain surrogate, reflecting and transmitting waves of similar magnitude, leading to a diffused wave front. This wave, which is curved by the head's curvature, propagates across the brain, decaying in magnitude. As the curvature inverts, the wave reflects on the inner cranium, leading to the observed tension in the posterior of the brain. Even though the overall wave dynamics are clearly more complex than in any of the other models, due to the geometry of the skull and brain surrogates, the principles that govern the dynamics in those

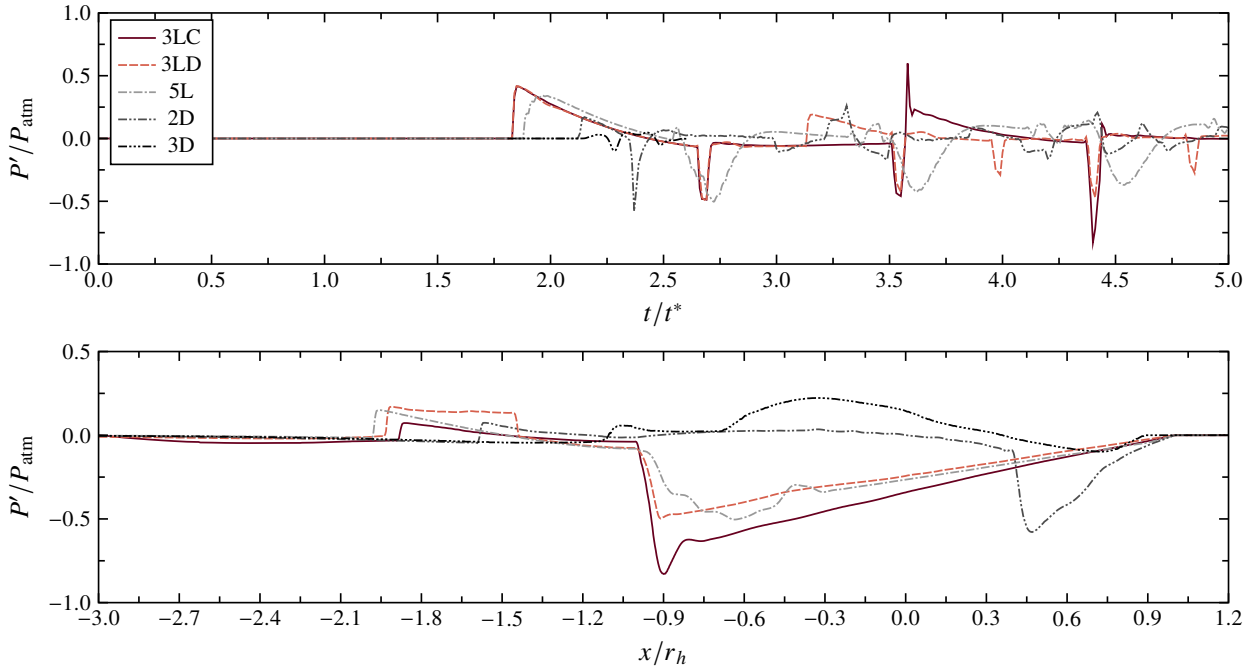


Figure 6.12: Pressure over time at the position of peak tension in each model (top) and center-line pressure at the time of peak tension for each model (bottom), where the three-layer model with constructive superposition of the second blast peak is denoted "3LC", the three-layer model with destructive superposition is "3LD", the five-layer model is "5L", the two-dimensional model is "2D", and the three-dimensional model is "3D".

simpler models hold and are directly applicable to this more complex system. Further, based on both three-dimensional and one- and two-dimensional results, we would expect that tension waves would manifest in the rear of the head for a front blast, with the pulse length and magnitude potentially sufficient to cause damaging cavitation.

In terms of cavitation, the pressure at the location of peak tension over time is useful for characterizing the pulse to which a microbubble may be subject. Figure 6.12 (top) plots $P(t; x = x(P = P_{\min}))$ for each model, where $P_{\min}$ is the minimum pressure (maximum tension) observed. Unsurprisingly, the one-dimensional models all produce relatively similar results, with significant tension emerging over multiple periods. The multidimensional models, however, exhibit a single primary tension event at the location of peak tension and only mild tension thereafter due to intracranial reflections.

Similarly, the centerline pressure at peak tension is useful for comparing the location of peak tension across models. Figure 6.12 (bottom) shows $P(x; t = t(P = P_{\min}))$ along the center line. As expected from the above results, the peak tension occurs in the posterior region of the brain surrogate for the two- and three-dimensional models, whereas it occurs in the anterior region for the one-dimensional models. While the tension magnitude is much lower in the skull case, this is

to be expected, as this is a relatively weak blast.

Regardless, in each model, the measured tensions are consistent with expectations from acoustic theory and are of sufficient duration and magnitude to suggest the potential for cavitation damage. As such, the presented sequence of models, alongside theoretical and numerical results, demonstrates that waves alone can lead to cavitation-relevant conditions in the brain.

Several features of the presented framework could be modified to account for higher-order effects. For example, the nonlinearities in wave deposition from the blast would be better modeled by a Riemann problem. As the blasts typically measured in experiments are not exceptionally strong [269], the results from such a calculation are not expected to differ substantially from the present analyses. Another consideration would be the changes to the blast front due to oblique interaction with the cylinder or skull, or from a curved blast front. Once again, however, the results considering such effects are not expected to differ substantially from the present results because intracranial wave propagation occurs on a shorter time scale than the exterior shock physics. Nonlinear geometric focusing, unlike the linear focusing considered here, is another effect expected to be insignificant (due to the small Mach number) but may increase fidelity. Finally, the most significant change to the modeling done herein would be to account for elastic (skull) and visco-elastic (brain) effects. This is out of the scope of the present work, but based on the discussion in §6.3, these effects are not expected to change the qualitative behavior demonstrated here.

6.6 Conclusions

In this work, it was demonstrated how purely acoustic effects induced by a blast can lead to conditions relevant to cavitation damage inside the brain. The theoretical discussion justified the use of acoustic dynamics to study bTBI in this context, and the fundamentals of wave dynamics and cavitation were presented. A sequence of one-, two-, and three-dimensional models was used to demonstrate the mechanisms governing wave dynamics that may lead to tension in the brain. The one-dimensional models were used to simply show the blast depositing the wave into the head, the generation of tension via reflections, the additive effect of a compound blast profile, and the changes induced by the addition of a layer with notably different impedance than the brain surrogate. The two-dimensional model was used to demonstrate geometric effects, such as focusing, and to resolve the unexpected peak tension location observed in the one-dimensional models. Finally, the three-dimensional model wrapped together the five-layer and two-dimensional models to simultaneously account for multiple layers, complex geometries, and three-dimensional dynamics. In all models, substantial tension was observed. The discussion of each model was paired with numerical simulations that quantified the pressures generated in each system for a moderate blast. Further, these numerical results were used alongside the Keller-Miksis equation

to verify that the observed conditions were indeed relevant to cavitation damage. Future work includes full finite-element modeling of blast-head interaction, accounting for elastic effects and more realistic geometries, and more detailed analyses of the potential for cavitation damage in the brain.

CHAPTER 7

Conclusion

This final chapter summarizes and discusses the key findings, addresses limitations, and suggests future research directions.

7.1 Summary of Key Findings

The preceding five chapters each represent novel research contributions. In Chapter 2, a new software is developed to study compressible N-material interfacial flows with consistent and conservative Phase-Field modeling. Built on the AMReX framework, `COMPAS` uses patch-based adaptive mesh refinement to resolve shocks and interfaces and is directly portable to CPU and GPU high-performance computing (HPC) systems. Key methods for conservation, reduction consistency, and interface regularization are implemented and adapted for adaptive mesh refinement (AMR). The modular inclusion of a variety of Phase-Field models makes `COMPAS` a useful tool for studying the effects of various regularization mechanisms, and the kernel-based programming streamlines further development. Verification, validation, and performance tests confirm both fidelity and efficiency, positioning `COMPAS` as a practical platform for future multi-physics studies of complex interfacial phenomena.

The consistent and conservative Phase-Field (PF) model is extended to a high-order discontinuous Galerkin (DG) scheme that incorporates non-conformal subcell finite-volume shock and interface capturing, as described in Chapter 3. This formulation preserves conservation and boundedness, maintains uniform interface thickness, and captures shocks while simultaneously mitigating numerical diffusion that leads to thick interfaces in conventional diffuse-interface implementations. Verification and validation against theory and experiments show that the scheme accurately represents both smooth fields and steep gradients across large density and pressure ratios, outperforming standard finite-volume schemes and DG approaches with limiters. The result is a robust, high-order accurate framework for simulating compressible interfacial flows with shocks.

Chapter 4 describes numerical simulations that reveal that the vortex cores generated by the

planar Richtmyer-Meshkov instability (RMI) can quickly become susceptible to vortex-line instabilities. The observed sinusoidal deformation and pinch-off reorganize nominally two-dimensional flows into three-dimensional structures, such as vortex rings, altering momentum transport and enhancing material mixing. Measured growth rates tend to classical Crow instability (CI) theory results as the Atwood number increases, establishing a previously unquantified pathway from shock-driven interface perturbations to large-scale coherent structures. These findings alter expectations for post-shock mixing and offer new metrics for assessing instability growth in environments such as supernovae and fusion energy experiments.

Three-dimensional simulations of shock-driven cylinder compression in Chapter 5 show that, similarly to the planar case in Chapter 4, Richtmyer-Meshkov-generated vortex cores experience the rapid onset of Crow-type vortex-line instabilities. Despite the geometric effects, including reshock, the measured perturbation growth is exponential and exhibits non-monotonic behavior with respect to wavenumber, as is characteristic of the CI, suggesting that vortex-pair instabilities are a probable mechanism for the observed structure formation. These results highlight that vortex-line instabilities may drive asymmetry in converging geometries (cylindrical and spherical), informing the design of double-shell inertial confinement fusion (ICF) experiments and similar systems.

In Chapter 6, a hierarchy of models shows how blast waves can generate tensions in the brain sufficient to trigger cavitation, offering a plausible explanation for blast-induced traumatic brain injury (bTBI) that earlier mechanisms under-predict. One-dimensional analyses identify the key reflection and superposition processes; two- and three-dimensional simulations, including a realistic cranial geometry, capture focusing, divergence, and layering effects that affect negative pressures. These results suggest intracranial cavitation as a potential damage mechanism and provide a simple framework to guide future experiments and numerical studies.

7.2 Discussion and Limitations

The code developed in Chapter 2 uniquely combines consistent and conservative Phase-Field models, GPU acceleration, and AMR into a single tool. Since the first version became available, `COMPAS` has been adopted by several researchers to address problems that extend well beyond the case studies presented in this dissertation. Application areas include high-speed droplet dynamics [37, 38], high-energy-density physics [256], astrophysical flows [259], hyper-elastic modeling of gas-liquid-solid flows [75], and methods development [104]. Peer-reviewed, conference, and under-review manuscripts utilizing `COMPAS` include:

1. T. M. Perez, S. C. Dick, R. F. Smith, P. M. Celliers, J. H. Eggert, S. J. Ali, E. Johnsen, and J. K. Wicks. Low viscosity of solid MgO at high pressures and strain rates measured using the laser-driven Richtmyer-Meshkov instability. *Physical Review B*, 111(14), April 2025.

2. W. J. White, Z. Huang, and E. Johnsen, A high-order discontinuous Galerkin method for compressible interfacial flows with consistent and conservative Phase Fields. *Journal of Computational Physics* 527:113830, 2025.

3. M. J. Wadas, W. J. White, H. J. LeFevre, C. C. Kuranz, A. Towne, and E. Johnsen. Hydrodynamic Mechanism for Clumping along the Equatorial Rings of SN1987A and Other Stars. *Physical Review Letters*, 132:111201, 2024.

4. M. Wadas, H. LeFevre, Y. Elmore, X. Xie, W. White, C. Kuranz, and E. Johnsen. Feasibility of an experiment on clumping induced by the Crow instability along a shocked cylinder. *Physics of Plasmas*, 31(6):062103, 2024.

5. W. J. White, Z. Huang, and E. Johnsen. A discontinuous Galerkin method for the five-equations multiphase model. In *AIAA SCITECH 2024 Forum*, Orlando, FL, 2024.

6. T. Huang, M. J. Zahr, K. Holst, J. Freels, AJ Baker, P. Batten, N. Bachchan, J. B. Maeng, T. Fisher, D. Kim, K. Yune, C. Kim, J. Kim, J. Park, S. K. Jee, W. J. White, and E. Johnsen. High-Fidelity CFD Verification Workshop 2024 Summary: Shock-Dominated Flows. In *AIAA AVIATION FORUM AND ASCEND 2024*. American Institute of Aeronautics and Astronautics, July 2024.

7. E. M. Burrell, Z. Huang, W.J. White, B. Perfect, and E. Johnsen. Resolution Considerations for Two-Dimensional High-Speed Droplet Impact Along a Rigid Wall. In *AIAA Aviation Forum and ASCEND co-located Conference Proceedings*, page 3822, Las Vegas, Nevada, USA, 2025.

8. Z. Huang, S. Mirjalili, M. Khanwale, S. Jain, and E. Johnsen. Investigation of Phase-Field models in compressible multiphase flows in a unified high-order and bound-preserving framework. Center for Turbulence Research, Proceedings of the Summer Program:327–337, 2024.

9. Z. Huang, W. J. White, and E. Johnsen. Consistent and conservative Phase-Field method for compressible two- and N-phase flows with adaptive mesh refinement. *Manuscript under review at Journal of Computational Physics*, 2025.

Several additional manuscripts leveraging `COMPAS` are also in preparation. Further, `COMPAS` has been used by other students for computational flow physics studies central to their PhD theses:

1. M. J. Wadas, *Hydrodynamics of Shocked Interfaces*, PhD thesis, University of Michigan, 2023.

2. S. C. Dick, *Hydrodynamic Instabilities in Condensed Matter at High Pressures*, PhD thesis, University of Michigan, 2024.

3. E. M. Burrell, *Numerical Investigation of High-Speed Droplet Impact Upon a Solid Surface*, PhD thesis, University of Michigan, 2025.

Numerous conference presentations have also relied on COMPAS, including talks delivered by researchers from four universities at APS DPP, APS DFD, ICNMMF, ICMF, AGU, AIAA SciTech, AIAA Aviation, MultiMat, and more. The widespread use of COMPAS illustrates the software's generality and usability: users can change models and methods while retaining a common, performance-portable back end. Ongoing work includes implementing magnetohydrodynamics, radiation hydrodynamics, and embedded boundaries to further increase functionality. This framework, however, currently lacks implementations of methods such as the recently introduced THINC (Tangent of Hyperbola for INterface Capturing) method for N-material compressible flows [283]. Implementing techniques such as THINC, ghost fluid [20], and VOF [16] alongside new physics models would make COMPAS a powerful platform for methods testing and physics studies of compressible interfacial flows. The development process, however, elucidated how challenging it is to simultaneously develop, maintain, and use a research code. Significant effort is needed for even just one of these, and attempting to do all three can slow progress. Utilizing available tools where possible can help alleviate this, with a clear example being AMReX, which radically reduced the effort required to achieve an AMR-enabled code.

The methods developed in Chapter 3 have been implemented into the COMPAS to provide a high-order accurate solver option for multi-material flows. The scheme's consistent, conservative, and high-order-accurate nature is uncommon [108, 241], and the general Phase-Field framework sets this method apart. Further, the subcell finite-volume method offers greater robustness and accuracy than many limiting techniques [66, 112, 126, 275]. The high-order solver, however, is limited by the lack of a robust DG implementation for patch-based AMR structures. Other researchers have studied DG with AMR [158, 213, 279], so combining these techniques with the present methodology presents an opportunity for a potent solver. Despite the advancements of the presented method, its development demonstrated the fragility of high-order solvers when applied to problems with discontinuities. The gains enabled by such a method should be weighed against the effort required to make it work for the problem of interest; sometimes, lower-order methods are sufficient if they enable rapid prototyping and discovery.

With the numerical framework verified and validated for interfacial and vortex instabilities, the dissertation shifts to physics studies. The generation of three-dimensional structure from nominally two-dimensional RMI is typically attributed to the growth of small transverse variations by the same mechanism that drives the dominant perturbation growth. However, such a theory would require

analyzing other studies of flows transitioning from two to three dimensions, calculating the RMI growth in the transverse direction, and determining whether the transition to three dimensions matches expectations. The simulations in Chapter 4 propose an alternate path to three-dimensional structure by demonstrating that a nominally two-dimensional RMI flow can generate significant three-dimensional structures (vortex rings) via vortex-pair instabilities, namely the CI. Unlike RMI-driven growth, which is generally linear, the CI classically grows exponentially, so that for close pairs of vortices or large perturbations, vortex-pair instabilities may dominate the dynamics and alter mixing rates. In simulations across three Atwood numbers, computed with `COMPAS`, exponential growth in the core perturbation amplitude has been observed. Further, the computed non-dimensional growth rates closely match the trends predicted from classical CI theory at the largest Atwood number. While other instabilities may be present, the close agreement with classical theory in both the optimal and cutoff wavenumbers suggests that the CI is the dominant instability mechanism, revealing a novel and potentially significant pathway to mixing in shocked systems with long-wavelength perturbations. Despite this, questions remain regarding the influence of mesh resolution, initial perturbation amplitudes, and plausibility in real systems. These uncertainties, coupled with the lack of a comprehensive theoretical framework, highlight the substantial work left to characterize this system.

Curvature adds new components to RMI flows, with geometric effects altering the already complex dynamics [140]. The simulations from Chapter 5 suggest that a similar mechanism of structure formation to that observed in the planar RMI case (Chapter 4) is active in converging cylindrical geometries. This result is significant for converging systems with long-wavelength or equatorial features, such as in double-shell ICF, where the assembly joint could be a source of long, converging pairs of vortex rings. The observed non-monotonic trend in growth rates and the exponential perturbation growth profile suggest that vortex-pair instabilities are responsible for the formation of secondary vortex rings, rather than other effects such as simple vortex reconnection. Quantifying this instability pathway in practical systems and assessing the potential impact on mixing models is a logical next step, with implications on ICF target design. To do this, however, requires a more detailed study of the shock, reshock, geometry, and materials, which would help assess the phenomenon's realistic importance. Additionally, this study demonstrated clearly how numerical inaccuracies in components of the method can lead to spurious features, such as small-scale RMI growth at the interface seeded by the Cartesian grid. Managing these features is crucial for accurately modeling interfacial physics.

Finally, Chapter 6 demonstrates the wave dynamics that may lead to conditions relevant to cavitation in bTBI. One-, two-, and three-dimensional simulations have been coupled with simplified wave mechanics to argue several points: (i) viscous, heat transfer, and elasticity effects are not substantial when considering the primary waves in the head, (ii) geometric focusing leads to a focal

point of peak tension, and (iii) the relative consistency of the acoustic impedance values of various biological structures inside the cranium suggests that modeling the skull and brain are sufficient to capture wave dynamics. Simulations conducted with COMPAS demonstrate that, for comparatively weak blasts, the tensions generated within the brain by blast interaction may result in damaging hoop stretches. Additionally, the hierarchy of models provides researchers with a simple framework for understanding the potential origin of cavitation damage in the brain. While useful, the effects of the simplifications in these models should be quantified to fully assess the accuracy of the presented analyses.

Throughout this research, it became clear that scientific progress depends on both developing new computational tools and methods and on persistent critical assessment of their potency, limitations, and practicality. Working closely with the methods/models, their software implementations, and applications, it was elucidated that many practical systems *require* advancements beyond state-of-the-art techniques to make discovery science a reality. Conversely, many systems are simple enough that cheap, robust schemes that solve relatively basic models can capture all the physics needed to enable discovery, removing the need for time-consuming method and model research. These ideas are clearly demonstrated in COMPAS, where, despite having state-of-the-art methods and models for compressible interfacial flows, users most commonly utilize robust finite-volume schemes without Phase Fields. Further, the process of implementing different physics capabilities, such as MHD, and comparing them with the radiation and laser capabilities of codes like FLASH demonstrated that trying to make a code that solves every problem of interest is a fool's errand; development should be pointed and purposeful, intended to address specific deficiencies. Finally, in the realm of Phase Field models, experiments with different regularization mechanisms made it clear how significant an impact interface treatments can have on the resulting flow physics, and that significant care needs to be taken when making modeling choices. Interface regularization techniques are very useful, but can also be challenging to implement robustly. The same goes for high-order methods and AMR; very effective, but potentially delicate. Compressible interfacial problems are a great way to stress-test these methods as they probe all of the weaknesses of each. Ultimately, the best method, model, or code is the one that enables accurate and useful science, even if it is not the latest or greatest. This philosophy strongly encourages a utilitarian approach to numerical work; implement or develop what is needed, when it is needed, but carefully assess the long-term impacts of the choices made in that process. Moving forward, continued collaboration across computationalists and experimentalists will be essential to realize the full potential of methods and codes. In summary, this work aims to illustrate the keen awareness of numerical techniques and the thoughtful recognition of the limits and possibilities inherent in modeling and the methods required to advance computational science.

7.3 Future Work

The findings described in this dissertation expose several directions for future research in flow physics and numerical methods. These areas include accurate, bound-preserving DG-AMR schemes, stability studies of converging and magnetized vortex ring pairs, and investigations of dominant instability modes at shocked equatorial joints.

7.3.1 Accurate and Bound-Preserving Refinement and Coarsening Procedures for h-Adaptive Discontinuous Galerkin Schemes

The present work provides a bound-preserving finite-volume framework with adaptive mesh refinement for compressible interfacial flows, as well as a high-order discontinuous Galerkin method for Phase-Field modeling, but does not discuss bound preservation in the context of DG-AMR schemes. Several authors have implemented DG methods within adaptive mesh frameworks with good results [158, 213, 279], but little work has examined the possibility of generating out-of-bounds states as a result of cell coarsening or refinement. As such, developing DG schemes that are robust and accurate during coarsening and refinement is a key future work to ensure the success of large-scale calculations, and we propose a pathway here.

In the context of finite volume discretizations on patch-based AMR grids, the process to refine a coarse cell into a set of fine cells is straightforward because the solution is a constant: The fine cells can be directly filled with that same constant, or neighbor cells can be used to reconstruct a polynomial in the coarse cell that conservatively adds higher-order variation. Ensuring conservation and bound preservation during refinement is straightforward, as there is no modal information to commute between mesh levels. With DG schemes, however, the high-order data make the refinement nontrivial as the boundedness of cell averages and quadrature points is paramount to stability. As such, ensuring accurate and well-behaved refinement is crucial.

When refining a single coarse DG cell in the context of conformal refinement, i.e., $\Omega^c = \bigcup_s \Omega_s^f$, the most obvious method is to directly interpolate the coarse cell onto the fine cells,

$$\int_{\Omega_s^f} q^c v_i^f \, d\Omega = \int_{\Omega_s^f} q_s^f v_i^f \, d\Omega \quad 0 \leq i \leq \mathcal{K} - 1, \tag{7.1}$$

where q^c is the coarse cell solution, v_j^f are the basis functions defined in the fine cell, q_s^f are the fine cell solutions, $s \in [0, R]$ where $R - 1$ is the number of refined cells, Ω^c is the coarse cell domain, Ω_s^f is the domain of refined cell s, and $\mathcal{K} - 1$ is the number of basis functions. This operation can

be written as a matrix-vector system, assuming an orthonormal basis,

$$\boldsymbol{A}^f = \begin{bmatrix} \boldsymbol{a}_0^f \\ \vdots \\ \boldsymbol{a}_R^f \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} a_{0,0}^f \\ \vdots \\ a_{0,\mathcal{K}}^f \end{bmatrix} \\ \vdots \\ \begin{bmatrix} a_{R,0}^f \\ \vdots \\ a_{R,\mathcal{K}}^f \end{bmatrix} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \int_{\Omega_0^f} v_0^c v_0^f d\Omega & \dots & \int_{\Omega_0^f} v_{\mathcal{K}}^c v_0^f d\Omega \\ \vdots & \ddots & \vdots \\ \int_{\Omega_0^f} v_0^c v_{\mathcal{K}}^f d\Omega & \dots & \int_{\Omega_0^f} v_{\mathcal{K}}^c v_{\mathcal{K}}^f d\Omega \end{bmatrix} \\ \vdots \\ \begin{bmatrix} \int_{\Omega_R^f} v_0^c v_0^f d\Omega & \dots & \int_{\Omega_R^f} v_{\mathcal{K}}^c v_0^f d\Omega \\ \vdots & \ddots & \vdots \\ \int_{\Omega_R^f} v_0^c v_{\mathcal{K}}^f d\Omega & \dots & \int_{\Omega_R^f} v_{\mathcal{K}}^c v_{\mathcal{K}}^f d\Omega \end{bmatrix} \end{bmatrix} \begin{bmatrix} a_0^c \\ \vdots \\ a_{\mathcal{K}}^c \end{bmatrix} = \boldsymbol{R}\boldsymbol{a}^c \tag{7.2}$$

where $\boldsymbol{A}^f \in \mathbb{R}^{R\mathcal{K}}$, $\boldsymbol{R} \in \mathbb{R}^{R\mathcal{K}\times\mathcal{K}}$, and $\boldsymbol{a}^c \in \mathbb{R}^{\mathcal{K}}$. By construction, this operation is conservative, as shown by integrating and summing the refined cells.

