

CRANFIELD UNIVERSITY

ANDREW DANIEL MOSEDALE

**MODELLING SHOCK-INDUCED
INSTABILITIES, TRANSITION
AND TURBULENT MIXING USING
HIGH-ORDER METHODS**



SCHOOL OF ENGINEERING

PhD

Cranfield University

School of Engineering

PhD

Year: 2008

Andrew Daniel Mosedale

Modelling Shock-Induced Instabilities,
Transition and Turbulent Mixing using
High-Order Methods.

Supervisor: Dimitris Drikakis

April 21, 2009

© Cranfield University, 2008.
All rights reserved. No part of this publication may be reproduced
without the written permission of the copyright holder.

Abstract

High-order numerical methods have been considered and implemented in order to assess their applicability in a range of complex flows centering on shock-induced turbulent mixing. Specifically, Weighted Essentially Non-Oscillatory (WENO) variable reconstruction schemes of fifth and ninth order accuracy have been investigated within the context of a finite volume Godunov solver. In addition to this there have been further numerical developments to assess the HLLC Riemann solver and various quasi-conservative multi-component models in conjunction with the high-order methods.

Understanding the physics of fundamental flow instabilities and turbulence is increasingly necessary to the development of a vast range of engineering applications with relation to fluid dynamics. It is desirable to develop numerical methods that possess sufficient accuracy to capture the detail of such flows while remaining robust and viable in terms of cost.

The WENO schemes have been tested on a number of cases in comparison with more traditional second-order MUSCL schemes. These include two and three dimensional, single and multi mode Richtmyer-Meshkov instabilities with differing initial perturbations, a cube of homogeneous decaying turbulence and two hypersonic geometry cases were simulated. The results from this research were consistent. The higher-order methods provided measurably greater resolution of small scale fluctuations. By conducting grid convergence studies it was seen that the effect of the higher-order methods was comparable to the effect of increasing the number of grid points. The cost analysis repeatedly showed that despite the additional cost of using a higher-order method they were much better value as they could resolve flow features on a significantly coarser grid.

The high-order methods were not only validated for a range of flow problems but shown to offer great value for their additional cost; they could potentially help advance understanding and development in a wide range of fields much faster than is currently the case.

Acknowledgements

“When one has much to put in to them, a day has a hundred pockets”

Nietzsche

Completing this thesis has not been a trivial task, and there are many people without whom it would not have been accomplished and others who achieved the vital goal of making the experience worthwhile and at times even enjoyable.

First is my supervisor, Professor Dimitris Drikakis, who never wavered in his belief that I could complete any piece of work in a timely fashion. His guidance on the direction of the research was as important as his knowledge on academic matters and his willingness to share it. I thank him for giving me the opportunity to work in this group on such interesting and varied projects - on which I should also thank Robin Williams and David Youngs of AWE for their support. My gratitude extends to the rest of the group, in particular Ben Thornber but also Marco Hahn, Evgeniy Shapiro and all the others with whom I discussed the code and who were as responsible for the development process as I, both in the sharing of ideas and in keeping me sane.

On this last point there are many more people around the university who deserve special mention for their role. Sue Richardson has been a constant and reassuring non-academic point of contact, providing endless distractions along with the other staff of the CSA and the many friends that passed through those doors when the bar was the best place to work. I must also express my deepest gratitude for all the Apocalyptic Cows I have known, both for putting up with me and for making my time here so rewarding and more fun than I could have possibly imagined. I must also single out Max Starr, who has not only put up with me but agreed to proof-read the thesis. It has been much appreciated.

Finally, I wish to thank my family. They have always been immensely supportive and given me the aspirations that have led to all I have achieved. None of this would have been possible without them, and their willingness to answer whenever I called. They helped keep everything in perspective, and in that regard I want to end with a mention for Samuel and Isabel - an unending source of joy and inspiration. Thank you.

Contents

Contents	i
List of figures	iii
List of tables	ix
Glossary	xiii
1 Introduction	1
2 Review, Implementation & Development	9
2.1 Godunov's Method	10
2.2 Riemann Solvers	12
2.2.1 The HLLC Riemann Solver	12
2.2.2 MUSTA	18
2.2.3 Introduction to the Sod problem	18
2.3 Reconstruction	21
2.3.1 Variable selection	22
2.3.2 MUSCL	22
2.3.3 WENO	26
2.3.4 Low Mach Number Correction	31
2.3.5 WAF	31
2.3.6 Further one-dimensional test cases	31
2.4 Multi-component models	38
2.4.1 Total Enthalpy Conservation Model (ThCM)	38
2.4.2 Quasi-conservative Johnsen and Colonius model for WENO (JC)	39

2.4.3	Quasi-conservative 5-equation model of Allaire	39
2.5	Time integration	40
2.5.1	Runge-Kutta	40
2.5.2	Implicit Time Integration	41
2.5.3	ADER	41
2.6	Axisymmetric problems	41
3	Two-dimensional single-mode Richtmyer-Meshkov Instability	43
3.1	Introduction	43
3.2	Early comparison of methods	45
3.3	‘Error’ accumulation	50
3.3.1	LES and the ‘error’	50
3.3.2	Errors in Highly-Resolved Flow	52
3.4	Growth and theoretical models	56
3.5	Conclusions	66
4	Homogeneous Decaying Turbulence	67
4.1	Background	67
4.2	Turbulence and initialisation	67
4.3	Measuring turbulence	68
4.3.1	Isotropy	68
4.3.2	The Decay of Kinetic Energy	70
4.3.3	Energy Spectra	71
4.3.4	The Structure of Turbulence	72
4.4	Conclusion	74
5	Three-dimensional Multi-mode Richtmyer-Meshkov Instability	89
5.1	Introduction	89
5.2	Planar RMI	89
5.2.1	Group work	90
5.2.2	The experiments of Vetter and Sturtevant	92
5.3	Large-scale Perturbation	94
5.3.1	Double-bump	95

5.3.2	Inverse Chevron	97
5.4	Conclusions	104
6	Hypersonic Problems	109
6.1	Introduction	109
6.2	Methodology	109
6.2.1	Test Case: Double Cone	110
6.2.2	Test Case: HB-2	110
6.2.3	Test Details	111
6.3	Results	112
6.3.1	Double Cone	112
6.3.2	HB-2	114
6.3.3	Efficiency and robustness	118
6.4	Conclusions	120
7	Concluding Remarks	135
7.1	Future Work	137
	References	138

List of Figures

1.1	Kelvin-Helmholtz clouds	2
1.2	Crab Nebula	3
1.3	Leonardo da Vinci's turbulent pool	4
1.4	The scales of turbulence	4
1.5	Waterfalls: Mean and instantaneous flow	5
2.1	Basic finite volume discretisation	11
2.2	Riemann problem wave structure	13
2.3	HLLC derivation	13
2.4	Comparison of Riemann solvers for the Sod problem	20
2.5	Illustration of effect of choice of variables in the reconstruction	23
2.6	WENO reconstruction stencils	28
2.7	Comparison of lower-order reconstruction methods for the Sod problem	32
2.8	Comparison of higher-order reconstruction methods for the Sod problem	33
2.9	Initial condition for advection of Gaussian pulse	34
2.10	Comparison of reconstruction methods for advection of Gaussian pulse with 16 grid points	34
2.11	Comparison of reconstruction methods for advection of Gaussian pulse with 32 grid points	35
2.12	Comparison of reconstruction methods for advection of Gaussian pulse with 64 grid points	35
2.13	Comparison of reconstruction methods for the entropy wave test cases	37

3.1	Schematic representation of the initial condition for the single-mode Richtmyer-Meshkov Instability	44
3.2	Experimental picture of single-mode Richtmyer-Meshkov Instability	45
3.3	Development of single-mode RMI over time based on a sharp initial interface	46
3.4	Comparison of reconstruction methods and grid resolution for single-mode RMI with sharp initial interface	48
3.5	Comparison of reconstruction methods and grid resolution for single-mode RMI with diffuse initial interface	49
3.6	Highly resolved single-mode RMI	51
3.7	9th-order WENO 16x16 simulation of Kelvin Helmholtz at Mach 0.2 and 0.02 with and without low Mach number (LM) scaling correction	53
3.8	Asymmetry in single-mode RMI initial condition	54
3.9	Diffusion in single-mode RMI initial condition	55
3.10	Sub-grid resolution in single-mode RMI initial condition	55
3.11	Original single-mode RMI growth comparison	57
3.12	Development of zero Atwood No. single-mode RMI over time	60
3.13	Comparison of simulated growth rates with models at zero Atwood No.	61
3.14	Development of high Atwood No. single-mode RMI over time	62
3.15	Comparison of simulated growth rates with experiment at high Atwood No.	63
3.16	Development of low Atwood No. single-mode RMI over time	64
3.17	Late time development of low Atwood number single-mode RMI	65
3.18	Comparison of simulated growth rates with models and experiment at low Atwood No.	65
4.1	Isosurfaces of vorticity in the 256 cube at early times	69
4.2	Comparison of decay of kinetic energy for 32 and 64 cubes with varying reconstruction method	75
4.3	Comparison of decay of kinetic energy for 128 and 256 cubes with varying reconstruction method	76
4.4	Comparison of decay of kinetic energy between low and high order methods	77

4.5	Comparison of development of enstrophy for 32 and 64 cubes with varying reconstruction method	78
4.6	Comparison of development of enstrophy for 128 and 256 cubes with varying reconstruction method	79
4.7	Comparison of development of enstrophy for ninth order WENO on all grids	80
4.8	Comparison of kinetic energy spectra for 32 and 64 cubes with varying reconstruction method	81
4.9	Comparison of kinetic energy spectra for 128 and 256 cubes with varying reconstruction method	82
4.10	Comparison of kinetic energy spectra for fifth order methods on 256 cube	83
4.11	Cross-section of cube showing vorticity for 32 grid with different reconstruction methods	84
4.12	Cross-section of cube showing vorticity for 64 grid with different reconstruction methods	85
4.13	Cross-section of cube showing vorticity for 128 grid with different reconstruction methods	86
4.14	Cross-section of cube showing vorticity for 256 grid with different reconstruction methods	87
5.1	Visualisation of multi-mode planar RMI	90
5.2	Development of the energy spectrum over time for the multi-mode planar RMI	91
5.3	Comparison of energy spectrum for original and improved WENO 9th-order schemes	91
5.4	Visualisation of Vetter and Sturtevant problem	92
5.5	Growth of the mixing layer in Richtmyer-Meshkov problem of Vetter and Sturtevant for large initial perturbation (Amplitude (m), Time (s))	93
5.6	Development of energy spectrum in Richtmyer-Meshkov problem of Vetter and Sturtevant	94
5.7	Initial condition for the ‘double-bump’ problem	95
5.8	Averaged 3D simulation 160x80x40 for the ‘double-bump’ problem	95
5.9	Approximate comparison of double-bump flow at different grid resolutions with experiment	96

5.10	Three-dimensional representation of development of flow density in ‘double-bump’ problem	97
5.11	Initial condition for the ‘inverse-chevron’ problem	97
5.12	Comparison of development of inverse-chevron case over time with experiment	98
5.13	Inverse-chevron density comparison with AWE code	99
5.14	Inverse-chevron volume fraction comparison with AWE code	100
5.15	Distribution of turbulent kinetic energy for inverse-chevron problem	100
5.16	Inverse-chevron turbulent kinetic energy comparison with AWE code	101
5.17	Development of feature positions over time for different codes	102
5.18	Development of mixing over time for different codes	103
5.19	Effect of grid resolution and reconstruction method on mixing	104
5.20	Development of turbulent kinetic energy over time for different codes	105
5.21	Development of enstrophy over time for different codes	106
5.22	Effect of grid resolution and reconstruction method on enstrophy	107
6.1	Schematic of double-cone geometry	110
6.2	The computational mesh for the double-cone, showing every 1 in 4 grid lines	111
6.3	Schematic of HB-2 geometry	112
6.4	The computational mesh for the HB-2, showing every 1 in 4 grid lines	112
6.5	Double-cone visualisation	113
6.6	Double-cone pressure and heat transfer for different methods	115
6.7	Comparison of separation zone size	116
6.8	Double-cone pressure and heat transfer showing grid convergence	117
6.9	Pressure distribution for HB-2 at Mach 5	118
6.10	Pressure distribution for HB-2 at Mach 5 at cylinder-flare transition	118
6.12	Pressure distribution for inclined HB-2 (Case 3)	119
6.11	Longitudinal velocity for inclined HB-2 (Case 3)	119

6.13	Comparison of pressure coefficient and Stanton number for run 2893 for grids 128x48 and 256x96 at first order	122
6.14	Comparison of pressure coefficient and Stanton number for run 2893 for grids 128x48 and 256x96 at second order	123
6.15	Comparison of pressure coefficient and Stanton number for run 2893 for grids 128x48 and 256x96 at ninth order	124
6.16	Comparison of pressure coefficient and Stanton number for run 2894 for grids 128x48 and 256x96 at first order	125
6.17	Comparison of pressure coefficient and Stanton number for run 2894 for grids 128x48 and 256x96 at second order	126
6.18	Comparison of pressure coefficient and Stanton number for run 2894 for grids 128x48 and 256x96 at ninth order	127
6.19	Pressure distribution for HB-2 at Mach 5 at first order for CNS3D and FLUENT at both grid resolutions.	128
6.20	Pressure distribution for HB-2 at Mach 5 at second order for CNS3D and FLUENT at both grid resolutions.	129
6.21	Pressure distribution for HB-2 at Mach 5 using WENO ninth-order in CNS3D.	130
6.22	Pressure distribution for HB-2 at Mach 7.5 at first order for CNS3D and FLUENT at both grid resolutions.	131
6.23	Pressure distribution for HB-2 at Mach 7.5 at second order for CNS3D and FLUENT at both grid resolutions.	132
6.24	Pressure distribution for HB-2 at Mach 7.5 using WENO ninth-order in CNS3D.	133

List of Tables

2.1	Order of accuracy of solver based on peak error convergence . . .	36
3.1	CPU runtime for schemes VL, W5 and W9 at grid resolution 20, 40 and 80 for a given length of simulated time	50
4.1	Estimate of exponent p in decay of kinetic energy for all simu- lations	71
4.2	Estimate of third order structure function for all simulations . . .	73
4.3	Estimate of fourth order structure function for all simulations . .	74
6.1	A matrix showing the twelve test runs simulated for the double- cone geometry.	113
6.2	A matrix showing the run conditions simulated for the HB-2 geometry.	113
6.3	Zero-lift Forebody Axial Force Coefficient	121
6.4	Normal Force Coefficient and pitching moment at 15°	121
6.5	Computational cost normalised by CNS3D first-order.	121

Glossary

ADER	An advanced high-order method
AMR	Adaptive-Mesh-Refinement (refers to results by Bates)
AWE	Atomic Weapons Establishment
CBM	Characteristics-Based Method
CFD	Computational Fluid Dynamics
CFL	Courant-Friedrichs-Lewy
CIR	Courant-Isaacson-Rees
CNS3D	In-house CFD code
CPU	Central Processing Unit
DNS	Direct Numerical Simulation
ENO	Essentially Non-Oscillatory
FLUENT	Commercial CFD code
HLL	Harten, Lax and van Leer (Riemann solver)
HLLC	Harten, Lax and van Leer, Contact (Improved Riemann solver of Toro)
HPC	High-Performance Computing
ILES	Implicit Large-Eddy Simulation
JC	The multi-component model of Johnsen and Colonius
LES	Large-Eddy Simulation
LM	Low Mach no. correction
M5	5th order MUSCL scheme
MM	MinMod MUSCL limiter
MUSCL	Monotone Upwind-centered Scheme for Conservation Laws
MUSTA	MUlti-STAge
PVRS	Primitive Variable Riemann Solver
RANS	Reynolds Averaged Navier Stokes
RMI	Richtmyer-Meshkov Instability
SF6	Sulphur hexafluoride, dense gas used in experiments
ThCM	Total enthalpy Conservation Model
TRRS	Two-Rarefraction Riemann Solver
TSRS	Two-Shock Riemann Solver
TURMOIL	AWE CFD code (refers to external results)
TVD	Total Variation Diminishing
VL	Van Leer MUSCL limiter
W5	5th order WENO scheme
W9	9th order WENO scheme
WAF	Weighted Average Flux
WENO	Weighted Essentially Non-Oscillatory

Mathematical Symbols

γ	Ratio of specific heats
γ, Γ	Vortex strength
ϵ	Vortex displacement
ϵ	Arbitrarily small number
λ	Wave speed
μ	Dynamic viscosity
ρ	Density
ω, Ω	Vorticity
ω	Weightings
a	Speed of sound
a, A	Amplitude
B	Bubble amplitude in vortex model
C_p	Pressure Coefficient
c_v	Specific heat capacity
e	Internal energy
\mathbf{F}	Flux vector
i, j, k	Indices
k	Wavenumber
p, P	Pressure
q	Total Velocity
r	Radius (in axisymmetric co-ordinates)
Re	Reynold's number
S	Wave Speed
St	Stanton number
t, T	Time
\mathbf{U}	vector of conserved variables
u, v, w	Cartesian components of velocity
V_0	Initial growth rate
x, y, z	Cartesian components of displacement
z	Colour function / volume fraction

*In places where this beauty
has already disappeared, we
will reconstruct it.*

Fritz Todt

1

Introduction

It is commonly conceived that turbulence may be the last and possibly the greatest challenge of classical physics. Fluid dynamics now spans such a range of applications that understanding fluids better could potentially advance the design, manufacture or operation of virtually every classically derived piece of technology. From flow/structure interaction in building construction to the obvious aerodynamic applications to all vehicles; from weather systems, atmospheric, oceanic and galactic motion to internal pipe flows of water, oil, gas and even blood. From combustion at every level - from vehicle engines to the Sun - to the manufacture of metals, plastics, composites, paper and glass; fluid dynamics impacts on everyone's life at every turn. Fluid processes and, in particular, turbulent instabilities (Figs. 1.1, 1.2) have always been a part of nature and that nature has long been recognised (Fig. 1.3) on an intuitive level, but the Navier-Stokes equations governing fluid flow are in turn both straightforward and highly complex. While a number of analytical solutions exist for simple and well-constrained cases, in more general, turbulence-driven flows the equations are essentially insoluble and may remain this way. In order to get the most out of them at the present time requires numerical simulation, and as computing power has grown in the past few decades, so has Computational Fluid Dynamics (CFD).

In the absence of suitably quantified experimental data on certain fundamental physical flows, CFD is finding an increasing role to play in driving theory and understanding. Turbulent theory in particular has not advanced significantly in over half a century, despite turbulence being prevalent in virtually every application of fluid dynamics, due to the difficulty in both controlling



(a) [12]

(b) [16]

Figure 1.1: Nature's approach to flow visualisation

and measuring the phenomenon experimentally. While this limits our ability to draw conclusions from the experiments themselves, it also restricts the extent to which CFD codes can be validated. Indeed a CFD simulation can be specified so precisely that, given the levels of accuracy now being obtained, it is increasingly viewed as the benchmark against which the experiments ought to be assessed. More than ever it is vital to understand the factors involved in these simulations and their strengths and limitations.

There are essentially three approaches to solving fluid flow problems numerically. The most obvious is Direct Numerical Simulation (DNS). This entails a numerical discretisation of the governing equations onto a grid which is smaller than the smallest scales present in the flow. These are the Kolmogorov microscales of turbulence and are typically several orders of magnitude smaller than the integral length scale (the typical size of the largest turbulent eddies) and many more orders of magnitude smaller than the large-scale flow features (Fig. 1.4). In order to represent all these scales in a region of any practical engineering size would require a phenomenal number of grid points and with current computing power it would take centuries to calculate the flow round an aircraft for example. DNS is useful in researching pure turbulent behaviour in a small region but currently has no practical application in design.

To reduce the computational cost of the calculations there is another method which simplifies the equations by performing statistical averaging on them. These Reynolds Averaged Navier Stokes (RANS) equations provide a useful way of understanding various flows, in particular those that are steady in time for which the mean flow from the ensemble average is equivalent to the time average which can be more efficiently solved than unsteady RANS. However the averaging does cause a loss of information and accuracy, and even still the equations are not closed. Any RANS approach requires a model for the turbulent fluctuation correlations. This method is most prevalent in industry due in part to its simplicity and robustness. With many tuneable parameters




Figure 1.2: The universe's approach to flow visualisation [14]



Figure 1.3: Leonardo da Vinci's approach to flow visualisation [15]



Figure 1.4: The scales of turbulence [61]



Figure 1.5: The photographer's approach to fluid modelling

for different problems it is hard to validate the results, but there has been so much research into RANS in the past 50 years that a lot of confidence has been gained. Turbulence is inherently chaotic, it cannot be predicted, and all the models have their strengths and flaws; by accepting this it has been possible to manufacture the 'correct' answer in many cases. For accurate representation of unsteady turbulent flows however, RANS is essentially incapable of coping with instabilities. In certain cases of steady flow the Reynolds averaging is analogous to simple time averaging as depicted in figure 1.5. In order to create pictures that mimic an 'artist's impression', rivers and waterfalls are often photographed with a long exposure so as to give an indication of the smooth mean flow rather than capturing the exact detail of each variation for a single moment in time.

The third approach is something of a combination of the previous two. Large Eddy Simulation (LES) seeks to accurately compute the large-scale motion and so is based on the original governing equations like DNS. However, to reduce the computational load and make the simulations of practical use, LES does not resolve the smaller turbulent eddies - the grid size is several orders of magnitude greater than the Kolmogorov scale. Typically, it is hoped to simulate enough scales to capture 90% of the energy, as most of the energy is contained within the large eddies although it is passed down to smaller and

smaller eddies until it is finally dissipated by viscosity. It is believed that these dissipative eddies have no influence on the large-scale motion, and similarly that viscosity has no effect on the upper end of this turbulent energy cascade. These are some of the key principles in determining the models to be used to calculate the dissipation at the sub-grid scales. Much as with the RANS, there are various models with differing characteristics which can be used to represent the complex small scales of turbulence. In essence, LES simply requires that all energy beyond a certain wavelength be dissipated appropriately and it can still accurately represent the flow field.

Beyond the problem of devising a sub-grid model that accurately deals with the dissipation, traditional LES has to theoretically achieve zero numerical dissipation in the scheme. This is significant as the truncation error in the straightforward second-order accurate discretisation of the equations is of the same order of magnitude as the dissipation required by the sub-grid scale. An alternative approach to LES is therefore to design a scheme with the correct amount of inherent dissipation so as to not require a separate sub-grid model - nor a base scheme that is perfectly without error. This is known as Implicit LES (ILES), and is the base area of research for this project. The objective is to develop high resolution schemes, those which by definition give a better than first-order resolution while restricting spurious oscillations, specifically addressing the order of accuracy in the reconstruction of the flow from averaged data.