While the above method is accurate, efficient, and applicable to arbitrary polynomials, it can encounter issues regarding boundedness. The first issue is that the boundedness of polynomial solutions in DG schemes is almost exclusively checked at quadrature points $\boldsymbol{\xi}$, while the polynomial solution may be unbounded at other points [67]. This issue is not problematic in the absence of refinement; however, after refinement, the solution is resampled at new quadrature points where the state may not be bounded, so the post-refinement boundedness must be considered. Assuming that the solution is bounded at all quadrature points before refinement, we propose two methods to address this issue: limiting the coarse solution before refinement based on the quadrature points of the fine cells, and limiting the fine solutions after refinement based on their own cell averages.

The first method can be written explicitly as,

$$\tilde{q}^c(\boldsymbol{x}) = \bar{q}^c + \theta(q^c(\boldsymbol{x}) - \bar{q}^c), \tag{7.3}$$

where $\tilde{q}$ is q scaled around its cell-average, the scaling factor θ is,

$$\theta = \min\left\{1, \frac{\bar{q}^c - \min q^c(\boldsymbol{\xi}^c)}{\bar{q}^c - \min q^c(\boldsymbol{\xi}^f)}\right\}, \tag{7.4}$$

$\boldsymbol{\xi}^c$ are the quadrature points of the coarse cell, and $\boldsymbol{\xi}^f = \bigcup_s \boldsymbol{\xi}_s^f$ is the set of all quadrature points of the refined cells. Figure 7.1 (left) shows how this scaling acts on a solution. Due to the assumption of the solution being bounded at all coarse quadrature points, the coarse cell-average is guaranteed to be bounded (by convexity of the quadrature rule) such that there is always a $\theta \in [0, 1]$ that can make the solution bounded at any point.

The second method for ensuring the refined solutions are bounded at quadrature points is to

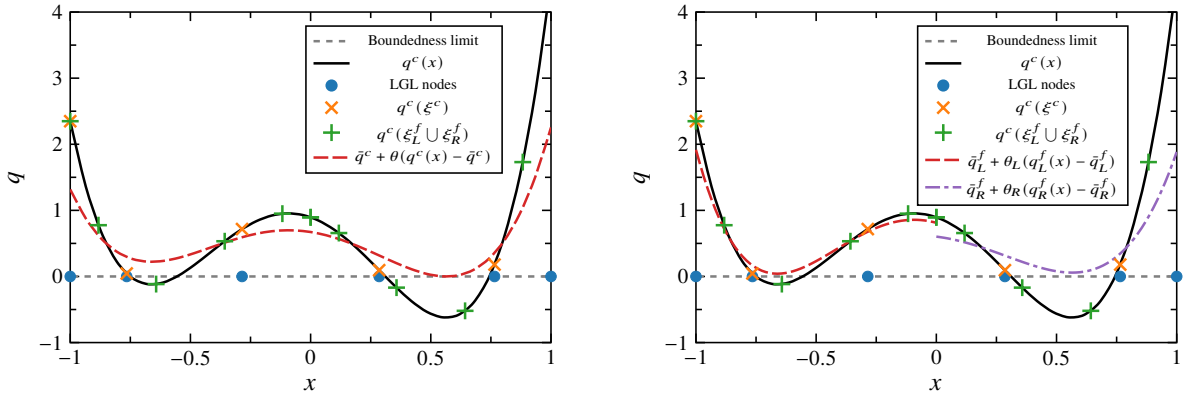


Figure 7.1: Coarse solutions scaled around the coarse cell-average value (left) and around the refined cell averages (right).

scale the refined cells around their own cell averages,

$$\tilde{q}_s^f(\boldsymbol{x}) = \bar{q}_s^f + \theta_s(q_s^f(\boldsymbol{x}) - \bar{q}_s^f), \quad 0 \le s \le R, \tag{7.5}$$

where θ_s is given by,

$$\theta_s = \min\left\{1, \frac{\bar{q}_s^f - \min q^c(\boldsymbol{\xi}_s^c)}{\bar{q}_s^f - \min q_s^f(\boldsymbol{\xi}_s^f)}\right\} \tag{7.6}$$

and $\boldsymbol{\xi}_s^c$ are the coarse quadrature points that are in Ω_s^f. Figure 7.1 (right) shows an example of the refined solutions that result from this method. While this bounding operation is more local, it exposes the possibility of the refined cell-average value being unbounded even if the coarse cell is bounded at all quadrature points, which cannot be resolved by a conservative limiter.

Suppose a polynomial $P(x) \in [-1, 1]$ of order p has $P(\xi_i) \ge \sigma \in \mathbb{R}$ at $N = p + 1$ Legendre-Gauss-Lobatto quadrature points ξ_i with $\sum_i \xi_i = 2$. Then $\exists P(x)$ such that one or both of $\int_{-1}^{0} P(x)dx < \sigma$ and $\int_{0}^{1} P(x)dx < \sigma$ could occur, despite $\bar{P} \ge \sigma$. Examples of this are readily found, suggesting that the boundedness of the refined cell-averages needs to be addressed before considering the boundedness of the refined quadrature points. As such, we propose an additional scaling that applies to the coarse solution that bounds the refined cell averages by the minimum value at any of the coarse quadrature points,

$$\check{q}^c(\boldsymbol{x}) = \bar{q}^c + \theta_C(q^c(\boldsymbol{x}) - \bar{q}^c), \tag{7.7}$$

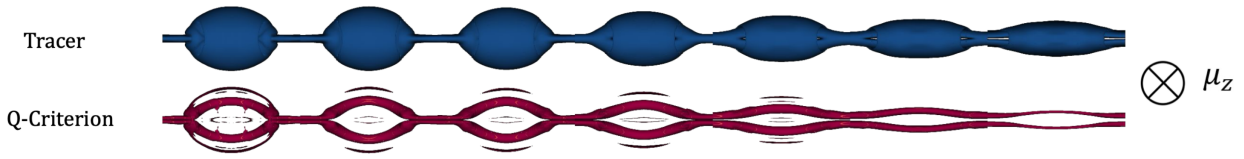


Figure 7.2: Contours of the material tracer (top) and Q-criterion (bottom) for varying magnetic field strengths μ_z, increasing from left (no field) to right.

where the scaling constant is given by,

$$\theta_C = \min\left\{1, \frac{\bar{q}^c - \min q^c(\boldsymbol{\xi}^c)}{\bar{q}^c - \min\{\bar{q}_s^f\}_{s=0}^R}\right\}. \tag{7.8}$$

Constricting the high-order data with θ_C ensures that all refined cells have bounded averages and values at quadrature points, which can be bounded via Eq. (7.6).

Further development and testing of the above bounding operations for systems, as well as the development of similar operations for accurate coarsening, can help improve the robustness of DG schemes in AMR implementations.

7.3.2 Linear Stability Theory of Converging Vortex Ring Pairs

As demonstrated in Chapter 5, converging, counter-rotating pairs of vortex rings appear to be unstable to vortex-pair instabilities. Recent work has shown that diverging vortex rings are predictably unstable to such mechanisms [255], with applications spanning vortex ring collisions in tabletop water tanks [142] to supernova remnants [259]. The divergent behavior of these systems dramatically alters the dynamics compared to planar instabilities [57], as the instability growth rates change over time. This behavior is a direct consequence of the zeroth-order motion of the rings, which continuously increases the arcwise wavelength of perturbations. Conversely, in the converging case, the arcwise wavelength decreases as the ring radii decrease, potentially exciting shorter-wavelength instabilities than those reported by Crow [123, 267]. Computing the theoretical linear growth rates for converging vortex-pair instabilities is thus of interest for characterizing flows in which these features may be present, with the analysis likely to follow that of [255].

In addition to the study of converging pairs of vortex rings, preliminary work has investigated the effect of magnetic fields on the classical Crow instability. Theoretical results suggest that magnetic fields tend to reduce the instability growth rate compared to the purely hydrodynamic case, but in some regimes, the growth is accelerated. Figure 7.2 shows contours of the material tracer and Q-criterion for several vortex pairs. These preliminary results also suggest that magnetic fields tend to suppress the Crow instability and delay or prevent the formation of vortex rings via

reconnection. Future studies should map the parameter space by the plasma β, Alfvénic Mach number, and field orientation to identify regimes where magnetic tension suppresses, neutralizes, or amplifies the instability growth. This work could be applied to systems ranging from laboratory fusion experiments to astrophysics, both in planar and cylindrical geometries.

7.3.3 Competition Between Instabilities in Shocked Gaps

While the results in Chapter 4 and Chapter 5 suggest that vortex-pair instabilities may be significant in RMI flows, future studies must clarify when, and under what flow parameters, large-scale vortex-line instabilities decisively outpace small-scale interfacial growth in shock-interface interactions. Consider a nominally planar shocked gap between two heavy fluids, filled with a light fluid; the interfaces are impulsively accelerated by the shock, seeding Richtmyer-Meshkov and Kelvin-Helmholtz roll-ups along the developing shear layer. If coherent counter-rotating cores form, they themselves may become susceptible to vortex-line instabilities, including the Crow [57] and elliptical [123] instabilities, depending on the flow conditions. These vortex lines may eventually pinch off into vortex rings, exponentially quickly in the case of the CI. In order for this to occur, however, the cores must be (i) coherent enough to move cohesively and (ii) perturbed enough initially to outpace the growth of other mixing mechanisms that would disrupt the core instability. The competition is therefore multiscale: the RMI amplifies short-wave interface perturbations and produces tightly packed cores, while the CI reorganizes vorticity into larger, more coherent structures (vortex rings) that can alter momentum transport and material mixing. The planar and cylindrical simulations in Chapter 4 and Chapter 5 demonstrated that, for gas-gas systems, Crow growth rates can be large enough to exhibit coherent instability dynamics before small-scale RMI perturbations destabilize the cores. This represents a hand-off from interfacial mixing to vortex-ring-dominated transport, but future studies should explore this with experiments and simulations of shocked gaps.

In double-shell ICF capsules, the assembly joint creates a narrow, shocked gap between the high-Z outer pushers and the low-Z foam and fuel layers. Depending on the dynamics, coherent counter-rotating vortex lines may form at the gap. The cylindrical results in Chapter 5 suggest that these vortex lines are poised to undergo Crow-type instabilities as the shells converge, potentially imprinting large-scale, low-mode distortions on the inner interfaces. A comprehensive study should use simulations to track the transition from small-scale interfacial growth to large-scale vortex-ring formation at the assembly joint under drive conditions matching National Ignition Facility or OMEGA double-shell designs [56, 168, 210]. If there is evidence for vortex instabilities, then mixing models and symmetry assumptions could be modified to account for this instability pathway.

APPENDIX A

Path-Conservative Scheme for the Divergence-Form Volume Fraction Equation

We contrast the proposed scheme for the volume fraction equation with the path-conservative-style scheme of [154, 200] and extended to multiphase flows in [91, 225]. Beginning with Eq. (3.17) and denoting $\alpha_1(1+K_1) = H$, we apply Dal Maso–LeFloch–Murat theory [59] to write,

$$\begin{aligned}\int_{\Omega_e} v_i \frac{\partial \alpha_{1,e}}{\partial t} d\Omega + \int_{\partial\Omega_e} v_i \mathbf{n}\cdot \widehat{(\alpha_1 \mathbf{u})} dS - \int_{\Omega_e} \alpha_1 \mathbf{u}\cdot\nabla v_i d\Omega - \int_{\Omega_e} v_i H \nabla\cdot\mathbf{u} d\Omega \\ + \int_{\partial\Omega_e} \frac{1}{2}(v_{i,L}+v_{i,R}) \left[\int_0^1 H(\psi(\tau;\mathbf{U}_L,\mathbf{U}_R)) \frac{\partial\psi}{\partial\tau}(\tau;\mathbf{u}_L,\mathbf{u}_R) d\tau\right]\cdot\mathbf{n} dS = 0, \quad 0 \le i < \mathcal{K}, \quad \forall\Omega_e \in \mathcal{E}_h,\end{aligned} \tag{A.1}$$

where ϕ_L and ϕ_R denote the left and right values of a variable ϕ on the interface between two cells, respectively. The bracketed integral in the fifth term is a path integral over τ that connects the discontinuous states across an element interface and approximates a surface contribution from the $H\nabla\cdot\mathbf{u}$ term, like a penalty. Choosing a linear path $\psi(\tau;\phi_L,\phi_L) = \phi_L + \tau(\phi_R - \phi_L)$, across $\tau \in [0,1]$ the path integral is evaluated as,

$$\int_0^1 H(\psi(\tau;\mathbf{U}_L,\mathbf{U}_R)) \frac{\partial\psi}{\partial\tau}(\tau;\mathbf{u}_L,\mathbf{u}_R) d\tau = (\mathbf{u}_R - \mathbf{u}_L)\int_0^1 H(\psi(\tau;\mathbf{U}_L,\mathbf{U}_R)) d\tau = (\mathbf{u}_R - \mathbf{u}_L) H^*. \tag{A.2}$$

Inserting this result into (A.1) gives a useful weak form,

$$\begin{aligned}\int_{\Omega_e} v_i \frac{\partial \alpha_{1,e}}{\partial t} d\Omega + \int_{\partial\Omega_e} v_i \mathbf{n}\cdot \widehat{(\alpha_1 \mathbf{u})} dS - \int_{\Omega_e} \alpha_1 \mathbf{u}\cdot\nabla v_i d\Omega - \int_{\Omega_e} v_i H \nabla\cdot\mathbf{u} d\Omega \\ + \int_{\partial\Omega_e} \frac{1}{2} H^* (v_{i,L}+v_{i,R})(\mathbf{u}_R - \mathbf{u}_L)\cdot\mathbf{n} dS = 0, \quad 0 \le i < \mathcal{K}, \quad \forall\Omega_e \in \mathcal{E}_h,\end{aligned} \tag{A.3}$$

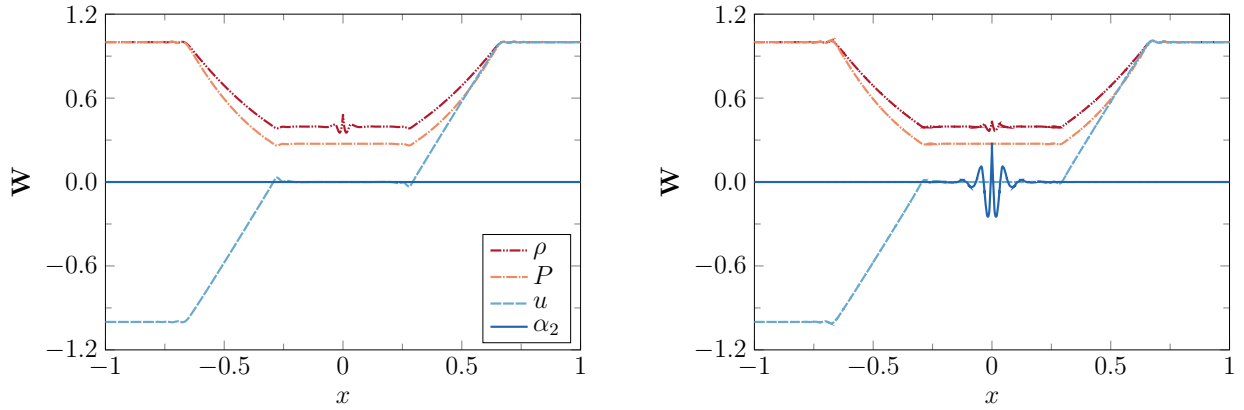


Figure A.1: Primitive variables $\mathbf{W}$ for the double rarefaction problem at $t = 0.3$. Left: Solution from the proposed well-balanced scheme. Right: Solution from the described path-conservative scheme [154, 200]. Results are generated with 64 $p = 3$ elements.

As with the proposed scheme, to test if this scheme is exact away from interfaces we set $\alpha_1 = \bar{\alpha}_1$. Without loss of generality, we set $\{K_m = 0\}_{m=1}^{N}$ such that $H^* = \frac{1}{2}(\alpha_{1,L} + \alpha_{1,R})$, leading to,

$$\begin{aligned} \int_{\Omega_e} v_i \frac{\partial \alpha_{1,e}}{\partial t} d\Omega + \bar{\alpha}_1 \int_{\partial\Omega_e} \mathbf{n}\cdot \left[v_i \hat{\mathbf{u}} + \frac{1}{2}(v_{i,L} + v_{i,R})(\mathbf{u}_R - \mathbf{u}_L) \right] dS \\ -\bar{\alpha}_1 \int_{\Omega_e} (\mathbf{u}\cdot\nabla v_i + v_i \nabla\cdot\mathbf{u})\, d\Omega = 0, \quad 0 \le i < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h. \end{aligned} \tag{A.4}$$

The spatial terms clearly do not vanish in the same manner as before, although it is possible that their sum is zero. To investigate further we set $p = 0$ and manipulate Eq. (A.4) to find,

$$\int_{\Omega_e} \frac{\partial \alpha_{1,e}}{\partial t} d\Omega + \bar{\alpha}_1 \int_{\partial\Omega_e} \mathbf{n}\cdot \left[\hat{\mathbf{u}} + (\mathbf{u}_R - \mathbf{u}_L) \right] dS = 0, \quad 0 \le i < \mathcal{K}, \quad \forall \Omega_e \in \mathcal{E}_h, \tag{A.5}$$

where the third integral vanishes because when $\mathbf{u}$ becomes constant $\bar{\mathbf{u}}$ its divergence is also zero, i.e., $\nabla\cdot\bar{\mathbf{u}} = 0$ under direct evaluation for $p = 0$. The remaining spatial term vanishes if the velocity is constant or $\alpha_1 = 0$. However, in the case of nonzero α_1 or a velocity gradient, the velocity from the Riemann solver must negate the difference between the face velocities, which is not generally true. The time step must go to zero to keep this term from modifying the cell average, but this result does not lead to the same time-step constraint as that for the Euler equations and thus does not *numerically* satisfy consistency of reduction. As such, the scheme is not exact away from interfaces and may introduce errors in the volume fraction. To demonstrate this error, we simulate a double rarefaction with the proposed well-balanced scheme and with the described path-conservative scheme. The initial conditions are $(\rho, u, P, \alpha_1) = (1, -\mathcal{H}(-x) + \mathcal{H}(x), 1, 1)$. The resulting numerical solutions are shown in Fig. A.1 for the proposed scheme and the path-conservative scheme. While the

proposed scheme maintains a constant volume fraction in the presence of a velocity gradient, the path-conservative scheme generates spurious oscillations in the volume fraction that are not bounded (leading to a non-convex combination of material properties) and does not conserve the volume fraction as we would expect for a pure material. This error manifests in violating volume fraction positivity and generating a fictitious phase. The spike in density observed in both simulations is well documented and is a product of the Riemann solver [68].

APPENDIX B

Pressure Fields Produced by Single-Bubble Collapse Near a Corner

The material of this chapter is adapted from W. White, S. A. Beig, and E. Johnsen, Pressure fields produced by single-bubble collapse in a corner, *Physical Review Fluids*. 8(2):023601, 2023. Simulations and preliminary analyses were conducted as an undergraduate under Professor Eric Johnsen, with analysis, writing, and publication completed during doctoral studies.

B.1 Abstract

Damage produced by repeated bubble collapse to neighboring rigid objects in hydraulic systems is an important consequence of cavitation. Although bubble collapse near a single wall has received significant attention in the past, few studies exist on the dynamics of bubbles collapsing near a corner, i.e., two flat rigid surfaces intersecting at a right angle. In this work, we quantify the pressure fields produced by a single bubble collapsing near two perpendicular rigid walls. Using a high-order accurate shock- and interface-capturing method to solve the three-dimensional compressible Navier-Stokes equations for gas/liquid flows, we simulate the dynamics of a single bubble collapsing at different initial stand-off distances from the two walls. By contrast to a bubble collapsing near a single wall, the collapse of bubbles within a critical stand-off distance is not symmetric about the bisecting plane due to the interaction between the bubble and the second wall. The second wall affects the pressure produced during the collapse in the following ways: (i) for bubbles initially located sufficiently close to both walls, the re-entrant jet produced during collapse no longer points in the direction normal to the closest wall but at an angle toward the corner, (ii) the part of the emitted shock with the highest amplitude propagates in line with the jet, and (iii) the bubble migrates in that same direction during its collapse with a dependence on the stand-off distance consistent with predictions made using Kelvin impulse. The location of maximum pressure along the walls is measured for the different initial stand-off distances. Using

acoustic arguments we find a semi-empirical relationship to predict the initial stand-off distances for which the maximum pressure occurs in the corner. We find that when the bubble is sufficiently close to equidistant from each boundary, the maximum pressure is observed in the corner due to the water-hammer and implosion shocks reflecting off the boundaries and intersecting in the corner. We also show that when bubbles are initially attached to either wall the wall pressure produced can be significantly increased compared to bubbles detached from either wall.

B.2 Introduction

Cavitation has long been known to cause structural damage to solids in a variety of hydrodynamic and acoustic applications [63, 187, 240, 254]. The central problem in cavitation-induced damage is bubble collapse. In the idealized case of an infinite medium, a single initially spherical bubble collapses symmetrically [198], thereby compressing its contents until the interior pressure is sufficiently large to halt the collapse; the high pressures produced release a shock and possibly visible light [32, 64, 73, 133, 143, 173, 187, 198, 228]. When collapsing near a rigid surface, the bubble no longer maintains a spherical shape; a re-entrant jet forms and impinges upon the opposite side of the bubble, generating a water-hammer shock [25, 171, 189, 254]. Past simulations [24] show that the strength of this shock is greatest in the jet direction and decreases azimuthally; the maximum wall pressure is therefore along the axis of symmetry (i.e., the jet direction). This maximum pressure is an important predictor of damage [125].

Although there have been many experimental [25, 33, 132, 143, 171, 188, 254] and computational [24, 115] studies of a bubble collapsing near a single, rigid, flat wall, the collapse of a bubble near a corner, i.e., two flat rigid surfaces intersecting at a right angle, has received far less interest. Such a geometry is common in internal and external flows comprising channels, ducts, cavities, etc., which can be represented by an infinite corner when the bubble is small compared to the geometry [35]. The collapse of a bubble near a corner was first investigated experimentally by [34]. Using lasers to induce cavitation bubbles and high-speed cameras to capture their dynamics, several collapse parameters were characterized. In particular, the angle of the jet and the bubble translation were measured. The work of Cui *et al.* [58] demonstrated that the bubble is primarily influenced by the nearer of the two walls. An analytical estimate for the jet angle using potential flow was proposed in [233] and it was shown through experiments that the jet angle can be accurately predicted by potential flow [167]. In simulations based on potential flow theory (e.g.,boundary integral or element methods) it was found that the bubble migration towards the near wall decreases with initial stand-off distance and that the jet velocity decreases as the initial stand-off distances decrease. Additionally, the inertial collapse of a vapor bubble near slotted walls has been explored numerically and experimentally [139, 244].

While accurately predicting the bubble dynamics, potential-flow-based methods cannot represent compressible-flow features such as shock waves. The pressure fields, which are essential to predicting potential damage, therefore cannot accurately be determined. Recently, shock-capturing methods have been used to investigate shock waves produced by bubbles collapsing near a wall [11, 24, 115, 204], though these methods have yet to be applied to bubbles collapsing near a corner.

The objective of our work is to determine the time-dependent pressure fields produced by a bubble collapsing near a corner. The location and magnitude of the maximum pressure produced are not trivially determined, as the presence of the second wall deflects the direction of the re-entrant jet. As a result, the strongest part of the emitted shock no longer impinges upon the closest wall surface at the projected bubble center on that wall, by contrast to collapse near a single wall. Using a high-order shock-/interface-capturing method, we simulate the collapse initially located at different stand-off distances from the two walls making up the corner. In addition to characterizing the collapse dynamics, we develop a geometric prediction of initial stand-of distances for which the maximum wall pressure occurs in the corner. The article is organized as follows. After a description of the problem and methods in § B.3, the dynamics of a detached bubble collapsing in a corner are investigated in § B.4. We compare our simulations to experimental results and further explore the kinematics in § B.4.3. The wall pressure and temperature produced are analyzed in § B.4.4 and § B.4.5, respectively. Then, the dynamics of a bubble initially attached to the wall are discussed in § B.5. Concluding remarks are provided in § B.6.

B.3 Problem Description and Physical Model

We examine the classical Rayleigh collapse [198] of an isolated bubble near two perpendicular, rigid boundaries. We assume that the bubble is initially spherical with radius $R_o = 500\,\mu$m, filled with non-condensible gas with water vapor properties. While in practice the bubble may not be perfectly spherical and stationary at maximum expansion, experimental images show a high degree of radial symmetry at the corresponding time and we assume any interfacial velocity is small [34, 233]. Initially, the water temperature T_∞ and the bubble temperature $T_{b,o}$ are 300 K. The initial gas pressure is set to 3.55 kPa, which is the saturation pressure of water at this temperature. The problem is initiated by instantaneously raising the liquid pressure $p_\infty = 3.55$ kPa by Δp. We set $\Delta p/p_\infty = 1400$ for all simulations, which is relevant to a variety of cavitation applications [76]. As illustrated in Fig. B.1, the bubble is initially located a distance H_x from a vertical rigid wall in the y-z plane and a distance H_z from a horizontal rigid wall in the x-y plane. The corresponding normalized initial stand-off distances are $\delta_x = H_x/R_o$ and $\delta_z = H_z/R_o$. As in [34] the angle between the line connecting the corner and the bubble center with the x-axis is θ_b. Cases bisected by $\theta_b = \pi/4$ are symmetric about the bisecting plane. Cases in which $\delta_x > 1$ and $\delta_z > 1$ represent

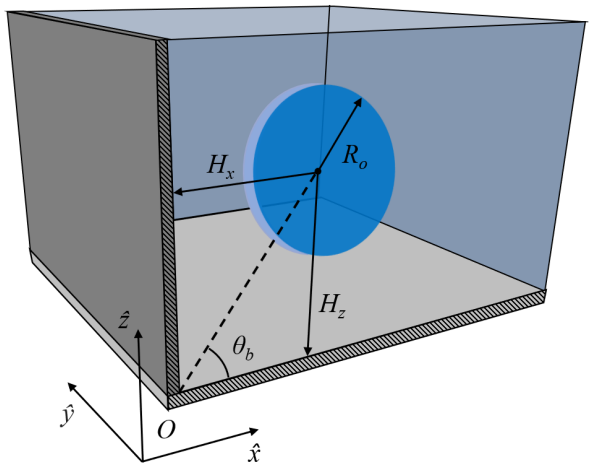


Figure B.1: Schematic of the Rayleigh collapse of a single bubble near two perpendicular rigid walls.

a bubble initially detached from both walls; $\delta_x \leq 1$ and/or $\delta_z \leq 1$ represent a bubble initially attached to both walls; $\delta_x \leq 1$ and $\delta_z > 1$ represent a bubble initially attached to only the y-z wall. For attached bubbles, the gas is in direct contact with the wall surface, i.e. there is no liquid film between the bubble and the wall; the initial bubble volume is adjusted such that the initial potential energy of the collapse $E_o = \Delta p V_o$ is the same as that in the detached cases. We also consider the limit of $\delta_x \to \infty$.

The compressible Navier-Stokes equations are solved in both the liquid and the gas:

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j}(\rho u_j) = 0, \tag{B.1a}$$

$$\frac{\partial (\rho u_i)}{\partial t} + \frac{\partial}{\partial x_j}(\rho u_i u_j + p\delta_{ij}) = \frac{\partial \tau_{ij}}{\partial x_j}, \tag{B.1b}$$

$$\frac{\partial E}{\partial t} + \frac{\partial}{\partial x_j}[u_j(E+p)] = \frac{\partial}{\partial x_j}(u_i \tau_{ij} - Q_j), \tag{B.1c}$$

where ρ is the density, u_i the velocity vector, p the pressure, $E = \rho e + \rho u_i u_i/2$ the total energy, e the internal energy, and δ_{ij} the identity tensor. The viscous stress tensor τ_{ij} and heat flux Q_j are given by:

$$\tau_{ij} = \mu\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \frac{2}{3}\frac{\partial u_k}{\partial x_k}\delta_{ij}\right) + \mu_B \frac{\partial u_k}{\partial x_k}\delta_{ij}, \qquad Q_j = -\kappa\frac{\partial T}{\partial x_j}, \tag{B.2}$$

where μ is the dynamic shear viscosity, μ_B is the bulk viscosity, and κ is the thermal conductivity. To capture material interfaces between different phases, we use the five-equations multiphase

Coefficients	Water vapor	Water
μ (Pa s)	1.8531×10^{-5}	8.3283×10^{-4}
μ_B (Pa s)	1.1119×10^{-5}	2.5818×10^{-3}
κ (W m^{-1} K^{-1})	0.026107	0.61497
n	1.47	1.19
B (MPa)	0	702.8
b (kg m^{-3})	0	6.61×10^{-4}
c (J kg^{-1} K^{-1})	955	3610
q (J kg^{-1})	2.1×10^6	-1.2 $\times 10^6$

Table B.1: Relevant gas and liquid properties.

model [169] and solve the following transport equations:

$$\frac{\partial(\rho^{(k)}\alpha^{(k)})}{\partial t} + \frac{\partial}{\partial x_j}\left(\rho^{(k)}\alpha^{(k)}u_j\right) = 0, \tag{B.3a}$$

$$\frac{\partial\alpha^{(k)}}{\partial t} + u_j\frac{\partial\alpha^{(k)}}{\partial x_j} - \alpha^{(k)}\alpha^{(k')}\frac{\rho^{(k')}(a^{(k')})^2 - \rho^{(k)}(a^{(k)})^2}{\alpha^{(k)}\rho^{(k')}(a^{(k')})^2 + \alpha^{(k')}\rho^{(k)}(a^{(k)})^2}\frac{\partial u_j}{\partial x_j} = 0, \tag{B.3b}$$

where a is the sound speed and k and k' represent each phase $\alpha^{(k)} = 1 - \alpha^{(k')}$. As done in relevant past studies [88, 115, 239], we neglect surface tension and mass transfer effects since the flow is inertia dominated. The Noble-Able Stiffened-Gas (NASG) equation of state [134] closes the system:

$$\rho(e-q) = \frac{p}{n-1}(1-\rho b) + \frac{nB}{n-1}(1-\rho b), \qquad \text{(pressure-wise)} \tag{B.4a}$$

$$= \rho cT + B(1-\rho b), \qquad \text{(temperature-wise)} \tag{B.4b}$$

where q, n, b, and B are experimentally determined parameters, given in table B.1. For ideal gases, $n = \gamma$ represents the specific heat ratio, $c = c_v$ is the specific heat at constant volume, and B, q, and b are zero.

The numerical method of Beig and Johnsen [11] is used to accurately and stably represent discontinuities in the flow. The solution is advanced in time with a third-order total variation diminishing Runge-Kutta scheme [85] with an adaptive step size satisfying advection and diffusion stability constraints. The spatial scheme is high-order accurate, in which fifth-order Weighted Essentially Non-Oscillatory [111] is used at shocks and interfaces [113] and fourth-order central difference elsewhere; the sensor of Henry de Frahan [91] discriminates between these regions. For stability purposes, the source term in (B.3b) is set to zero for the first 1000 time steps [24] and

interfaces are initially smoothed over three computational cells. The validity of this numerical scheme was thoroughly evaluated by Beig and Johnsen [11].

Owing to the problem symmetry with respect to the x-z plane, we simulate only half of the domain with a symmetry boundary condition along this plane. The walls are adiabatic and perfectly reflecting with no slip. Non-reflecting boundary conditions are used along the remaining boundaries. A uniform Cartesian grid with 192 cells per initial bubble radius is used and the largest domain size being $8R_o \times 8R_o \times 4R_o$. Beig and Johnsen [11] examined the convergence of the spherical case in comparison with the Keller-Miksis [122] solution and demonstrated the chosen resolution is satisfactory. To understand the dependence of the bubble dynamics on the proximity of the walls, we conduct simulations in which the stand-off distances are varied as follows: $0.25 \leq \delta_x \leq 4.0$ and $0.25 \leq \delta_z \leq 4.0$. Due to the infinite nature of the walls, the dynamics of cases with initial stand-off distances $(\delta_x, \delta_z) = (a, b)$ are the same as those with initial stand-off distances $(\delta_x, \delta_z) = (b, a)$ so we only simulate cases with $\delta_z \geq \delta_x$. As a reference, we compare our results to bubble collapse near a single wall.

B.4 Detached Bubbles

B.4.1 Qualitative Dynamics

We start by examining the dynamics and resulting pressure fields of detached bubbles. Figures B.2 and B.3 show contours of (numerical) Schlieren, pressure, and velocity for simulations with $(\delta_x, \delta_z) = (1.5, 3.0)$ and $(\delta_x, \delta_z) = (2.5, 3.0)$, respectively. Time is normalized by the Rayleigh collapse time $t_c = 0.915R_0\sqrt{\rho_l/\Delta p}$. These two cases are chosen to contrast the dynamics that occur when $\delta_x/\delta_z \approx 1$ and when δ_x/δ_z is farther from unity. As the collapse begins, a rarefaction wave propagates outward while a transmitted shock wave converges towards the bubble center. The rarefaction wave impinges upon the closest wall first, then the farther wall. Upon each interaction, the relevant part of the wave is reflected back towards the bubble. Until the time when the reflected wave hits the bubble, the bubble is unaffected by the walls, so the collapse remains spherical [24]. The interaction of the reflected rarefactions with the bubble gives rise to a non-uniform pressure distribution around the bubble, thus causing the interface to accelerate asymmetrically. The pressure fields resulting from these wave dynamics can be understood as the effects of the image bubbles, whose influence is inversely proportional to their distance from the original bubble. However, by contrast to potential flow calculations [233] in which these effects instantaneously take place, the present simulations account for the finite propagation speed of the waves. Following early contraction a region of high-pressure and high-speed liquid is visible on the far side of the bubble. A re-entrant jet forms in a direction that is normal to neither wall and subsequently impinges upon

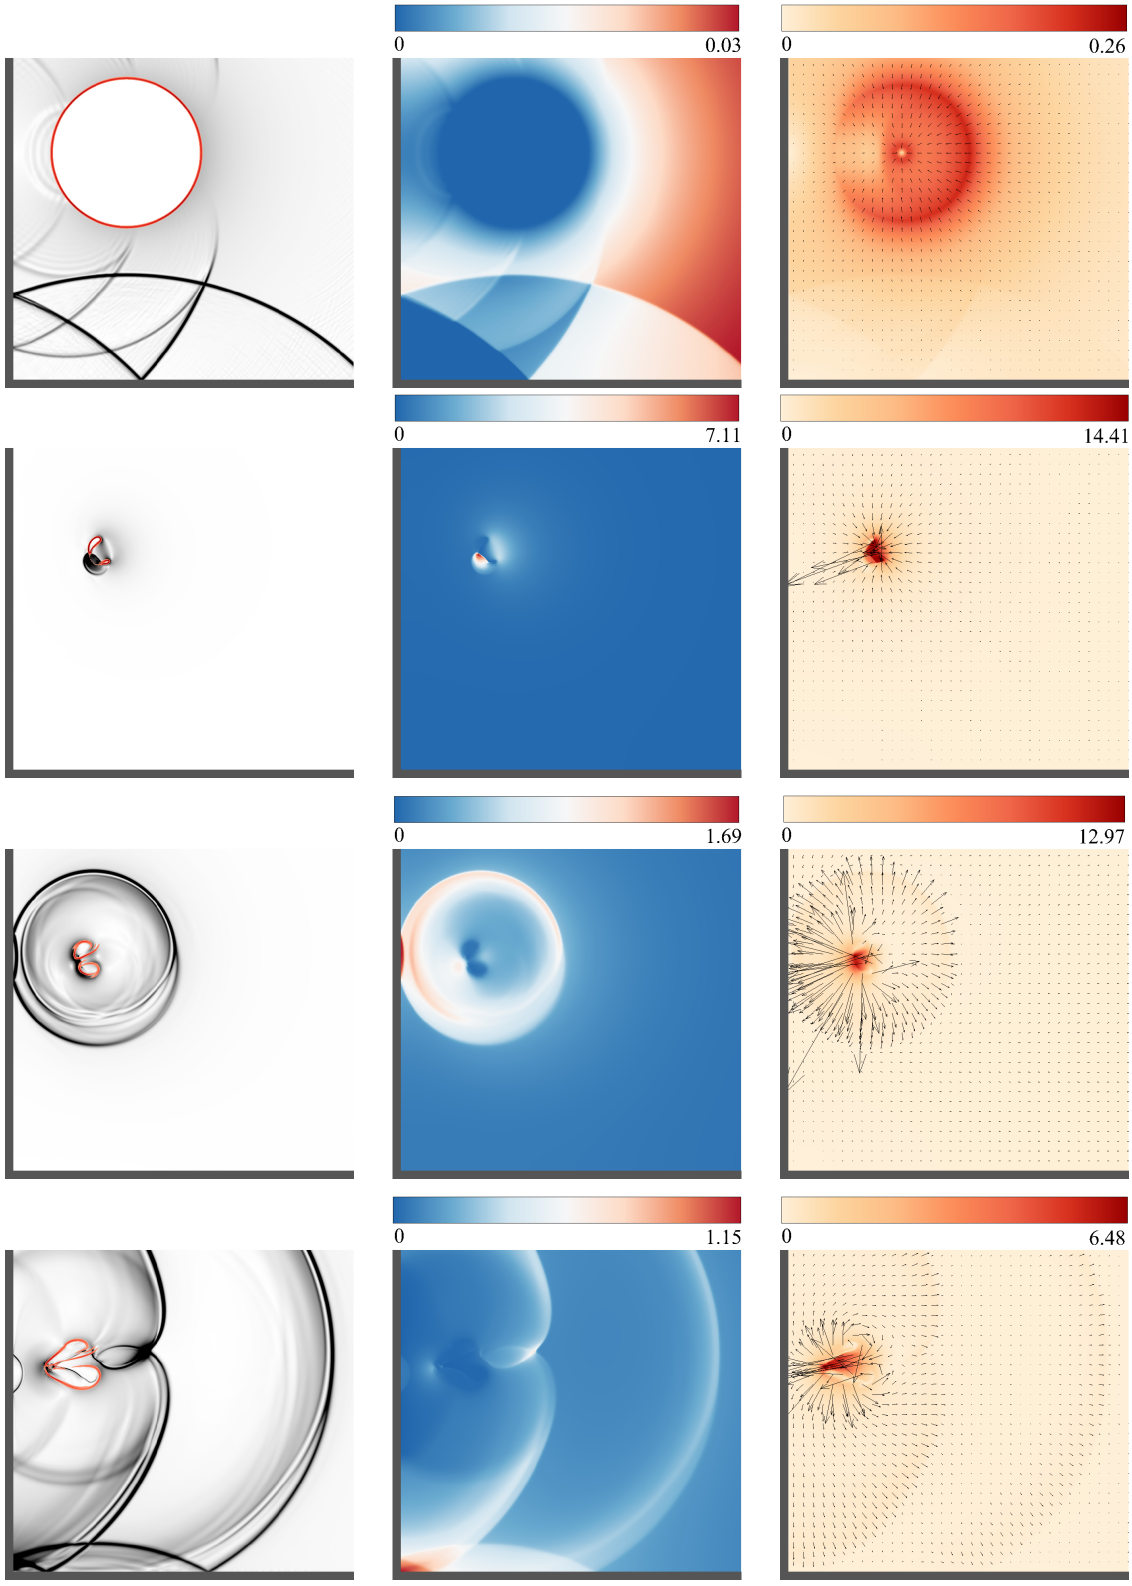


Figure B.2: Contours of (numerical) Schlieren, pressure normalized by $\rho_l a_l \sqrt{\Delta p/\rho_l}$, and velocity magnitude normalized by $\sqrt{\Delta p/\rho_l}$ for $(\delta_x, \delta_z) = (1.5, 3.0)$. Contours are of times $t = 0.163$, $t = 1.155$, $t = 1.199$, and $t = 1.294$, normalized by the Rayleigh collapse time $0.915 R_0 \sqrt{\rho_l/\Delta p}$.

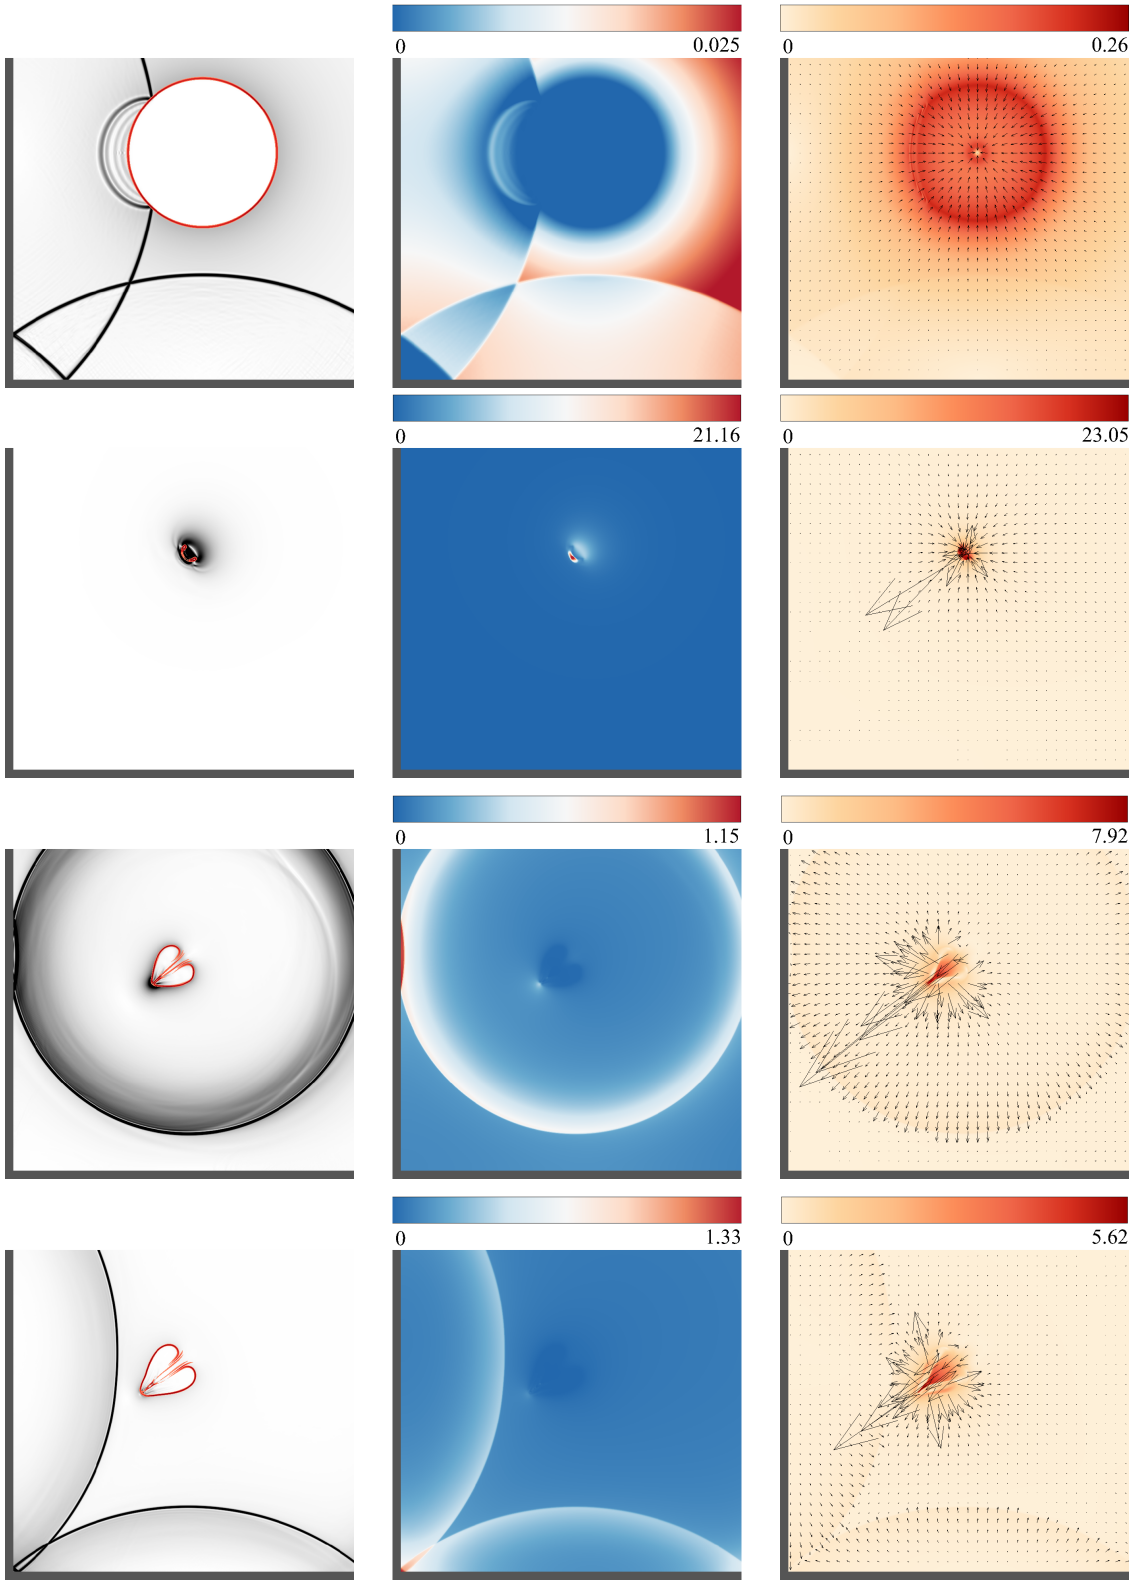


Figure B.3: Contours of (numerical) Schlieren, pressure normalized by $\rho_l a_l \sqrt{\Delta p/\rho_l}$, and velocity magnitude normalized by $\sqrt{\Delta p/\rho_l}$ for $(\delta_x, \delta_z) = (2.5, 3.0)$. Contours are of times $t = 0.163$, $t = 1.102$, $t = 1.198$, and $t = 1.261$, normalized by the Rayleigh collapse time $0.915 R_0 \sqrt{\rho_l/\Delta p}$.

the opposite side of the bubble. This impact generates a water-hammer shock and causes the bubble to take the shape of a vortex ring (Fig. B.2(b) and Fig. B.3(b)). The water-hammer shock subsequently interacts with the bubble, thus causing it to collapse and emit another shock. The pressure of the resulting shocks varies azimuthally and appears to be highest in the direction of the jet [24].

For $(\delta_x, \delta_z) = (2.5, 3.0)$, the collapse is close to symmetric as the initial location is close the angle $\theta_b = \pi/4$. The wave front impinging upon the closer wall interacts with the bubble slightly before the wave front impinging upon the farther wall. This small difference in time between the two wave fronts gives rise to an asymmetry that is visible only at late times in the collapse. Though not shown here, a bubble initially located along the bisection plane exhibits symmetry in its collapse. For the case with $(\delta_x, \delta_z) = (1.5, 3.0)$, the non-uniformity of the collapse is clear, particularly at the point of the jet impact impacting the bubble wall and late times.

The topology of the bubble throughout the collapse varies significantly with different initial stand-off distances, which dictate the jet direction. For $(\delta_x, \delta_z) = (2.5, 3.0)$, the upper and lower portions of the bubble at the point of jet impact are visibly very similar, whereas for $(\delta_x, \delta_z) = (1.5, 3.0)$ the upper region of the bubble is noticeably larger. Also different between these two cases are the shock waves that intersect in the corner. For the case with $(\delta_x, \delta_z) = (1.5, 3.0)$, the water-hammer and implosion shocks are distinct when they intersect in the corner. The substantial distance from the corner in the case with $(\delta_x, \delta_z) = (2.5, 3.0)$ results in a single shock by the time it impinges upon the walls. When the ratio δ_x/δ_z is small or large the dynamics of the collapse are more similar to those near a single wall. For cases where $\delta_x \approx \delta_z$, the reflected waves intersect the bubble at the same time and the collapse is symmetric about $\theta_b = \pi/4$. The observed bubble shapes match well with those from experimental studies, and the use of simulations in the present work affords the ability to examine the kinematics of the collapse in detail and to quantify the pressure and temperature fields.