The aims for this research were as follows:

- Research and implement very-high-order methods, such as Weighted Essentially Non-Oscillatory (WENO) schemes into the implicit LES framework;
- Perform comparative studies between second- and higher-order methods for shock-induced instabilities and turbulent mixing;
- Investigate the range of improvement techniques and modifications established in the literature in conjunction with very high-order methods;
- Perform ILES computations on very fine grids using the High-Performance Computing (HPC) facilities at Cranfield to investigate the methods behaviour in comparison with AWE experimental data and shed light on the physics of turbulent mixing.
- Conduct theoretical analysis to develop models for the growth of the instability in the light of the computational results

The prime area of interest is shock-induced instabilities in multi-component gases. This introduces two other elements to our code - a finite volume formulation for effective shock-capturing and multi-component models to handle

different gases within the flow. These and other basic elements of the code will be expanded on in Chapter 2 along with the development of methods to improve these results to much higher-order accuracy to deal with complex smooth flow features - without reducing the efficacy of the shock-capturing scheme. It will be seen how this work builds upon current best practise in a range of areas to achieve unprecedented accuracy and computational efficiency in the specific flow situations we are attempting to resolve, ultimately moving further towards simulations that can confidently be used to replace live experiments. The subsequent chapters cover a range of such flows, each extending the complexity and level of validation of our methods, culminating in the fully three-dimensional simulation of multiple mode instabilities. Chapter 6 also demonstrates the flexibility of the code when it was turned to hypersonic applications.

2

Review, Implementation & Development

In this chapter the equations and numerical methods researched and implemented will be presented, along with discussion of additional literature and, where appropriate, of developments made. The task of the fluid dynamics modeller is to make viable assumptions in order to simplify a given problem and render it soluble, while not invalidating the results. The first step in the process is to determine which equations need to be solved. There are many established systems of equations for treating certain classes of problems, the majority of the work done here is based on solution of the Euler equations. These are effectively the Navier-Stokes equations with the assumption of zero viscosity and comprise equations for the conservation of mass, momentum and energy.

In reality there is always finite viscosity even in high Reynolds number flows and this is mirrored numerically by the numerical errors which act to dissipate energy at a very low level. This dissipation is in fact necessary to the modelling of turbulence. The concept of using inviscid equations means only that viscosity is thought not to affect the larger-scale motions and need only be considered at the smallest scales, where it is sufficient for it to exist without further mention of magnitude or form. The extent to which this can be assumed in numerical simulation shall be considered in Chapter 4. The following are the full three-dimensional governing equations in Cartesian coordinates in the absence of any source terms:

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{E}}{\partial x} + \frac{\partial \mathbf{F}}{\partial y} + \frac{\partial \mathbf{G}}{\partial z} = 0,$$

where

$$\begin{aligned} \mathbf{U} &= [\rho, \rho u, \rho v, \rho w, e]^T, \\ \mathbf{E} &= [\rho u, \rho u^2 + p, \rho uv, \rho uw, (e + p)u]^T, \\ \mathbf{F} &= [\rho v, \rho uv, \rho v^2 + p, \rho vw, (e + p)v]^T, \\ \mathbf{G} &= [\rho w, \rho uw, \rho vw, \rho w^2 + p, (e + p)w]^T, \end{aligned}$$

$$e = \rho i + 0.5 \rho q^2$$

and ρ, i, u, v, w are the density, internal energy and cartesian velocity components respectively. In general terms the variable \mathbf{U} changes over time as determined by the spatial distribution of the related flux \mathbf{F} . The first step to solving these equations computationally is discretisation. As mentioned previously the need to capture shockwaves is greatly facilitated through use of a finite-volume formulation to ensure conservation. The governing equations are written in integral form for an arbitrary control volume, which in simple one-dimensional form gives the continuous equation,

$$\int \mathbf{U} dt + \int \mathbf{F} dx = 0,$$

which in discrete form can be expressed as

$$\frac{\mathbf{U}_i^{n+1} - \mathbf{U}_i^n}{\Delta t} + \frac{\mathbf{F}_{i+\frac{1}{2}}^n - \mathbf{F}_{i-\frac{1}{2}}^n}{\Delta x} = 0.$$

If we consider a computational cell, this is merely saying that the change in the variables over time within the cell are given by the sum of the fluxes through the cell boundaries (Fig. 2.1). This is how we guarantee conservation of the variables. The cell averages of the variables are stored for each cell as the cell-centre value. If the fluxes at each cell interface are known then the solution is exact, however this is not the case. As can be seen, the only data gained at the next time level is the cell-averaged variables. We need some method of calculating the appropriate fluxes.

2.1 Godunov's Method

Godunov achieved the first extension of the basic CIR flux scheme [13] to systems of equations. Godunov's method is at first glance remarkably straightforward. The flow variables are updated for each computational cell by the net flux into the cell by the (one-dimensional) equation




Figure 2.1: Basic finite volume discretisation

$$\mathbf{U}_i^{n+1} = \mathbf{U}_i^n + \frac{\Delta t}{\Delta x} (\mathbf{F}_{i-\frac{1}{2}} - \mathbf{F}_{i+\frac{1}{2}})$$

but there are two more critical steps to allow this to become a physically stable scheme. Godunov used the assumption that the flow variables were piece-wise constant in each cell and then dictated that the value used for the flux calculation at each interface be the solution of the resulting Riemann problem between cells

$$\begin{aligned} \mathbf{U}_t + \mathbf{F}(\mathbf{U})_x &= \mathbf{0}, \\ \mathbf{U}(x, 0) &= \begin{cases} \mathbf{U}_L & \text{if } x < 0, \\ \mathbf{U}_R & \text{if } x > 0, \end{cases} \end{aligned}$$

evaluated at $x/t = 0$. This then captures the basic physics of any flow, allowing for a stable and conservative solution of shocked flows. There are two major drawbacks to Godunov's method, one is the exact solution of the Riemann problems which has been largely overcome and will be discussed next. The other is the loss of information caused by the assumption of piece-wise constant data, and is the main subject of this research.

2.2 Riemann Solvers

The characteristics-based (CBM) Riemann Solver of Eberle [23] has been used extensively in the code. It makes use of the Riemann invariance across characteristic waves to infer the variable values at the interface at a given time from the values at a determined point of space at known time. It is possible to derive equations to calculate the flux based on these variables from the information contained in the related Riemann problem. This is not straight-forward and requires pressure derivatives to be calculated from the equation of state, making extensions of the system of equations being solved quite involved. Further details on the implementation of this method can be found in [20, 78, 3] which are not reproduced here however this represents the baseline from which the code was developed.

Another well-known approximate Riemann solver is the Roe scheme [70, 85]. It has been used extensively and had a number of modifications proposed, many of which aim to provide an ‘entropy fix’ as the original scheme violated basic physics. This is indicative of the problems surrounding Riemann solvers in general. They each have their strengths and weaknesses and the appropriate choice may ultimately depend on the flow regime to be studied or other special considerations, but they can usually be made to work with little difference between the ensuing results. However the HLLC scheme of Toro [85] provides simplicity and robustness across all problems of interest to this work and has not yet been shown to require particular treatment [5].

2.2.1 The HLLC Riemann Solver

The HLLC Riemann solver developed by Toro is an improvement upon the basic HLL solver first proposed by Harten, Lax and van Leer [35]. The concept is to approximate the wave structure resulting from a discontinuity, all contained within an arbitrary control volume, and evaluate the changes in the fluxes across each wave. The original HLL solver reduced the Riemann problem to a shock and expansion wave pair, with a single average state in between. While simple, this method is highly dissipative. Toro addressed this by restoring a contact discontinuity to the theoretical wave structure, as shown in figure 2.2.

The HLLC solver has two intermediate states, and three waves corresponding to all eigenvalues. This makes it significantly more accurate than the HLL solver. The extra level of complexity is negligible as the same principles are used in deriving the equations.

A control volume is placed around the waves structure to an arbitrary time T (Fig. 2.3). The flux through the interface of these two cells, that is at $x/t = 0$,




Figure 2.2: Riemann problem wave structure




Figure 2.3: Control volume derivation of HLLC scheme

is required over this period of time. As seen in the diagram this is constant. The left and right states ($\mathbf{U}_L, \mathbf{U}_R$) are known and the control volume is defined such that these are also constant.

Integrating over an arbitrary area gives

$$\int_{x_L}^{x_R} \mathbf{U}(x, T) dx = \int_{x_L}^{x_R} \mathbf{U}(x, 0) dx + \int_0^T \mathbf{F}(\mathbf{U}(x_L, t)) dt + \int_0^T \mathbf{F}(\mathbf{U}(x_R, t)) dt,$$

$$\int_{x_L}^{x_R} \mathbf{U}(x, T) dx = x_R \mathbf{U}_R - x_L \mathbf{U}_L + T(\mathbf{F}_L - \mathbf{F}_R),$$

based on

$$\mathbf{U}_t + \mathbf{F}(\mathbf{U})_x = \mathbf{0}.$$

Reducing this to the control volume,

$$\int_{x_L}^{x_R} \mathbf{U}(x, T) dx = \int_{x_L}^{TS_L} \mathbf{U}(x, T) dx + \int_{TS_L}^{TS_R} \mathbf{U}(x, T) dx + \int_{TS_R}^{x_R} \mathbf{U}(x, T) dx,$$

$$\int_{x_L}^{x_R} \mathbf{U}(x, T) dx = \int_{TS_L}^{TS_R} \mathbf{U}(x, T) dx + (TS_L - x_L) \mathbf{U}_L + (x_R - TS_R) \mathbf{U}_R,$$

$$\int_{TS_L}^{TS_R} \mathbf{U}(x, T) dx = T(S_R \mathbf{U}_R - S_L \mathbf{U}_L + \mathbf{F}_L - \mathbf{F}_R).$$

At this point no assumptions have been made. In order to obtain an expression for the variables at $x/t = 0$ the HLL solver assumes a constant profile and takes an average of the integral to give

$$\frac{1}{T(S_R - S_L)} \int_{TS_L}^{TS_R} \mathbf{U}(x, T) dx = \mathbf{U}^{hll} = \frac{S_R \mathbf{U}_R - S_L \mathbf{U}_L + \mathbf{F}_L - \mathbf{F}_R}{S_R - S_L},$$

which, given the wave speeds, results in

$$\mathbf{F}^{hll} = \mathbf{F}_L + S_L(\mathbf{U}^{hll} - \mathbf{U}_L),$$

$$\mathbf{F}^{hll} = \begin{cases} \mathbf{F}_L & \text{if } 0 \leq S_L \\ \frac{S_R \mathbf{F}_L - S_L \mathbf{F}_R + S_L S_R (\mathbf{U}_R - \mathbf{U}_L)}{S_R - S_L} & \text{if } S_L \leq 0 \leq S_R \\ \mathbf{F}_R & \text{if } 0 \geq S_R. \end{cases}$$

The HLLC approach allows for two average states, with the sign of the contact wave determining in which region the interface lies. Similar integration can be carried out to evaluate these states, but there is an even simpler and more elegant approach. The Rankine-Hugoniot conditions, which define the change in flux across a wave as $\Delta \mathbf{F} = S \Delta \mathbf{U}$ as an expression of conservation, can be used directly to give the fluxes

$$\mathbf{F}^{hllc} = \begin{cases} \mathbf{F}_L & \text{if } 0 \leq S_L \\ \mathbf{F}_L + S_L(\mathbf{U}_{*L} - \mathbf{U}_L) & \text{if } S_L \leq 0 \leq S_* \\ \mathbf{F}_R + S_R(\mathbf{U}_{*R} - \mathbf{U}_R) & \text{if } S_* \leq 0 \leq S_R \\ \mathbf{F}_R & \text{if } 0 \geq S_R. \end{cases}$$

By assuming constant normal velocity and pressure across the contact discontinuity, while allowing the tangential velocities to be constant across the left and right waves, and by assuming the contact surface moves with speed inside the star regions we can give expressions for the *_L and *_R variables

$$\mathbf{U}_{*K} = \rho_K \begin{pmatrix} \frac{S_K - u_K}{S_K - S_*} \\ 1 \\ S_* \\ v_K \\ w_K \\ \frac{E_K}{\rho_K} + (S_* - u_K) \left[S_* + \frac{p_K}{\rho_K(S_K - u_K)} \right] \end{pmatrix}.$$

Results published for the HLLC Riemann solver indicate it to be both effective and robust and worth implementing. In terms of the code structure, the Riemann solver subroutine receives the reconstructed conserved variables for the left and right of the interface in question. These are called *RRL*, *UUL*, *VVL*, *WWL* & *EEL* in the code (2.1) for density, components of momentum and energy respectively - similarly for the right-hand side (*RRR*, *UUR*, *VVR*, *WWR* & *EER*). A number of other variables are set at this point, primarily in order to efficiently calculate the wave speeds which are the most demanding part of the scheme.

```

CCC  RHO contains the total density after the reconstruction step;
CCC  the left and right values are denoted by the L and R suffix;
CCC  the inverse is calculated once here for efficiency;
CCC  C1O2 is the coefficient 1/2 - for clarity and efficiency

      RIL=1./RHOL
      RIR=1./RHOR

CCC  GAMMA variables determined for the allaire model

      GR=(ETAR+1.)/ETAR
      GL=(ETAL+1.)/ETAL
      G1R=GR-1.
      G1L=GL-1.

CCC  Flag to cut velocity if solid wall

      LEFT=ILN
      IRIG=IRN

CCC  KINETIC ENERGY
      RQ2HL=C1O2*(UUL*UUL+VVL*VVL+WWL*WWL)*RIL
      RQ2HR=C1O2*(UUR*UUR+VVR*VVR+WWR*WWR)*RIR

CCC  PRESSURE
      HPPL= G1L*(EEL-RQ2HL)
      HPPR= G1R*(EER-RQ2HR)

CCC  SPEED OF SOUND variables
      SL2=GL*G1L*(EEL-RQ2HL)*RIL
      SR2=GR*G1R*(EER-RQ2HR)*RIR
      HAL=SQRT(SL2)
      HAR=SQRT(SR2)
      HALI=1./HAL
      HARI=1./HAR

CCC  Averages for PVRs approximation
      HCR=C1O2*(RHOL+RHOR)
      HCA=C1O2*(HAL+HAR)

CCC  Velocity in flux direction (eg. XI)
      UL=XITI+(UUL*XX+VVL*YY+WWL*ZZ)*RIL
      UR=XITI+(UUR*XX+VVR*YY+WWR*ZZ)*RIR

```

Code Fragment 2.1: Quantities required for the HLLC Riemann Solver

There are several approaches to estimating the wave speeds. Naturally the simplest is to make a direct estimate using the flow velocity and speed of sound already calculated. A number of different averaging and limiting processes

have been suggested to refine this estimate. Another option is to consider the pressure internal to the Riemann structure and use this to determine the wave speeds. Again there are different ways this can be applied, the most complete of which is the hybrid scheme put forward by Toro which was implemented in this case. We determine the wave speeds to be

$$S_L = u_L - a_L q_L, \quad S_R = u_R - a_R q_R,$$

where,

$$q_K = \begin{cases} 1 & \text{if } p_* \leq p_K \\ [1 + \frac{\gamma+1}{2\gamma}(p_*/p_K - 1)]^{\frac{1}{2}} & \text{if } p_* > p_K \end{cases},$$

and the middle wave can be determined from these exactly,

$$S_* = \frac{p_R - p_L + \rho_L u_L (S_L - u_L) - \rho_R u_R (S_R - u_R)}{\rho_L (S_L - u_L) - \rho_R (S_R - u_R)}.$$

Initial estimates of the pressure in the star region are made based on the Primitive Variable Riemann Solver (PVRS) approach. In theory this is an iterative procedure however in practise it was noted by Toro that one iteration was sufficient. The pressure is given as

$$p_* = \frac{1}{2}(p_L + p_R) - \frac{1}{2}(u_R - u_L)(\bar{\rho}\bar{a}),$$

which is then used to determine the likely wave structure. If the star pressure is estimated to be less than the pressure on either side then we assume both waves are rarefactions. We can then use exact relations to better determine the star values for pressure and velocity - although it should be noted that the pressure ratios are an estimate and of course the concept of a single ‘star value’ is itself an approximation. The new values as taken from the two-rarefaction Riemann solver (TRRS) are

$$u_* = \frac{P_{LR}u_L/a_L + u_R/a_R + 2(P_{LR} - 1)/(\gamma - 1)}{P_{LR}/a_L + 1/a_R},$$

$$p_* = \frac{1}{2} \left\{ p_L \left[1 + \frac{\gamma - 1}{2a_L}(u_L - u_*) \right]^{\frac{2\gamma}{\gamma-1}} + p_R \left[1 + \frac{\gamma - 1}{2a_R}(u_* - u_R) \right]^{\frac{2\gamma}{\gamma-1}} \right\},$$

where $P_{LR} = \left(\frac{p_L}{p_R}\right)^{\frac{\gamma-1}{2\gamma}}$. Similarly if the initial estimate for pressure is greater than that found on either side of the interface or there is a significant jump in

pressure between the left and right values then we can improve our estimate by assuming the methodology of the two-shock Riemann solver (TSRS). In this case,

$$p_* = \frac{g_L p_L + g_R p_R - (u_R - u_L)}{g_L + g_R},$$

where

$$g_K = \left[\frac{2}{\rho_K(p_*(\gamma + 1) + p_K(\gamma - 1))} \right]^{\frac{1}{2}}.$$

In most instances the Riemann problem will be in a more conventional form and the initial PVRS estimate will be used, however this adaptive procedure allows a more sophisticated solver to come into effect in critical regions of high gradients and extrema. The code implementation (2.2) illustrates how this approach is used in practise.

It is worth noting in particular that the velocity of the central ‘star’ wave is utterly dependant on the estimates for the outer wave speeds, which in turn are based on the estimated pressure. It shall be seen shortly that the central wave is used to determine whether the left or right values are used, the upwinding direction. This is a crucial aspect of the solver which is why special attention is appropriate if the gradients are large and there is a dominant direction. Having calculated the wave speeds it only remains to evaluate the fluxes. As the code works on a curvilinear basis, the equations have been implemented in curvilinear form [19]. For convenience this form of the equations is reproduced here for the star region, along with the code implementation (2.3).

$$\mathbf{U}_{*K} = \frac{1}{S_K - S_*} \begin{pmatrix} \rho_K(S_K - \lambda_K) \\ \rho u_K(S_K - \lambda_L) + (p^* - p_K)\tilde{\xi}_x \\ \rho v_K(S_K - \lambda_L) + (p^* - p_K)\tilde{\xi}_y \\ \rho w_K(S_K - \lambda_L) + (p^* - p_K)\tilde{\xi}_z \\ e_K(S_* - \lambda_K) + S_* p^* - p_K \lambda_K \end{pmatrix}$$

$$\mathbf{F}_{*K} = \begin{pmatrix} \rho_{*K} S_* \\ \rho u_{*K} S_* + p^* \tilde{\xi}_x \\ \rho v_{*K} S_* + p^* \tilde{\xi}_y \\ \rho w_{*K} S_* + p^* \tilde{\xi}_z \\ (e_K + p^*) S_* \end{pmatrix}$$

$$p^* = \rho_K(\lambda_K - S_K)(\lambda_K - S_*) + p_K$$

```

CCC  CALCULATE PVRS PRESSURE
PSTAR=C1O2*((HPPL+HPPR)-(UR-UL)*HCR*HCA)

CCC  DETERMINE WHAT APPROXIMATION TO USE FOR STAR VALUES
HPMAX=MAX(HPPL,HPPR)
HPMIN=MIN(HPPL,HPPR)
QUSER=2.

  IF((HPMAX.GE.QUSER*HPMIN).OR.(PSTAR.LT.HPMIN).OR.
    (PSTAR.GT.HPMAX))THEN
    IF(PSTAR.LT.HPMIN)THEN      !Switch to TRRS
      HPLR=(HPPL/HPPR)**(C1O2*G1L/GL)
      USTAR=(HPLR*UL*HALI+UR*HARI+2.*(HPLR-1)*ETAL)/
        (HPLR*HALI+HARI)
      PSTAR=C1O2*(HPPL*(1.+(C1O2*G1L*HALI)*
        (UL-USTAR))**2.*GL*ETAL)
        +HPPR*(1.+(G1R*C1O2*HARI)*
        (USTAR-UR))**2.*GR*ETAL)
    ELSE                          !Switch to TSRS
      HGL=SQRT((2.*RIL)/(PSTAR*(GL+1.)+(HPPL*G1L)))
      HGR=SQRT((2.*RIR)/(PSTAR*(GR+1.)+(HPPR*G1R)))
      PSTAR=(HGL*HPPL+HGR*HPPR-(UR-UL))/(HGL+HGR)
    ENDIF
  ENDIF

CCC  ESTIMATE WAVE SPEEDS FROM PSTAR

  IF(PSTAR.LE.HPPL)THEN
    HQL=1.
  ELSE
    HQL=SQRT(1.+C1O2*(GL+1.)*(PSTAR/HPPL-1.)/GL)
  ENDIF
  IF(PSTAR.LE.HPPR)THEN
    HQR=1.
  ELSE
    HQR=SQRT(1.+C1O2*(GR+1.)*(PSTAR/HPPR-1.)/GR)
  ENDIF

  HSL=UL-HAL*HQL
  HSR=UR+HAR*HQR
  HSDL=HSL-UL
  HSDR=HSR-UR

  SSTAR=(HPPR-HPPL+RHOL*UL*HSDL-RHOR*UR*HSDR)/
    (RHOL*HSDL-RHOR*HSDR)

```

Code Fragment 2.2: Estimating wave speeds for the HLLC Riemann Solver

2.2.2 MUSTA

Although not strictly a Riemann solver, the MUSTA approach [86] was also implemented into the code. This is a multi-stage solution procedure which has the advantage of not requiring a known equation of state. The code is given in [86] and while it was found to perform acceptably in our research it was not as accurate as the Riemann solvers proper and so not appropriate given the focus on perfect gas cases.