B.4.2 Characterization of the Bubble Non-Spherical Shape

For bubbles lying off the bisecting plane, spherical symmetry is broken when the first portion of the reflected rarefaction front interacts with the bubble. Bubbles sufficiently far from the walls, i.e., distances such that the instant of collapse (minimum volume) occurs before the time when the reflected rarefaction interacts again with the bubble, remain spherical until collapse. Comparing this time-of-flight with the collapse time, it follows that bubbles within the following stand-off

distance from either wall collapse non-spherically:

$$\delta_{cr} \approx 0.5\left(1+\sqrt{\frac{\rho_l a_l^2}{\Delta p}+\left[\frac{p_\infty}{(\gamma-1)\Delta p}\right]^{\frac{1}{3(\gamma-1)}}}\right), \tag{B.5}$$

which yields $\delta_{cr} \approx 10.5$ for the driving pressure used here. All simulations, with the exception of single-wall cases for comparison, have both δ_x and δ_z smaller than δ_{cr}, and as such the dynamics are affected by both walls.

The collapse properties are characterized by the two geometric parameters δ_x and δ_z for certain geometrical properties (e.g., jet angle, migration angle) over a range of initial stand-off distances. Other quantities however depend on the absolute distance from the walls; e.g. the bubble reaches a smaller volume before the reflected rarefaction waves can return and inform the bubble of the geometry. As such, the characterization of quantities associated with the interfacial velocity (collapse time, jet speed, etc.) require additional information beyond the ratio δ_x/δ_z. For this purpose we calculate the deviation from spherical shape of the bubble, $\xi = \pi\lambda/\mathcal{P}$ where λ is the film thickness and $\mathcal{P}$ is the perimeter of the bubble's outline along the center-plane. The film thickness is the shortest distance between the front and back sides of the bubble at any given time, as illustrated in Fig. B.4. In the limit of a spherical collapse $\lambda = 2R(t)$ and $\xi = 1$ at all times during the collapse. When the jet forms, the shortest distance between the front and back sides is at the tip of the jet. The area of the bubble cross-section in this case is always greater than that of a circle with diameter equal to the film thickness, so for non-spherical collapse $0 \leq \xi \leq 1$. For all bubbles that are initially spherical, the film thickness decreases from $2R_0$ at $t = 0$ to zero at the point at which the jet impinges the opposite side; we calculate the film thickness and perimeter of the bubble along the center-plane until this point in time.

To appreciate the non-spherical evolution of the bubble, outlines of the interface along the centerplane up to the point at which $\lambda \approx 0.05$ are shown in Fig. B.4 for $(\delta_x, \delta_z) = (1.5, 3.0)$ and $(\delta_x, \delta_z) = (2.5, 3.0)$. Figure B.5 shows ξ for different initial stand-off distances with λ chosen to be around 0.05, as at this value the jet has nearly penetrated the bubble wall. A different value of λ close to zero may be chosen to calculate ξ, however the resultant trend in ξ will be similar, just scaled. The collapse is most non-spherical for the bubbles closest to the wall, as evidenced by values of ξ close to zero. As expected from the geometry, the values of ξ are symmetric about the line $\delta_x = \delta_z$. The data nearly collapse onto a single line when plotted versus $\min(\delta_x, \delta_z)$, indicating the closest wall has the greatest effect on the collapse. The more spherical collapses occur furthest from the corner where the walls have the least effect. The strength of the initial rarefaction is inversely proportional to its original distance from the bubble, so increasing δ_x or δ_z decreases its effect on the collapse. For two cases with the same values of δ_x/δ_z, the case with the greater value

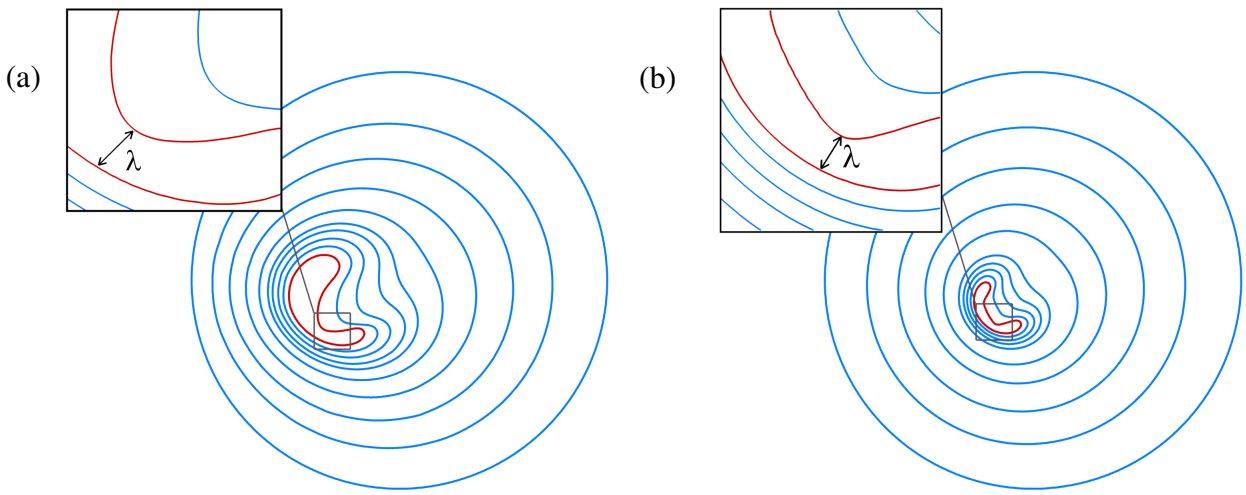


Figure B.4: Outlines for (a) $(\delta_x, \delta_z) = (1.5, 3.0)$ and (b) $(\delta_x, \delta_z) = (2.5, 3.0)$, and the measurement of λ.

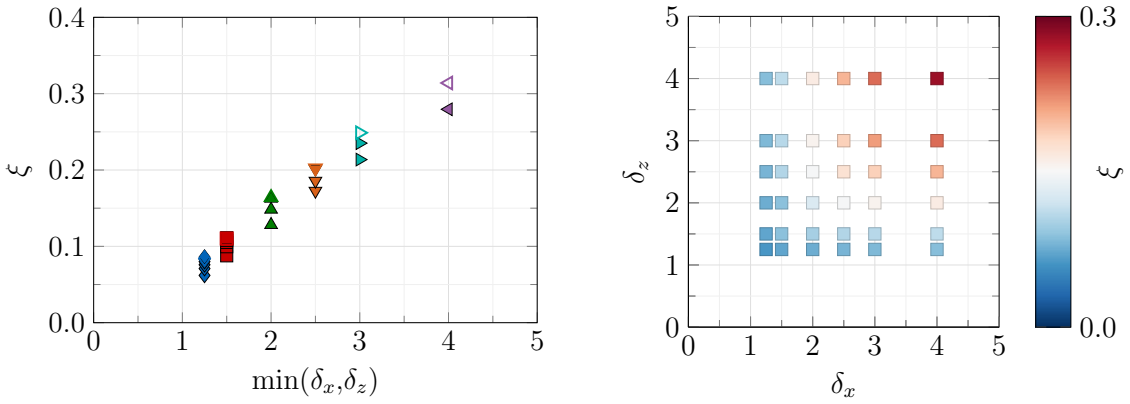


Figure B.5: Non-sphericity ξ vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty.

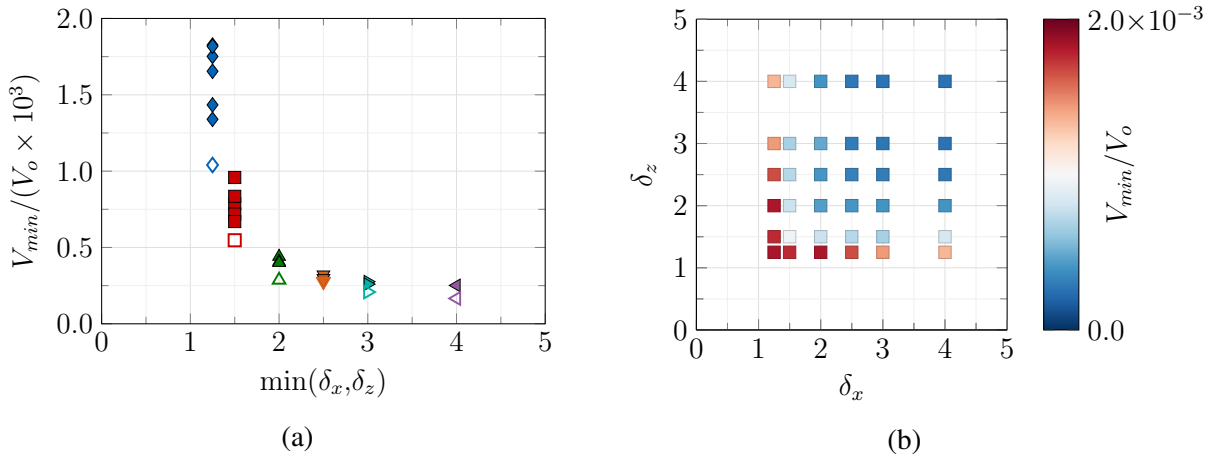


Figure B.6: Minimum volumes vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty.

of $\sqrt{\delta_x^2 + \delta_z^2}$ collapses more spherically for the same reason. The strength of all subsequent waves that interact between the bubble and the walls decreases in magnitude while still contributing to the asymmetric collapse by modifying the driving pressure and acceleration of the interface. When $\delta_x \neq \delta_z$, the non-spherical dynamics are most affected by the waves reflected off the closest wall, until the wave reflected off the other wall returns. At this point, the collapse becomes more non-spherical due to the enhanced interactions of the additional sinks. The dynamics of two bubbles with the same values of δ_x/δ_z but different values of δ_x and δ_z are distinct because of the different reflected rarefaction strengths and because the bubble/wave interactions begin at different points in the collapses. Those bubbles initially closest to the corner experience the least spherical collapses because the asymmetry is imparted on the bubble interface earliest in the collapse. For a given δ_x, increasing δ_z decreases the deviation from spherical shape until the point at which $\delta_z > \delta_{cr}$, in which case the dynamics are those of a bubble near a single wall, as the rarefaction reflected off of the additional wall does not reach the bubble before it collapses.

The increased non-sphericity due to the bubble-wall interactions affects key collapse quantities. The minimum volume V_{min} and collapse time $t_{collapse}$ (time to minimum volume) are shown in Fig. B.6 and Fig. B.7, respectively, for different initial stand-off distances. The minimum bubble volume is a measure of the strength of the emitted shock; a smaller minimum volume gives rise to a higher bubble pressure and thus stronger emitted shock. Holding either stand-off distance constant and increasing the other leads to a smaller minimum volume. This behavior is consistent with collapse near a single wall, where the minimum volume achieved at collapse is smaller for bubbles increasingly far from the wall, because their collapse is more spherical. As with ξ the data

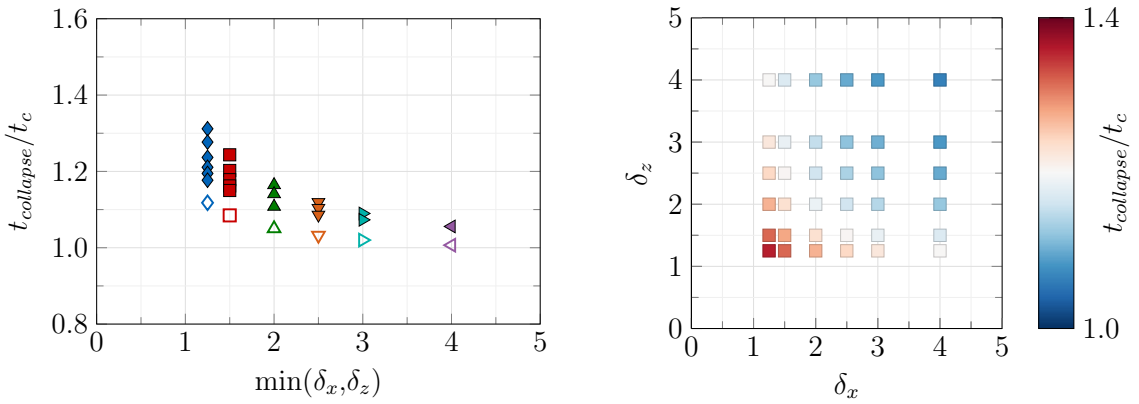
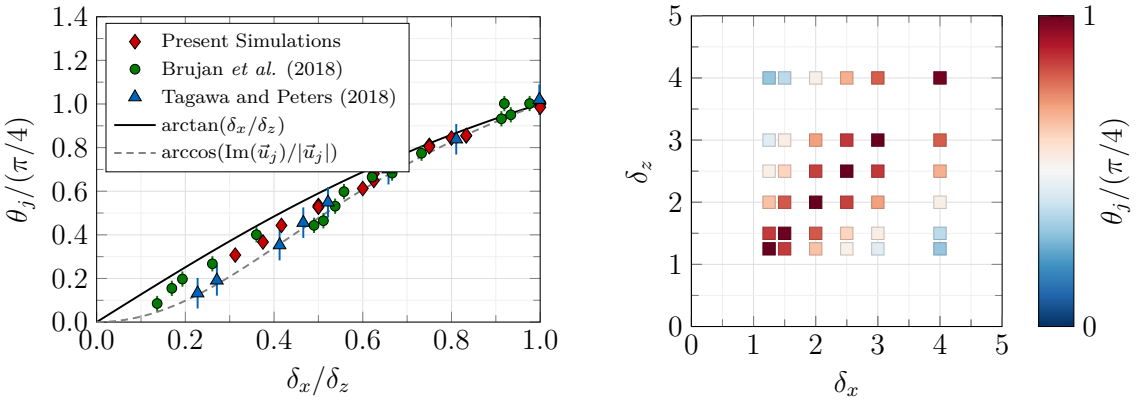


Figure B.7: Collapse times scaled by the Rayleigh collapse time vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$; gray circles: other single-wall cases. Corner cases are filled and single-wall cases are empty.

collapse when V_{min} is plotted against $\min(\delta_x, \delta_z)$. For a given δ_x however, the minimum volume is larger in the cases near a corner than those near a single wall. These larger minimum volumes can be explained by the dynamics leading to more intricate bubble morphologies of the collapses when compared to those near a single wall, which results in less effective focusing of the liquid kinetic energy. The impact of this decreased focus can also be seen in the longer collapse times for bubbles near a corner than those in collapse near a single wall. The increased collapse time is a consequence of the two additional image sinks, whose induced velocity reduces the inward interfacial velocity. Bubbles initially closest to the wall have the longest collapse times because the initially emitted rarefaction wave reduces the local driving pressure sooner in the collapse than in the case of bubbles located farther away from the wall; a lower driving pressure slows down the dynamics and thus give rise to a longer collapse time [198]. Under the influence of both walls making up the corner, this effect is increased as the second wall reflects an additional portion of the initial rarefaction wave that then returns to the bubble and further decreases the pressure field, although to a lesser degree than the portion of the wave that returns from the closest wall. This behavior is supported by the observation that bubbles collapsing near a corner have longer collapse times than a bubble collapsing near a single wall. For all cases near the corner, the collapse time is notably longer than the Rayleigh collapse time for these reasons.



(a) (b)

Figure B.8: Jet angle θ_j as a function of (a) the ratio of the initial stand-off distances, δ_x/δ_z; red diamonds: direct simulations; green circles: experimental data from [34]; blue triangles: experimental data from [233]; black line (arctan(δ_x/δ_z)): Eq. (B.6); gray dashed line (arccos(Im($\vec{u}_j$)/|$\vec{u}_j$|)): Eq. (B.8) and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$; gray circles: other single-wall cases. Corner cases are filled and single-wall cases are empty.

B.4.3 Bubble Kinematics and Dynamics

The kinematics and dynamics of bubbles collapsing near a corner have been the focus of past experimental studies [34, 58, 233], so we use this section to validate our results and examine specific features of the collapse. As illustrated in Fig. B.2 and Fig. B.3, the direction of the re-entrant jet formed in bubble collapse near a corner is dependent on the initial stand-off distances from the two walls; additionally, the part of the emitted shock with the highest amplitude appears to propagate in the direction of the jet. In cases where the bubble is sufficiently far from the second wall ($\delta_z > \delta_{cr}$), the jet is perpendicular to the closest wall. However, for bubbles sufficiently close to both walls ($(\delta_x, \delta_z) < \delta_{cr}$), the jet points in a direction that is neither at the closest initial location along either wall, but somewhere in between, which is one of the ways in which the shock wave dynamics are affected by the presence of the second wall.

To calculate the jet angle, we define the jet tip as the point along the interface on the far side of the bubble that is closest to the interface on the far side, once the bubble starts to involute, i.e., is no longer purely concave. The jet angle as a function of time is then calculated as the angle between jet positions at successive time steps. The jet angle is averaged over the interval starting from the time at which the jet is first detected and the time at which the jet impacts the far side of the bubble; the angle is nearly constant throughout this interval. For a range of initial stand-off

distances similar to ours, experimental data were found to obey the following power law [34],

$$\theta_j = \frac{\pi}{2}(\delta_x/\delta_z)^m \approx \arctan(\delta_x/\delta_z), \quad m = 0.981 \pm 0.035, \tag{B.6}$$

where m is a parameter fit to experimental results. Another measure of jet angle is from Tagawa and Peters [233], who used potential flow to determine the jet angle based on the jet velocity $\vec{u}_j$,

$$\frac{\vec{u_j}}{C} = \frac{e^{i\theta_b} - e^{i(2\pi-\theta_b)}}{[1-\cos(2\theta_b)]^{3/2}} + \sum_{k=1}^{g-1}\left[\frac{e^{i\theta_b} - e^{i(2\pi k/g-\theta_b)}}{[1-\cos(2\theta_b - 2\pi k/g)]^{3/2}} + \frac{e^{i\theta_b} - e^{i(2\pi k/g+\theta_b)}}{[1-\cos(2\pi k/g)]^{3/2}}\right], \tag{B.7}$$

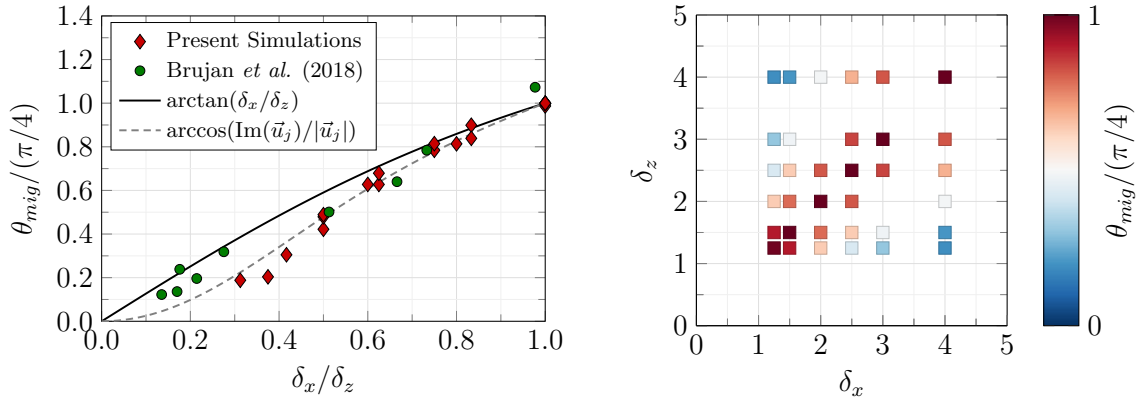
where $C = p/(\pi\sqrt{128}R_0^2)$, g the integer index of the corner geometry, p the sink strength, and $g = 2$ describes two perpendicular walls. From equation (B.7), the jet angle is:

$$\theta_j = \arccos(\mathrm{Im}(\vec{u}_j)/|\vec{u}_j|). \tag{B.8}$$

Figure B.8 shows the jet angle as a function of δ_x/δ_z, comparing our numerical simulations to the experimental data from [34, 233], as well as with theoretical results from equations (B.6) and (B.8). Our simulation results fit well within the bounds of the error bars of Brujan *et al.* [34] and Tagawa and Peters [233], and follow trends similar to those proposed in each respectively. For most of the δ_x/δ_z range under consideration, the models provide upper and lower bounds to the experimental and simulations data. This agreement indicates that our simulations also accurate represent the jet dynamics governed by interactions between the bubble and its images as described by potential flow theory.

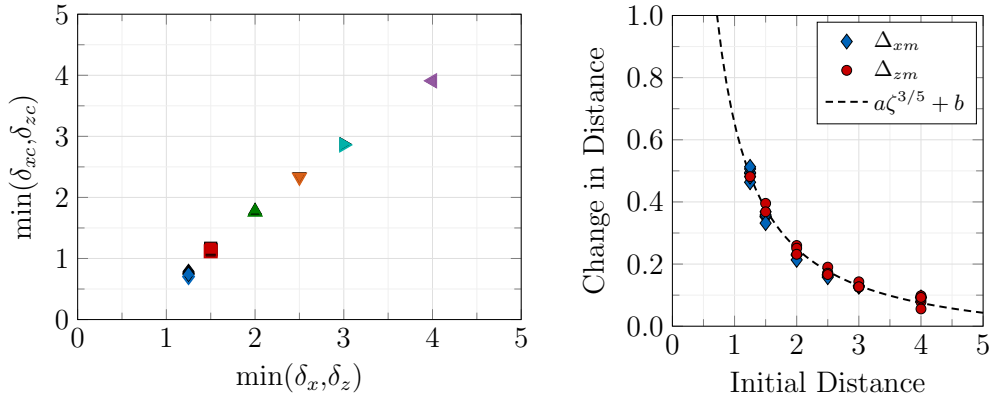
Over the course of the collapse, the bubble migrates toward both walls to varying degrees. We calculate the volumetric centroid of the bubble by taking moments of the volume fraction field, throughout the collapse. Bubbles with one of δ_x or δ_z greater than δ_{cr} follow a straight line toward the closest wall. Bubbles with initial stand-off distances such that they are affected by both walls follow trajectories somewhere in-between. One method of characterizing the trajectory is utilized by Brujan *et al.* [34], who quantified the general direction of the bubble motion with a parameter defined as the migration angle θ_{mig}, measured between the initial centroid location and the location at collapse. Figure B.9 compares this migration angle between our simulations, Brujan's experiments, and equations (B.6) and (B.8). Again, the agreement between simulations and experiments is reasonable; discrepancies may be explained by the deviation in angle between the first and subsequent collapses in [34]. Our numerical simulations agree well with the potential flow model from [233], while the experiments show better agreement with equation (B.6), which was developed based on these experiments.

The location of the bubble centroid at collapse is important as it is a good approximation of the



(a) (b)

Figure B.9: Bubble migration angle θ_j as a function of (a) the ratio of the initial stand-off distances, δ_x/δ_z; red diamonds: direct simulations; green circles: experimental data from [34]; black line (arctan(δ_x/δ_z)): Eq. (B.6); gray dashed line (arccos(Im($\vec{u}_j$)/|$\vec{u}_j$|)): Eq. (B.8) and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$; gray circles: other single-wall cases. Corner cases are filled and single-wall cases are empty.



(a) (b)

Figure B.10: (a) min(δ_{xc}, δ_{zc}) vs. min(δ_x, δ_z); blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty. (b) Migration from initial stand-off distances to collapse location vs. initial stand-off distance; red circles: Δ_{xm} vs. δ_x; blue squares: Δ_{zm} vs. δ_z. The black dashed curve is a direct fit of Δ_{xm} and Δ_{zm} to $a\zeta^{3/5} + b$ with $a \approx 1.9$ and $b \approx -0.06$.

origin of the shock emitted upon collapse. Figure B.10(a) shows the distance $\min(\delta_{xc}, \delta_{zc})$ between of the bubble centroid at collapse and the near wall, where δ_{xc} and δ_{zc} are the normalized x and z positions of the bubble centroid at the point of collapse.

While the centroid positions at collapse are very similar for collapse near a corner and collapse near a single wall, the latter are systematically closer to the wall than the former. The difference in centroid locations between the single wall and corner cases is largest near the corner, and monotonically decreases as $\min(\delta_x, \delta_z)$ increases. In other words, the location of the bubble at collapse is slightly farther from the wall in collapse near a corner due to the effect of the second wall. This behavior is consistent with the trend in the maximum wall pressure observed in the next section, in that the difference between the maximum wall pressure in the single wall and corner cases decreases with increasing $\min(\delta_x, \delta_z)$.

The distance traveled from the initial position is calculated as $(\Delta_{xm}, \Delta_{zm}) = (\delta_x - \delta_{xc}, \delta_z - \delta_{zc})$. This motion is important when predicting impact loads as the centroid location at collapse is a reasonable measure of the location where the shock is released. Figure B.10(b) shows the distance traveled in each direction, measured from the initial stand-off distances to the position at collapse.

Bubbles that are initially furthest from either wall experience the smallest change in the position of their centroids until the collapse. The change in distance in x and in z follow a similar trend. Supponen *et al.* [229] defined an anisotropy parameter $\zeta = 0.195(\delta/R_0)^{-2}$, which is a dimensionless version of the Kelvin impulse and showed that it can be used to predict the dynamics of bubbles collapsing near a single rigid wall. They further suggested that for strong jets the displacement of the centroid scales with $\zeta^{\frac{3}{5}}$ when the bubble location is the point where the jet impacts the opposite side ($\zeta > 0.1$ is considered strong, present simulations have $0.5856 \leq \zeta \leq 0.0572$). Our simulations exhibit a similar dependence of the migration distance on the stand-off distance, which can be seen as the fit curve in Fig. B.10(b). The bubble centroid at collapse can be then predicted by $\delta_{jc} = \delta_j - a\zeta_j^{3/5} - b$, where $j = x, z$. The necessity to scale and add a constant to $\zeta^{\frac{3}{5}}$ is likely due to different measurements of the bubble centroids and to the effects of the second wall. With the introduction of a constant scaling and offset to account for the different physics in the corner compared to near a single wall, the Kelvin impulse accurately predicts the migration behavior observed in each direction. The motion towards both walls is a key difference in the dynamics between those of collapse near a single wall.

The speed of the jet is indicative of the strength of the water-hammer shock produced upon impact with the opposite side of the bubble, as its pressure scales as $\rho_l a_l |u_j|$ where ρ_l is the liquid density, a_l is the sound speed in the liquid, and $|u_j|$ is the velocity magnitude [227]. Figure B.11 shows the maximum jet speed for different initial stand-off distances. The results are consistent with our observations of minimum volume and collapse time, with the fastest jets occurring furthest from the corner and the slowest occurring closest. Based on the observation that non-sphericity increases

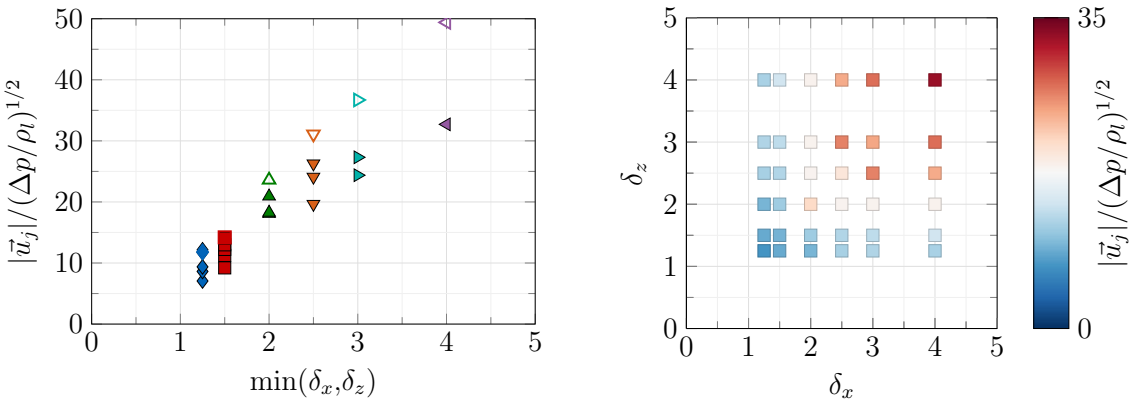


Figure B.11: Maximum jet speed vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty.

as bubbles start their collapse closer to the walls, the maximum jet speed by the deviation from spherical shape indicates that more spherical collapses give rise to increased speeds. Additionally, when affected by two walls, the jet speeds are not as high as those achieved in collapse near a single wall, consistent with the fact that the velocity is reduced due to the flow induced by the additional two image sources. As such, the water hammer pressure is lower for bubble collapse near a corner than near a single wall.

B.4.4 Wall Pressure

The maximum pressure produced by a collapsing bubble along a wall is the quantity characterizing impact loads in cavitation erosion. If only δ_x or δ_z is greater than δ_{cr}, the problem corresponds to collapse near a single wall, in which case the maximum pressure decreases with increasing initial stand-off distance [24, 115]. In the present work, all cases near the corner have initial stand-off distances within δ_{cr}. Since the dynamics are affected by both walls, the pressure produced by collapse near a corner is different than that near a single wall in the following ways: (i) the maximum re-entrant jet velocity is affected by the proximity of the second wall, (ii) the part of the emitted shock with the highest amplitude propagates in the direction of the jet, which no longer points toward the nearest wall but is deflected toward the corner, and (iii) the location of the bubble centroid at collapse is attracted toward the second wall. For bubbles whose collapse is affected by both walls, the maximum pressure is unlikely to be produced at the location along the wall closest to the initial bubble position, as would be the case in collapse near a single wall. To better understand

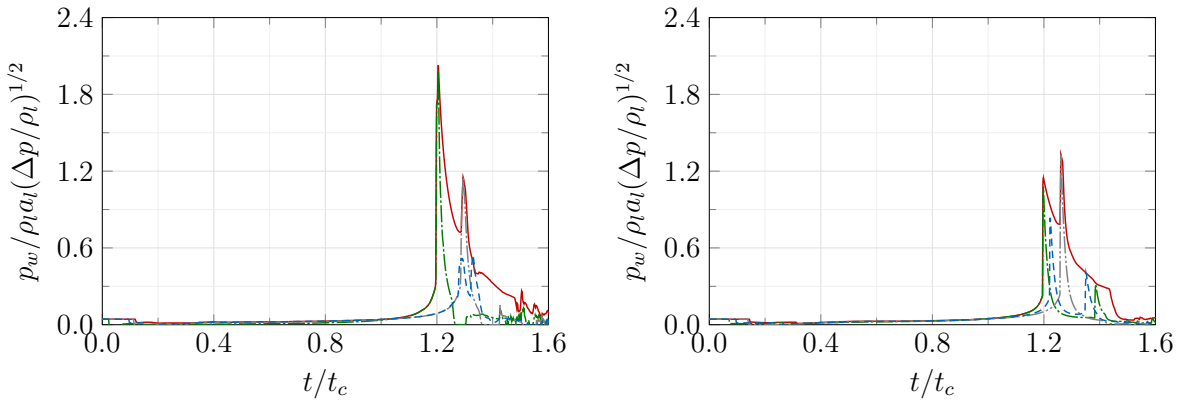


Figure B.12: Time history of wall pressure at the location of the maximum wall pressure (red solid line), at the corner (green dashed-dotted line), $(x, y, z) = (\delta_x, 0, 0)$ (gray dash-dot-dot), and $(x, y, z) = (0, 0, \delta_z)$ (blue dashed line) for (a) $(\delta_x, \delta_z) = (1.5, 3.0)$ and (b) $(\delta_x, \delta_z) = (2.5, 3.0)$.

these effects at play, Fig. B.12 shows the time history of the wall pressure at $(x, y, z) = (\delta_x, 0, 0)$, $(x, y, z) = (0, 0, \delta_z)$, the pressure at the corner, and the maximum wall pressure over the course of the simulation, for two sets of stand-off distances. Interpreting the data in conjunction with Fig. B.2 and Fig. B.3, the wall pressure initially decreases due to the rarefaction reflecting off the closest wall and then does not change significantly until the water-hammer shock and implosion shock interact with the closest wall, which causes a rapid increase in the pressure. For the case $(\delta_x, \delta_z) = (1.5, 3.0)$, the maximum wall pressure is achieved by this process, though at a location along the wall closer to the corner. A similar phenomenon occurs along the farther wall, though the peak pressure is smaller. The time-delay between the water-hammer and implosion shocks can be seen in the first peak of Fig. B.12(a), with the maximum occurring only after the implosion shock impinges upon the wall. In the case $(\delta_x, \delta_z) = (2.5, 3.0)$, the portion of the shock front reflecting off the closest wall and that reflecting off the farther wall intersect close to the corner, thereby giving rise to the highest pressure at that location, even if the shocks do not exactly intersect there. In this case, there is no delay between the water-hammer and implosion shocks in this case, as expected from Fig. B.3.

The above observations suggest that there are geometrical configurations for which the pressure generated in the corner is larger than that produced along either wall. Despite the complex bubble morphologies and shock dynamics in these collapses, acoustic arguments can be used to predict the initial stand-off distances for which the maximum wall pressure is larger in the corner than along either wall. In the case of $\delta_x > \delta_{cr}$ and $\delta_z > \delta_{cr}$, the collapse is spherical and the shock pressure is

given by,

$$\frac{p_s(r)}{p_\infty} = \left[\gamma^{-1/(\gamma-1)}\left(6\gamma^2\frac{p_\infty}{\rho_l a_l^2}\right)^{1/(2\gamma-1)}\right]^{-(\gamma-0.38)}\left(\frac{\Delta p}{p_\infty}\right)^{(\gamma-0.38)/(2\gamma-1)}\left(\frac{r}{R_0}\right)^{-1.13}, \tag{B.9}$$

where γ is the polytropic index, and r is the radius of the radially propagating shock wave [24]. We assume that the shock pressure scales in the same way for collapse in a corner. Neglecting bubble migration during collapse and assuming the jet impacts the bubble wall at the centroid, the distances from the bubble to the wall in the direction of travel of the jet, and the distance from the bubble to the corner are $L_w = \delta_x/\cos\theta_j$ and $L_c = \sqrt{\delta_x^2 + \delta_z^2}$ respectively. The pressures at the wall and corner, given the shock pressure is greatest in line with the jet [24] are then $p_w = 2p_{bc}L_w^{-1.13}$ and $p_c = 4p_{bc}L_c^{-1.13}$ where p_{bc} is the shock pressure immediately after the collapse. The coefficients 2 and 4 come from the image bubbles doubling and quadrupling the shock pressure at the wall and corner respectively, and the power of -1.13 is based on underwater explosion scalings [50]. The ratio of pressures at the wall and corner is then,

$$\frac{p_c}{p_w} = 2\left[\frac{L_c}{L_w}\right]^{-1.13} = 2\left[\frac{\delta_x}{\cos\theta_j\sqrt{\delta_x^2+\delta_z^2}}\right]^{1.13}. \tag{B.10}$$

All of the constants divide out such that the ratio of pressures depends only on the initial stand-off distances. This result directly predicts the initial stand-off distances for which the maximum pressure occurs directly in the corner. As observed in experiments [34] and in the present simulations, the bubble migrates during its collapse due to the presence of the walls. Given that the shock emitted upon collapse decays as it propagates outward, closer proximity to the wall(s) at the instant of collapse results in higher pressures than if that same shock had been released at the initial centroid location. The degree of motion towards each wall is straightforward to calculate as described in §B.4.3. As such, we replace δ_x and δ_z with $\delta_{xc} = \delta_x - \Delta_{xm}$ and $\delta_{zc} = \delta_z - \Delta_{zm}$ to find,

$$p_r = 2\left[\frac{\delta_{xc}}{\cos\theta_j\sqrt{\delta_{xc}^2+\delta_{zc}^2}}\right]^{1.13} \approx 2\left[\frac{\delta_x - a\zeta_x^{3/5} - b}{\cos\theta_j\sqrt{(\delta_x - a\zeta_x^{3/5} - b)^2 + (\delta_z - a\zeta_z^{3/5} - b)^2}}\right]^{1.13}. \tag{B.11}$$

This semi-empirical expression predicts the region in $\delta_x - \delta_z$ space for which the pressure in the corner exceeds that on either wall, taking into consideration the motion of the bubble centroid. The jet angle can be predicted using only the initial stand-off distances with equation (B.6) or (B.8). Comparing this expression to our simulation results, the wall location δ_{zp} at which the maximum pressure is achieved is shown in Fig. B.13. The region of initial stand-off distances where equa-

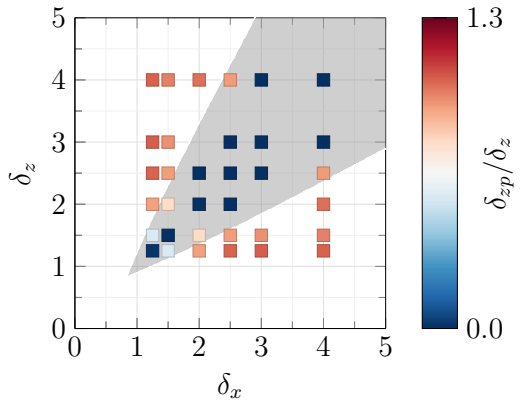


Figure B.13: Location of maximum wall pressure δ_{zp}/δ_z; gray region: $p_r \geq 1$.

tion (B.11) predicts $p_r \geq 1$, i.e. the maximum pressure occurs in the corner, is shown in gray. The maximum pressure is achieved in the corner for cases that lie along the $\delta_x = \delta_z$ line due to the double reflection of the shock at that location, as well as for those cases within a certain range of this line. In all other cases, the maximum pressure occurs within $0.5R_o$ from the closest initial wall location, in the direction of the corner, consistent with the jet angle. Overall, the agreement between our simulations and our theoretical model (B.11) is good, indicating that the location of the maximum wall pressure near a corner is well described using these shock interactions. If the desired effect is to reduce damage in a corner of a component undergoing repeated cavitation, this result suggests that a design that minimizes nucleation sites with $\delta_x/\delta_z \approx 1$ is desirable. Conversely, if high pressure in the corner is desirable (as may be the case for ultrasonic cleaners) then maximizing nucleation sites with $\delta_x/\delta_z \approx 1$ would be optimal.

The maximum wall pressure itself is a key quantity of interest in cavitation damage. Figure B.14 shows the maximum wall pressures for different initial stand-off distances. The highest wall pressures occur for collapse closest to the walls, as expected. The maximum wall pressure decreases quickly with initial stand-off distance for both collapse near a corner and near a single wall because of the decay of the emitted shock with propagation distance. When parameterized by ξ, the pressures produced by collapse near a single wall are consistently larger than those produced by collapse near two walls. This follows from (i) the decreased jet speeds for collapse near a corner, which result in a lower water-hammer pressure, and (ii) the larger minimum volumes of collapse near a corner, which result in a weaker implosion shock. These two differences in the dynamics are a direct consequence of the second wall making up the corners. Those cases for which the wall pressures are higher for collapse near a corner than for collapse near a single wall are cases where the reflected shocks intersect at the corner. The difference between the pressures in collapse near a corner and collapse near a single wall decreases with increasing distance from the walls as the effect the walls becomes small. The asymmetry of the collapse due to the second wall generally

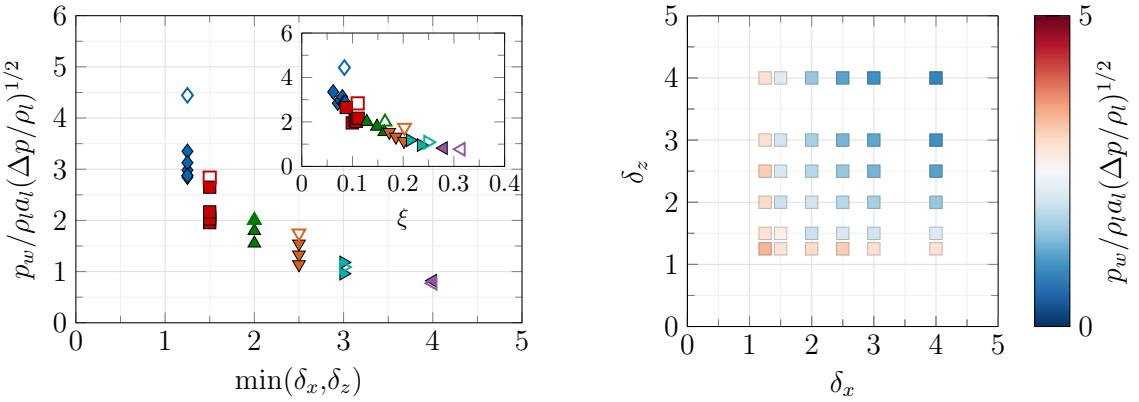


Figure B.14: Maximum wall pressure vs. (a) $\min(\delta_x, \delta_z)$ and ξ (inset), and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty.

manifests in less effective focusing of the energy in the collapse, leading to weaker shocks and therefore lower wall pressures than those produced in collapse near a single wall. It follows that the presence of the second wall is generally expected to decrease the damage potential of bubble collapse compared to collapse near a single wall.

B.4.5 Wall Temperatures

Certain materials may be sensitive to temperatures produced during the bubble collapse [63], in which case the maximum wall temperature is an important quantity of interest. As illustrated in [24], bubble collapse near a single wall can increase the wall temperature via two mechanisms: (i) as it reflects off the wall, the shock emitted upon collapse raises the temperature of the water adjacent to the wall for a short period; and (ii) bubbles sufficiently close to the wall come into close proximity with the wall at collapse, at which point their contents can reach a high temperature and lead to heat conduction to the wall. The resulting wall temperature can be calculated by solving the corresponding heat conduction problem. Though both mechanisms are short-lived, the relatively large temperature difference can give rise to significant changes in wall temperatures. The maximum wall temperature achieved during collapse near a corner is shown in Fig. B.15 for different initial stand-off distances assuming properties corresponding to ultra-high-molecular-weight polyethylene (UHMWPE) as it is relevant to applications of interest [24]. Most cases follow mechanism (i), with the shock wave impingement creating a transient increase in the wall temperature. Bubbles initially closest to one or both walls (i.e. with the smallest ratio δ_x/δ_z) follow mechanism (ii), whereby the

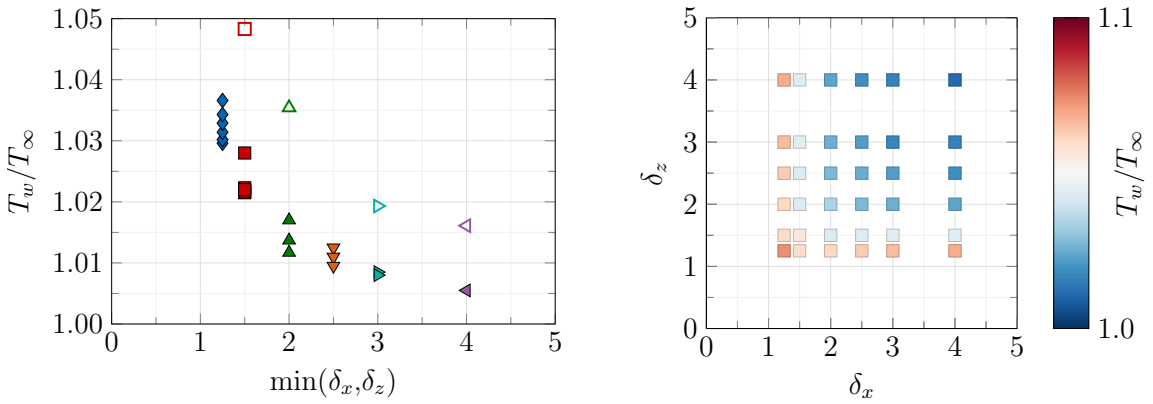


Figure B.15: Maximum wall temperature (UHMWPE) vs. (a) $\min(\delta_x, \delta_z)$ and (b) initial stand-off distances; blue diamonds: $\delta_x = 1.25$; red squares: $\delta_x = 1.5$; green deltas: $\delta_x = 2.0$; orange gradients: $\delta_x = 2.5$; teal right triangles: $\delta_x = 3.0$; purple left triangles: $\delta_x = 4.0$. Corner cases are filled and single-wall cases are empty.

bubble comes in contact with the wall. The importance of contact is visible by considering the cases with $\delta_x = 1.25$ and increasing δ_z; the wall temperature increases despite the maximum wall temperature of all cases considered occurring for $\delta_x = \delta_z = 1.25$.

Though not shown here for conciseness, the locations of the maximum wall temperatures exhibit a behavior similar to the locations of the maximum wall pressures. Cases for which δ_x/δ_z is close to unity experience the maximum wall temperature directly in, or very close to, the corner. Other cases give rise to the maximum wall temperature along the wall, away from the corner. Equation (B.11) can also be used to determine initial stand-off distances for which the maximum wall temperature occurs in the corner rather than along the wall. Initial stand off distances for which the maximum temperature occurs in the corner exactly match those for which the maximum pressure occurs in the corner.

The temperatures calculated here do not however reach the melting point of the wall, and thus significant damage is unlikely caused as a direct result of the heat produced; instead, experimentally observed damage is likely primarily due to increased wall pressure. This finding is consistent with that of Beig *et al.* [24] for UHMWPE. Despite the inability of a single bubble to reach melting temperatures, repeated collapses may transfer sufficient heat to cause damage. Materials with lower melting temperatures may also be susceptible to melting.

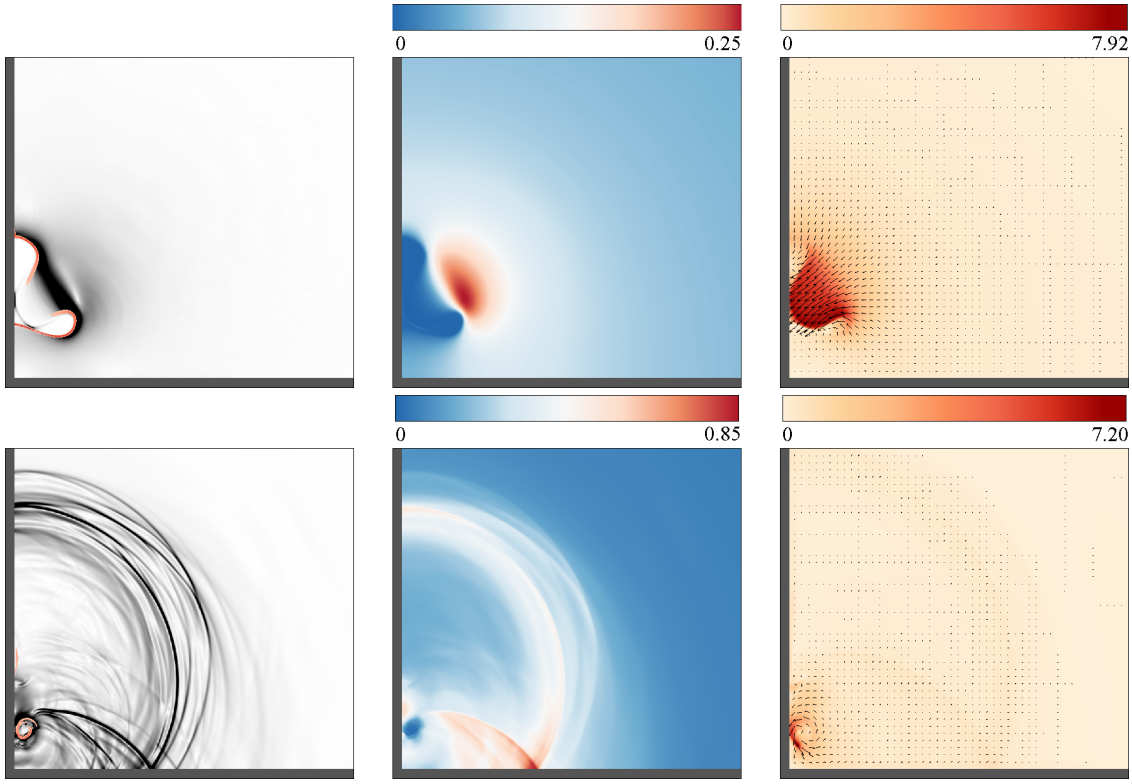


Figure B.16: Contours of (numerical) Schlieren, pressure normalized by $\rho_l a_l \sqrt{\Delta p/\rho_l}$, and velocity magnitude normalized by the characteristic interfacial speed $\sqrt{\Delta p/\rho_l}$ for $(\delta_x, \delta_z) = (0.75, 1.25)$ at $t = 1.387$ and $t = 1.607$.

B.5 Attached Bubbles

Bubbles attached to a rigid surface have been shown to produce large pressures [240] as the re-entrant jet directly impinges upon the surface. Additionally, given the direct contact of the hot gas at minimum volume, high wall temperatures may be produced. Although the contact line dynamics cannot be represented in the present numerical framework, we assume like Trummler *et al.* [244] that their effect is negligible given the present interest in inertia-dominated collapse.

Figures B.16 and B.17 show contours of numerical Schlieren, pressure, and velocity magnitude for simulations with $(\delta_x, \delta_z) = (0.75, 1.25)$ and $(\delta_x, \delta_z) = (0.25, 0.75)$. The bubble is immediately aware of the wall(s) to which it is attached. A re-entrant jet develops, again at an angle that is not normal to the attached wall, and penetrates the bubble. The jet directly impinges upon the wall, thereby generating an intricate set of shocks that further interact with the bubble. In the $(\delta_x, \delta_z) = (0.75, 1.25)$ case, the bubble appears to take a toroidal shape with one part attached to the wall, while for $(0.25, 0.75)$ the bubble converges toward the corner and eventually detaches from one side after the shock is generated.