2.2.3 Introduction to the Sod problem

The modified Sod shock tube problem is a commonly used one-dimensional test case and is a good basis for comparison with other schemes [73]. Starting from two regions of flow at rest, it is essentially a large-scale Riemann problem,

```

CCC   DETERMINE FLUXES

      IF(SSTAR.GT.0.)THEN      !use left values
        RFLS=(UUL*XIX+VVL*XIY+WWL*XIZ)
        L=1-LEFT-IRIG
        LRIL=L*RIL
        U=UUL*LRIL
        V=VVL*LRIL
        W=WWL*LRIL

        IF(HSL.GT.0.)THEN      !interface lies outside Riemann structure

          XHS=U*RFLS+HPPL*XIX
          YHS=V*RFLS+HPPL*XIY
          ZHS=W*RFLS+HPPL*XIZ
          EFLS=RFLS*(HPPL+EEL)*RIL-HPPL*XITI

CCC   Variables for solution of volume fraction advection
      RIN=U*B
      QIN=MML*RFLS*RIL

      FFLHS=FFL*RFLS*RIL
      RFLS=RRL*RFLS*RIL

      ELSE                      !interface lies inside Riemann structure
        HUSL=SSTAR*B/(HSL-SSTAR)
        HPS=RHOL*(-HSDL)*(UL-SSTAR)+HPPL
        HPSMHPPL=HPS-HPPL
        RRLHSDL=RHOL*HSDL

        RFLS=HUSL*RRL*HSDL
        FFLHS=HUSL*FFL*HSDL
        XHS=HUSL*(RRLHSDL*U+HPSMHPPL*XX)+HPS*XIX
        YHS=HUSL*(RRLHSDL*V+HPSMHPPL*YY)+HPS*XIY
        ZHS=HUSL*(RRLHSDL*W+HPSMHPPL*ZZ)+HPS*XIZ

        EFLS=HUSL*(EEL*HSDL+HPS*SSTAR-HPPL*UL)
1      +SSTAR*B*HPS

      RIN=SSTAR*B
      QIN=MML*RIN

      ENDIF

      ELSE                      !use right values

```

Code Fragment 2.3: Evaluating fluxes for the HLLC Riemann Solver

creating an expansion, contact and shock wave. It is a simple enough problem that a fair comparison of methods can be made. In particular there is a lot of interest in minimizing the width of the contact wave, which in theory should be infinitely thin as should the shock wave. In practise, obtaining a contact discontinuity that covers four computational cells is considered the practical limit. The initial conditions for the variable-gamma case are given as

$$(\rho, u, p, \gamma)_L = (1, 0, 1, 1.4) ,$$

$$(\rho, u, p, \gamma)_R = (0.125, 0, 0.1, 1.2) ,$$

while a single-species version simply sets $\gamma = 1.4$ in both sections of the fluid. Figure 2.4 shows the resulting density distribution for this single-species case

with the initial fluid regions defined either side of $x = 0.5$ with standard second-order MUSCL reconstruction comparing the CBM solver to the HLLC solver along with the exact solution. While it is reassuring to note that the results are for the most part identical, there is a clear flaw in the CBM solution at the base of the expansion wave generated from the initial condition and maintained as the waves travel away from the center. It should be noted this oscillation appears even with first-order reconstruction. It is a relatively small issue but set against the smoother result generated by the HLLC solver it is significant. The Sod problem will be revisited along with other one-dimensional test cases in investigating other developments of the code.

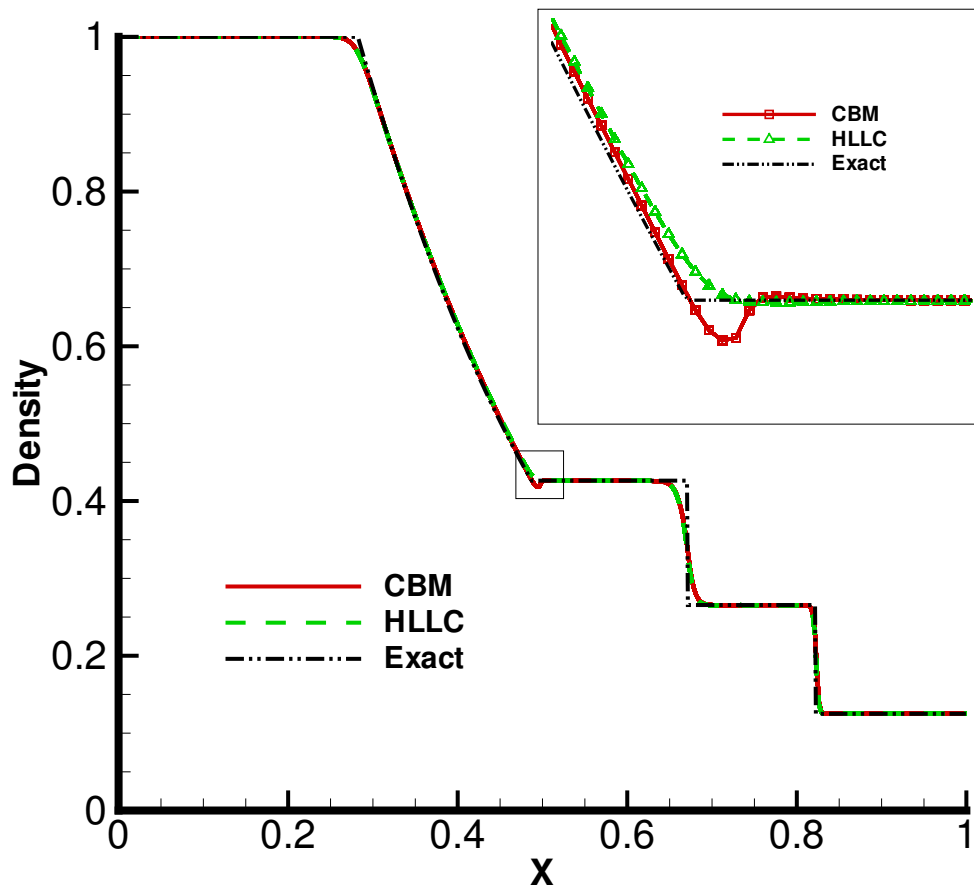


Figure 2.4: Comparison of Riemann solvers for the Sod problem

In most simulations there is no evident difference between the solvers, however in addition to the greater ease of implementation and development offered by the HLLC solver and the arguable improvement to results, it was noted as being slightly faster in computational time and more robust to providing

solutions in more extreme problems where the CBM failed to maintain positive energy. The overall impression was that the HLLC solver was a better choice to carry forward in development whenever there was any distinction between the two methods. As shall be seen in the following sections this also facilitated development of the multi-component extension of the code as well as in the area of increasing order of accuracy. Unless otherwise specified, all subsequent results were obtained with the HLLC solver.

2.3 Reconstruction

The second part of Godunov’s legacy, and by far the more challenging, is laid down in his own theorem.

“There are no monotone, linear schemes. . . of second or higher order accuracy.”

Monotonicity is the requirement that cell variable updates are bounded by the neighbouring cells. In the scheme

$$u_i^{n+1} = \sum_{k=-l_L}^{l_R} b_k u_{i+k}^n,$$

monotonicity can be interpreted as positive or zero coefficients

$$b_k \geq 0, \forall k.$$

It is essentially what we mean by removing the spurious oscillations that occur around shock waves - no new minima or maxima can be created. Godunov’s theorem explicitly states that monotonicity and accuracy are contradictory aims. Naturally people have been trying to circumvent this for years, and the key to this effort is in the assumption of a linear method. It is perhaps not surprising that in order to fully engage with the highly non-linear governing equations, a non-linear solver be used. In many schemes this non-linearity is introduced through artificial viscosity (the high-order accurate central schemes have dissipation added artificially where necessary to damp out the inherent oscillations). Thus far the Godunov method is inherently free of oscillations, but restricted to first order. To improve this requires a non-linear extension that will increase the order of accuracy in smooth parts of the flow but not interfere with the region of any discontinuity - quite the opposite effect of artificial viscosity. There has been a lot of work on this, however there must first be consideration of which variables to interpolate.

2.3.1 Variable selection

The need for a conservative scheme defines the variables used in the flux calculations. There is freedom however in the reconstruction phase. Interpolating the conserved variables can give very poor results in both accuracy and stability. Using the primitive variables can essentially remove these concerns and is a simple enough operation. As with the discussion of Riemann solvers it is not always easy to establish a clear difference between different options, never mind a preferred option. However there is a test case that is suitably sensitive to illustrate the potential improvement in results that a change in variables can provide. Figure 2.5 shows part of the pressure profile following the passage of a shock wave through a slab of helium with second-order reconstruction using the HLLC solver. The definition for this problem is

$$(\rho, u, p, \gamma, c_v)_{x < 0.25} = (1.3765, 0.3948, 1.57, 1.4, 717.2) ,$$

$$(\rho, u, p, \gamma, c_v)_{0.25 < x < 0.4} = (1.0, 0.0, 1.0, 1.4, 717.2) ,$$

$$(\rho, u, p, \gamma, c_v)_{0.4 < x < 0.6} = (0.138, 0.0, 1.0, 1.67, 3114.9) ,$$

$$(\rho, u, p, \gamma, c_v)_{0.6 < x} = (1.0, 0.0, 1.0, 1.4, 717.2) ,$$

and the resulting flow has a complex series of reflected and transmitted shocks making it a more demanding case than the Sod problem but also more relevant to practical applications.

There is a clear argument for the use of primitive variables to avoid the pressure oscillations seen in the figure. It is well worth remembering that this is highly coupled to both the reconstruction method itself and the multi-component scheme employed which provides some of the variables in question. Many papers suggest, for the higher-order methods in particular, that it is necessary to perform the characteristic transformation [67]. While this may perhaps prove fruitful in certain flows with complex shock structures it does not offer any particular treatment of the multi-component equations which are typically the most challenging element of variable choice and in terms of stability the characteristic variables have not been required in the work done.

2.3.2 MUSCL

The MUSCL (Monotone Upwind-centred Scheme for Conservation Laws) approach of Van Leer [87, 18] was an obvious extension to Godunov's method.




Figure 2.5: Pressure in the helium slab for different choice of variables in the reconstruction

Instead of assuming piece-wise constant data, van Leer assumed piece-wise linear data thereby achieving second order accuracy. The values used for either side of the Riemann problem are a much better approximation of the real interface value. The non-linearity is introduced through ‘slope limiters’, functions that restrict the slope imposed on each cell in the vicinity of the discontinuity in order to maintain monotonicity. These were developed as Total-Variation Diminishing (TVD) schemes, which for one-dimensional linear advection is synonymous with preserving monotonicity. The basic aim of these restrictions was that maxima would decrease in value and minima increase, thereby avoiding the creation of spurious oscillations. The basic consequence was that in the region of a shock-wave the scheme defaulted to first-order accuracy. Various limiters have been developed with varying degrees of dissipation; there is a clear correlation between improved resolution of discontinuities and in-

crease in oscillations. Sweby [75] covers and standardises all the traditional limiters. However even the most dissipative of these is a definite improvement over the first-order solutions. In the past decade there have been many attempts to push the limit of the TVD concept [65, 48, 59, 96]. Another paper which accepts that TVD is limited to one dimension [44] explores an equivalent multi-dimensional concept of Local Extremum Diminishing (LED). Higher order interpolations such as piece-wise quadratic can improve accuracy further but are all restricted by the quality of the limiting functions available.

One alternative to slope limiting is what is known as flux limiting [20, 24]. The idea is the same, with a limiting function controlling the ‘mix’ of a first and second order flux depending on the nature of the local solution. The option to combine two different fluxes with different characteristics has some merits, but has yet to prove particularly more effective than using first and second order of the same flux as slope limiters do.

In terms of our code, a range of limiters were previously available and much early work centered around evaluating these and implementing alternatives such as the piecewise quadratic [76] scheme mentioned. This was ongoing as part of the familiarisation with the code and consideration of Riemann solvers. The only limiters carried forward here are Minmod - the most dissipative and therefore robust traditional option - and the van Leer limiter which provided best accuracy reliably. For an interface at $i + \frac{1}{2}$ the left and right values for the general MUSCL reconstruction are given as

$$\mathbf{U}_{i+1/2}^L = \mathbf{U}_i + \frac{1}{2}(\mathbf{U}_i - \mathbf{U}_{i-1})\phi(r_L^{lim}) ,$$

$$\mathbf{U}_{i+1/2}^R = \mathbf{U}_{i+1} - \frac{1}{2}(\mathbf{U}_{i+2} - \mathbf{U}_{i+1})\phi(r_R^{lim}) ,$$

where

$$r_L^{lim} = \frac{\mathbf{U}_{i+1} - \mathbf{U}_i}{\mathbf{U}_i - \mathbf{U}_{i-1}} , r_R^{lim} = \frac{\mathbf{U}_{i+1} - \mathbf{U}_i}{\mathbf{U}_{i+2} - \mathbf{U}_{i+1}} ,$$

and $\phi(r^{lim})$ is the limiting function. A negative value for r^{lim} implies an extremum in which case the limiting function automatically reverts to first order. Otherwise, for Minmod (MM) and van Leer (VL) type limiting respectively [85],

$$\phi_{MM} = \min(1, r^{lim}) ,$$

$$\phi_{VL} = \frac{2r^{lim}}{1 + r^{lim}} .$$

```

!MUSCL
C   CALCULATE LEFT AND RIGHT DIFFERENCES
    DLFT=HL-HLL
    DLOC=HR-HL
    DRGT=HRR-HR
    DLFTI=1./(DLFT+EPS)
    DRGTI=1./(DRGT+EPS)
C   CALCULATE R FOR THE LIMITER FUNCTIONS
    RLIMR =DLOC*DRGTI
    RLIML =DLOC*DLFTI
    .
    .
    ELSEIF(LIMLOCAL.EQ.2)THEN      !MM
      DO I=1,7
C   RIGHT      ELIMR (I)=MAX(0.,MIN(1.,RLIMR (I)))
C   LEFT      ELIML (I)=MAX(0.,MIN(1.,RLIML (I)))
      END DO
    .
    .
    ELSEIF(LIMLOCAL.EQ.4)THEN      !VL
      ELIMR =(RLIMR +ABS(RLIMR))/(1.+RLIMR )
      ELIML =(RLIML +ABS(RLIML))/(1.+RLIML )
    .
    .
    MUSCLR=ELIMR
    MUSCLL=ELIML
    RRL=HL(1)+MUSCLL(1)*DLFT(1)*C1O2
    FFL=HL(2)+MUSCLL(2)*DLFT(2)*C1O2
    MML=HL(3)+MUSCLL(3)*DLFT(3)*C1O2
    UUL=HL(4)+MUSCLL(4)*DLFT(4)*C1O2
    VVL=HL(5)+MUSCLL(5)*DLFT(5)*C1O2
    WWL=HL(6)+MUSCLL(6)*DLFT(6)*C1O2
    EEL=HL(7)+MUSCLL(7)*DLFT(7)*C1O2
    RRR=HR(1)-MUSCLR(1)*DRGT(1)*C1O2
    FFR=HR(2)-MUSCLR(2)*DRGT(2)*C1O2
    MMR=HR(3)-MUSCLR(3)*DRGT(3)*C1O2
    UUR=HR(4)-MUSCLR(4)*DRGT(4)*C1O2
    VVR=HR(5)-MUSCLR(5)*DRGT(5)*C1O2
    WWR=HR(6)-MUSCLR(6)*DRGT(6)*C1O2
    EER=HR(7)-MUSCLR(7)*DRGT(7)*C1O2

```

Code Fragment 2.4: The pre-existing code to reconstruct interface values using MUSCL

These equations form part of the reconstruction subroutine in the code 2.4 where the variables to be reconstructed are held in the arrays prefixed by H with suffixes denoting the cell they belong to relative to the interface and EPS is a small number to avoid division by zero.

As piecewise-linear methods these two reconstruction methods are limited to a nominal second-order. While it is not instructive to look at the implementation of the majority of alternatives considered in the early part of this research, there is a notable exception that was added by others in the group during the

project that belongs to the MUSCL category and will be used extensively in subsequent results. It is a fifth-order scheme [51, 52] which requires more data points and is notably more complex however it is a convenient extension of the MUSCL framework to high order. The limiting functions in this case are different for the left and right constructs and are given by

$$\phi_{M5,K}^{lim} = \max(0, \min(2, 2r_{K,i}^{lim}, \tilde{\phi}_{M5,K}^{lim})) ,$$

where

$$\tilde{\phi}_{M5,L}^{lim} = \frac{11 - \frac{2}{r_{L,i-1}^{lim}} + 24r_{L,i}^{lim} - 3r_{L,i}^{lim}r_{L,i+1}^{lim}}{30} ,$$

$$\tilde{\phi}_{M5,R}^{lim} = \frac{11 - \frac{2}{r_{R,i+2}^{lim}} + 24r_{R,i+1}^{lim} - 3r_{R,i+1}^{lim}r_{R,i}^{lim}}{30} ,$$

with

$$r_{L,i}^{lim} = \frac{\mathbf{U}_{i+1} - \mathbf{U}_i}{\mathbf{U}_i - \mathbf{U}_{i-1}} , \quad r_{R,i}^{lim} = \frac{\mathbf{U}_i - \mathbf{U}_{i-1}}{\mathbf{U}_{i+1} - \mathbf{U}_i} .$$

Monotonicity is maintained in theory with this approach, however in smooth regions of the flow the accuracy can be expected to be improved. It is worth noting at this point that the scheme still reverts to first order in the presence of extrema. It is possible to increase the notional order of accuracy further in MUSCL schemes but they offer little return for their increased cost. These will prove to be significant details when this fifth-order MUSCL (M5) reconstruction is compared against other high-order approaches.

2.3.3 WENO

The main advance since MUSCL in this area has been the increasing development of Essentially Non-Oscillatory (ENO) schemes [34, 8, 26] and more recently Weighted-ENO (WENO) schemes [45, 56]. ENO and WENO aim to provide very high order accuracy without reducing the integrity of the shock-wave region in finite volume schemes [83]. They are effectively complex polynomial interpolations to the data values at the interface based on a large support stencil - finite difference ENO and WENO interpolated the fluxes themselves to the cell interfaces. Instead of limiting functions to dissipate undesirable oscillations, ENO and WENO actually maintain the same level of accuracy in all areas of the flow by use of an adaptive stencil. ENO selects its support such that it does not cross a discontinuity, whereas WENO uses a convex combination of all available stencils (with discontinuous stencils being afforded zero

weight). WENO has several other advantages - it gives a smoother, more predictable flux; can obtain the optimal accuracy from its support base ($(2r - 1)$ where r is the size of the stencil); and removes the logical statements necessary for ENO, which makes it twice as fast on vector supercomputers. However to achieve a proper multi-dimensional high-order interpolation in finite volume schemes on unstructured meshes can require an excessively large number of calculations. It is not necessarily practical to use these schemes for problems where there is little return for supposedly increased accuracy, but using them in split dimensions on cartesian grids such as those used here, while limited in its effect, is fairly straightforward to implement. The weaknesses of these very high-order schemes, apart from additional computational time, are detecting the discontinuities, particularly when two discontinuities are close together. These methods are not perfect, a fact acknowledged in their name as *essentially* non-oscillatory. They are a move away from the more traditional thinking of TVD schemes, yet there are cases where further monotonicity restraints are required for stability [2, 74], and various approaches to artificial compression are currently being investigated to counter unwanted diffusion [92]. For all these supposed flaws, the ability to calculate relatively stable and clear results at fifth, ninth or potentially higher orders of accuracy has made these methods very attractive as shall be seen in the results. The extra cost of evaluating the large stencils is more than covered by the reduction in grid resolution that can be made while maintaining an accurate and well-resolved solution. Comparisons of the different high-order methods can be found in [55, 72] and alternative approaches to improving the WENO concept in [69, 66, 17].

Recently a number of improvements have been proposed in the literature geared towards releasing the optimal performance by reconsidering the weightings in WENO schemes and their implementation into the code has proved valuable. In order to understand these changes, it is necessary to consider the underlying equations of the WENO scheme which shall be illustrated here for the fifth-order case - extension to ninth- (and higher) order follows similarly.

The basic principle of a WENO interpolation in finite volume form is to interpolate an interface value for each variable from all available stencils, and to average them with appropriate weighting to achieve the final result which is then entered into the calculation of flux. Through this a range of stencils three cells in length can theoretically produce a fifth-order accurate reconstruction. Non-linearity is introduced through the weightings, which are made dependant on the ‘smoothness’ of each candidate stencil - those stencils adjudged to contain a discontinuity are given an effective zero weighting. This can be expressed mathematically as

$$U_{j+\frac{1}{2}} = \sum_{k=1}^3 \omega_k \hat{U}_{j+\frac{1}{2}}^k,$$

where $\hat{U}_{j+\frac{1}{2}}^k$ is the interpolated value for each stencil, calculated as in [83], and ω_k is the weight.

In the work of Shu, these weights were originally formulated as

$$\omega_k = \frac{\alpha_k}{\sum_{i=1}^3 \alpha_i}, \text{ where } \alpha_k = \frac{\bar{\omega}_k}{(\epsilon + \beta_k)^p}.$$

ϵ is a small parameter to keep α bounded, β_k are the ‘smoothness indicators’ as laid out in [2], p is a free parameter and $\bar{\omega}_k$ are the ideal weights required to reduce to a fifth-order central difference scheme in smooth flow. Interpolating for the left-hand side, the indices $i, k = 1, 2, 3$ refers to the three stencils (S_k , Fig. 2.6) for the $j + \frac{1}{2}$ interface which are

$$S_1 = (j, j + 1, j + 2), S_2 = (j - 1, j, j + 1), S_3 = (j - 2, j - 1, j).$$




Figure 2.6: Stencils for left-sided interpolation of interface values

In our code (2.5) the j^{th} cell is identified as *HL* with $j + 1$ being *HR*, $j - 1$ as *HLL* and so on. The ideal weights are given as

$$\bar{\omega}_1 = \frac{3}{10}, \bar{\omega}_2 = \frac{6}{10}, \bar{\omega}_3 = \frac{1}{10}.$$

The smoothness indicators are then calculated for each variable,

$$\beta_1 = \frac{13}{12}(U_j - 2U_{j+1} + U_{j+2})^2 + \frac{1}{4}(3U_j - 4U_{j+1} + U_{j+2})^2,$$

$$\beta_2 = \frac{13}{12}(U_{j-1} - 2U_j + U_{j+1})^2 + \frac{1}{4}(U_{j-1} - U_{j+1})^2,$$

$$\beta_3 = \frac{13}{12}(U_{j-2} - 2U_{j-1} + U_j)^2 + \frac{1}{4}(U_{j-2} - 4U_{j-1} + 3U_j)^2,$$

```

!WENO 3RD

!LEFT

CCC Ideal weightings
    DEE(1)=0.3
    DEE(2)=0.6
    DEE(3)=0.1

CCC Calculate smoothness indicators
    DO I=1,7
    BEE(1)=(13./12.)*(HL(I)-2.*HR(I)+HRR(I))**2
    .      +0.25*(3.*HL(I)-4.*HR(I)+HRR(I))**2
    BEE(2)=(13./12.)*(HLL(I)-2.*HL(I)+HR(I))**2
    .      +0.25*(HLL(I)-HR(I))**2
    BEE(3)=(13./12.)*(HLLL(I)-2.*HLL(I)+HL(I))**2
    .      +0.25*(HLLL(I)-4.*HLL(I)+3.*HL(I))**2

    EPS=1.E-36

CCC Relative indicator to determine if flow is discontinuous
CCC or just highly variable
    RSIK=MAX(BEE(1),BEE(2),BEE(3))
    .      /(EPS+MIN(BEE(1),BEE(2),BEE(3)))

    IF(RSIK.LT.10.)BEE=0.

CCC Calculation of weights
    ALPH=DEE/(EPS+BEE) !**2

    ALPHATOT=1./(ALPH(1)+ALPH(2)+ALPH(3)+ALPH(4)+ALPH(5))
    OMEGA=ALPH*ALPHATOT

CCC Mapping of weights
    GEE=OMEGA*(DEE+DEE**2-3.*DEE*OMEGA
    .      +OMEGA**2)/(DEE**2+OMEGA*(1.-2.*DEE))

    GEETOT=1./(GEE(1)+GEE(2)+GEE(3))
    OMEGA=GEE*GEETOT

CCC Final reconstruction
    WENOL(I)=OMEGA(1)*(-1.*HRR(I)+5.*HR(I)+2.*HL(I))/6.
    .      +OMEGA(2)*(-1.*HLL(I)+5.*HL(I)+2.*HR(I))/6.
    .      +OMEGA(3)*(2.*HLLL(I)-7.*HLL(I)+11.*HL(I))/6.
    ENDDO

    RRL=WENOL(1)
    FFL=WENOL(2)
    MML=WENOL(3)
    UUL=WENOL(4)
    VVL=WENOL(5)
    WWL=WENOL(6)
    EEL=WENOL(7)

```

Code Fragment 2.5: Implementation of modified fifth-order WENO reconstruction for the left side of the interface

and combined with the previous formulae to give the weights.