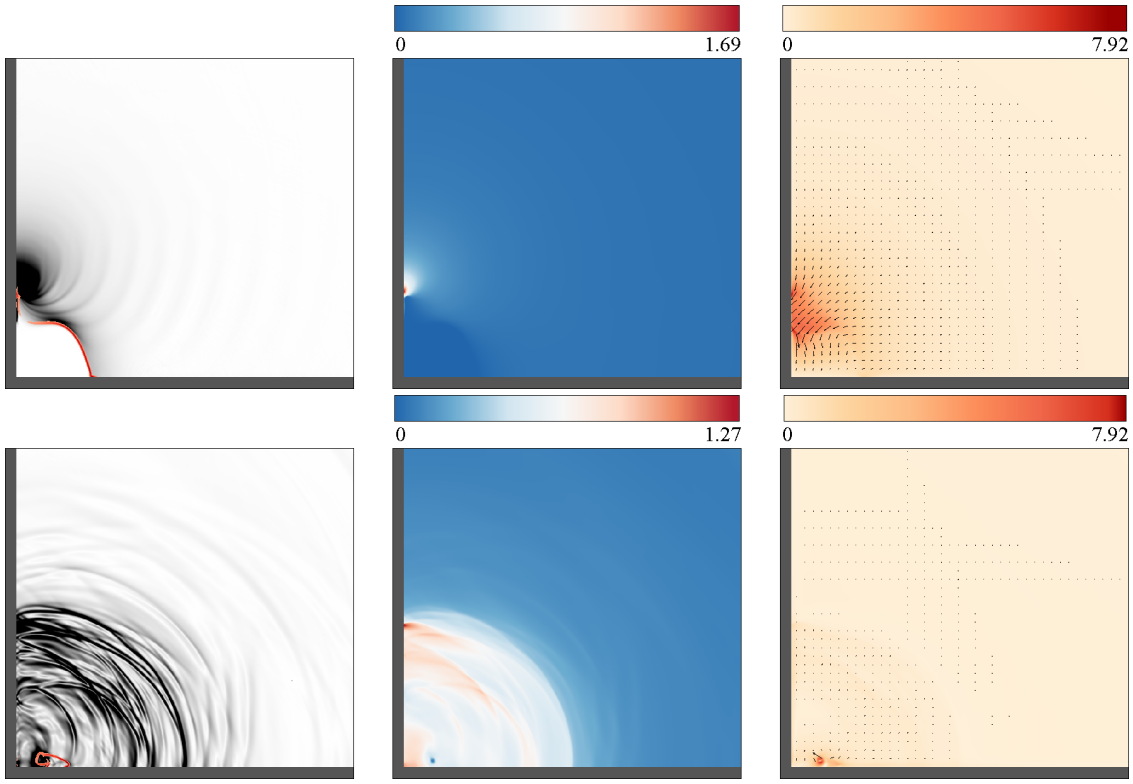


Figure B.17: Contours of (numerical) Schlieren, pressure normalized by $\rho_l a_l \sqrt{\Delta p/\rho_l}$, and velocity magnitude normalized by the characteristic interfacial speed $\sqrt{\Delta p/\rho_l}$ for $(\delta_x, \delta_z) = (0.25, 0.75)$ at $t = 1.319$ and $t = 1.493$.

As compared to collapse of a bubble initially detached from the walls, the shock waves produced by collapse of a bubble initially attached to one or both walls are more numerous and complex. This is a consequence of the impact of the re-entrant jet on the wall, which reflects the generated shock back into the portions of the bubble attached to the wall. This reflected shock propagates through the bubble and impinges upon the air-water interface of the bubble wall and generates a reflected rarefaction which returns to the wall and further contributes to the number of waves present in the flow field. The direct impact of the re-entrant jet on the wall produces large pressures along the wall, and the strength of the shock does not decay significantly as they travel to the walls due to their proximity with the walls. As such, the wall pressures produced by collapse of a bubble initially attached to one or both walls are expected to be larger than those produced by collapse of a bubble initially detached from both walls.

Conversely, the strong effects of the wall are expected to reduce the intensity of collapse for attached bubbles. Figure B.18 shows the collapse times for different values of δ_z. The collapse times decrease with increasing δ_z and are slower than in the detached case, as expected due to the more rapid reduction in driving pressure following wave reflection off the wall. The lack of strong dependence on δ_x indicates that the collapse time depends most on the shock reflecting off the far wall. As previously discussed, a longer collapse time implies a less violent collapse, which is consequential for the jet speed and maximum wall pressure.

The jet angles point almost directly towards the corner for small δ_z and decrease with increasing δ_z. Consistent with the collapse times, the data do not show variability with δ_x; they also approach the expected value of $\theta_j = 0$ as δ_z is increased towards δ_{cr}. The jet angle of bubbles initially attached to one or either wall do not appear to follow (B.6) or (B.8), which implies potential flow does not accurately approximate the collapse dynamics, likely due to the significant compressibility effects.

The ambivalence towards δ_x in both the collapse time and jet angle of the attached cases (particularly those attached to only one wall) is likely because the initial rarefaction travels normal to the bubble surface, so the effect of the wall is not felt on the bubble interface until the radially converging transmitted shock wave reflects off of the walls and returns to the bubble interface. The rarefaction returning from the far wall breaks the axismmetry and non-uniformly changes the driving pressure around the bubble–hence the strong dependence on δ_z.

Consistent with the collapse times, the jet speeds are lowest for collapse closest to the corner. While for each value of δ_x the jet speed does not increase monotonically with increasing distance from the corner, the jet speeds are generally increasing. The maximum jet speed depends on local dynamics, which are complex due to the wave-interface interactions, such that their dependence on δ_x is not straightforward and we observe this non-monotonic behavior. As for fully detached bubbles, the jet speeds are lower than in collapse near a single wall. Compared to the jet speeds observed

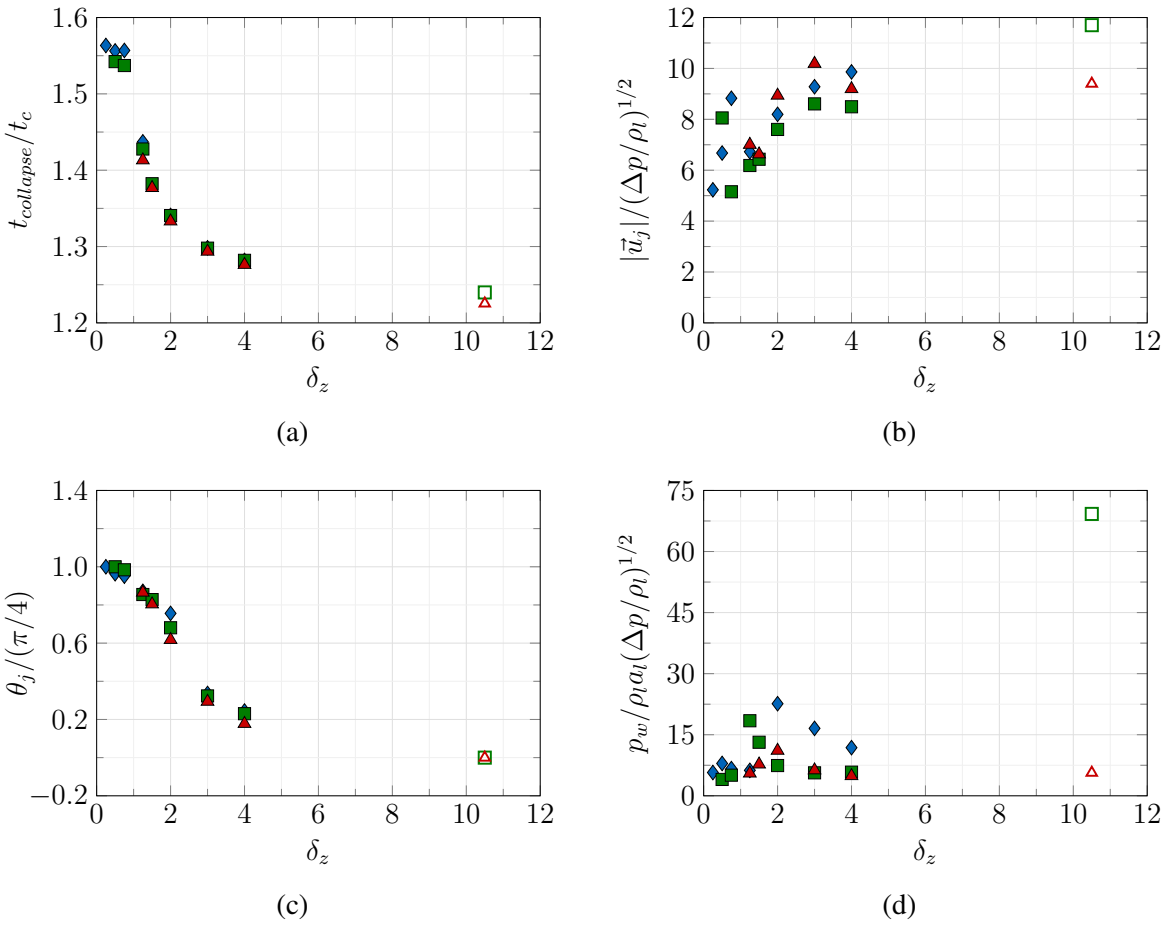


Figure B.18: (a) Collapse time, (b) jet speed, (c) jet angle, and (d) maximum wall pressure vs. δ_z; blue diamonds: $\delta_x = 0.25$; green squares: $\delta_x = 0.5$; red deltas: $\delta_x = 0.75$. Corner cases are filled and single-wall cases are empty.

for fully detached bubbles, the jet speeds observed here are generally lower, again suggesting that the collapse of bubbles initially attached to either wall are less violent than their initially detached counterparts.

The maximum wall pressure occurs where the jet impinges upon the wall, and is thus dependent on the jet angle. The resulting maximum wall pressure depends on two competing factors: on one hand, the jet speeds are lower for attached bubbles than for detached bubbles, so the water-hammer pressure is expected to be lower; on the other hand, the bubble is closer, if not touching the wall, such that the wall feels the full effect of the jet impact.

There are two regions of behavior for each δ_x value. For $\delta_x = 0.25$ and $\delta_x = 0.75$, when $\delta_z < 2$ there is no clear trend in the pressure. The same is noted for $\delta_x = 0.5$ and $\delta_z < 1.25$, and similar to the maximum jet speed, in these cases the dynamics are governed by the complex shock-interface interactions and characterization is not straightforward. For $\delta_x = 0.25$ and $\delta_x = 0.75$, when $2 \leq \delta_z \leq \delta_c$, and for $\delta_x = 0.5$, when $1.25 \leq \delta_z \leq \delta_c$, we observe a monotonic decrease in the maximum wall pressure with increasing δ_z. The maximum wall pressure appears to converge to the limiting case of bubble collapse near a single wall for $\delta_x = 0.75$; however collapse near a single wall with $\delta_x = 0.5$ shows a substantially higher pressure. This behavior results from a faster re-entrant jet which causes a strong water hammer. We expect that as δ_z approaches δ_{cr} the maximum wall pressures for $\delta_x = 0.5$ and $\delta_x = 0.25$ increases to match those in collapse near a single wall.

B.6 Conclusions

In this article we conduct high-resolution numerical simulations of Rayleigh collapse of a gas bubble near a corner for varying initial stand-off distances to understand the role of the second (perpendicular) wall in Rayleigh collapse near a single wall. We are particularly interested in the maximum pressure produced along the wall and the phenomena governing it, as this quantity is the most consequential hydrodynamic metric for cavitation erosion.

Compared to collapse near a single wall, a bubble is affected by the second wall only if it is within a critical distance from the wall based on the time it takes for the initially released rarefaction to reflect back to the bubble. This reflected rarefaction wave alters the pressure distribution around the bubble, thereby modifying the pressure field by reducing it locally and giving rise to an asymmetric distribution. In the case of collapse near a corner, the different parts of the rarefaction interact with the bubble at different times depending on the initial stand-off distance. This reflection off the second wall communicates the presence of the extra image sinks and breaks the axisymmetry of the converging flow, leading to a deflection of the re-entrant jet toward the corner. As such, the angle of the jet can be parameterized solely by the ratio of stand-off distances. Other key collapse properties (e.g., collapse time, minimum volume, maximum jet speed) can also be described by

a single parameter – the deviation from spherical shape of the bubble, which can be related to the Kelvin impulse. Overall, the presence of the second wall inhibits the collapse, i.e., collapse times are slower, minimum volumes are larger, jet speeds are slower when the bubble collapses in a corner compared to collapse near a single wall. These dynamics lead to shock emission with properties different from those observed in collapse near a single wall, in the following ways: (i) the maximum re-entrant jet speed (and thus water-hammer shock magnitude) is reduced by the proximity to the second wall, (ii) the highest magnitude along the emitted shock front is in the direction of the jet and thus not in the direction normal to either wall, and (iii) the migration of the bubble toward the closest wall is hindered by the second wall, such that the bubble collapses farther from the closest wall than in collapse near a single wall. The combination of these effects leads to, in general, wall pressures that are lower than it would in the case of collapse near a single wall. An additional difference is that the second wall causes an additional reflected shock, which interacts with the shock reflected off the original wall. A consequence is the doubling of pressure in the corner and enhanced wall pressure. Using theoretical arguments, we develop expressions for the maximum wall pressure produced by bubble collapse near a corner and determine which initial stand-off distances lead to higher pressures in the corner than along the closest wall. We also consider heat transfer during this process, as well as the collapse of attached bubbles.

This work sets the stage for studies in more intricate geometries, as well as collapse of a larger number of bubbles. In such situations, bubbles may form along rigid surfaces, such that contact line dynamics may need to be investigated.

BIBLIOGRAPHY

[1] S. I. Abarzhi, D. L. Hill, K. C. Williams, J. T. Li, B. A. Remington, D. Martinez, and W. D. Arnett. Fluid dynamic mathematical aspects of supernova remnants. *Physics of Fluids*, 35(3):034106, 2023.

[2] R. Abgrall. How to Prevent Pressure Oscillations in Multicomponent Flow Calculations: A Quasi Conservative Approach. *Journal of Computational Physics*, 125(1):150–160, 1996.

[3] R. Abgrall and S. Karni. Computations of Compressible Multifluids. *Journal of Computational Physics*, 169(2):594–623, 2001.

[4] R. Abgrall and S. Karni. A comment on the computation of non-conservative products. *Journal of Computational Physics*, 229(8):2759–2763, 2010.

[5] D. Adalsteinsson and J. A. Sethian. A Fast Level Set Method for Propagating Interfaces. *Journal of Computational Physics*, 118(2):269–277, 1995.

[6] M. Adams, P. Colella, D. T. Graves, J. N. Johnson, N. D. Keen, T. J. Ligocki, D. F. Martin, P. W. McCorquodale, D. Modiano, P. O. Schwartz, T. D. Sternberg, and B. Van Straalen. Chombo Software Package for AMR Applications – Design Document. Technical Report LBNL-6616E, Lawrence Berkeley National Laboratory, 2014.

[7] U. Adhikari, A. Goliaei, and M. L. Berkowitz. Nanobubbles, cavitation, shock waves and traumatic brain injury. *Physical Chemistry Chemical Physics*, 18(48):32638–32652, 2016.

[8] S. Adjerid, K. D. Devine, J. E. Flaherty, and L. Krivodonova. A posteriori error estimation for discontinuous Galerkin solutions of hyperbolic problems. *Computer Methods in Applied Mechanics and Engineering*, 191(11-12):1097–1112, 2002.

[9] S. Adjerid and T. C. Massey. Superconvergence of discontinuous Galerkin solutions for a nonlinear scalar hyperbolic problem. *Computer Methods in Applied Mechanics and Engineering*, 195(25-28):3331–3346, 2006.

[10] J. AHRENS, B. GEVECI, and C. LAW. *ParaView: An End-User Tool for Large-Data Visualization*, page 717–731. Elsevier, 2005.

[11] S. Alahyari Beig and E. Johnsen. Maintaining interface equilibrium conditions in compressible multiphase flows using interface capturing. *Journal of Computational Physics*, 302:548–566, 2015.

[12] H. Alfvén. Existence of Electromagnetic-Hydrodynamic Waves. *Nature*, 150(3805):405–406, 1942.

[13] G. Allaire, S. Clerc, and S. Kokh. A Five-Equation Model for the Simulation of Interfaces between Compressible Fluids. *Journal of Computational Physics*, 181(2):577–616, 2002.

[14] S. M. Allen and J. W. Cahn. A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening. *Acta Metallurgica*, 27(6):1085–1095, 1979.

[15] A. S. Almgren, V. E. Beckner, J. B. Bell, M. S. Day, L. H. Howell, C. C. Joggerst, M. J. Lijewski, A. Nonaka, M. Singer, and M. Zingale. CASTRO: A NEW COMPRESSIBLE ASTROPHYSICAL SOLVER. I. HYDRODYNAMICS AND SELF-GRAVITY. *The Astrophysical Journal*, 715(2):1221–1238, 2010.

[16] M. Ancellin, B. Després, and S. Jaouen. Extension of generic two-component VOF interface advection schemes to an arbitrary number of components. *Journal of Computational Physics*, 473:111721, 2023.

[17] W. Aniszewski, T. Ménard, and M. Marek. Volume of Fluid (VOF) type advection methods in two-phase flow: A comparative study. *Computers & Fluids*, 97:52–73, 2014.

[18] R. E. A. Arndt. Cavitation in fluid machinery and hydraulic structures. *Annu. Rev. Fluid Mech.*, 13:273–328, 1981.

[19] M. Baer and J. Nunziato. A two-phase mixture theory for the deflagration-to-detonation transition (ddt) in reactive granular materials. *International Journal of Multiphase Flow*, 12(6):861–889, 1986.

[20] X. Bai and M. Li. Simulating compressible two-phase flows with sharp-interface discontinuous Galerkin methods based on ghost fluid method and cut cell scheme. *Journal of Computational Physics*, 459:111107, 2022.

[21] A. Bakhsh, S. Gao, R. Samtaney, and V. Wheatley. Linear simulations of the cylindrical Richtmyer-Meshkov instability in magnetohydrodynamics. *Physics of Fluids*, 28(3), 2016.

[22] E. Bar-Kochba, M. T. Scimone, J. B. Estrada, and C. Franck. Strain and rate-dependent neuronal injury in a 3D in vitro compression model of traumatic brain injury. *Scientific Reports*, 6(1), 2016.

[23] C. Barajas and E. Johnsen. The effects of heat and mass diffusion on freely oscillating bubbles in a viscoelastic, tissue-like medium. *J. Acoust. Soc. Am.*, 141:908–918, 2017.

[24] S. A. Beig, B. Aboulhasanzadeh, and E. Johnsen. Temperatures produced by inertially collapsing bubbles near rigid surfaces. *Journal of Fluid Mechanics*, 852:105–125, 2018.

[25] T. B. Benjamin and A. T. Ellis. The collapse of cavitation bubbles and the pressure thereby produced against solid boundaries. *Phil. Trans. R. Soc. Lond.*, 260:221–240, 1966.

[26] M. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82(1):64–84, 1989.

[27] M. J. Berger and J. Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53(3):484–512, 1984.

[28] J. P. Best and A. Kucera. A numerical investigation of non-spherical rebounding bubbles. *J. Fluid Mech.*, 245:137–154, 1992.

[29] D. Biskamp. *Nonlinear Magnetohydrodynamics*. Cambridge University Press, 1993.

[30] J. R. Blake and D. C. Gibson. Cavitation bubbles near boundaries. *Ann. Rev. Fluid Mech.*, 19:99–123, 1987.

[31] M. Brassel and E. Bretin. A modified phase field approximation for mean curvature flow with conservation of the volume. *Mathematical Methods in the Applied Sciences*, 34(10):1157–1180, 2011.

[32] M. P. Brenner, S. Hilgenfeldt, and D. Lohse. Single-bubble sonoluminescence. *Rev. Mod. Phys.*, 74:425–484, 2002.

[33] E. A. Brujan, G. S. Keen, A. Vogel, and J. R. Blake. The final stage of the collapse of a cavitation bubble close to a rigid boundary. *Phys. Fluids*, 14:85–92, 2002.

[34] E. A. Brujan, T. Noda, A. Ishigami, T. Ogasawara, and H. Takahira. Dynamics of laser-induced cavitation bubbles near two perpendicular rigid walls. *J. of Fluid Mech.*, 841:28–49, 2018.

[35] E.-A. Brujan, A.-M. Zhang, Y.-L. Liu, T. Ogasawara, and H. Takahira. Jetting and migration of a laser-induced cavitation bubble in a rectangular channel. *Journal of Fluid Mechanics*, 948:A6, 2022.

[36] S. H. Bryngelson, K. Schmidmayer, V. Coralic, J. C. Meng, K. Maeda, and T. Colonius. MFC: An open-source high-order multi-component, multi-phase, and multi-scale compressible flow solver. *Computer Physics Communications*, 266:107396, 2021.

[37] E. Burrell, Z. Huang, W. White, B. Perfect, and E. Johnsen. Resolution Considerations for Two-Dimensional High-Speed Droplet Impact Along a Rigid Wall. In *AIAA Aviation Forum and ASCEND co-located Conference Proceedings*, page 3822. Las Vegas, Nevada, USA, 2025.

[38] E. M. Burrell. *Numerical Investigation of High-Speed Droplet Impact Upon a Solid Surface*. Ph.D. thesis, University of Michigan, 2025.

[39] J. W. Cahn and J. E. Hilliard. Free Energy of a Nonuniform System. I. Interfacial Free Energy. *The Journal of Chemical Physics*, 28(2):258–267, 1958.

[40] W. Cao, Z. Zhang, and Q. Zou. Superconvergence of Discontinuous Galerkin Methods for Linear Hyperbolic Equations. *SIAM Journal on Numerical Analysis*, 52(5):2555–2573, 2014.

[41] A. Carreon, S. Zhang, S. Sharma, J. Singh, and V. Raman. GPU Performance Modeling and Assessment of High-Speed Combustion Simulations Using Adaptive Mesh Refinement. In *AIAA SCITECH 2025 Forum*. American Institute of Aeronautics and Astronautics, 2025.

[42] J. I. Castor. *Radiation Hydrodynamics*. Cambridge University Press, 2004.

[43] C. E. Castro and E. F. Toro. A Riemann solver and upwind methods for a two-phase flow model in non-conservative form. *International Journal for Numerical Methods in Fluids*, 50(3):275–307, 2006.

[44] M. J. Castro Díaz, A. Kurganov, and T. Morales de Luna. Path-conservative central-upwind schemes for nonconservative hyperbolic systems. *ESAIM: Mathematical Modelling and Numerical Analysis*, 53(3):959–985, 2019.

[45] H. Chaves, M. Knapp, A. Kubitzek, F. Obermeier, and T. Schneider. Experimental Study of Cavitation in the Nozzle Hole of Diesel Injectors Using Transparent Nozzles. *SAE Paper*, page 950290, 1995.

[46] J. Cheng and F. Zhang. A Bound-Preserving and Positivity-Preserving Path-Conservative Discontinuous Galerkin Method for Solving Five-Equation Model of Compressible Two-Medium Flows. *SIAM Journal on Scientific Computing*, 44(4):B1195–B1220, 2022.

[47] H. Childs, E. Brugger, B. Whitlock, J. Meredith, S. Ahern, K. Bonnell, M. Miller, G. Weber, C. Harrison, D. Pugmire, T. Fogal, C. Garth, A. Sanderson, E. W. Bethel, M. Durant, D. Camp, J. Favre, O. Rübel, P. Navratil, and F. Vivodtzev. VisIt: An End-User Tool for Visualizing and Analyzing Very Large Data. *Proceed SciDAC*, pages 1–16, 2011.

[48] P.-H. Chiu and Y.-T. Lin. A conservative phase field method for solving incompressible two-phase flows. *Journal of Computational Physics*, 230(1):185–204, 2011.

[49] B. Cockburn. Discontinuous Galerkin methods. *Journal of Applied Mathematics and Mechanics*, 83(11):731–754, 2003.

[50] R. H. Cole. *Underwater Explosions*. Princeton University Press., 1948.

[51] A. J. Coleman, J. E. Saunders, L. A. Crum, and M. Dyson. Acoustic cavitation generated by an extracorporeal shockwave lithotripter. *Ultrasound Med. Biol.*, 13:69–76, 1987.

[52] B. D. Collins and J. W. Jacobs. PLIF flow visualization and measurements of the Richtmyer–Meshkov instability of an air/SF_6 interface. *Journal of Fluid Mechanics*, 464:113–136, 2002.

[53] S. S. Collis. Discontinuous Galerkin methods for turbulence simulation. In *Proceedings of the Summer Program 2002*, page 13. Center for Turbulence Research, 2002.

[54] A. W. Cook. Artificial fluid properties for large-eddy simulation of compressible turbulent mixing. *Physics of Fluids*, 19(5):055103, 2007.

[55] V. Coralic and T. Colonius. Shock-induced collapse of a bubble inside a deformable vessel. *European Journal of Mechanics, B/Fluids*, 40:64–74, 2013.