More recent work has discovered that the above formulation is not sufficient to maintain the maximum possible accuracy around critical points [37]. This paper suggests that a simple mapping of the weights calculated as before to bring them closer to the ideal weights, while still retaining the required behaviour away from these regions resolves this issue. The modified weights then are calculated as

$$\omega_k^{(M)} = \frac{\alpha_k^*}{\sum_{i=0}^2 \alpha_i^*},$$

$$\alpha_k^* = g_k(\omega_k),$$

$$g_k(\omega_k) = \frac{\omega_k(\bar{\omega}_k + \bar{\omega}_k^2 - 3\bar{\omega}_k\omega_k + \omega_k^2)}{\bar{\omega}_k^2 + \omega_k(1 - 2\bar{\omega}_k)},$$

thus the mapping is relatively cheap. In implementation it has also been seen to improve the results, gaining greater convergence for a coarse grid. The paper also made an additional point regarding the value of ϵ , which it illustrated should be much smaller than the parameter used in Shu's work. Without the mapping corrected scheme however, reducing ϵ led to an unacceptable loss of formal accuracy. The modified scheme is robust to the change and so can take advantage of making ϵ as small as the machine accuracy will allow.

Another paper has been written recently with an aim to optimising WENO methods for turbulent flows [77]. Again the aim has been to reduce the influence of the previously over-constrictive weights to allow turbulent flow to develop to a high-order approximation. The first simple change is to consider the value of the exponent p in the original calculations. It has been set at 2 in the work of Shu and Toro, but the paper by Taylor et al. identifies that a value of 1 can significantly reduce the non-linear limiting throughout the flow without any great loss of stability in regions where the adaptation is necessary. Again this modification was quickly shown to improve the results given by our implementation of the WENO code.

The main development of the paper however was in the use of 'relative limiters'. These are in effect a relatively crude assessment of the smoothness indicators which help differentiate between a single discontinuity and a rapidly varying but smooth turbulent flow. If the smoothness of all the stencils under consideration is relatively comparable, then the flow is deemed smooth and the indicators adjusted as such. A very small change to the overall code, taking advantage of the information given by the basic but large WENO calculations, again reduces unnecessary dissipation in the method allowing for the weights to more closely match the ideal weighting of the central difference scheme. The paper found the scheme to be robust for different problems if these changes were made relative, and a threshold of one order of magnitude between the biggest and smallest smoothness indicators within a set of stencils has proved acceptable in our code as suggested in the paper.

2.3.4 Low Mach Number Correction

During the course of this research an issue was highlighted within the group concerning the dissipation of energy in low mach number flow. This was established to be caused by the reaveraging inherent in the Riemann solution process [81]. This Mach number dependency could be removed with significant improvement in results as shall be seen in later chapters by modifying the reconstructed variables. Having interpolated left and right interface values, the two are simply blended together linearly with Mach number such that at zero-speed the resulting interface value is the mean of the two reconstructed values. This is expressed as

$$\mathbf{U}_L^{LM} = \frac{1}{2}((1 + M)\mathbf{U}_L + (1 - M)\mathbf{U}_R) ,$$

$$\mathbf{U}_R^{LM} = \frac{1}{2}((1 + M)\mathbf{U}_R + (1 - M)\mathbf{U}_L) ,$$

where $M = \min(1, \text{Mach No})$.

2.3.5 WAF

Before looking at the relative performance of the MUSCL and WENO reconstruction methods there is again an alternative approach that was considered. The Weighted-Average-Flux (WAF) [6, 84] method provides an intriguing way to improve the accuracy of a scheme. It is distinct from the flux splitting methods considered previously, in that it is more an extension of the HLLC Riemann solver. Rather than selecting which of the four states in the Riemann problem corresponds to the interface position, the WAF method takes a weighted average of all four states. It offers increased accuracy as it can be applied in addition to reconstruction of the variables, although it is still dependent on limiting functions similar to those used in the MUSCL scheme. For the higher-order methods there is not necessarily much of a gain and there is notable extra computational expense so the work was not carried forward to the more complex applications however it was included in early test cases.

2.3.6 Further one-dimensional test cases

Although consideration of the order of accuracy of reconstruction forms the bulk of work done in later chapters it is useful to look at the simpler test cases involved in the development of the code to assess the veracity of the concepts gained from the literature. We will first revisit the Sod problem introduced in the previous section. Figure 2.7 shows the relative improvements offered

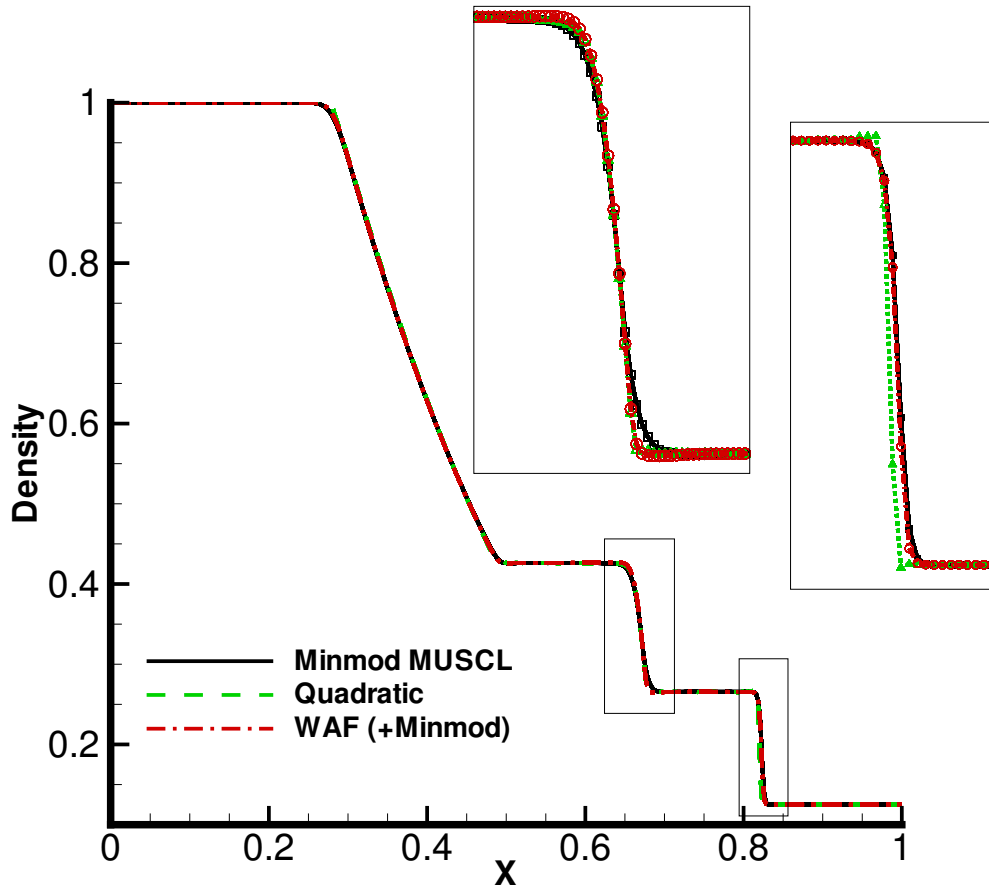


Figure 2.7: Comparison of lower-order reconstruction methods for the Sod problem

by the piecewise-quadratic interpolation and the WAF extension to the basic Minmod-limited second-order MUSCL scheme. As expected the effect of improving accuracy is to sharpen the contact discontinuity - the quadratic reconstruction also gives a marginally sharper shockwave. As mentioned previously these options were not carried forward beyond the early period of research, the reason of which is partly illustrated in figure 2.8. Here the fifth and ninth order WENO reconstructions are used in their original form. While there are other coding issues affecting the quality of results, the improvement in resolution of the discontinuities is immediately apparent and this is achieved with less complexity than the other options. The apparent difference in shock position for the 9th order scheme is simply due to output time. The lack of strict monotonicity is clearly visible as ‘overshoots’ in the WENO results, and there is a range of developments aimed at limiting this problem in the literature, some of which were implemented in the code though not carried forward - for

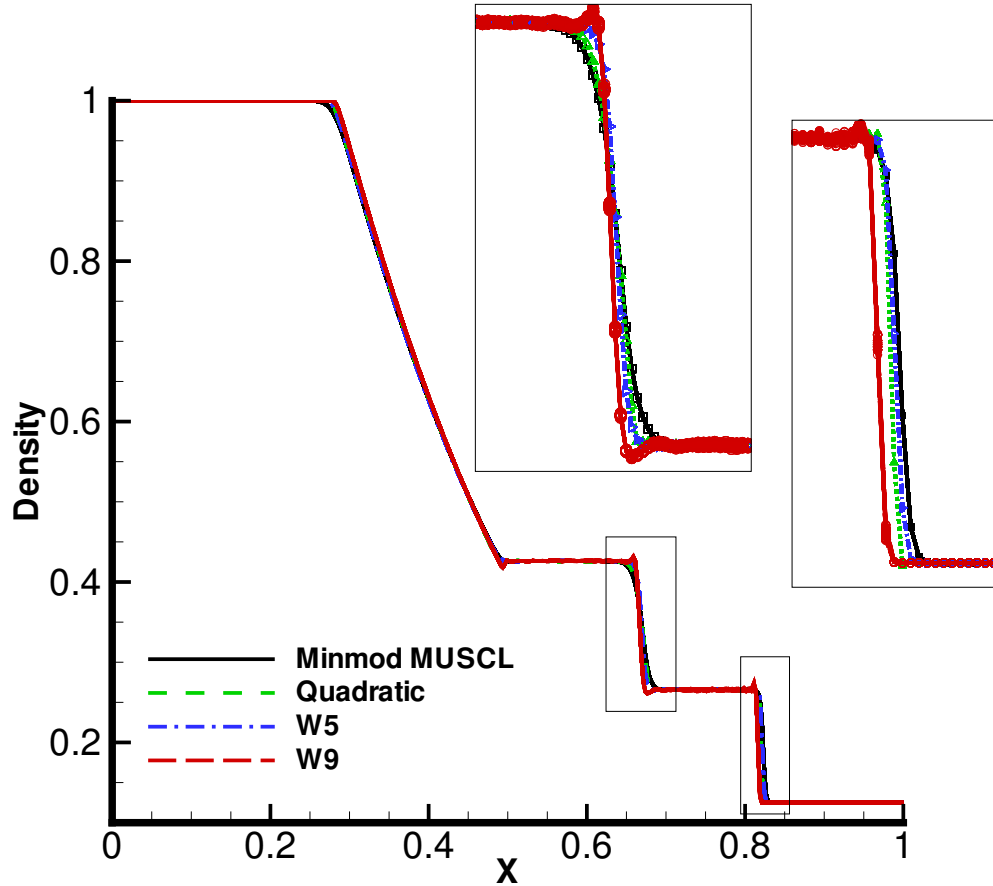


Figure 2.8: Comparison of higher-order reconstruction methods for the Sod problem

practical application in our turbulence-dominated flows these oscillations were not significant enough to outweigh the gains in overall accuracy.

More recent work with the improved WENO code helps illustrate some important considerations concerning the dissipation inherent in different reconstruction methods. A simple test case is the advection of a Gaussian pulse [69]. The current multi-component version of the code involves advection of the volume fraction which can be specified as a simple colour function (z) when the two components are identical. The function is defined as

$$z = \frac{1}{4}(1 + e^{(-30x-0.5)^2}) ,$$

and is shown in figure 2.9. The flow is given a uniform velocity of $1ms^{-1}$ and two periods are allowed to pass at very low CFL to reduce time inaccuracies. As this is an inviscid problem the shape of the pulse ought to be maintained




Figure 2.9: Initial condition for advection of Gaussian pulse with constant uniform velocity $1ms^{-1}$




Figure 2.10: Comparison of reconstruction methods for advection of Gaussian pulse with 16 grid points

without any dissipation as it is advected with the flow. The case was run at 3 different grid resolutions - 16, 32 and 64 points - in order to gain convergence information in addition to simply observing how well the different methods preserved the pulse. Figures 2.10, 2.11, 2.12 show the results of this test case




Figure 2.11: Comparison of reconstruction methods for advection of Gaussian pulse with 32 grid points




Figure 2.12: Comparison of reconstruction methods for advection of Gaussian pulse with 64 grid points

for the nominal second-order van Leer limiter (VL), the fifth-order MUSCL (M5) and ninth-order WENO (W9).

There are several interesting things to note. The most obvious perhaps is that increasing the number of grid points reduces the dissipation. This is

to be expected as the cell averaging in the finite volume method will better represent the detail if the cells are smaller. It is also immediately apparent that increasing the order of reconstruction of this averaged data has a similar effect in maintaining the structure of the pulse. Furthermore while the WENO method appears to be grid converged with the correct solution at 64 points, the MUSCL methods seem to be converging to a different solution. This is most likely a manifestation of the first-order treatment of extrema discussed in the MUSCL implementation. This introduces an important theoretical point - the nominal order of accuracy of these reconstruction methods only refers to their maximum order in smooth flows and does not represent the overall accuracy of the code [29]. This can be illustrated informally by consideration of the rate of error reduction in this case. Looking simply at the difference in peak values between a given solution and the initial condition we can gain a measure of the error and how it changes with grid resolution. Table 2.1 shows the peak values for each simulation, the relative error and the associated ‘order’ of improvement over the previous grid resolution.

Table 2.1: Order of accuracy of solver based on peak error convergence

	Grid Size	Peak Value	Relative Error	Order of accuracy (p), ($\frac{ERR_{16}}{ERR_{32}} = 2^p$)
VL	16	0.416	16.8%	-
	32	0.462	7.6%	1.15
	64	0.486	2.8%	1.45
M5	16	0.442	11.6%	-
	32	0.471	5.8%	1.0
	64	0.487	2.6%	1.15
W9	16	0.474	5.2%	-
	32	0.495	1.0%	2.4
	64	0.4995	0.1%	3.3

While only an approximate and variable measure this not only enables the difference in the reconstruction methods to be quantified but also demonstrates to some extent why the high-order WENO schemes are able to give physical solutions when ILES theory requires that the dissipation represented by the inherent errors in the scheme be around second order [58]. It also introduces the concept that the ninth-order WENO scheme can obtain results comparable to lower-order schemes while using half the number of grid points. It shall be seen in later chapters that the WENO reconstruction is correspondingly twice as expensive to compute, however in multi-dimensional problems the reduction in grid points dominates. Finally the table would suggest that the M5 scheme is of lower order than the van Leer scheme, yet the results are more accurate. This is partly due to the first-order limit leaving little room for improvement,

and further suggests that quantifying the accuracy of these schemes absolutely in theoretical terms is a considerable challenge.

There is one further well-documented test case which shall be presented to further emphasise the limits of these MUSCL formulations in comparison with the adaptability and preserved high order of WENO. The case is a shock/entropy waves interaction as simulated in [2]. The initial conditions are defined as

$$(\rho, P, u_x) = (3.857143, 10.3333, 2.629369) \quad x < -0.8$$

$$(\rho, P, u_x) = (1 + 0.2 \text{sine}(5\pi x), 1, 0) \quad x > -0.8 .$$




Figure 2.13: Comparison of reconstruction methods for the entropy wave test cases

It is clear from these results (Fig. 2.13) that for this problem there is a much wider gap in resolution of the high frequency waves between WENO

and MUSCL in our implementation than might be expected. The first-order points in the MUSCL approach are severely limiting, so much so that the WENO method has notably better accuracy with half the number of grid points.

2.4 Multi-component models

It has been mentioned several times prior to this point that the code includes multi-component models necessary to handle the more complex flows considered in this research. There are a number of options to consider however there is no one preferred method for all cases. Indeed the results presented in the previous section use all three of implemented models with no significant impact on the comparison of methods. The first of these models is the Total Enthalpy Conservation Model (ThCM) [89] which was already in the code. As part of the current research the quasi-conservative models proposed by Johnsen and Colonius [46] and by Allaire et al. [1] have been implemented to improve the behaviour of the higher-order methods in keeping the volume fraction bounded. A more complete consideration of the relative merits of these different models can be found in [78]. While part of the code development, it is only of tangential interest to the problems at hand and so the three models are discussed briefly here.

2.4.1 Total Enthalpy Conservation Model (ThCM)

This model was originally in the code as a development of the standard mass-fraction model. It requires an additional two equations for total enthalpy conservation of multi-component flow as proposed by Wang et al.. The additional variables being conserved are $\frac{\rho\chi}{M}$, $\frac{\rho}{M}$, where M is molecular weight, and χ is defined as

$$\chi = \frac{\rho^i}{p} + 1$$

The total enthalpy approach had been found to be better for the current range of problems than the mass fraction model commonly found in industrial codes. However it still created pressure oscillations around contact discontinuities which degraded the overall reliability of the results.

2.4.2 Quasi-conservative Johnsen and Colonius model for WENO (JC)

It has been noted that the WENO methods are particularly sensitive to the choice of interpolated variables due to the lack of enforced monotonicity. With the implementation of HLLC and WENO into the code, a new model designed for such solvers developed by Johnsen and Colonius seemed an obvious candidate to be investigated. For the range of problems considered it appears to significantly improve upon the ThCM model in the pressure field, although as it is a quasi-conservative approach the gamma can fluctuate at high Mach numbers. The model requires only one additional equation, plus a small modification to the HLLC solver to solve the main Euler equations in a suitably quasi-conservative fashion. In effect, the method merely employs a stiffened equation of state, which reduces to

$$\Gamma P = E - \frac{1}{2}\rho u^2, \text{ where } \Gamma = \frac{1}{\gamma - 1}$$

for perfect gases. While convenient and a notable improvement over the ThCM for higher-order methods, there was still some concern over the efficacy of the multi-component part of the code.

2.4.3 Quasi-conservative 5-equation model of Allaire

The quasi-conservative 5-equation model of Allaire et al. works over a wide range of conditions and maintains a good interface with higher-order methods. Each fluid has its own continuity equation and there is an advection equation for the volume fraction of one of the fluids. Along with total momentum and energy these make up the system,

$$\begin{aligned} \frac{\partial z_1 \rho_1}{\partial t} + \nabla \cdot (z_1 \rho_1 \mathbf{u}) &= 0, \\ \frac{\partial z_2 \rho_2}{\partial t} + \nabla \cdot (z_2 \rho_2 \mathbf{u}) &= 0, \\ \frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \times \mathbf{u} + P) &= 0, \\ \frac{\partial \rho e}{\partial t} + \nabla \cdot (\rho \mathbf{u} (e + \frac{P}{\rho})) &= 0, \\ \frac{\partial z}{\partial t} + \mathbf{u} \nabla z &= 0. \end{aligned}$$

In their work they considered different closures for this under-specified system and determined that the isobaric approach was preferable. It is a relatively

simple model which is easily adapted to the HLLC solver. The only complication is advection of the volume fraction, however as mass of each gas is already conserved there is some freedom here. As suggested in the original paper, the transport equation for volume fraction, z , can be written

$$\frac{\partial z}{\partial t} + \nabla \cdot (z\mathbf{u}) - z\nabla \cdot \mathbf{u} = 0$$

which can be calculated as with the conservative equations but with a non-conservative correction. On balance this model has proven to be preferable for the majority of flows and was used as the standard model in the latter stages of this research.

2.5 Time integration

Compared to the efforts in spatial discretisation, the time development of a flow has long been a simple matter. The stability requirements are well characterised and there has been little reason to be concerned with accuracy. As discussed previously, the spatial error remains large even with the higher-order schemes. Varying the order of the time-stepping scheme has not been seen to significantly affect the results, and use of a third order scheme ensures that the temporal error is dominated by a spatial error thought to be no better than second order. The basic equation for Godunov,

$$\mathbf{U}_i^{n+1} = \mathbf{U}_i^n + \frac{\Delta t}{\Delta x} (\mathbf{F}_{i-\frac{1}{2}} - \mathbf{F}_{i+\frac{1}{2}}),$$

shows the first-order forward in time algorithm. There are many alternatives and extensions but the use of these is largely determined in practise by the application.

2.5.1 Runge-Kutta

The optimal option for these shock-driven unsteady flows is generally accepted to be the third-order Runge-Kutta method [20, 43]. A Total-Variation-Diminishing (TVD) formulation,

$$\begin{aligned} \mathbf{U}_j^1 &= \mathbf{U}_j^n + \frac{\Delta t}{\Delta x} \mathbf{F}(\mathbf{U}_j^n), \\ \mathbf{U}_j^2 &= \mathbf{U}_j^n + \frac{1}{4} \frac{\Delta t}{\Delta x} [\mathbf{F}(\mathbf{U}_j^n) + \mathbf{F}(\mathbf{U}_j^1)], \\ \mathbf{U}_j^{n+1} &= \mathbf{U}_j^n + \frac{1}{6} \frac{\Delta t}{\Delta x} [\mathbf{F}(\mathbf{U}_j^n) + \mathbf{F}(\mathbf{U}_j^1) + 4\mathbf{F}(\mathbf{U}_j^2)], \end{aligned}$$

is used as preferable for the higher-order methods. Splitting each time-step into stages is the basic mechanism for increasing accuracy and stability. In principle Runge-Kutta methods reduce at first-order to the basic forward method. This third-order method is actually more stable than the first-order method, but increasing the accuracy further reduces the stability. It is the standard approach for such fluctuating flows and well-established from early work in the aeronautics industry.

2.5.2 Implicit Time Integration

The code is also equipped to solve steady-state problems with an implicit time integration method that can take advantage of large iteration steps without concern for accuracy (for example [30]). While considerable effort was employed to make this part of the code functional with the latest version to facilitate the work on hypersonics presented later, there was little scope for considering the underlying equations. It is mentioned here due to its use in producing some results, however it has not been the subject of any development during this research.

2.5.3 ADER

Finally an intriguing possibility which has yet to reach wider application is the ADER scheme developed by Titarev and Toro [82]. They have shown that the higher-order spatial discretisation can be used to determine similarly arbitrary high-order terms in time. It actually requires a different framework to be adopted which has not as yet been done here, however it may prove a very useful direction of further development. Recent work by the authors demonstrates how this can all be combined with the high-order methods with very promising results [22].

2.6 Axisymmetric problems

While the majority of cases simulated in this research take place in a box-geometry with a cartesian mesh there is scope within the code to simulate more complex geometries in a curvilinear framework as implied previously. To efficiently and accurately simulate axisymmetric problems it was necessary to implement additional terms to account for the difference in the formation of equations from cartesian to cylindrical geometries [33, 21, 71]. Using an r-z co-ordinate system

$$r \frac{\partial \mathbf{U}}{\partial t} + \tilde{\nabla} \cdot ((\mathbf{F}(\mathbf{U}) + \mathbf{V}(\mathbf{U}))r) = \mathbf{S}(\mathbf{U}) .$$

Here \mathbf{F} represents the convective part of the flux and \mathbf{V} the viscous fluxes as necessary for the hypersonic problem. The axi-symmetric source terms appear as $\mathbf{S}(\mathbf{U}) = (0, S_r, S_z, S_E)^T$,

$$S_r = p - \tau_{\theta\theta} - \frac{2}{3} \frac{\partial(\mu u_r)}{\partial r} ,$$

$$S_z = -\frac{2}{3} \frac{\partial(\mu u_r)}{\partial z} ,$$

$$S_E = \frac{2}{3} \mu \left(\frac{u_r^2}{r} \right) + r \left(\frac{\partial}{\partial r} \left(\frac{2}{3} \mu \frac{u_r^2}{r} \right) + \frac{\partial}{\partial z} \left(\frac{2}{3} \mu \frac{u_r u_z}{r} \right) \right) ,$$

where

$$\tau_{\theta\theta} = \frac{2}{3} \mu \left(-\nabla \cdot \mathbf{u} + 2 \frac{u_r}{r} \right) .$$

3

Two-dimensional single-mode Richtmyer-Meshkov Instability

3.1 Introduction

The Richtmyer-Meshkov Instability (RMI) derives from the work of Richtmyer [68] where he considered the case of a shock wave passing from light fluid into heavy fluid with a perturbed interface as an extension of the theory of Taylor regarding the growth of such irregularities. Meshkov [60] established the more general case where the heavy fluid can be accelerated into the light fluid. Figure 3.1 shows the initial condition schematically - the flow can be assumed to be inviscid and the timescales are such that gravity is not considered (the orientation of the fluids is arbitrary). The significance of this test case will be more apparent when considering more realistic problems in Chapter 5 though it is not difficult to conceive of situations involving the interaction of shock waves with fluid interfaces.

Handling the theory around such complex flow is not straightforward, however some insight into the physics of the instability can be achieved through consideration of the two-dimensional vorticity (ω) equation,

$$\frac{D\omega}{Dt} = \frac{1}{\rho}(\nabla\rho \times \nabla p) .$$

In the absence of viscosity we are left with this source term which gives rise to vorticity when the local pressure gradient and density gradient are not parallel.

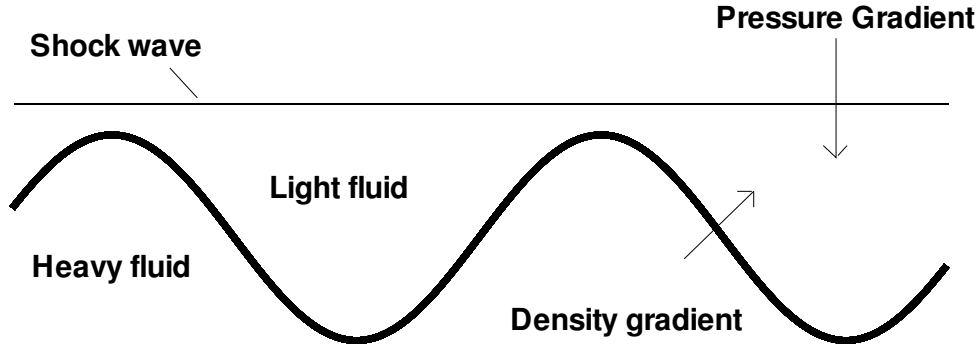


Figure 3.1: Schematic representation of the initial condition for the single-mode Richtmyer-Meshkov Instability

As can be seen in figure 3.1 this is the situation created by a perturbed fluid interface with the passage of a planar shock wave. The extent to which the vorticity so deposited by the shock wave can account for the growth of the instability will be considered later in the chapter in the section on modelling growth. The earlier sections address the effect of numerics in simulating a two-dimensional RMI.

The Richtmyer-Meshkov Instability modelled here is taken from the experimental work of Jacobs [41]. Similar work simulating this problem has also been conducted by Latini et al. [53]. In the current research the simulation was first initialised as perfect - pure gases on either side of the interface with a pre-shock Atwood number of 0.692. Due to necessity the experiment has some diffusion as a result of which air is mixed with the SF6. The numerical interface between the two gases is initialised with a simple sinusoidal perturbation over a wavelength of 59mm, and a shock of strength Mach number 1.3 is passed through this interface. The experimental and numerical problems are similar enough to bear comparison, not only in the visual appearance of the instability (Fig. 3.2) but also the growth rates of the bubble and spike (to distinguish between the peaks and troughs in RMI flow the spike refers to the more extensive penetration of the heavier fluid while the bubble is the alternate extrema), which are appropriately non-dimensionalised in any case. Simulations with a modified initial condition to allow for the experimental diffusion were then carried out with use of the JC two-fluid model, rather than the ThCM. It is expected that increasing the thickness of the interface will have a depressing effect on the growth rate of the instability - an effect seen experimentally [7].

3.2 Early comparison of methods

The typical growth of the instability is shown in Figure 3.3 for a reasonably well resolved scheme, where the dark region represents the SF6 in a plot of volume fraction. After the initial linear growth phase the vorticity deposited by the shock wave can be seen to be gathering at the head of the instability as it starts to pull material round the sides. This roll-up continues, establishing clear vortical structures which then start to manifest the secondary instability seen in experiments as a series of small Kelvin-Helmholtz features which disturb the smooth lines of the vortex leading to an eventual breakdown into turbulence. Comparison with experiment will quickly reveal however that the curvature at the tip of the spike is flawed in this simulation. This is due to how the initial perturbation is resolved on the Cartesian grid, indeed there are many factors to consider in setting the initial conditions as the high-order schemes are very sensitive to small differences.

While the vortex dynamics dominate the flow and can destabilise the upper surface, there is no evidence that the shape of the very tip of the instability has any major impact on the formation of the vortices nor the breakdown to turbulence. Indeed, the comparison of methods conducted by Liska and Wendroff [55] in which they consider the Rayleigh-Taylor instability, which is closely related to Richtmyer-Meshkov, shows a host of unphysical and unpredictable features at the head of the instability depending on the method used. This is a manifestation of how different methods respond to “errors” in different ways. Should the initial conditions be perfectly specified as part of a continuum then one would expect a properly rounded tip to the instability but this is one area that demands very careful approximations to be simulated accurately.

Figure 3.4 shows the comparison of different order accuracy methods over a range of coarse regular grids with 20, 40 or 80 cells per wavelength. At the




Figure 3.2: Experimental picture [41] of single-mode Richtmyer-Meshkov Instability




Figure 3.3: Plot of Volume Fraction of SF6 showing the development of the instability over time, using 5th order WENO on 80x240 grid, based on a sharp initial interface

coarsest resolution very little can be seen, though it is well worth noting that the overall amplitude of the instability at this time is comparable with the more highly resolved simulations. With less than 20 cells this soon ceases to be the case, so this point can be considered to be when the primary instability is resolved. It is hard to see much difference between the schemes, however the higher-order WENO methods do give an indication that some genuine vorticity is at work at the top of the spike causing the shear layer to roll up. At 40

cells across, features are much more clearly visible. The plots (b) and (h) in Figure 3.4 still appear very similar at the large scale but the 9th-order WENO interpolation shows clear vortical roll-up at this stage and the ‘mushroom’ shape compares favourably to the experimental results. Yet, increasing the resolution further allows the second-order method (plot (c) in Figure 3.4) to start showing some kind of vortex development, but still it is not qualitatively the same as the much more advanced vortex coil for the 5th-order scheme (plot (f) in Figure 3.4) nor even particularly comparable to the coarser WENO simulations (plots (e) and (h)). The most resolved picture (i) has already passed the onset of secondary instability within the vortex coil (visible as ‘waviness’ in the 5th order simulation) and has broken down into a turbulently mixed area. Also quite clearly seen is some separated fluid drawn down from the vortices which is caused by small Kelvin-Helmholtz instabilities that form on the stem. At this resolution only the effects of distortion can be seen. They are not found in the experiment due to the slightly diffuse nature of the initial perturbation. By contrast these simulations had a sharp discontinuity (to the level of resolution of the grid) and the higher-order schemes do not carry sufficient dissipation to damp out the numerical errors.

Time-equivalent representations of the volume fractions within the flow are shown in figure 3.5 for the same range of methods as seen previously but with a diffuse initial interface - the volume fraction being given a linear profile over a 5mm region about the interface. The Atwood number has also been adjusted to 0.605 to reflect that actually measured in the experiments. There is little difference in how the various reconstruction methods behave relative to each other, however there is a noticeable difference in the visible level of vorticity and resulting shape of the head of the instability. A critical issue in dealing with this case is that as the roll-up of the vortices becomes established the two-dimensional vortex dynamics begin to dominate the flow. This affects the growth and stability of the stem in later times, as well as having a large impact on the shape of the overall instability. Until the solution has converged it is to be expected that different shapes will emerge dependant on the scheme as the vortices remain under-resolved. However it is possible to see qualitatively that the resolution of features is greater in the higher-order schemes, particularly comparing the fifth-order WENO with the second-order MUSCL. It must be borne in mind that the WENO schemes can only be as accurate as the number of cells in any continuous region for stability, which means for the ninth-order scheme to be truly effective, nine cells between discontinuities is required. This is not the case with the stem in these coarser grids, so there are regions where the order of accuracy is limited yielding little improvement overall on the fifth-order results. It should also be noted that this is a very approximate approach to considering the effects of diffusion in the initial condition.

Table 3.1 shows the relative CPU cost for a given length of simulated time for

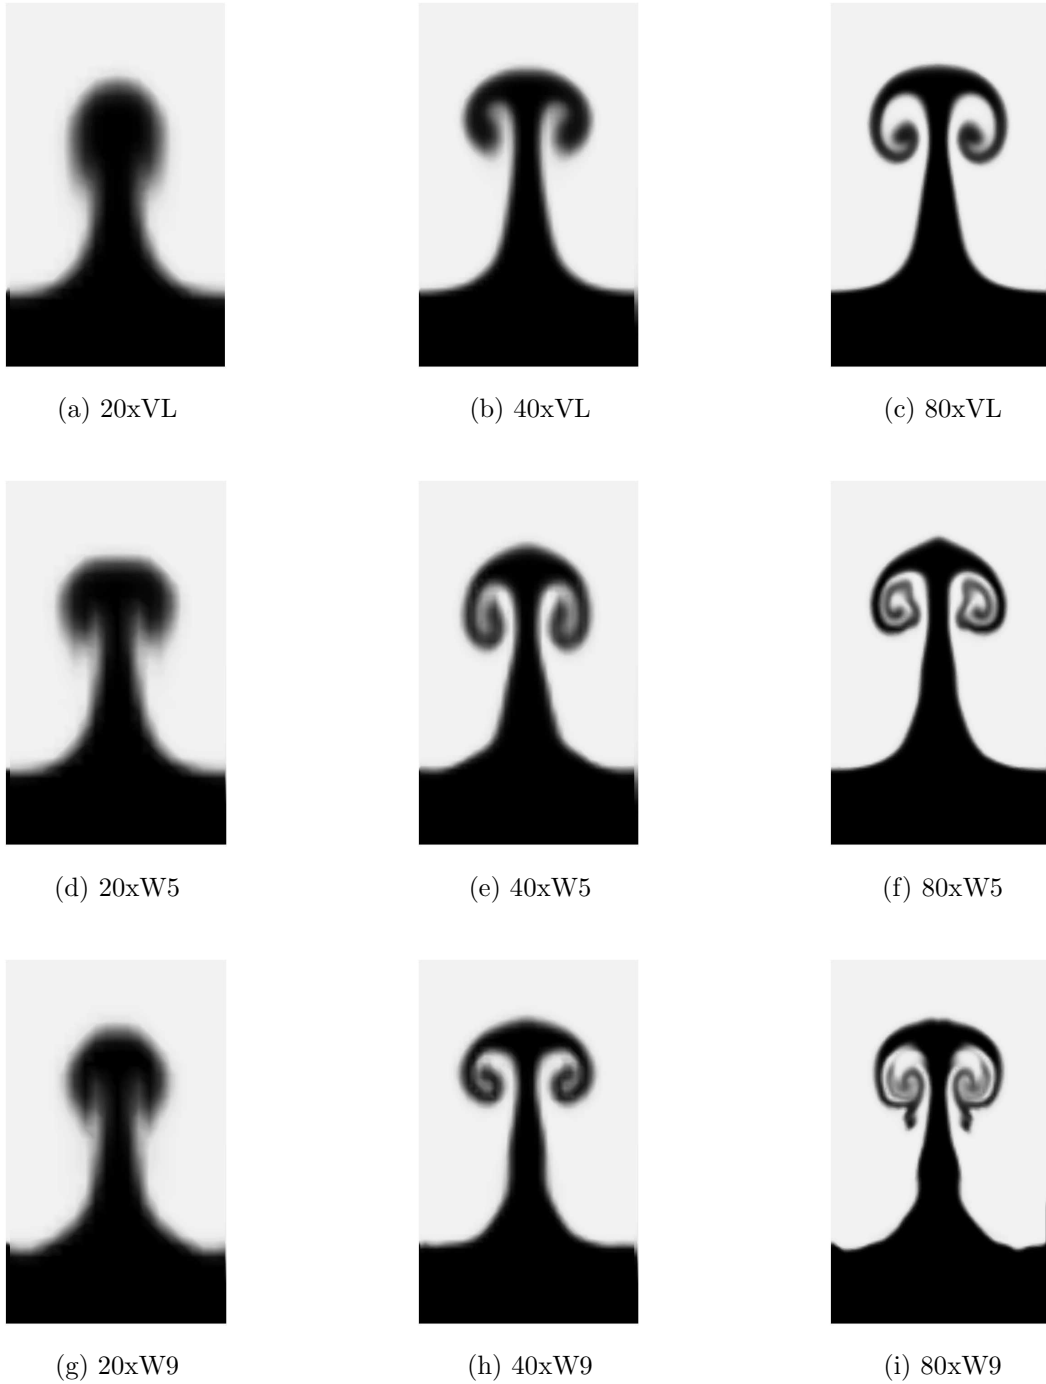


Figure 3.4: Volume fraction plots of Richtmyer-Meshkov Simulations with no. of cells per wavelength and reconstruction method (VL:2nd-order van Leer, W5:5th-order WENO, W9:9th-order WENO) for sharp initial interface and Atwood No. of 0.692




Figure 3.5: Richtmyer-Meshkov Simulations with no. of cells per wavelength and reconstruction method (VL:2nd-order van Leer, W5:5th-order WENO, W9:9th-order WENO) for diffuse initial interface and Atwood No. of 0.605

the different schemes used, normalised by the runtime for the 20 cell van Leer scheme. For the very coarse grid the ‘fixed’ costs start to dominate such that there is less difference between the different order schemes. However what is important to note is the relative cost of, for example, the ninth-order scheme on a 40 cell grid with the second-order scheme on an 80 cell grid. The higher order WENO methods are more expensive, but are better value for the level

of fine-scale detail they provide - and this comparison ought yield better results with the latest developments of the code. More refined grids and further comparisons with experimental results are required to assess whether such detail is physically correct, however it has been noticed that the higher-order schemes are particularly sensitive to very small differences in initial conditions and there is a limit to how well the experiment can be modelled to achieve highly comparable results with simulations. Indeed it may not be possible with this problem to identify whether the scheme is introducing spurious errors or merely reflecting discrepancies in other parts of the model. Thus far the behaviour has appeared physical, and with coarse simulations the WENO methods do well at mimicking the results of the the second-order schemes on finer grids, but it has not been definitively seen that both approaches converge to the same solution for this problem, and there are a number of issues that need to be addressed.

Resolution	20	40	80
VL	1	4.1	18.2
W5	1.02	10.4	45.1
W9	1.70	11.5	53.5


Table 3.1: CPU runtime for schemes VL, W5 and W9 at grid resolution 20, 40 and 80 for a given length of simulated time

In order to help validate these results much finer grids were used. Figure 3.6 shows the fifth-order WENO result for 400 grid cells, as well as how it compares overlayed on the 80 cell image. The overall amplitude is comparable as expected, as is the general bubble shape and position of the vortices. Although the coarser simulation does not evidence the secondary instability that leads to breakdown of the vortices, it is up to this point a very good approximation of the flow. It may be noted that the ninth-order results in figure 3.5 do not appear to be converging to this same shape, and indeed that has been seen in other simulations as the resolution is increased. It has been identified that this is due to sensitivity to spurious perturbations, which it has previously been noted can significantly alter the development of the flow.

3.3 ‘Error’ accumulation

3.3.1 LES and the ‘error’

With or without a subgrid model when simulating complex flows, and in particular transitional flows, there is a minimum resolution required to capture the structure of the instabilities which drive these flows as has been seen earlier in this chapter. Interpreting these physically requires understanding of



(a) Highly resolved fine grid

(b) Fine grid result superimposed on coarse grid

Figure 3.6: Fine grid (400 cell) 5th-order WENO simulation compared to coarse 80 cell

effective viscosity and damping. A better approach to appreciate the relation between grid resolution and order of reconstruction is to view it numerically. In a cell-averaged finite volume scheme the information is limited. Reducing the number of cells or using a lower order reconstruction is a basic loss of information. In this instance if the main vortices are not resolved, what begins as a range of similar flows from one initial condition can diverge to significantly different results in the later time development of the instability. In reality this effect generally occurs on a much greater scale, any practical application of an LES is not going to be able to afford the grid cells to fully resolve each spike. Supposed ‘effective order of accuracy’ is not a sufficient measure when features are severely under-resolved and the basic physics of the problem differs from that expected.

At greater detail, there are secondary instabilities which will affect the overall growth rate of the mixing region as the flow progresses. These appear to be a new flow feature altogether, witnessed in the experiments but not so often captured in any simulation at a useful resolution. These features change the late-time behaviour, altering the growth rates. Taken from this point of view the analysis would indicate that there is little difference between the different methods and grids, all capturing the overall growth accurately within

the experimental range of scatter. Without pursuing the simulations to very late times, the differences so evident when the flow is visualised would not be identified. This then is a human source of error combining with our basic level of error. These transitional, inherently unstable flows are heavily reliant on resolution of the driving features, which are often very small.

Another issue is that these features are often at low speed, and compressible (Godunov-type) methods are well known for increasing ‘dissipation’ at low Mach number. This is shown in the simple case of a single-mode Kelvin-Helmholtz instability (Figure 3.7). The upper and lower boundaries are periodic, pressure and density are initially constant throughout the flow and a velocity of 0.5 m/s is applied vertically up in the left-hand half of the domain and down in the right-hand half. This is a fundamental inviscid instability created at shear layers where the simple action of fluids passing in opposing directions causes the layer to roll up into a vortex. Excessive dissipation suppressing such instabilities would make it effectively impossible to resolve certain flows, particularly wall-bounded flows. However, the Mach number dependency can be removed through consideration of the structure of the Godunov scheme, allowing the compressible method to be used with confidence for a much wider range of flows. It is worth noting again that without the high-order methods, the vortex would be unresolved at this grid-level, and given the increasingly common presumption in turbulence simulations that the spectrum be resolved all the way to cut-off frequency, the number of cells required to capture this instability is quite significant.

These are the errors that are in some part necessary, they are a function of the solution approach and have to be understood in order to harness their strengths as an approach to modelling.

3.3.2 Errors in Highly-Resolved Flow

Having considered the importance of methods being able to resolve certain features in order to properly consider the flow to be physical, it is useful to look at how higher-resolution methods can reveal problems previously unnoticed. Many of these could be considered human error, but it is still important to recognise them as such. The single-mode Richtmyer-Meshkov problem is surprisingly intractable as an initial condition when you remove the convenient ‘damping’ that can suppress small perturbations before they become genuine instabilities. In particular, the ninth-order WENO method is quite capable of delivering appreciable changes in result when provided with errors on the level of the machine (for example 10^{-16}). Often this is still insignificant, and at times such behaviour is decidedly beneficial when handling turbulent flows - sensitivity being a key concept. However in the interests of assessing the quality of LES problems, possible sources of error have to be addressed and,




Figure 3.7: 9th-order WENO 16x16 simulation of Kelvin Helmholtz at Mach 0.2 and 0.02 with and without low Mach number (LM) scaling correction

in the case of these methods, that extends right down to the smallest errors - if only to discount them.

Figure 3.8 shows one example of how a small error in the initialisation of a well-resolved (MUSCL 5th order) simulation can have an effect. Failing to align the initial condition symmetrically with the grid creates ‘errors’ of the order 10^{-6} yet the consequences are clearly visible. It is perhaps easy when schemes tend towards the dissipative, or subgrid models specify what the physical behaviour at small scales is, to overlook the issue of sensitivity, and the importance of initialising the correct problem. This is illustrated again in Figure 3.9 - the

difference between an initially sharp interface and one which has diffused over a number of cells. The experiment had some diffusion prior to the passage of the shock. On the other hand, the small Kelvin-Helmholtz roll-ups are to be expected in a high- Re ($Re \rightarrow \infty$), inviscid simulation of this kind - provided there is sufficient resolution to capture them. Running a lower-order scheme with a sharp interface can give much the same result as the higher-order scheme that is initially diffuse. Both can reasonably approximate the experiment, but only one can be considered an accurate simulation of the problem.





Figure 3.8: Effect of misaligning initial condition with grid by 0.01% with 5th-order MUSCL

These smaller instabilities are not just a quirk of inviscid physics, it is possible to make mistakes during the initialisation which produce more such artifacts, yet they are indistinguishable from the physical features. Furthermore, representing the initial sinusoidal perturbation on a Cartesian grid provides seeding points for such instabilities. Such factors can be improved upon by better resolving the initial condition, using a higher-order approximation. Figure 3.10 shows the effect of this on a relatively coarse grid (this also shows the effect of the latest developments of the code in comparison with Fig. 3.4), but ultimately the flow is absolutely unstable and the more that is resolved, the more features appear. Identifying which are genuine and which erroneous can seem unnecessary, but while they may not be fully controlled, they can be assessed.