[56] R. S. Craxton, K. S. Anderson, T. R. Boehly, V. N. Goncharov, D. R. Harding, J. P. Knauer, R. L. McCrory, P. W. McKenty, D. D. Meyerhofer, J. F. Myatt, A. J. Schmitt, J. D. Sethian, R. W. Short, S. Skupsky, W. Theobald, W. L. Kruer, K. Tanaka, R. Betti, T. J. B. Collins, J. A. Delettrez, S. X. Hu, J. A. Marozas, A. V. Maximov, D. T. Michel, P. B. Radha, S. P. Regan, T. C. Sangster, W. Seka, A. A. Solodov, J. M. Soures, C. Stoeckl, and J. D. Zuegel. Direct-drive inertial confinement fusion: A review. *Physics of Plasmas*, 22(11):110501, 2015.

[57] S. C. Crow. Stability theory for a pair of trailing vortices. *AIAA Journal*, 8(12):2172–2179, 1970.

[58] J. Cui, Z.-P. Chen, Q. Wang, T.-R. Zhou, and C. Corbett. Experimental studies of bubble dynamics inside a corner. *Ultrasonics Sonochemistry*, 64:104951, 2020.

[59] G. Dal Maso, P. H. Le Floch, and F. Murat. Definition and weak stability of nonconservative products. *Journal de Mathématiques Pures et Appliquées*, 74:483–548, 1995.

[60] M. De Lorenzo, M. Pelanti, and P. Lafon. HLLC-type and path-conservative schemes for a single-velocity six-equation two-phase flow model: A comparative study. *Applied Mathematics and Computation*, 333:95–117, 2018.

[61] F. De-Xun, M. Yan-Wen, and L. Xin-Liang. Direct Numerical Simulation of Three-Dimensional Richtmyer–Meshkov Instability. *Chinese Physics Letters*, 25(1):188, 2008.

[62] M. J. Del Razo and R. J. LeVeque. Computational study of shock waves propagating through air-plastic-water interfaces. *Bulletin of the Brazilian Mathematical Society, New Series*, 47(2):685–700, 2016.

[63] T. Deplancke, O. Lame, J. Y. Cavaille, M. Fivel, M. Riondet, and J. P. Franc. Outstanding cavitation erosion resistance of Ultra High Molecular Weight Polyethylene (UHMWPE) coatings. *Wear*, 328:301–308, 2015.

[64] Y. T. Didenko, W. B. McNamara, and K. S. Suslick. Molecular emission from single-bubble sonoluminescence. *Nature (London)*, 407:877–879, 2000.

[65] J. Ding, T. Si, J. Yang, X. Lu, Z. Zhai, and X. Luo. Measurement of a Richtmyer-Meshkov Instability at an Air-SF_6 Interface in a Semiannular Shock Tube. *Physical Review Letters*, 119(1), 2017.

[66] M. Dumbser, O. Zanotti, R. Loubère, and S. Diot. A posteriori subcell limiting of the discontinuous Galerkin finite element method for hyperbolic conservation laws. *Journal of Computational Physics*, 278:47–75, 2014.

[67] T. Dzanic. Continuously bounds-preserving discontinuous Galerkin methods for hyperbolic conservation laws. *Journal of Computational Physics*, 508:113010, 2024.

[68] B. Einfeldt, C. Munz, P. Roe, and B. Sjögreen. On Godunov-type methods near low densities. *Journal of Computational Physics*, 92(2):273–295, 1991.

[69] X. Escaler, E. Egusquiza, M. Farhat, F. Avellan, and M. Coussirat. Detection of cavitation in hydraulic turbines. *Mech. Syst. Signal Process.*, 20:983 – 1007, 2006.

[70] J. B. Estrada, H. C. Cramer, M. T. Scimone, S. Buyukozturk, and C. Franck. Neural cell injury pathology due to high-rate mechanical loading. *Brain Multiphysics*, 2:100034, 2021.

[71] S. Fechter and C.-D. Munz. A discontinuous Galerkin-based sharp-interface method to simulate three-dimensional compressible two-phase flow. *International Journal for Numerical Methods in Fluids*, 78(7):413–435, 2015.

[72] E. Fievisohn, Z. Bailey, A. Guettler, and P. VandeVord. Primary Blast Brain Injury Mechanisms: Current Knowledge, Limitations, and Future Directions. *Journal of Biomechanical Engineering*, 140(2):020806, 2018.

[73] D. J. Flannigan, S. D. Hopkins, C. G. Camara, S. J. Putterman, and K. S. Suslick. Measurement of pressure and density inside a single sonoluminescing bubble. *Phys. Rev. Lett.*, 96:204301, 2006.

[74] N. Fleischmann, S. Adami, and N. A. Adams. A shock-stable modification of the HLLC Riemann solver with reduced numerical dissipation. *Journal of Computational Physics*, 423:109762, 2020.

[75] B. Fonkwa Kamga and E. Johnsen. Numerical investigation of the deformation of solid surfaces due to bubble collapse. In *APS Division of Fluid Dynamics Meeting Abstracts*, APS Meeting Abstracts, page J07.004. 2024.

[76] J. P. Franc, M. Riondet, A. Karimi, and G. L. Chahine. Impact Load Measurements in an Erosive Cavitating Flow. *J Fluids Engng.*, 133:121301, 2011.

[77] J. P. Freidberg. Ideal magnetohydrodynamic theory of magnetic fusion systems. *Reviews of Modern Physics*, 54(3):801–902, 1982.

[78] F. G. Friedlander. The diffraction of sound pulses I. Diffraction by a semi-infinite plane. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 186(1006):322–344, 1946.

[79] B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo. FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. *The Astrophysical Journal Supplement Series*, 131(1):273–334, 2000.

[80] F. Gibou, R. Fedkiw, and S. Osher. A review of level-set methods and some recent applications. *Journal of Computational Physics*, 353:82–109, 2018.

[81] S. K. Godunov and I. Bohachevsky. Finite difference method for numerical computation of discontinuous solutions of the equations of fluid dynamics. *Matematicheskii Sbornik*, 47(89):271–306, 1959.

[82] J. Goeller, A. Wardlaw, D. Treichler, J. O'Bruba, and G. Weiss. Investigation of Cavitation as a Possible Damage Mechanism in Blast-Induced Traumatic Brain Injury. *Journal of Neurotrauma*, 29(10):1970–1981, 2012.

[83] D. O. Gómez, S. M. Mahajan, and P. Dmitruk. Hall magnetohydrodynamics in a strong magnetic field. *Physics of Plasmas*, 15(10), 2008.

[84] S. Goodarzi, I. Sagert, J. P. Sauppe, P. A. Keiter, E. N. Loomis, R. F. Sacks, Z. L. Mohamed, S. Palaniyappan, E. C. Merritt, B. M. Haines, B. M. Patterson, D. D. Meyerhofer, D. S. Montgomery, and D. W. Schmidt. Characterizing the effects of drive asymmetries, component offsets, and joint gaps in double shell capsule implosions. *Physics of Plasmas*, 31(8), 2024.

[85] S. Gottlieb and C.-W. Shu. Total variation diminishing Runge-Kutta schemes. *Mathematics of Computation*, 67(221):73–85, 1998.

[86] F. H. Harlow. The particle-in-cell method for numerical solution of problems in fluid dynamics. Technical report, Los Alamos Scientific Lab., N. Mex., 1962.

[87] F. H. Harlow and J. E. Welch. Numerical Calculation of Time-Dependent Viscous Incompressible Flow of Fluid with Free Surface. *The Physics of Fluids*, 8(12):2182–2189, 1965.

[88] N. A. Hawker and Y. Ventikos. Interaction of a strong shockwave with a gas bubble in a liquid medium: a numerical study. *J. Fluid Mech.*, 701:59–97, 2012.

[89] Helmholtz. XLIII. On discontinuous movements of fluids. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 36(244):337–346, 1868.

[90] D. Henneaux, P. Schrooyen, B. Ricardo Barros Dias, A. Turchi, P. Chatelain, and T. Magin. Extended Discontinuous Galerkin Method for Solving Gas-Liquid Compressible Flows with Phase Transition. In *AIAA AVIATION 2020 FORUM*. American Institute of Aeronautics and Astronautics, VIRTUAL EVENT, 2020.

[91] M. T. Henry de Frahan, S. Varadan, and E. Johnsen. A new limiting procedure for discontinuous Galerkin methods applied to compressible multiphase flows with shocks and interfaces. *Journal of Computational Physics*, 280:489–509, 2015.

[92] C. Hirt and B. Nichols. Volume of fluid (VOF) method for the dynamics of free boundaries. *Journal of Computational Physics*, 39(1):201–225, 1981.

[93] Z. Huang and E. Johnsen. A consistent and conservative Phase-Field method for compressible multiphase flows with shocks. *Journal of Computational Physics*, 488:112195, 2023.

[94] Z. Huang and E. Johnsen. A consistent and conservative phase-field method for compressible N-phase flows: Consistent limiter and multiphase reduction-consistent formulation. *Journal of Computational Physics*, 501:112801, 2024.

[95] Z. Huang, G. Lin, and A. Ardekani. A mixed upwind/central WENO scheme for incompressible two-phase flows. *J. Comput. Phys.*, 387:455–480, 2019.

[96] Z. Huang, G. Lin, and A. Ardekani. Consistent and conservative scheme for incompressible two-phase flows using the conservative Allen-Cahn model. *J. Comput. Phys.*, 420:109718, 2020.

[97] Z. Huang, G. Lin, and A. Ardekani. Consistent, essentially conservative and balanced-force Phase-Field method to model incompressible two-phase flows. *J. Comput. Phys.*, 406:109192, 2020.

[98] Z. Huang, G. Lin, and A. Ardekani. A consistent and conservative model and its scheme for N-phase-M-component incompressible flows. *J. Comput. Phys.*, 434:110229, 2021.

[99] Z. Huang, G. Lin, and A. Ardekani. A consistent and conservative volume distribution algorithm and its applications to multiphase flows using Phase-Field models. *Int. J. Multiphase Flow*, 142:103727, 2021.

[100] Z. Huang, G. Lin, and A. Ardekani. A consistent and conservative Phase-Field method for multiphase incompressible flows. *J. Comput. Appl. Math.*, 408:114116, 2022.

[101] Z. Huang, G. Lin, and A. Ardekani. A consistent and conservative Phase-Field model for thermo-gas-liquid-solid flows including liquid-solid phase change. *J. Comput. Phys.*, 449:110795, 2022.

[102] Z. Huang, G. Lin, and A. Ardekani. Implementing contact angle boundary conditions for second-order Phase-Field models of wall-bounded multiphase flows. *J. Comput. Phys.*, 471:111619, 2022.

[103] Z. Huang, S. Mirjalili, M. Khanwale, S. Jain, and E. Johnsen. Investigation of Phase-Field models in compressible multiphase flows in a unified high-order and bound-preserving framework. *Center for Turbulence Research*, Proceedings of the Summer Program:327–337, 2024.

[104] Z. Huang, W. J. White, and E. Johnsen. Consistent and conservative Phase-Field method for compressible two- and N-phase flows with adaptive mesh refinement, 2025. Manuscript under review at Journal of Computational Physics.

[105] A. Huerta, E. Casoni, and J. Peraire. A simple shock-capturing technique for high-order discontinuous Galerkin methods. *International Journal for Numerical Methods in Fluids*, 69(10):1614–1632, 2012.

[106] IT'IS Foundation. Tissue Properties Database V5.0, 2025.

[107] S. S. Jain. Accurate conservative phase-field method for simulation of two-phase flows. *Journal of Computational Physics*, 469:111529, 2022.

[108] S. S. Jain, M. C. Adler, J. R. West, A. Mani, P. Moin, and S. K. Lele. Assessment of diffuse-interface methods for compressible multiphase fluid flows and elastic-plastic deformation in solids. *Journal of Computational Physics*, 475:111866, 2023.

[109] S. S. Jain, A. Mani, and P. Moin. A conservative diffuse-interface method for compressible two-phase flows. *Journal of Computational Physics*, 418:109606, 2020. ArXiv:1911.03619 [physics].

[110] S. S. Jain and P. Moin. A kinetic energy–and entropy-preserving scheme for compressible two-phase flows. *Journal of Computational Physics*, 464:111307, 2022.

[111] G.-S. Jiang and C.-W. Shu. Efficient Implementation of Weighted ENO Schemes. *Journal of Computational Physics*, 126(1):202–228, 1996.

[112] Z.-H. Jiang, X. Deng, F. Xiao, C. Yan, J. Yu, and S. Lou. Hybrid Discontinuous Galerkin/Finite Volume Method with Subcell Resolution for Shocked Flows. *AIAA Journal*, 59(6):2027–2044, 2021.

[113] E. Johnsen and T. Colonius. Implementation of WENO schemes in compressible multicomponent flow problems. *Journal of Computational Physics*, 219(2):715–732, 2006.

[114] E. Johnsen and T. Colonius. Shock-induced collapse of a gas bubble in shockwave lithotripsy. *The Journal of the Acoustical Society of America*, 124(4):2011–2020, 2008.

[115] E. Johnsen and T. Colonius. Numerical simulations of non-spherical bubble collapse. *J. Fluid Mech.*, 629:231–264, 2009.

[116] E. Johnsen and F. Ham. Preventing numerical errors generated by interface-capturing schemes in compressible multi-material flows. *Journal of Computational Physics*, 231(17):5705–5717, 2012.

[117] P. E. Johnson. *A Recovery-Assisted Discontinuous Galerkin Method for Direct Numerical Simulation of Compressible Turbulence*. Ph.D. thesis, University of Michigan, 2018.

[118] S. J. Kamkar. *Mesh Adaption Strategies for Vortex-Dominated Flows*. Ph.D. thesis, Stanford University, 2011.

[119] J. Kane, R. P. Drake, and B. A. Remington. An Evaluation of the Richtmyer-Meshkov Instability in Supernova Remnant Formation. *The Astrophysical Journal*, 511(1):335–340, 1999.

[120] W. Kang, A. Adnan, T. O'Shaughnessy, and A. Bagchi. Cavitation nucleation in gelatin: Experiment and mechanism. *Acta Biomaterialia*, 67:295–306, 2018.

[121] A. K. Kapila, R. Menikoff, J. B. Bdzil, S. F. Son, and D. S. Stewart. Two-phase modeling of deflagration-to-detonation transition in granular materials: Reduced equations. *Physics of Fluids*, 13(10):3002–3024, 2001.

[122] J. B. Keller and M. Miksis. Bubble oscillations of large amplitude. *The Journal of the Acoustical Society of America*, 68(2):628–633, 1980.

[123] R. R. Kerswell. Elliptical Instability. *Annual Review of Fluid Mechanics*, 34(1):83–113, 2002.

[124] J. Kim. Phase-Field Models for Multi-Component Fluid Flows. *Communications in Computational Physics*, 12(3):613–661, 2012.

[125] K. H. Kim, G. Chahine, J. P. Franc, and A. Karimi. Advanced experimental and numerical techniques for cavitation erosion prediction. *Springer*, 2014.

[126] L. Krivodonova. Limiters for high-order discontinuous Galerkin methods. *Journal of Computational Physics*, 226(1):879–896, 2007.

[127] F. Kummer. Extended Discontinuous Galerkin methods for multiphase flows: the spatial discretization. In *Annual Research Briefs 2013*, page 15. Center for Turbulence Research, 2013.

[128] F. Kummer, J. Weber, and M. Smuda. BoSSS: A package for multigrid extended discontinuous Galerkin methods. *Computers & Mathematics with Applications*, 81:237–257, 2021.

[129] C. C. Kuranz, H.-S. Park, C. M. Huntington, A. R. Miles, B. A. Remington, T. Plewa, M. R. Trantham, H. F. Robey, D. Shvarts, A. Shimony, K. Raman, S. MacLaren, W. C. Wan, F. W. Doss, J. Kline, K. A. Flippo, G. Malamud, T. A. Handy, S. Prisbrey, C. M. Krauland, S. R. Klein, E. C. Harding, R. Wallace, M. J. Grosskopf, D. C. Marion, D. Kalantar, E. Giraldez, and R. P. Drake. How high energy fluxes may affect Rayleigh–Taylor instability growth in young supernova remnants. *Nature Communications*, 9(1), 2018.

[130] J. Lang, R. Nathan, D. Zhou, X. Zhang, B. Li, and Q. Wu. Cavitation causes brain injury. *Physics of Fluids*, 33(3), 2021.

[131] M. Latini and O. Schilling. A comparison of two- and three-dimensional single-mode reshocked Richtmyer–Meshkov instability growth. *Physica D: Nonlinear Phenomena*, 401:132201, 2020.

[132] W. Lauterborn and H. Bolle. Experimental investigations of cavitation-bubble collapse in the neighbourhood of a solid boundary. *J. Fluid Mech.*, 72(2):391–399, 1975.

[133] W. Lauterborn and T. Kurz. Physics of bubble oscillations. *Rep. Prog. Phys.*, 73:106501, 2010.

[134] O. Le Métayer and R. Saurel. The Noble-Abel Stiffened-Gas equation of state. *Phys. Fluids*, 28:046102, 2016.

[135] T. Le Reun, B. Favier, A. J. Barker, and M. Le Bars. Inertial Wave Turbulence Driven by Elliptical Instability. *Physical Review Letters*, 119(3), 2017.

[136] S. K. Lele. Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics*, 103(1):16–42, 1992.

[137] R. J. LeVeque. *Numerical methods for conservation laws*. Lectures in mathematics ETH Zürich. Birkhäuser Verlag, Basel ; Boston, 2nd ed edition, 1992.

[138] R. J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2002.

[139] S.-M. Li, A.-M. Zhang, Q. X. Wang, and S. Zhang. The jet characteristics of bubbles near mixed boundaries. *Phys. Fluids*, 31(10):107105, 2019.

[140] X. Li, Y. Fu, C. Yu, and L. Li. Statistical characteristics of turbulent mixing in spherical and cylindrical converging Richtmyer–Meshkov instabilities. *Journal of Fluid Mechanics*, 928, 2021.

[141] Y. Liang and X. Luo. Hydrodynamic instabilities of two successive slow/fast interfaces induced by a weak shock. *Journal of Fluid Mechanics*, 955, 2023.

[142] T. T. Lim and T. B. Nickels. Instability and reconnection in the head-on collision of two vortex rings. *Nature*, 357(6375):225–227, 1992.

[143] O. Lindau and W. Lauterborn. Cinematographic observation of the collapse and rebound of a laser-produced cavitation bubble near a wall. *J. Fluid Mech.*, 479:327–348, 2003.

[144] J. D. Lindl, P. Amendt, R. L. Berger, S. G. Glendinning, S. H. Glenzer, S. W. Haan, R. L. Kauffman, O. L. Landen, and L. J. Suter. The physics basis for ignition using indirect-drive targets on the National Ignition Facility. *Physics of Plasmas*, 11(2):339–491, 2004.

[145] J. Liu, T. Tolle, D. Zuzio, J.-L. Estivalèzes, S. M. Damian, and T. Marić. Inconsistencies in unstructured geometric volume-of-fluid methods for two-phase flows with high density ratios. *Computers & Fluids*, 281:106375, 2024.

[146] W. H. Liu, C. Ping Yu, W. Hua Ye, L. Feng Wang, and X. Tu He. Nonlinear theory of classical cylindrical Richtmyer-Meshkov instability for arbitrary Atwood numbers. *Physics of Plasmas*, 21(6), 2014.

[147] X.-D. Liu, S. Osher, and T. Chan. Weighted Essentially Non-oscillatory Schemes. *Journal of Computational Physics*, 115(1):200–212, 1994.

[148] Y. X. Liu, L. F. Wang, Z. Chen, Z. Y. Li, J. F. Wu, J. Q. Dong, S. Y. Zou, Z. Yan, J. Li, Z. Lei, W. H. Ye, and Y. J. Li. Coupling dynamics of capsule interior defects and its impact on hydrodynamic instabilities at ablation fronts for inertial confinement fusion implosions. *Physics of Plasmas*, 31(3), 2024.

[149] M. Lombardini. *Richtmyer–Meshkov Instability in Converging Geometries*. Ph.D. thesis, California Institute of Technology, 2008.

[150] M. Lombardini, D. I. Pullin, and D. I. Meiron. Turbulent mixing driven by spherical implosions. Part 1. Flow description and mixing-layer growth. *Journal of Fluid Mechanics*, 748:85–112, 2014.

[151] M. Lombardini, D. I. Pullin, and D. I. Meiron. Turbulent mixing driven by spherical implosions. Part 2. Turbulence statistics. *Journal of Fluid Mechanics*, 748:113–142, 2014.

[152] T. Long, J. Cai, and S. Pan. A fully conservative sharp-interface method for compressible multiphase flows with phase change. *Journal of Computational Physics*, 493:112501, 2023.

[153] X. Luo, B. Guan, T. Si, Z. Zhai, and X. Wang. Richtmyer-Meshkov instability of a three-dimensional SF_6-air interface with a minimum-surface feature. *Phys. Rev. E*, 93:013101, 2016.

[154] Y. Lv. Development of a nonconservative discontinuous Galerkin formulation for simulations of unsteady and turbulent flows. *International Journal for Numerical Methods in Fluids*, 92(5):325–346, 2020.

[155] G. Malamud, E. Leinov, O. Sadot, Y. Elbaz, G. Ben-Dor, and D. Shvarts. Reshocked Richtmyer-Meshkov instability: Numerical study and modeling of random multi-mode experiments. *Physics of Fluids*, 26(8), 2014.

[156] V. Maltsev, M. Skote, and P. Tsoutsanis. High-order hybrid DG-FV framework for compressible multi-fluid problems on unstructured meshes. *Journal of Computational Physics*, 502:112819, 2024.

[157] J. Markert, G. Gassner, and S. Walch. A Sub-element Adaptive Shock Capturing Approach for Discontinuous Galerkin Methods. *Communications on Applied Mathematics and Computation*, 5(2):679–721, 2023.

[158] J. Markert, S. Walch, and G. Gassner. A Discontinuous Galerkin Solver in the FLASH Multi-Physics Framework. *Monthly Notices of the Royal Astronomical Society*, 511(3):4179–4200, 2022. ArXiv:2112.11318 [astro-ph].

[159] J. L. Marsh and S. A. Bentil. Cerebrospinal Fluid Cavitation as a Mechanism of Blast-Induced Traumatic Brain Injury: A Review of Current Debates, Methods, and Findings. *Frontiers in Neurology*, 12, 2021.

[160] D. A. McClintock, B. W. Riemer, P. D. Ferguson, A. J. Carroll, and M. J. Dayton. Initial observations of cavitation-induced erosion of liquid metal spallation target vessels at the Spallation Neutron Source. *J. Nucl. Mater.*, 431:147–159, 2012.

[161] R. McKeown, R. Ostilla-Mónico, A. Pumir, M. P. Brenner, and S. M. Rubinstein. Cascade leading to the emergence of small structures in vortex ring collisions. *Physical Review Fluids*, 3(12), 2018.

[162] E. E. Meshkov. Instability of the interface of two gases accelerated by a shock wave. *Fluid Dynamics*, 4(5):101–104, 1972.

[163] D. Mihalas, B. W. Mihalas, A. Fu, and W. D. Arnett. Foundations of Radiation Hydrodynamics. *Physics Today*, 39(6):90–92, 1986.

[164] S. T. Miller, C. F. Cooper, P. Elsbernd, J. Kerwin, R. Mejia-Alvarez, and A. M. Willis. Localizing Clinical Patterns of Blast Traumatic Brain Injury Through Computational Modeling and Simulation. *Frontiers in Neurology*, 12, 2021.

[165] S. Mirjalili, C. B. Ivey, and A. Mani. A conservative diffuse interface method for two-phase flows with provable boundedness properties. *Journal of Computational Physics*, 401:109006, 2020.

[166] M. Mohaghar, J. McFarland, and D. Ranjan. Three-dimensional simulations of reshocked inclined Richtmyer-Meshkov instability: Effects of initial perturbations. *Phys. Rev. Fluids*, 7:093902, 2022.

[167] L. Molefe and I. R. Peters. Jet direction in bubble collapse within rectangular and triangular channels. *Phys. Rev. E*, 100:063105, 2019.

[168] D. S. Montgomery, W. S. Daughton, B. J. Albright, A. N. Simakov, D. C. Wilson, E. S. Dodd, R. C. Kirkpatrick, R. G. Watt, M. A. Gunderson, E. N. Loomis, E. C. Merritt, T. Cardenas, P. Amendt, J. L. Milovich, H. F. Robey, R. E. Tipton, and M. D. Rosen. Design considerations for indirectly driven double shell capsules. *Physics of Plasmas*, 25(9), 2018.