To conclude, there are many potential sources for error in relatively simple unstable problems, and through the nature of high-resolution methods these cannot be ignored. Many are manifested in a physical sense, making it difficult to separate them out from genuine instabilities and thus requiring great care when setting up a simulation. Such resolution is, however, necessary in order



(a) Diffuse initial interface

(b) Sharp initial interface

Figure 3.9: Effect of initialising interface over finite region



(a) Standard initialisation

(b) Well-resolved initialisation

Figure 3.10: Using a cell-averaged approximation to the initial condition vs. resolving the initial condition on 10,000 ‘sub-cells’.

to avoid the greater risk of ‘error’ posed by failing to resolve sufficient details of the flow.

3.4 Growth and theoretical models

As mentioned previously, the growth rates seen in all the simulations are very similar. Figure 3.11 shows the non-dimensionalised plot of amplitude - defined as half the vertical distance from bubble to spike tip using the 0.1/0.9 volume fraction contours - over time. In line with the work done by Jacobs we non-dimensionalise the time by the initial growth rate V_0 and both amplitude and time by the wavenumber k . The points represent the experimental measurements taken from Jacobs’ paper and the lines show the results of the simulations for all resolutions and schemes presented in fig. 3.4. It is expected that if the simulations were extended to much later time there would be increasing deviation as the non-linearity of the growth becomes dominated by the turbulent mixing - a phenomenon not captured at the lower-resolutions.

Work has been done recently to further examine the factors affecting the growth of the instability to later times for comparison with theoretical models. In the experimental work [41], the authors try to fit several established models [97, 98] to their late-time measurements after earlier experiments [42, 62] did not provide sufficient data to be conclusive. The theoretical models considered show occasional passing similarity to observed results, but only one seems to show credible form - the vortex method proposed by Jacobs and Sheeley. The theory merely assumes the flow can be represented as a series of discrete vortices of alternating sign, and calculates the resulting velocity at spike and bubble tips using potential theory. Our own simulations show that these discrete vortices are in evidence after relatively short times and do seem to drive the flow. This suggests that this is a reasonable approach, and is in keeping with the earlier suggestion that the instability is rooted in the baroclinic vorticity production term in the inviscid vorticity equation. The theory is extended in more recent work [54] and also considered in the three-dimensional case [9] which has not been simulated in depth in here. A potential flow model was created to observe the development of line of vortices varying sinusoidally in strength. While limited in scope and unable to account for finite Atwood number it did yield a credible representation of the large-scale vortex formation associated with Richtmyer-Meshkov instabilities. A more sophisticated model along the lines developed elsewhere [25] may prove more useful however the preliminary results pointed to the possibility of initialising the LES with potential flow. To test this theory further a new set of simulations were set up. In order to simplify the flow and allow better comparison between the simulations and the theory, a new initial condition was devised.




Figure 3.11: Growth of instability, as predicted by different methods (lines), compared to experimental measurements (circles relate to equivalent problem, triangles to a weaker shock case), non-dimensionalised

Based on the assumption that the Richtmyer-Meshkov flow develops from the deposited vorticity on the interface due to the mis-aligned pressure and density gradients as the shock wave passes through the interface, an initial velocity perturbation was calculated from an arbitrary series of discrete vortices varying sinusoidally in strength analogous to the potential flow model. This could then be applied to an initially flat interface, removing ambiguities caused by differing initial amplitude between experiments, simulations and theory. It also removes the need for a shock wave, and the subsequent estimations of post-shock amplitude and Atwood number, any considerations of drift velocity or reflected waves and spurious oscillations. Furthermore this method is decoupled from the simulation itself such that the theoretical ‘zero-Atwood No.’ case can be simulated. Working from the assumption of similarity provided by the non-dimensionalisations, we need only extract the peak velocity in this imposed perturbation to know the initial growth rate and it can then be applied to a range of flows in the knowledge that the Atwood No. is absolutely the only variable.

Before looking at the results of these simulations we shall look further at the vortex model proposed by Likhachev and Jacobs. Earlier work had identified that the simple form for the amplitude based on equally spaced alternating vortices,

$$A = \sinh^{-1}(T) ,$$

where A is the non-dimensional amplitude and T the non-dimensional time, appeared to represent the shape of the growth observed in experiments. They extended this to an infinite row of vortices displaced by $\pm\epsilon$ to represent the effect of finite Atwood number. This equation is somewhat more complex, the non-dimensional spike (B_+) and bubble (B_-) interface positions relative to the vortex line being defined by

$$\frac{dB_{\pm}}{dT} = \pm \frac{\cos \epsilon}{\cosh B_{\pm} \mp \sin \epsilon} - \frac{1}{2} \tan \epsilon ,$$

with the absolute position then being given by

$$A_{\pm} = |B_{\pm}| \pm \frac{1}{2} \tan(\epsilon)T .$$

This can be solved implicitly, or simply integrated numerically as done for this research. The equation yields two parameters to be determined, the displacement ϵ and the vortex strength γ which is hidden in the non-dimensionalisation of time. Based in part on the earlier work the value $\gamma = \frac{\pi}{2}$ had been observed to give a good fit with results without any real theoretical basis provided. As it will impact on later discussion it is appropriate to offer an explanation for this value here.

First we consider the semi-infinite row of alternating vortices extending in non-dimensional space to one side of a given spike or bubble. For zero Atwood No. the two are interchangeable. If we allow each discrete vortex to have vorticity Ω_d then the velocity induced at any point by a given vortex is $V = \frac{\Omega_d}{r}$ where r is the distance from the vortex. Looking at the point where the tip of the spike/bubble would be we can see the first vortex is at $r = \frac{\pi}{2}$ and the velocity induced by the entire series at the tip can be expressed as

$$V_i = \frac{\Omega_d}{\pi/2} (1 - 1/3 + 1/5 - 1/7 + 1/9 \dots) ,$$

to infinity. The infinite series has a well-known value, $\frac{\pi}{4}$, which enables us to considerably simplify the expression. Now taking into account the second semi-infinite series extending from the other side of the spike we get the non-dimensional expression for total initial spike velocity,

$$V_s = \Omega_d .$$

Substituting circulation for vorticity, as $\Omega = \frac{\Gamma}{2\pi}$, and taking the estimation for circulation made by Jacobs and Sheeley [42], $\Gamma = 4\gamma V_s$, we can extract an expression for γ ,

$$\begin{aligned} \Gamma &= 2\pi\Omega_d = 2\pi V_s = 4\gamma V_s , \\ \gamma &= \frac{\pi}{2} , \end{aligned}$$

which is as expected. It is clear from this derivation that the extension to displaced vortices and the introduction of ϵ also requires a reassessment of the strength of the vortices represented in γ as the semi-infinite series will no longer be so easily formulated and will appear differently to spikes and bubbles.

Before considering this, we shall look at the simple case of zero Atwood No. Fig 3.12 shows the development over a long period of time of the volume fraction in the latest simulations using a reasonably well-resolved grid and fifth-order MUSCL scheme with the velocity perturbation initialisation. In this case as the two fluids are identical the volume fraction merely acts as a colour function. The grid used is well-resolved in the central section where the instability develops but coarsens towards the ends of the ‘shock tube’ resulting in spurious features at later time.

We can see the expected symmetric vortex roll-up with bubble and spike indistinguishable from each other as the vortices grow and become fully mixed. From this view the vortex model would seem very sensible, and this proves to be the case. Figure 3.13 shows the growth of the bubble, spike and overall amplitude over time. There is a slight discrepancy between these lines of the order of a grid cell most likely generated in the output rather than the flow itself. They are essentially the same, as expected. The vortex model with $\epsilon = 0$ is also shown, assuming $\gamma = \frac{\pi}{2}$. Naturally both bubble and spike give the same results, which follows the simulated flow very well over this long period. Also shown is the zero-Atwood No. case for the Sadot model - an empirical equation acknowledged to be one of the best fits for available data thus far and sharing some similar terms with the vortex model. In non-dimensional form it simplifies in the case to

$$\tilde{a} = \frac{1}{2} \left[\ln(\tilde{t}^2 + \tilde{t} + 1) + \frac{2}{\sqrt{3}} \tan^{-1} \left(\frac{1 + 2\tilde{t}}{\sqrt{3}} \right) \right] .$$

The late time behaviour seems to be captured well for this fundamental case by these equations, something which is not true of the bulk of theoretical models.

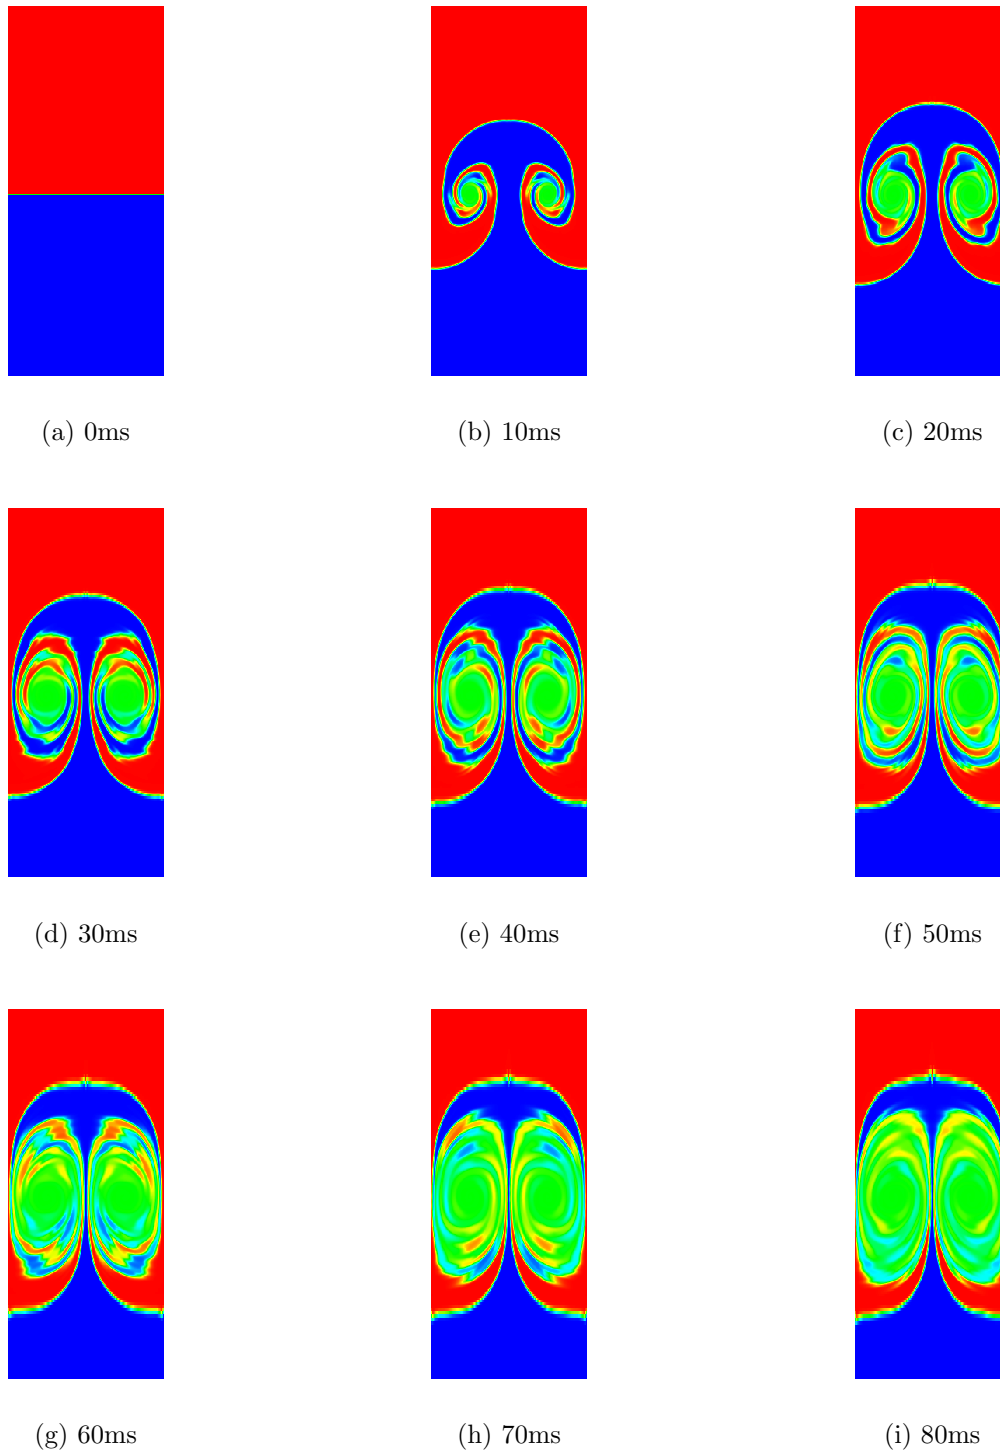


Figure 3.12: Volume fraction plots of zero Atwood number (single-fluid) Richtmyer-Meshkov corresponding to non-dimensional time $\approx 0 - 80$

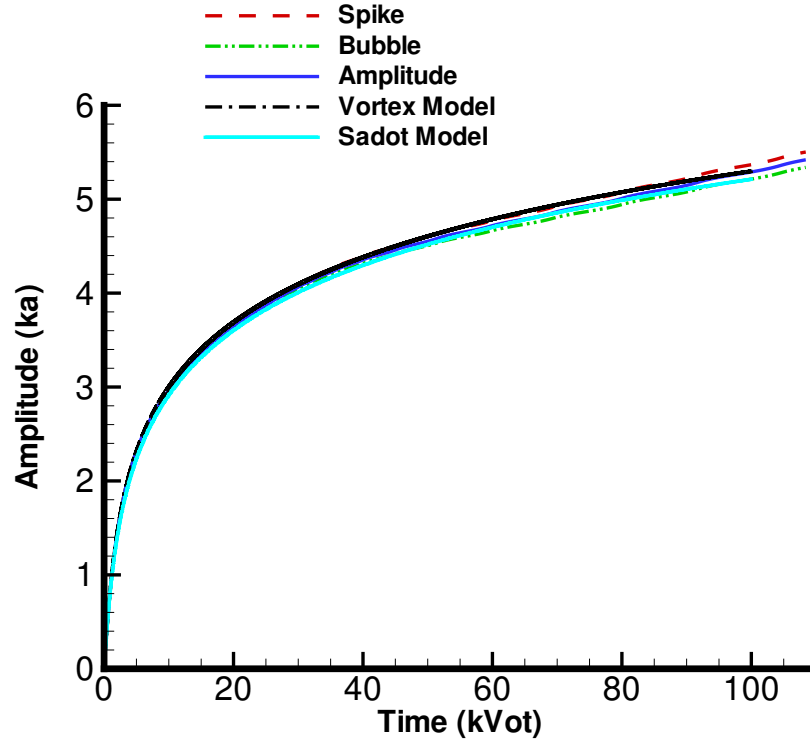


Figure 3.13: Simulated growth of instability amplitude (expressed as the mean of the individual bubble and spike amplitudes) at zero Atwood No. compared to models using the non-dimensionalisation where k is wavenumber and V_0 is initial growth rate

It is reassuring that the evolution of a vortex row can be simulated and modelled to such late times however this is not in itself a validation of this approach for true Richtmyer-Meshkov Instabilities and does not answer the challenge of finite Atwood No. Returning to the previous work simulating the high Atwood number case, the velocity perturbation method was applied to fluids of differing density to compare with experiment and previous simulation. The flow is visualised in Fig. 3.14 to later times. The early development does look very much like that seen in previous ‘authentic’ simulations.

This simulation is not run as late as the zero-Atwood No. case for the obvious reason that the domain would need to be much larger. A jet of fluid can be seen ejected from the tip of the spike as it reaches the coarser part of the grid. Because of this the overall growth results for late time should be viewed sceptically as the vortex pair can be seen to become fairly stationary in comparison. In any event the models fail to agree well in the early development making an accurate assessment of the late-time behaviour a moot point at

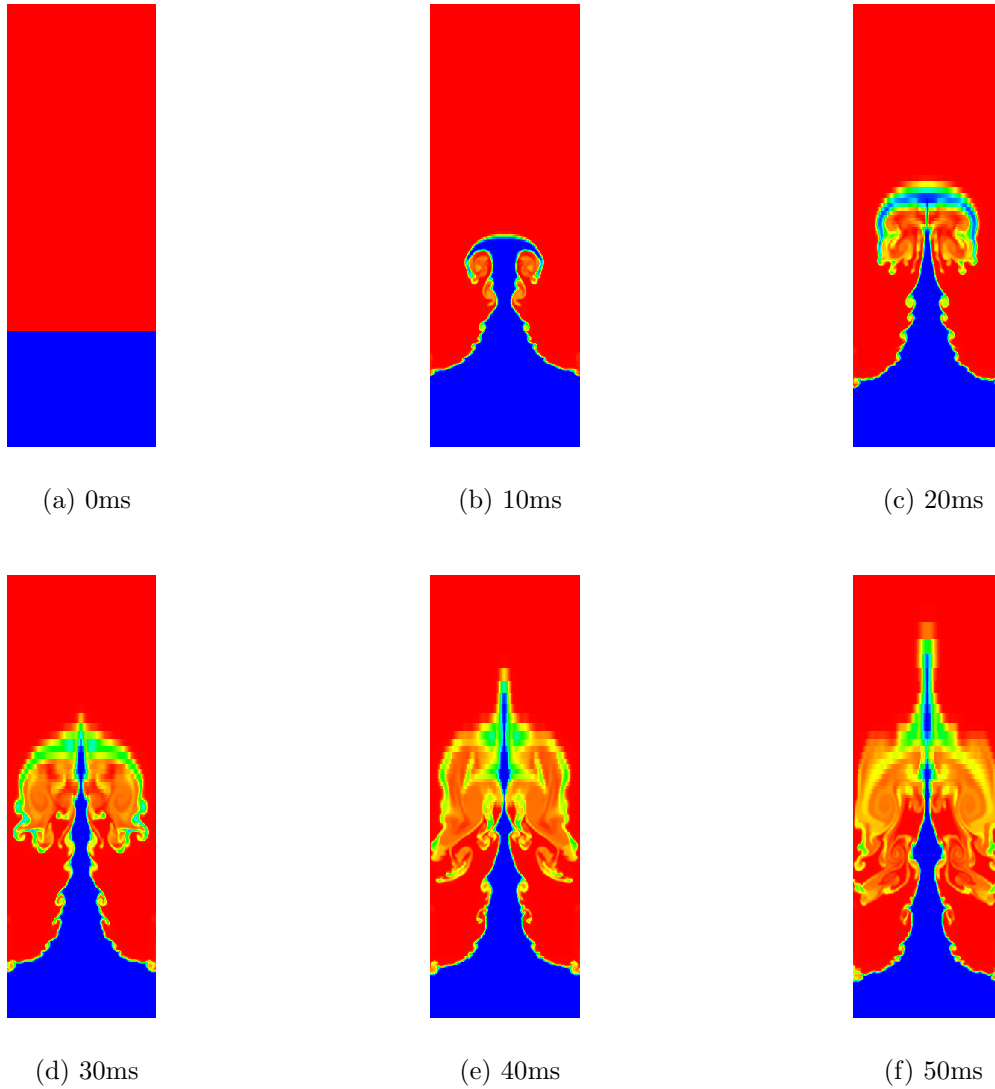


Figure 3.14: Volume fraction plots of Atwood No. of 0.692 over non-dimensional time $\approx 0 - 50$: SF6 is at bottom(in blue), Air is at top (in red)

present. The growth compared to the experimental results in Fig. 3.15 shows the simulation is physical in the region of interest but further work would need to be done to gain confidence in the late time results.

Although the vortex model has not been designed for high Atwood number cases it is possible that it could work in principle, should the vortices be shown to still drive the flow. One thing of note is that the bubble has not been seen to stagnate at any point, unlike the prediction of the vortex model, however the shape of the equation is still promising and it may simply require further consideration. This is illustrated with the final case of low Atwood number

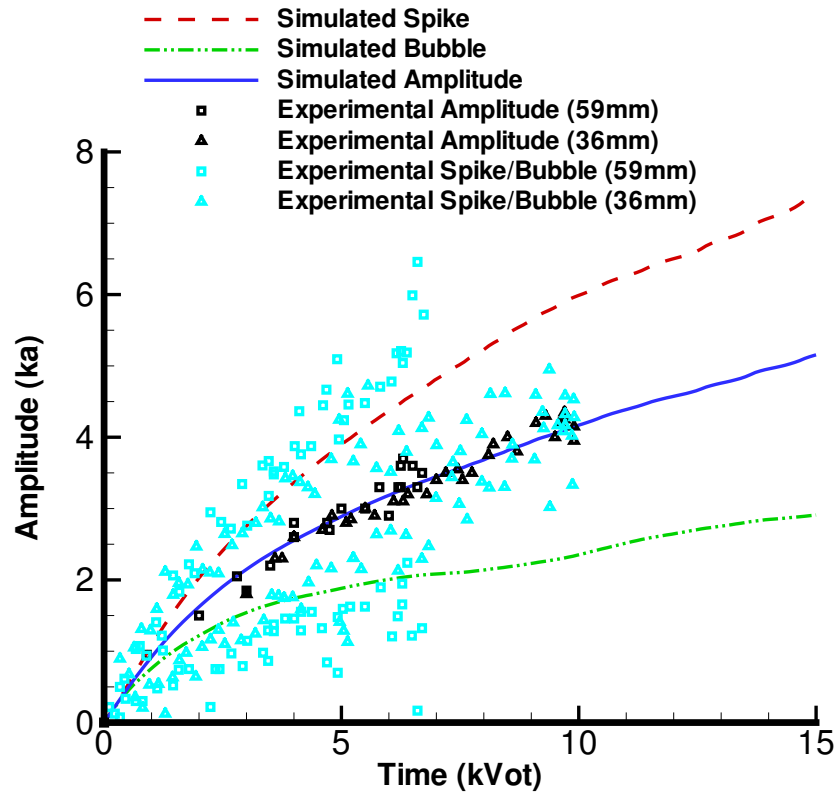


Figure 3.15: Simulated growth of instability at Atwood No. of 0.692 compared to experiments at different wavelengths

flow. Using the same setup as previously and taking the Atwood No. of 0.155 as in the original experiments on which the vortex model was created the flow was simulated as shown in Fig. 3.16.

A similar flow to the zero-Atwood No. case is seen initially, before the spike starts to grow in preference and the flow becomes more complex. Again there is strong mixing in the vortices. The late time behaviour here is somewhat unexpected as well, with a pair of counter-rotating vortices forming and the flow being pushed back towards the bubble (Fig. 3.17). Again it is unclear whether this is physical or an effect of the grid. It is also possible there are weak acoustic waves travelling through the domain which have not been entirely eliminated. The earlier times are well-behaved and as can be seen in Fig. 3.18 the features follow the experimental results well.

There is a fair degree of spread within the experimental results, which has previously led to the observation that some theoretical and empirical models seem to capture the correct behaviour. In comparison with the simulated results to later times this is not the case. The vortex model is presented

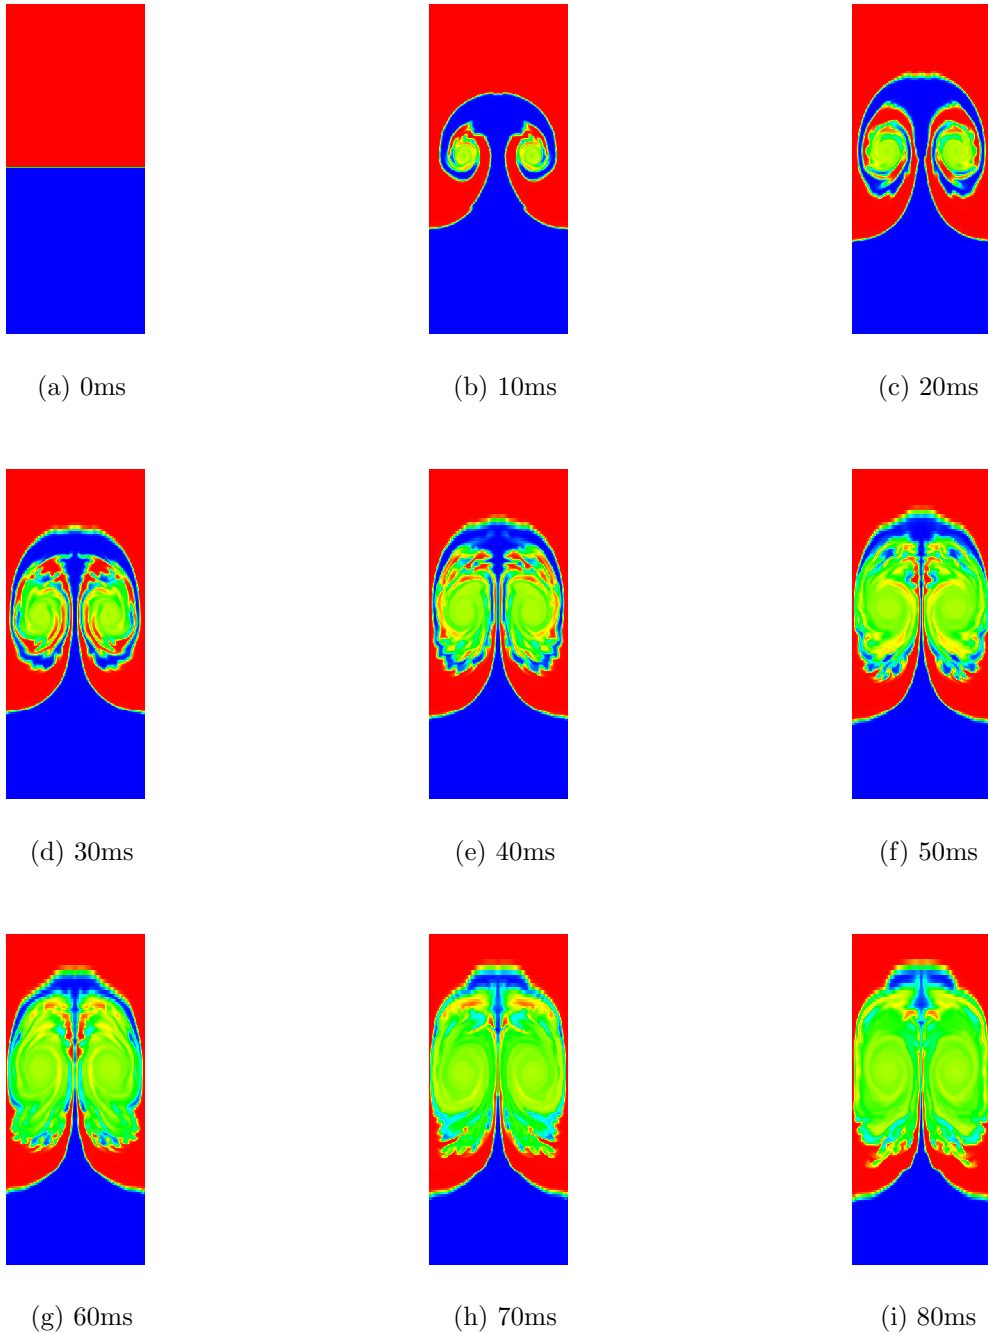


Figure 3.16: Volume fraction plots of Atwood No. of 0.155 over non-dimensional time $\approx 0 - 80$ SF6 is at bottom (in blue), Air is at top (in red)

with a finite value of ϵ and adjusted time scaling to allow for different γ as suggested earlier. This does give some room to fit the curves but no clear basis for determining these parameters has yet been found. There is a third consideration, in that a finite time must pass for the centre of vorticity to move




Figure 3.17: Volume fraction plots of Atwood No. of 0.155 at late time

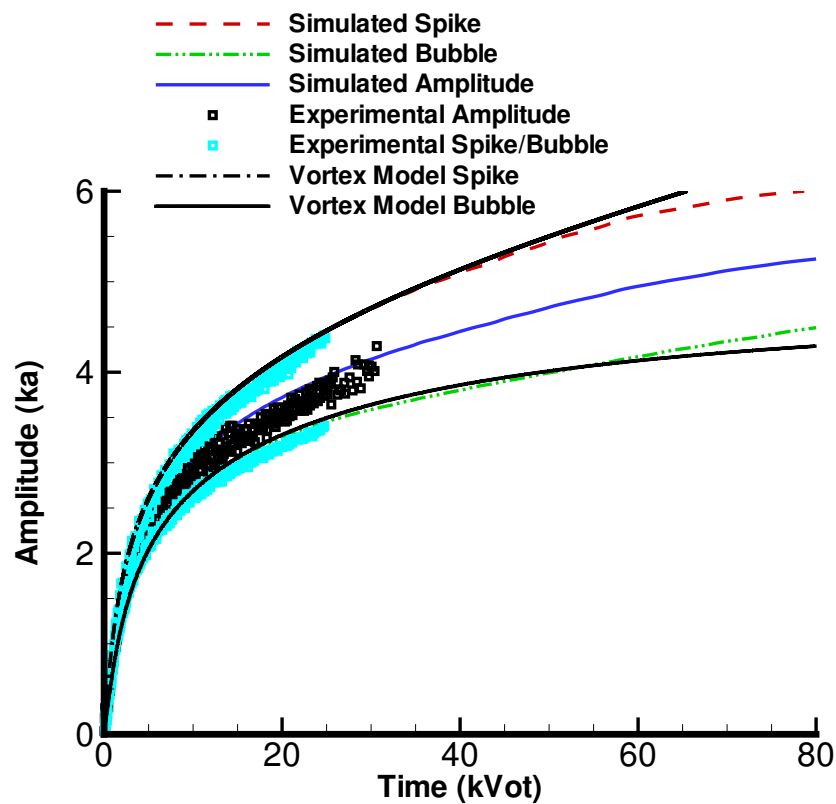


Figure 3.18: Simulated growth of instability at Atwood No. of 0.155 compared to experiment and models

from the initial points equidistant between spike and bubble to their displaced positions. It is difficult to identify the core so precisely from the simulations as to be able to confirm what the correct value for ϵ might be, and the adjusting of the strength is only approximate based on that displacement. It is difficult

to know at this point whether further work on the sophistication of this model will make it more credible but it would appear worth pursuing. It should also be noted that the overall amplitude predicted tends to be very good as the discrepancies in bubble and spike calculations are well-matched.

3.5 Conclusions

It has been seen that simulating the single-mode Richtmyer-Meshkov Instability is by no means simple. There are many factors to be taken in to account to accurately model the experiment, however the large-scale behaviour is well captured. The higher-order methods again prove their worth in resolving small-scale detail, though whether it is physical is harder to say. The theoretical growth models do not relate well to the simulated results, particularly at later times, although the Vortex Model appears to have potential.

4

Homogeneous Decaying Turbulence

4.1 Background

Turbulence remains the greatest challenge in fluid modelling. While stable and accurate capturing of shocks and interfaces is important for the development of mixing, most practical applications will involve low mach number turbulence and so it is important to test the Godunov method in this respect. Simulating the decay of a homogenous cube of turbulent flow is a classic test case and appropriate progression into three-dimensional problems. Much work has been done both experimentally and with Direct Numerical Simulation on such cases [64], against the various theories that have been developed regarding turbulence - notably the well-established work of Kolmogorov [49, 50]. Results have been mixed, with no real consensus on many key parameters. It is instructive therefore to see where ILES falls within the range for a number of these measures of flow, however the main concern as ever is to compare the different methods of reconstruction within the code.

4.2 Turbulence and initialisation

It has never been easy to classify turbulent flow, often it is said we know it when we see it. What we see then is a mass of chaotic eddies of magnificent complexity. The ways to measure and assess the characteristics of such a flow will be presented in the next section but it is first necessary to consider how to initialise an ideal case with no natural transition from laminar flow.

It is prudent to take advantage of work already done [80] and use the same method established in other simulations to facilitate comparison of the results. In this case the method used was derived by Youngs [93, 94]. As we are using a compressible code we need to consider this in the initialisation. By constructing the velocity components from a vector potential,

$$\mathbf{U} = \nabla \times \mathbf{A} ,$$

we guarantee the initial field is non-divergent. The amplitude and phases of the vector potential are set with a Gaussian distribution and scaled such that the mean turbulent kinetic energy is 0.5 and the initial mean turbulent Mach number is 0.1. This does not guarantee the flow will remain incompressible however the previous work [80] has shown that the compressible effects are negligible.

The initial kinetic energy spectrum is given by

$$E(k) = u'^2 \frac{k^4}{k_p^4} \sqrt{\frac{8}{k_p^2 \pi}} e^{-2(\frac{k}{k_p})^2} ,$$

as explained in [78], where k is the wavenumber and k_p defines the peak in the energy spectrum. To maintain comparability with the previous work $k_p = 4$.

4.3 Measuring turbulence

There are a number of quantities we can extract from turbulent flow for which theories have been put forth. For some there is no definite agreement and they only give a general indication as to whether a given flow is homogeneous and turbulent. Others allow a simple comparison between the methods to see which is resolving more of the flow. This test case has been run on 4 mesh resolutions: 32^3 , 64^3 , 128^3 and 256^3 . The simulations have also been carried out using the Minmod (MM) and van Leer (VL) second order limiters, along with the fifth order MUSCL (M5) and the WENO fifth (W5) and WENO ninth (W9) order reconstructions. The domain size remains fixed at a 1m cube and a non-dimensional time of $t = 5$ was reached to allow adequate statistics to be taken. This is approximately 8 eddy turnover times.

4.3.1 Isotropy

The first thing to assess is the claim of isotropy. This is one of the assumptions of turbulence theory and if not met then the subsequent analysis will be flawed. By using the longitudinal and transverse energy spectra to calculate separate




Figure 4.1: Isosurfaces of vorticity in the 256 cube at early times

estimates of the integral length for a given direction and then comparing the results, a measure of the level of isotropy can be achieved. These two measures of integral length are given by

$$L_1 = \frac{\pi}{u_{rms}^2} E_{11}(k=0) ,$$

$$L_1 = \frac{2\pi}{u_{rms}^2} E_{22}(k=0) ,$$

where

$$E_{11}(k_1) = \frac{1}{\pi} \langle u_1^2 \rangle \int_0^\infty dx_1 f(x_1) \cos k_1 x_1 ,$$

$$E_{22}(k_1) = \frac{1}{\pi} \langle u_1^2 \rangle \int_0^\infty dx_1 g(x_1) \cos k_1 x_1 ,$$

and f and g are the longitudinal and lateral second-order correlation functions respectively. A study of the numbers show that while all the simulations begin as essentially isotropic, anisotropy grows with time. This is particularly noticeable on the coarsest grid which had as much as a 30% difference in the two integral length scales. This rapidly improved to within 10% on finer grids. The main exception is the ninth-order WENO which maintained good isotropy on the coarser grids but became marginally worse on the finer grids. This would indicate that the subsequent results may not fit the theory too well.

4.3.2 The Decay of Kinetic Energy

The significance of turbulence in Large-Eddy Simulations is the transfer and dissipation of energy. Figures 4.2 and 4.3 show the decay of kinetic energy over time for each cube resolution comparing the different reconstruction methods. Physically there is a period of early time while the initialised flow breaks down into smaller vortices during which no energy is lost. Once features reach the size of the sub-grid scale the energy starts to be dissipated. Using DNS would require this scale to be comparable with the smallest physical scales of turbulence, the Kolmogorov lengthscale, however LES assumes that these small scales do not affect the larger motions so long as the energy is removed appropriately from the system. The figures show quite clearly that increasing the grid resolution extends the period before that smallest resolved scale is reached as expected. There is also a clear trend with order of reconstruction - the higher order methods resolving more on the coarser grids than the traditional second-order methods. This difference is illustrated in figure 4.4 where more energy is maintained for longer in the ninth-order WENO simulation of the 64 cube than in the second-order minmod simulation of the 256 cube. By this measure better resolution can be achieved with the higher-order method at a fraction of the cost. It can also be seen that the form of the decay is comparable.

It is commonly assumed that the decay of energy can be written as a power law,

$$\mathbf{K.E.} = A(t - t_0)^{-p} ,$$

where A and t_0 are constants and p is the determining exponent. The figures shown do not appear to fit this trend at higher resolution however we can still determine an approximate value for the decay exponent p . There is a good deal of uncertainty regarding the constants in the expression and so a non-linear regression script in Mathematica has been used to try and fit the data over the appropriate range where the turbulence has been established. The

results are shown in Table 4.1. There is no fixed value from theory although it is expected to be in the region of $p \approx 1.2 - 1.4$. Experiments on turbulence in wind tunnels has measured decay exponents $p \approx 1.2 - 1.3$. Based on this it is fair to say that the 128 and 256 grids produce very convincing results, with no great difference between the reconstruction methods. At the coarsest grid the more dissipative methods have a notably greater decay rate as expected, however all the results lie in the physically realistic range.

Resolution	MM	VL	M5	W5	W9
32	1.78	1.58	1.62	1.50	1.44
64	1.37	1.42	1.46	1.03	1.01
128	1.20	1.13	1.18	1.27	1.23
256	1.25	1.17	1.29	1.17	1.21

Table 4.1: Estimate of exponent p in decay of kinetic energy for all simulations

A closely linked quantity to the energy decay is the enstrophy. As the initial flow field breaks down into smaller and smaller vortices the enstrophy increases until the critical point where dissipation begins to dominate. All but the coarsest simulations exhibit this peak (Figs. 4.5 & 4.6) which corresponds to the commencement of decay in the previous figures. Again there is a clear trend in method of reconstruction which can be made analogous to increasing grid resolution. The less dissipative (higher-order) schemes have a greater peak enstrophy before losing energy. It is to be expected that the entropy calculated will increase as the grid resolution increases and smaller scales are captured. This can be seen in figure 4.7 for the ninth order WENO reconstruction.

4.3.3 Energy Spectra

Taking the analysis a step further, it is worth looking at the kinetic energy spectra. This is where Kolmogorov formalised the concept of an energy cascade proportional to $k^{-5/3}$. We define the three dimensional energy spectrum as

$$E(k) = 2\pi k^2 \phi_{ii}(k) ,$$

where

$$\phi_{ij}(k) = \frac{1}{(2\pi)^3} \int_{-\infty}^{\infty} Q_{ij}(r) e^{-ikr} dr ,$$

and Q_{ij} is the second order velocity correlation tensor. The energy spectra for each simulation at late time are shown in figures 4.8 & 4.9. Previous work

did not correct for the low mach number effects, resulting in a severe drop-off at high wavenumbers which is now significantly changed. As expected the amount of energy in the smallest scales (high K) increases with increasing order of reconstruction. However the WENO methods do not show a simple cascade, particularly on the coarser grids. Instead energy drops off sharply while piling up at the smallest scales. The MUSCL methods are all largely well-behaved and are arguably following a trend in the region of $k^{-5/3}$ over much of the range, though rarely matching it exactly. Figure 4.10 shows a comparison of the two fifth order methods on the 256 cube which actually shows a better agreement with theory from the WENO method up to the kink which seems prevalent on all grids and for all times after the cascade is established. This is a difficult case to draw too many conclusions from this quantity as it is the large scales we are ultimately interested in, however it does offer a way of looking at how the energy is distributed within the flow.

These figures do however suggest a very significant consideration, that in certain cases the high-order methods have insufficient dissipation. This would seem to indicate that there is an optimal order of accuracy for the reconstruction and that it is possible to go too far in trying to reduce dissipation. More concerning is the implication that the level of dissipation is not fixed for a given method, and that a scheme which is overly dissipative in solving one problem may not dissipate enough energy in another. This may in fact point towards a different cause. The excessive energy at high wave numbers only appeared when the higher-order methods were run in conjunction with the low mach number correction. The turbulent cube only contains low mach number flow, for which the correction provides a near average of the left and right states. This reduces the role of the reconstruction and the Riemann solver, making it difficult to provide a mechanism for dissipating the energy at the small scales. The low mach theory is incomplete for the higher-order methods, and it may be that this is an area which needs to be developed to address these results.

4.3.4 The Structure of Turbulence

There are certain characteristics of turbulent flow that can be observed. To illustrate this, two-dimensional cross-sections of the vorticity distribution have been plotted for each case. Figure 4.11 shows the results at early time for the 32 cube. The red areas show regions of high vorticity as the vortex tubes intersect the plane, while the blue areas have very little vorticity. Although not much time has elapsed from the initial condition, the difference in both the magnitude and occurrence of vorticity between the different reconstruction methods is clearly visible. Increasing the grid resolution (Figs. 4.12, 4.13) allows for greater detail and smaller, more concentrated vortices. It is interesting to note that the fifth-order WENO scheme is more visibly comparable

to the ninth-order WENO than the fifth-order MUSCL, while the two MUSCL schemes represented share similarities. At the finest resolution (Fig. 4.14) this distinction disappears as the ninth-order WENO has substantially larger peak values of vorticity than the other methods. Nonetheless the pattern is comparable when one considers the sparse regions on each plot, one of the properties of turbulence.

Velocity correlations are commonly used to statistically determine if a flow is turbulent. Although initialised with a Gaussian distribution it is expected that the third-order correlation, or skewness, will fall to approximately -0.33 and the fourth-order correlation (kurtosis) will reach about 4. The correlations are defined as

$$S_n = (-1)^n \frac{\langle (\partial u / \partial x)^n \rangle}{\langle (\partial u / \partial x)^2 \rangle^{n/2}}$$

There is no absolute agreement on what these values should be for turbulent flow. DNS simulations have consistently produced higher results than other sources, with $S_3 \approx -0.5$ and $S_4 \approx 6$. Conventional LES has produced results with typically $S_3 \approx -0.4$ and $S_4 \approx 3 - 4$ while experiments have a similar range but tend to measure $S_3 \approx 0.34$. The results for the current simulations are shown in Table 4.2 for the third order function and Table 4.3 for the fourth order function. There is a steady trend as the grid size is increased as both values grow, indicating more developed turbulence. There is little difference between the reconstruction methods except for the ninth-order WENO which is always significantly lower than the other methods. While not following any obvious explanation, it is not unexpected given the discrepancies seen with this method in the energy spectra and measurement of isotropy. It would be reasonable to surmise this is not proper turbulence, however the decay of energy and enstrophy fit the expected pattern very well in comparison to the other methods. What may be being seen is the insensitivity to small scale details of the overall flow.

Resolution	MM	VL	M5	W5	W9
32	-0.26	-0.24	-0.23	-0.20	-0.17
64	-0.28	-0.29	-0.28	-0.29	-0.20
128	-0.30	-0.32	-0.30	-0.32	-0.24
256	-0.33	-0.31	-0.31	-0.32	-0.26


Table 4.2: Estimate of third order structure function for all simulations

Resolution	MM	VL	M5	W5	W9
32	3.17	3.25	3.18	3.17	3.10
64	3.44	3.48	3.45	3.52	3.30
128	3.66	3.66	3.71	3.71	3.48
256	4.03	4.04	4.08	4.04	3.82


Table 4.3: Estimate of fourth order structure function for all simulations

4.4 Conclusion

While quantifying turbulence is difficult, there is clear indication that ILES can capture the essential physics. Good agreement with experiment and other conventional simulations on the kinetic energy decay rate and structure functions support this view. The energy spectra present a mixed picture of the distribution of energy but are notable for a lack of any drop-off at high wavenumbers. They do raise issues concerning the importance of correctly capturing the energy cascade for simulating the large-scale motion. The development of enstrophy and overall dissipation of energy over time show clear trends that support the notion of higher-order methods resolving features ordinarily only captured on finer grids.




(a) 32




(b) 64

Figure 4.2: Comparison of decay of kinetic energy for 32 and 64 cubes with varying reconstruction method



(a) 128



(b) 256

Figure 4.3: Comparison of decay of kinetic energy for 128 and 256 cubes with varying reconstruction method






Figure 4.4: Comparison of decay of kinetic energy between low and high order methods




(a) 32




(b) 64

Figure 4.5: Comparison of development of enstrophy for 32 and 64 cubes with varying reconstruction method



(a) 128



(b) 256

Figure 4.6: Comparison of development of enstrophy for 128 and 256 cubes with varying reconstruction method






Figure 4.7: Comparison of development of enstrophy for ninth order WENO on all grids




(a) 32




(b) 64

Figure 4.8: Comparison of kinetic energy spectra for 32 and 64 cubes with varying reconstruction method



(a) 128



(b) 256

Figure 4.9: Comparison of kinetic energy spectra for 128 and 256 cubes with varying reconstruction method







Figure 4.10: Comparison of kinetic energy spectra for fifth order methods on 256 cube




(a) VL



(b) M5





(c) W5




(d) W9

Figure 4.11: Cross-section of cube showing vorticity for 32 grid with different reconstruction methods







(a) VL



(b) M5





(c) W5




(d) W9

Figure 4.12: Cross-section of cube showing vorticity for 64 grid with different reconstruction methods







(a) VL



(b) M5





(c) W5




(d) W9

Figure 4.13: Cross-section of cube showing vorticity for 128 grid with different reconstruction methods







(a) VL



(b) M5




(c) W5



(d) W9

Figure 4.14: Cross-section of cube showing vorticity for 256 grid with different reconstruction methods



5

Three-dimensional Multi-mode Richtmyer-Meshkov Instability

5.1 Introduction

In practical application the Richtmyer-Meshkov instability has multiple modes generated from a random perturbation of the given interface. These modes grow at different rates before breaking down into a turbulent mixing zone. Successful simulation of such problems requires good capturing of the smaller modes to resolve the growth of the mixing region correctly, as well as being able to simulate the resulting turbulent flow. In some cases the shock wave is reflected to ‘reshock’ the flow, creating even more complexity and requiring good shock-capturing. The first section considers essentially planar problems with random small-scale perturbations, while the second section looks in more detail at simulating problems with an imposed large-scale ‘perturbation’.