[169] A. Murrone and H. Guillard. A five equation reduced model for compressible two phase flow problems. *J. Comput. Phys.*, 202:664–698, 2005.

[170] M. Natarajan, R. Chiodi, M. Kuhn, and O. Desjardins. An all-Mach multiphase flow solver using block-structured AMR. Institute of Liquid Atomization and Spray Systems, 2019.

[171] C. F. Naudé and A. T. Ellis. On the mechanism of cavitation damage by nonhemispherical cavities collapsing in contact with a solid boundary. *J. Basic Engng.*, 83:648–656, 1961.

[172] M. L. Norman, L. Smarr, M. D. Smith, and J. R. Wilson. Hydrodynamic formation of twin-exhaust jets. *The Astrophysical Journal*, 247:52, 1981.

[173] C. D. Ohl, O. Lindau, and W. Lauterborn. Luminescence from Spherically and Aspherically Collapsing Laser Induced Bubbles. *Phys. Rev. Lett.*, 80:393–396, 1998.

[174] N. A. OKONG'O and J. BELLAN. Direct numerical simulation of a transitional supercritical binary mixing layer: heptane and nitrogen. *Journal of Fluid Mechanics*, 464:1–34, 2002.

[175] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.

[176] M. Owkes and O. Desjardins. A discontinuous Galerkin conservative level set scheme for interface capturing in multiphase flows. *Journal of Computational Physics*, 249:275–302, 2013.

[177] H. Paillère, C. Corre, and J. García Cascales. On the extension of the AUSM+ scheme to compressible two-fluid models. *Computers & Fluids*, 32(6):891–916, 2003.

[178] S. Pan, L. Han, X. Hu, and N. A. Adams. A conservative interface-interaction method for compressible multi-material flows. *Journal of Computational Physics*, 371:870–895, 2018.

[179] A. K. Pandare, J. Waltz, and J. Bakosi. A reconstructed discontinuous Galerkin method for multi-material hydrodynamics with sharp interfaces. *International Journal for Numerical Methods in Fluids*, 92(8):874–898, 2020.

[180] A. K. Pandare, J. Waltz, W. Li, H. Luo, and J. Bakosi. On the design of stable, consistent, and conservative high-order methods for multi-material hydrodynamics. *Journal of Computational Physics*, 490:112313, 2023.

[181] M. B. Panzer, B. S. Myers, B. P. Capehart, and C. R. Bass. Development of a Finite Element Model for Blast Brain Injury and the Effects of CSF Cavitation. *Annals of Biomedical Engineering*, 40(7):1530–1544, 2012.

[182] T. Paula, S. Adami, and N. A. Adams. A robust high-resolution discrete-equations method for compressible multi-phase flow with accurate interface capturing. *Journal of Computational Physics*, 491:112371, 2023.

[183] M. Pelanti. Arbitrary-rate relaxation techniques for the numerical modeling of compressible two-phase flows with heat and mass transfer. *International Journal of Multiphase Flow*, 153:104097, 2022.

[184] M. Pelanti and K.-M. Shyue. A mixture-energy-consistent six-equation two-phase numerical model for fluids with interfaces, cavitation and evaporation waves. *Journal of Computational Physics*, 259:331–357, 2014.

[185] P.-O. Persson and J. Peraire. Sub-Cell Shock Capturing for Discontinuous Galerkin Methods. In *44th AIAA Aerospace Sciences Meeting and Exhibit*. American Institute of Aeronautics and Astronautics, Reno, Nevada, 2006.

[186] P.-O. Persson and B. Stamm. A discontinuous Galerkin method for shock capturing using a mixed high-order and sub-grid low-order approximation space. *Journal of Computational Physics*, 449:110765, 2022.

[187] A. Philipp and W. Lauterborn. Cavitation erosion by single laser-produced bubbles. *J. Fluid Mech.*, 361:75–116, 1998.

[188] M. S. Plesset and R. B. Chapman. Collapse of an initially spherical vapour cavity in the neighbourhood of a solid boundary. *J. Fluid Mech.*, 47:283–290, 1971.

[189] M. S. Plesset and A. Prosperetti. Bubble dynamics and cavitation. *Annual Review of Fluid Mechanics*, 9(1):145–185, 1977.

[190] S. Popinet. A quadtree-adaptive multigrid solver for the Serre–Green–Naghdi equations. *Journal of Computational Physics*, 302:336–358, 2015.

[191] S. V. Poplavski, A. V. Minakov, A. A. Shebeleva, and V. M. Boyko. On the interaction of water droplet with a shock wave: Experiment and numerical simulation. *International Journal of Multiphase Flow*, 127:103273, 2020.

[192] P. Prentice, A. Cuschieri, K. Dholakia, M. Prausnitz, and P. Campbell. Membrane disruption by optically controlled microbubble cavitation. *Nature Phys.*, 1:107–110, 2005.

[193] M. Probyn, R. Williams, B. Thornber, D. Drikakis, and D. Youngs. 2D single-mode Richtmyer–Meshkov instability. *Physica D: Nonlinear Phenomena*, 418:132827, 2021.

[194] A. Przekwas, M. R. Somayaji, and R. K. Gupta. Synaptic Mechanisms of Blast-Induced Brain Injury. *Frontiers in Neurology*, Volume 7 - 2016, 2016.

[195] T. Qin, C.-W. Shu, and Y. Yang. Bound-preserving discontinuous Galerkin methods for relativistic hydrodynamics. *Journal of Computational Physics*, 315:323–347, 2016.

[196] P. A. Quinto-Su, K. Y. Lim, and C. D. Ohl. Cavitation bubble dynamics in microfluidic gaps of variable height. *Phys. Rev. E*, 80:6–9, 2009.

[197] Rayleigh. Investigation of the Character of the Equilibrium of an Incompressible Heavy Fluid of Variable Density. *Proceedings of the London Mathematical Society*, s1-14(1):170–177, 1882.

[198] L. Rayleigh. On the pressure developed in a liquid during the collapse of a spherical cavity. *Philos. Mag.*, 34:94–98, 1917.

[199] W. H. Reed and T. R. Hill. Triangular mesh methods for the neutron transport equation, 1973.

[200] S. Rhebergen, O. Bokhove, and J. van der Vegt. Discontinuous Galerkin finite element methods for hyperbolic nonconservative partial differential equations. *Journal of Computational Physics*, 227(3):1887–1922, 2008.

[201] R. D. Richtmyer. Taylor instability in shock acceleration of compressible fluids. *Communications on Pure and Applied Mathematics*, 13(2):297–319, 1960.

[202] W. J. Rider and D. B. Kothe. Reconstructing Volume Tracking. *Journal of Computational Physics*, 141(2):112–152, 1998.

[203] A. Roccon, L. Enzenberger, D. Zaza, and A. Soldati. MHIT36: A phase-field code for GPU simulations of multiphase homogeneous isotropic turbulence. *Computer Physics Communications*, 316:109804, 2025.

[204] M. Rodriguez and E. Johnsen. A high-order accurate five-equations compressible multiphase approach for viscoelastic fluids and solids with relaxation and elasticity. *J. Comput. Phys.*, 379:70–90, 2019.

[205] P. Roe. A Simple Explanation of Superconvergence for Discontinuous Galerkin Solutions to $u_t + u_x = 0$. *Communications in Computational Physics*, 21(4):905–912, 2017.

[206] A. M. Rueda-Ramírez, S. Hennemann, F. J. Hindenlang, A. R. Winters, and G. J. Gassner. An entropy stable nodal discontinuous Galerkin method for the resistive MHD equations. Part II: Subcell finite volume shock capturing. *Journal of Computational Physics*, 444:110580, 2021.

[207] A. M. Rueda-Ramírez, W. Pazner, and G. J. Gassner. Subcell limiting strategies for discontinuous Galerkin spectral element methods. *Computers & Fluids*, 247:105627, 2022. ArXiv:2202.00576 [physics].

[208] R. Samtaney and N. J. Zabusky. Circulation deposition on shock-accelerated planar and curved density-stratified interfaces: models and scaling laws. *Journal of Fluid Mechanics*, 269:45–78, 1994.

[209] M. A. Sandoval, W. R. Hix, O. E. B. Messer, E. J. Lentz, and J. A. Harris. Three-dimensional Core-collapse Supernova Simulations with 160 Isotopic Species Evolved to Shock Breakout. *The Astrophysical Journal*, 921(2):113, 2021.

[210] J. Sauppe, I. Sagert, T. Day, K. Flippo, J. Kline, L. Kot, S. Palaniyappan, R. Roycroft, and D. Schmidt. Double cylinder implosion experiments at the National Ignition Facility. *High Energy Density Physics*, 53:101156, 2024.

[211] R. Saurel and C. Pantano. Diffuse-Interface Capturing Methods for Compressible Two-Phase Flows. *Annual Review of Fluid Mechanics*, 50(1):105–130, 2018.

[212] R. Saurel, F. Petitpas, and R. A. Berry. Simple and efficient relaxation methods for interfaces separating compressible fluids, cavitating flows and shocks in multiphase mixtures. *Journal of Computational Physics*, 228(5):1678–1712, 2009.

[213] K. Schaal, A. Bauer, P. Chandrashekar, R. Pakmor, C. Klingenberg, and V. Springel. Astrophysical hydrodynamics with a high-order discontinuous Galerkin scheme and adaptive mesh refinement. *Monthly Notices of the Royal Astronomical Society*, 453(4):4279–4301, 2015.

[214] K. Schmidmayer, S. H. Bryngelson, and T. Colonius. An assessment of multicomponent flow models and interface capturing schemes for spherical bubble dynamics. *Journal of Computational Physics*, 402:109080, 2020.

[215] K. Schmidmayer, F. Petitpas, S. Le Martelot, and E. Daniel. ECOGEN: An open-source tool for multiphase, compressible, multiphysics flows. *Computer Physics Communications*, 251:107093, 2020.

[216] D. W. Schmidt, P. M. Donovan, S. L. Edwards, F. Fierro, B. M. Haines, C. E. Hamilton, P. A. Keiter, E. N. Loomis, T. Morrow, S. Palaniyappan, B. M. Patterson, R. B. Randolph, H. F. Robey, J. P. Sauppe, D. J. Stark, and D. R. Vodnik. Los Alamos National Laboratory Double Shell Program Target Development. *Fusion Science and Technology*, 79(7):754–760, 2023.

[217] L. I. Sedov. *Similarity and Dimensional Methods in Mechanics*. CRC Press, 2018.

[218] S. Sembian, M. Liverts, N. Tillmark, and N. Apazidis. Plane shock wave interaction with a cylindrical water column. *Physics of Fluids*, 28(5), 2016.

[219] J. A. Sethian and P. Smereka. Level Set Methods for Fluid Interfaces. *Annual Review of Fluid Mechanics*, 35(1):341–372, 2003.

[220] V. Sharma, A. H. Rauch, and V. Raman. Accelerating CFD Simulations With Super-Resolution Feedback-Informed Adaptive Mesh Refinement. In *AIAA SCITECH 2025 Forum*. American Institute of Aeronautics and Astronautics, 2025.

[221] Z. Shen, J. Ren, and X. Cui. On the carbuncle instability of the HLLE-type solvers. *Journal of Physics: Conference Series*, 1290(1):012026, 2019.

[222] R. K. Shukla, C. Pantano, and J. B. Freund. An interface capturing method for the simulation of multi-phase compressible flows. *Journal of Computational Physics*, 229(19):7411–7439, 2010.

[223] K.-M. Shyue and F. Xiao. An Eulerian interface sharpening algorithm for compressible two-phase flow: The algebraic THINC approach. *Journal of Computational Physics*, 268:326–354, 2014.

[224] L. Smarr, J. R. Wilson, R. T. Barton, and R. L. Bowers. Rayleigh-Taylor overturn in supernova core collapse. *The Astrophysical Journal*, 246:515, 1981.

[225] W. Sollie, O. Bokhove, and J. Van Der Vegt. Space–time discontinuous Galerkin finite element method for two-fluid flows. *Journal of Computational Physics*, 230(3):789–817, 2011.

[226] S. C. Spiegel, H. Huynh, and J. R. DeBonis. A Survey of the Isentropic Euler Vortex Problem using High-Order Methods. In *22nd AIAA Computational Fluid Dynamics Conference*. American Institute of Aeronautics and Astronautics, 2015.

[227] V. Streeter, E. Wylie, and K. Bedford. *Fluid Mechanics*. Civil Engineering Series. WCB/McGraw Hill, 1998.

[228] O. Supponen, D. Obreschkow, P. Kobel, M. Tinguely, N. Dorsaz, and M. Farhat. Shock waves from nonspherical cavitation bubbles. *Phys. Rev. Fluids*, 2:093601, 2017.

[229] O. Supponen, D. Obreschkow, M. Tinguely, P. Kobel, N. Dorsaz, and M. Farhat. Scaling laws for jets of single cavitation bubbles. *J. Fluid Mech.*, 802:263–293, 2016.

[230] M. Sussman, P. Smereka, and S. Osher. A Level Set Approach for Computing Solutions to Incompressible Two-Phase Flow. *Journal of Computational Physics*, 114(1):146–159, 1994.

[231] S. Sutar and S. Ganpule. Investigation of wave propagation through head layers with focus on understanding blast wave transmission. *Biomechanics and Modeling in Mechanobiology*, 19(3):875–892, 2019.

[232] P. K. Sweby. High Resolution Schemes Using Flux Limiters for Hyperbolic Conservation Laws. *SIAM Journal on Numerical Analysis*, 21(5):995–1011, 1984.

[233] Y. Tagawa and I. R. Peters. Bubble collapse and jet formation in corner geometries. *Phys. Rev. Fluids*, 3:081601(R), 2018.

[234] M. S. Taljanovic, L. H. Gimber, G. W. Becker, L. D. Latt, A. S. Klauser, D. M. Melville, L. Gao, and R. S. Witte. Shear-Wave Elastography: Basic Physics and Musculoskeletal Applications. *RadioGraphics*, 37(3):855–870, 2017.

[235] G. Taylor. The Instability of Liquid Surfaces when Accelerated in a Direction Perpendicular to their Planes. I. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 201(1065):192–196, 1950.

[236] B. Thornber and Y. Zhou. Numerical simulations of the two-dimensional multimode Richtmyer-Meshkov instability. *Physics of Plasmas*, 22(3), 2015.

[237] V. Titarev and E. Toro. Finite-volume WENO schemes for three-dimensional conservation laws. *Journal of Computational Physics*, 201(1):238–260, 2004.

[238] A. Tiwari, J. B. Freund, and C. Pantano. A diffuse interface model with immiscibility preservation. *Journal of Computational Physics*, 252:290–309, 2013.

[239] A. Tiwari, C. Pantano, and J. B. Freund. Growth-and-collapse dynamics of small bubble clusters near a wall. *J. Fluid Mech.*, 775:1–23, 2015.

[240] Y. Tomita and A. Shima. Mechanisms of impulsive pressure generation and damage pit formation by bubble collapse. *J. Fluid Mech.*, 169:535–564, 1986.

[241] N. Tonicello and M. Ihme. A high-order diffused-interface approach for two-phase compressible flow simulations using a discontinuous Galerkin framework. *Journal of Computational Physics*, 508:112983, 2024.

[242] E. F. Toro. *Riemann solvers and numerical methods for fluid dynamics: a practical introduction*. Springer, Dordrecht ; New York, 3rd ed edition, 2009. OCLC: ocn401321914.

[243] E. F. Toro, M. Spruce, and W. Speares. Restoration of the contact surface in the HLL-Riemann solver. *Shock Waves*, 4(1):25–34, 1994.

[244] T. Trummler, S. H. Bryngelson, K. Schmidmayer, S. J. Schmidt, T. Colonius, and N. A. Adams. Near-surface dynamics of a gas bubble collapsing above a crevice. *J. Fluid Mech.*, 899:A16, 2020.

[245] G. Tryggvason, B. Bunner, A. Esmaeeli, D. Juric, N. Al-Rawahi, W. Tauber, J. Han, S. Nas, and Y.-J. Jan. A Front-Tracking Method for the Computations of Multiphase Flow. *Journal of Computational Physics*, 169(2):708–759, 2001.

[246] G. Tryggvason, R. Scardovelli, and S. Zaleski. *Direct Numerical Simulations of Gas–Liquid Multiphase Flows*. Cambridge University Press, 2001.

[247] J. S. Turner. *Buoyancy Effects in Fluids*. Cambridge University Press, 1973.

[248] S. O. Unverdi and G. Tryggvason. A Front-Tracking Method for Viscous, Incompressible, Multi-fluid Flows. *Journal of Computational Physics*, 1992.

[249] B. Van Leer. Towards the ultimate conservative difference scheme III. Upstream-centered finite-difference schemes for ideal compressible flow. *Journal of Computational Physics*, 23(3):263–275, 1977.

[250] B. van Leer. Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method. *Journal of Computational Physics*, 32(1):101–136, 1979.

[251] B. van Leer and H. Nishikawa. Towards the ultimate understanding of MUSCL: Pitfalls in achieving third-order accuracy. *Journal of Computational Physics*, 446:110640, 2021.

[252] M. Vandenboomgaerde, P. Rouzier, D. Souffland, L. Biamino, G. Jourdan, L. Houas, and C. Mariani. Nonlinear growth of the converging Richtmyer-Meshkov instability in a conventional shock tube. *Physical Review Fluids*, 3(1), 2018.

[253] G. Vivarelli, N. Qin, and S. Shahpar. A Review of Mesh Adaptation Technology Applied to Computational Fluid Dynamics. *Fluids*, 10(5):129, 2025.

[254] A. Vogel, W. Lauterborn, and R. Timm. Optical and acoustic investigations of the dynamics of laser-produced cavitation bubbles near a solid boundary. *J. Fluid Mech.*, 206:299–338, 1989.

[255] M. Wadas, S. Balakrishna, H. LeFevre, C. Kuranz, A. Towne, and E. Johnsen. On the stability of a pair of vortex rings. *Journal of Fluid Mechanics*, 979, 2024.

[256] M. Wadas, H. LeFevre, Y. Elmore, X. Xie, W. White, C. Kuranz, and E. Johnsen. Feasibility of an experiment on clumping induced by the Crow instability along a shocked cylinder. *Physics of Plasmas*, 31(6):062103, 2024.

[257] M. J. Wadas and E. Johnsen. Interactions of two bubbles along a gaseous interface undergoing the Richtmyer–Meshkov instability in two dimensions. *Physica D: Nonlinear Phenomena*, 409:132489, 2020.

[258] M. J. Wadas, L. H. Khieu, G. S. Cearley, H. J. LeFevre, C. C. Kuranz, and E. Johnsen. Saturation of Vortex Rings Ejected from Shock-Accelerated Interfaces. *Physical Review Letters*, 130(19):194001, 2023.

[259] M. J. Wadas, W. J. White, H. J. LeFevre, C. C. Kuranz, A. Towne, and E. Johnsen. Hydrodynamic Mechanism for Clumping along the Equatorial Rings of SN1987A and Other Stars. *Phys. Rev. Lett.*, 132:111201, 2024.

[260] Q. Wang, M. Mahmud, J. Cui, W. R. Smith, and A. D. Walmsley. Numerical investigation of bubble dynamics at a corner. *Physics of Fluids*, 32(5):053306, 2020.

[261] Z. J. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary, H. Deconinck, R. Hartmann, K. Hillewaert, H. T. 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*, 2012.

[262] A. Wardlaw and J. Goeller. Cavitation as a possible traumatic brain injury (TBI) damage mechanism. In *26th Southern Biomedical Engineering Conference SBEC 2010, April 30-May 2, 2010, College Park, Maryland, USA*, pages 34–37. Springer, 2010.

[263] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Comput. Phys.*, 12(6):620–631, 1998.

[264] W. White, S. A. Beig, and E. Johnsen. Pressure fields produced by single-bubble collapse near a corner. *Physical Review Fluids*, 8(2):023601, 2023.

[265] W. J. White, Z. Huang, and E. Johnsen. A discontinuous Galerkin method for the five-equations multiphase model. In *AIAA SCITECH 2024 Forum*. American Institute of Aeronautics and Astronautics, Orlando, FL, 2024.

[266] W. J. White, Z. Huang, and E. Johnsen. A high-order discontinuous Galerkin method for compressible interfacial flows with consistent and conservative Phase Fields. *Journal of Computational Physics*, 527:113830, 2025.

[267] S. E. Widnall, D. B. Bliss, and C.-Y. Tsai. The instability of short waves on a vortex ring. *Journal of Fluid Mechanics*, 66(1):35–47, 1974.

[268] S. Wiri, T. Massow, J. Reid, J. Whitty, C. Dunbar, W. Graves, A. Gonzales, D. Ortley, J. Longwell, C. E. Needham, A. Ziegle, V. Phan, F. Leonessa, and J. L. Duckworth. Dynamic monitoring of service members to quantify blast exposure levels during combat training using BlackBox Biometrics Blast Gauges: explosive breaching, shoulder-fired weapons, artillery, mortars, and 0.50 caliber guns. *Frontiers in Neurology*, 14, 2023.

[269] S. Wiri, C. Needham, D. Ortley, J. Duckworth, A. Gonzales, T. Walilko, and T. B. Bentley. Development of a Fast-Running Algorithm to Approximate Incident Blast Parameters Using Body-Mounted Sensor Measurements. *Military Medicine*, 187(11–12):e1354–e1362, 2021.

[270] F. D. Witherden and A. Jameson. Future Directions in Computational Fluid Dynamics. In *23rd AIAA Computational Fluid Dynamics Conference*. American Institute of Aeronautics and Astronautics, 2017.

[271] S. E. Woosley, S. Wunsch, and M. Kuhlen. Carbon Ignition in Type Ia Supernovae: An Analytic Model. *The Astrophysical Journal*, 607(2):921–930, 2004.

[272] J. Wu, H. Liu, and Z. Xiao. Refined modelling of the single-mode cylindrical Richtmyer–Meshkov instability. *Journal of Fluid Mechanics*, 908, 2020.

[273] J. Wu and W. L. Nyborg. Ultrasound, cavitation bubbles and their interaction with cells. *Adv. Drug Deliv. Rev*, 60:1103 – 1116, 2008.

[274] F. Xiao, Y. Honma, and T. Kono. A simple algebraic interface capturing scheme using hyperbolic tangent function. *International Journal for Numerical Methods in Fluids*, 48(9):1023–1040, 2005.

[275] Z. Xu, Y. Liu, and C.-W. Shu. Hierarchical reconstruction for discontinuous Galerkin methods on unstructured grids with a WENO-type linear reconstruction and partial neighboring cells. *Journal of Computational Physics*, 228(6):2194–2212, 2009.

[276] Z. Yan, Y. Fu, L. Wang, C. Yu, and X. Li. Effect of chemical reaction on mixing transition and turbulent statistics of cylindrical Richtmyer–Meshkov instability. *Journal of Fluid Mechanics*, 941, 2022.

[277] H. Yee, N. Sandham, and M. Djomehri. Low-Dissipative High-Order Shock-Capturing Methods Using Characteristic-Based Filters. *Journal of Computational Physics*, 150(1):199–238, 1999.

[278] D. L. Youngs. Time-dependent multi-material flow with large fluid distortion. In *Numerical methods for fluid dynamics*, pages 27–51. Academic Press, New York, 1982.

[279] O. Zanotti, F. Fambri, M. Dumbser, and A. Hidalgo. Space–time adaptive ADER discontinuous Galerkin finite element schemes with a posteriori sub-cell finite volume limiting. *Computers & Fluids*, 118:204–224, 2015.

[280] F. Zhang and J. Cheng. Analysis on physical-constraint-preserving high-order discontinuous Galerkin method for solving Kapila's five-equation model. *Journal of Computational Physics*, 492:112417, 2023.

[281] L. Zhang, R. Makwana, and S. Sharma. Brain Response to Primary Blast Wave Using Validated Finite Element Models of Human Head and Advanced Combat Helmet. *Frontiers in Neurology*, 4, 2013.

[282] W. Zhang, A. Almgren, V. Beckner, J. Bell, J. Blaschke, C. Chan, M. Day, B. Friesen, K. Gott, D. Graves, M. Katz, A. Myers, T. Nguyen, A. Nonaka, M. Rosso, S. Williams, and M. Zingale. AMReX: a framework for block-structured adaptive mesh refinement. *Journal of Open Source Software*, 4(37):1370, 2019.

[283] W. Zhang, T. Paula, A. Bußmann, S. Adami, and N. A. Adams. Extension of the hybrid WENO5IS-THINC scheme to compressible multiphase flows with an arbitrary number of components. *Journal of Computational Physics*, 524:113702, 2025.

[284] E. Zwaan, L. G. Séverine, K. Tsuji, and C. D. Ohl. Controlled cavitation in microfluidic systems. *Phys. Rev. Lett.*, 98:22–25, 2007.