5.2 Planar RMI

The simulations in this chapter make use of a random small-scale perturbation on the interface to generate multiple modes for the Richtmyer-Meshkov instability. Further details can be found in work by Thornber [78, 79]. Modes are excited in accordance with a power spectrum inversely proportional to the square of the wavenumber using a method devised by Youngs [95]. While this

offers a convenient level of realism in modelling experimental flows it was originally developed to consider as the sole source of instability in an idealised planar case.

5.2.1 Group work




Figure 5.1: Isosurfaces of volume fraction after development of the multi-mode planar Richtmyer-Meshkov instability for the original and modified WENO 9th-order scheme

Much work has been carried out elsewhere in the group on the basic planar multi-mode Richtmyer-Meshkov problem to consider the effect of the initial condition. It has not been investigated fully in the context of the present research, however some simulations have been run which demonstrate the developments implemented in the WENO reconstruction of the code. Figure 5.1 shows the development of the flow some time after the shock has passed for the original WENO implementation. Individual ‘mushrooms’ can be seen growing from the initial interface from the dominant modes as seen in the single-mode simulations. The figure also shows the same point in time for the simulation with the modified WENO reconstruction. Rather than long, thin spikes there is more development in the vortices and the flow in the turbulent region is visibly more mixed with greater fine-scale detail.

It is difficult to draw too many conclusions on the physical reality from such representations although it is something of a truism that resolving more small scales is the ambition of any Large-Eddy simulation. We can instead look a bit more at the statistics of the flow in line with what was done in the simulations of decaying turbulence. Figure 5.2 shows the later time spectra of




Figure 5.2: Development of the energy spectrum (m^3s^{-2}) against K (m^{-1}) over time for the multi-mode planar Richtmyer-Meshkov instability with the improved WENO 9th-order scheme set against the $k^{5/3}$ line





Figure 5.3: Comparison of energy spectrum (m^3s^{-2}) against K (m^{-1}) for original and improved WENO 9th-order schemes at a given time for the original planar multi-mode case


the kinetic energy with it clearly falling into a strong cascade as expected in theory. Indeed this is a stronger case in some respects for assessing turbulent behaviour and none of the undesired behaviour seen in the decaying turbulence results can be found here. The comparison between the original and improved

WENO reconstructions is visible in figure 5.3, confirming that the improved scheme redistributes more energy to the smaller scales in what appears to be a physical manner in comparison with Kolmogorov theory.

5.2.2 The experiments of Vetter and Sturtevant



(a) Large initial amplitude



(b) Small initial amplitude

Figure 5.4: Isosurfaces of volume fraction after development of the Richtmyer-Meshkov problem of Vetter and Sturtevant for differing initial perturbation (dimensions in metres)

Finally the development of the energy spectra after the reshock (Fig. 5.6)




Figure 5.5: Growth of the mixing layer in Richtmyer-Meshkov problem of Vetter and Sturtevant for large initial perturbation (Amplitude (m), Time (s))

shows similar behaviour to that seen previously. The flow has developed identifiable turbulent properties.

While a useful test case for simulations, the previous problem did not allow for any real experimental validation. This is a standing difficulty as little experimental work has been carried out to the requisite standard to be of use in multi-mode simulation. One such experiment which has attracted a number of simulated comparisons [38, 11] is the work of Vetter and Sturtevant [88]. Unfortunately the bulk of their work involved a set up in which the membranes maintaining the interface prior to the arrival of the shock disrupted the resulting flow, typically suppressing growth of instabilities. There appear to have been a number of simulations which do not allow for this and yet claim good results - an indication of the excess level of dissipation present in many codes. In this experiment, the membrane was held in place by a fine wire grid. It is generally assumed that the grid size determines the small-scale perturbations, but that there is also a larger-scale deflection of the grid which ought be considered. These details are important as they significantly affect the result. Assuming the real perturbation to be too small to resolve, previous simulations have approximated the problem with larger perturbations. The results of this using our code are shown in Figure 5.4. The initial interface is shown along with the flow shortly after the reflected shock has passed through. The impact of considering a more realistic perturbation size is clear and illustrates




Figure 5.6: Development of energy spectrum (m^3s^{-2}) against K (m^{-1}) after reshock in Richtmyer-Meshkov problem of Vetter and Sturtevant against the $k^{5/3}$ line

the importance of correctly capturing the details of any given experiment in validation work.

There was one experiment carried out with a different rearrangement of the membrane to effectively remove it as a consideration. What the few results reported show is a significant difference to the other experiments. While the images are not terribly clear, the reported growth of the mixing layer is available for before and after the ‘reshock’. The comparison with simulation is shown in figure 5.5. Typically there is little correlation expected prior to reshock, however if the initial perturbation is resolved correctly at a small enough scale there is no obvious reason this cannot be simulated accurately. Although the data is sparse, the growth rates appear comparable and further work on the initial amplitude of the perturbation ought improve the correlation.

5.3 Large-scale Perturbation

It is of interest to properly investigate how the addition of a large-scale feature to the multi-mode problem affects the flow, and there is good experimental work and other simulations to compare results with.

5.3.1 Double-bump




Figure 5.7: Initial condition for the ‘double-bump’ problem

This case is based on experimental work [40] carried out at the Atomic Weapons Establishment (AWE). The ‘double-bump’ initial condition (Fig. 5.7) consists of a slab of SF6 with air either side. The shock passes from left to right in the figure before reflecting off the end wall and returning through the SF6 region. Both interfaces have a small-scale random perturbation in line with that used in the planar multi-mode problems, while the right-hand interface has two bumps as shown.





Figure 5.8: Averaged 3D simulation 160x80x40 for the ‘double-bump’ problem


Two-dimensional simulations carried out on a number of different grids have been compared to experiment in figure 5.9. The salient features of this flow are the two inverted ‘bumps’ which become large recirculation regions and the central jet which is ejected from the flow. These figures are not all at the exact same time, accounting for slight discrepancies, but the effect of improving grid resolution is clear. The turbulent regions are more developed and well mixed - mixing is typically the property of interest in these problems.




(a) Experimental 2ms




(b) Experimental 4ms




(c) 80x40 2ms




(d) 80x40 4ms




(e) 160x80 2ms



(f) 160x80 4ms



(g) 320x160 2ms



(h) 320x160 4ms

Figure 5.9: Density comparison of different grid resolutions at approximate times for ninth-order WENO simulation with experiment




Figure 5.10: Three-dimensional representation of development of flow density in ‘double-bump’ problem

Turbulence only exists in three-dimensions so it is necessary to simulate the third dimension to get a physical measure of the mixing region (Fig. 5.8). The two-dimensional simulations are appreciably asymmetric, however looking at the region of 3D flow, the results are symmetric on average. Figure 5.10 gives an indication of how the random perturbations manifest themselves across the third dimension. This problem has been considered more thoroughly in [31]; the same methodology applies to the second case-study, the ‘inverse-chevron’, which is discussed in more detail here.

5.3.2 Inverse Chevron




Figure 5.11: Initial condition for the ‘inverse-chevron’ problem

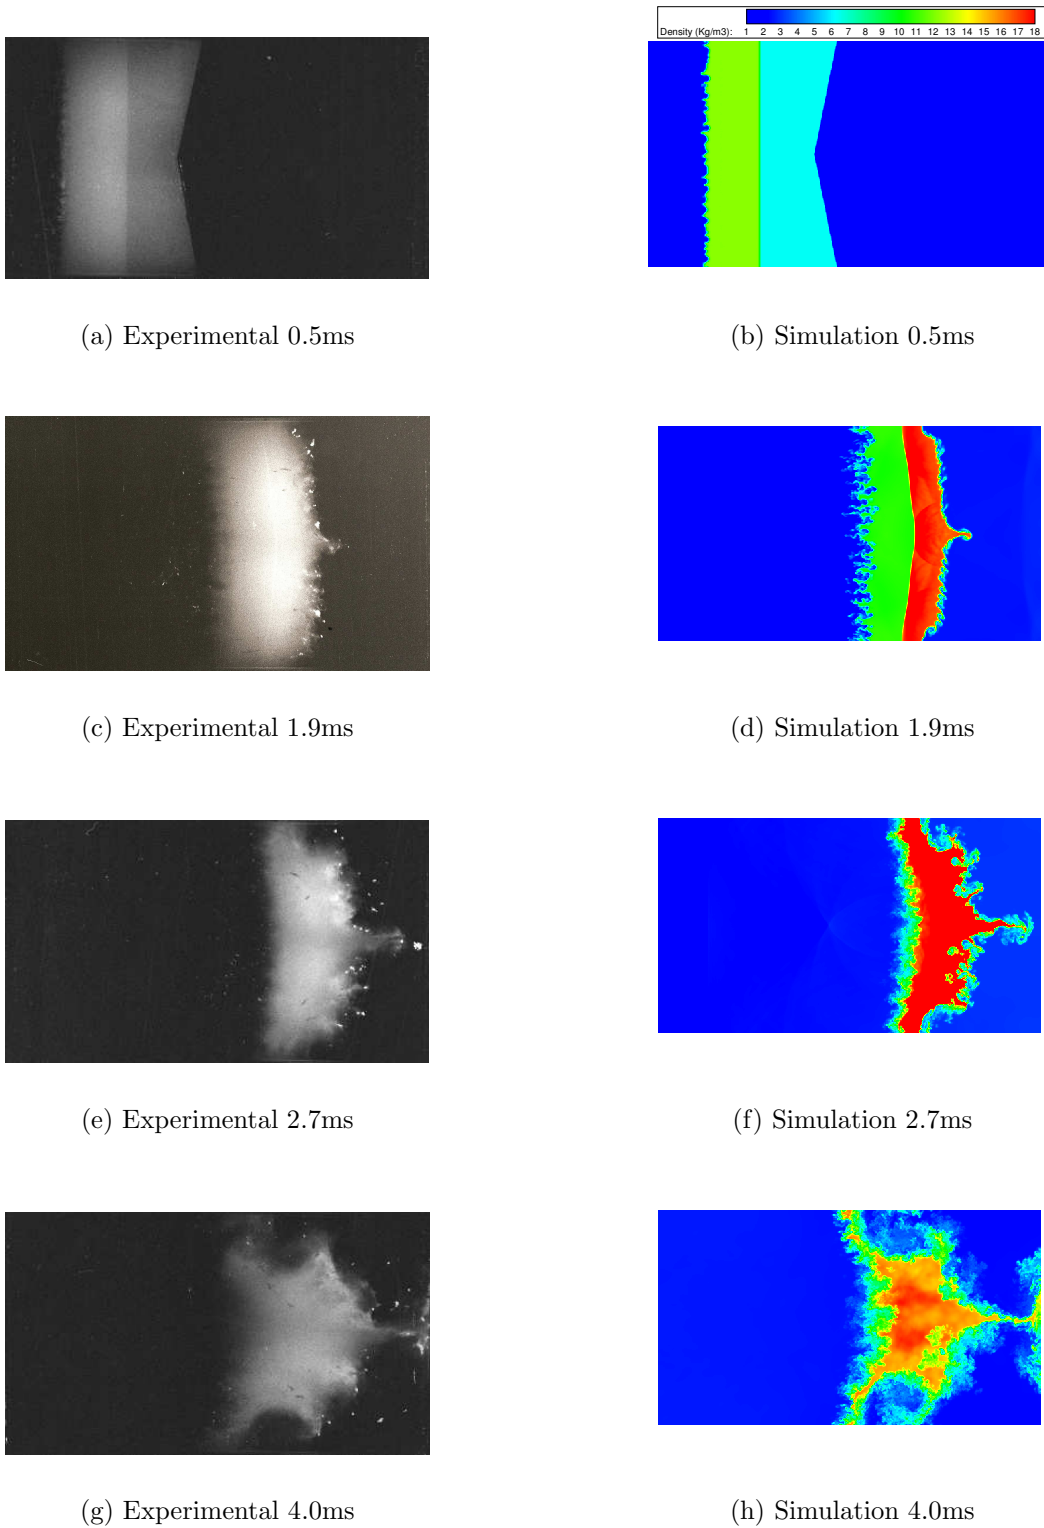


Figure 5.12: Development of density of experiment and ninth-order WENO simulation

The most recent work on the multi-mode problems using the latest code developments involves the ‘inverse chevron’ perturbation (Fig. 5.11). The experimental set-up is the same as for the double-bump, but with the chevron replacing the two bumps as the large-scale feature.

The development of the flow is shown in detail in figure 5.12 for the experiment alongside two-dimensional slices of the most-resolved simulation. The experimental pictures are taken with a laser sheet and are subject to scattering which gives a partially averaged view through the depth of the domain. The most notable feature of this flow is the central jet accelerated towards the end wall. There are also two main ‘bubbles’ identifiable at either side. These features shall form the basis for visual comparison.

Besides direct comparison with the experiment, this work was intended to be an inter-code validation. The results presented here include TURMOIL, a semi-lagrangian code, provided by AWE [94]. There will also be comparison with results from Bates [4]. Figure 5.13 shows the plot of density at the point where the reflected shock is passing back through the SF6 region. The first requirement of the simulations is to correctly place the shock, and this has been achieved. Further qualitative comparison is provided in figure 5.14 where contour bands of the volume fraction of SF6 have been visualised at the end of the simulation showing the same general flow structure.




Figure 5.13: Comparison of density at 1.9ms between codes

A more interesting quantity to consider is the turbulent kinetic energy. The distribution of this shown in figure 5.15 for our code shows how the activity is concentrated in the two side bubbles and the jet. Comparing this with TURMOIL results (Fig. 5.16) shows a similar pattern, though the position




Figure 5.14: Comparison of volume fraction of SF6 at 4ms between codes : with 0.1, 0.5 and 0.9 contours

within the side vortices where the turbulent energy is most concentrated is notably different. This may lead to differences in development of the flow at even later times however it may prove to be not significant as the TURMOIL code is known to over-estimate turbulent kinetic energy when under-resolved.




Figure 5.15: Distribution of turbulent kinetic energy in CNS3D simulation at 4ms




Figure 5.16: Comparison of turbulent kinetic energy at 4ms between codes : with 25, 50, 100, 200 kgm^2s^{-2} contours

It is desirable to gain a more quantitative comparison of how the codes compare. Figure 5.17 shows the growth of the features identified earlier in comparison with the experiment for the three codes discussed. The jet position represents the distance from the end-wall, and as such decreases until the point where the jet impinges on the wall. The jet width meanwhile slowly increases until that point where it spreads rapidly along the end-wall. The bubble position is harder to capture and is taken as an average of the two bubbles. All the simulations provide similar results, although there are spurious points picked up by the output algorithm in CNS3D. There is some distinct deviation from the experimental results however this is to be expected as these results are simply measured off the photographs which struggle to show full intensity near the wall and do not adhere to strict contour levels for measurement.

As stated previously, mixing is the fundamental consideration for much of this work. The measured mixing parameter is given by

$$M = \int \rho^2 m_1 m_2 dV,$$

where m_1 , m_2 are the mass fractions of SF6 and air respectively, ρ is the total density and the product is integrated over the whole volume. This parameter is plotted for each code (Fig. 5.18). The passage of the shock is visible as sudden jumps in the level of mixing. There are two interesting points to note from this figure. Most importantly is that the level of mixing at the end of the simulation is essentially identical for all three codes. This indicates they are all capturing the essential large-scale flow well-enough in comparison




Figure 5.17: Development of feature positions over time for different codes

to each other. The other point to note is the early stages where the finite-volume codes exhibit greater mixing due to numerical diffusion than the semi-lagrangian TURMOIL. It is a good illustration of how measured mixing does not guarantee well-resolved turbulent flow.

Considering just CNS3D we can see the effect of grid resolution and reconstruction method on this parameter. The early time numerical diffusion is largely dependant on the grid size as expected, while the late-time level of mixing depends on resolution of the small scales. This is most appreciable when considering the low mach number correction - without which even the ninth-order scheme struggles to resolve enough turbulence to give what appears to be the converged level of mixing. The fifth-order MUSCL scheme requires the finer grid to reach this level, whereas the improved ninth-order WENO scheme appears to resolve the flow equally well on the coarser grid.

The overall turbulent kinetic energy in the domain is a useful measure of how much energy is being dissipated. At this grid resolution both the AMR




Figure 5.18: Development of mixing over time for different codes

code of Bates and CNS3D have converged, whereas TURMOIL overestimates the value as noted earlier. Previous work on similar problems has shown that the turbulent kinetic energy in TURMOIL continues to fall as the grid is refined well beyond the current size. Simulations on a coarser grid with CNS3D indicate this is a converged solution, as with the mixing.

The final parameter considered here is enstrophy (Fig. 5.21). This quantity has no converged value, and so is expected to increase as the grid resolution increases. As in the case of the homogeneous decaying turbulence, high levels of enstrophy indicate many small energetic vortices, and as such can be used as a measure of resolution of small scales for a scheme. The figure shows a significant difference between the codes, suggesting that the ninth-order WENO scheme in CNS3D captures far more smaller scales than Bates' second-order AMR method, as expected. It is worth noting that the amount of enstrophy can vary greatly, as evidenced by observations that the second-order AMR result is comparable with second-order results in CNS3D. Figure 5.22 gives an




Figure 5.19: Effect of grid resolution and reconstruction method on mixing

indication of how differing grid sizes and reconstruction methods can affect the amount of enstrophy measured, again illustrating the significance of the low mach number correction in the code.

5.4 Conclusions

Current ILES codes can simulate more complex multi-mode problems which compare well when experimental data is available. The results presented support the thesis, in that highly-accurate reconstruction provides significant and physical improvement in capturing small scales in turbulent and unstable flows. There is also an indication supporting the argument for LES as a whole that resolution of the small scales does not necessarily affect the large-scale structure and development of the flow.




Figure 5.20: Development of turbulent kinetic energy over time for different codes




Figure 5.21: Development of enstrophy over time for different codes




Figure 5.22: Effect of grid resolution and reconstruction method on enstrophy

6

Hypersonic Problems

6.1 Introduction

There has been much work, both experimentally [10, 90, 39] and computationally [36, 28, 21, 57, 91, 32, 63], in the past eight years to identify the factors affecting the complex flows that surround hypersonic objects. In particular, it is non-trivial to simulate both the strong shock interactions and boundary layer separation. This is a comparable problem to that of solving compressible turbulence, although the flow regime is significantly different and other elements have to be considered. Most notably the time integration needs to be considered in order to make the problem stable. CNS3D has been applied to two hypersonic flow configurations and associated test cases are presented in the first section 6.2. The results from the simulations are then compared with experimental data in section 6.3.

6.2 Methodology

Several cases have been simulated in the hypersonic regime under assumptions of steady flow and ideal gas behaviour. Both of these assumptions are currently open to debate, and there is a growing body of work, including work within our own group, addressing the extent to which observed results can be better understood by adopting more realistic models; however, this is beyond the scope of this thesis. In both cases, the wall temperature is assumed to be constant.

6.2.1 Test Case: Double Cone

The 25° – 55° double-cone configuration (Fig. 6.1) was originally run at the Calspan University Buffalo Research Center (CUBRC) providing experimental data for a series of CFD code validations. The same configuration has since been tested in the Arnold Engineering Development Center (AEDC) Hypervelocity Wind Tunnel No.9 [10], and it is these flow conditions which have been simulated in the current work. It has thus far generally been considered appropriate to restrict investigations to a sufficiently low Reynolds number such that the effects of unsteadiness or turbulence are insignificant. Leaving aside the issue of whether that has been successful, the experimental runs that have been modelled here are Run 2983 and Run 2984 as in the published work. The Mach number is approximately 12.7 in both cases; Run 2983 is calculated to be at Reynolds No. 2.92×10^4 , while Run 2984 was at Reynolds No 2.03×10^4 based on the length of the first cone.




Figure 6.1: Schematic of double-cone geometry (dimensions in inches)

The features observed in the double-cone flow are well-documented. It is a challenging case, with shock-shock and shock-boundary interactions providing positive feedback via a separation zone. Typically the pressure coefficient (C_p) and heat transfer (Stanton number, St) are measured along the wall. Plotting either coefficient is sufficient to identify the point of separation, and the reattachment peak where the transmitted shock impinges on the body. Figure 6.2 shows the computational domain used for this problem.

6.2.2 Test Case: HB-2

The second geometry considered is the blunted-cone-cylinder-flare designated as HB-2 (Fig. 6.3). This geometry has been extensively used in aerodynamic test facilities. While it is not as extreme a case as the double-cone, it does form a convenient test for validation of CFD methods. There is documented




Figure 6.2: The computational mesh for the double-cone, showing every 1 in 4 grid lines

data for numerous experimental conditions, at the present time only three of these have been simulated. Case 1 is at the apparent limit of Reynolds number dependence, at Mach 5 with a $Re_D = 2.32 \times 10^6$ [47]. Case 2 is above this limit at Mach 7.5 with $Re_D = 0.13 \times 10^6$ [27]. Both these cases are run at zero angle of attack. Case 3 has the same flow parameters as Case 1 but is inclined at 15° to the incoming flow. The resulting pressure distribution is normalised by the post-shock stagnation pressure and plotted against longitudinal position. In the inclined case this pressure is reported at three circumferential stations. The grid for the inclined case is represented in fig. 6.4.

6.2.3 Test Details

In addition to running the test cases over the range of reconstruction methods, they have also been run on different grid resolutions. Table 6.1 below outlines the full details of the simulations performed for the double-cone, while Table 6.2 shows the data for the HB-2 cases. The cases have also been run in FLUENT for comparison, the outcome of which is presented in the results section.




Figure 6.3: Schematic of HB-2 geometry




Figure 6.4: The computational mesh for the HB-2, showing every 1 in 4 grid lines

6.3 Results

6.3.1 Double Cone

The flow structure typical of the double-cone geometry can be seen in the visualisation of pressure contours in Figure 6.5, illustrating the shock structure.

Table 6.1: A matrix showing the twelve test runs simulated for the double-cone geometry.

Run	2893		2894	
M_∞	12.73		12.62	
Re_L	2.92×10^4		2.03×10^4	
P_∞ (Pa)	8.07		5.04	
ρ_∞ (kgm^{-3})	5.9×10^{-4}		3.9×10^{-4}	
T_∞ (K)	46.1		42.7	
<i>Mesh</i>	128x48	256x96	128x48	256x96

Table 6.2: A matrix showing the run conditions simulated for the HB-2 geometry.

α	0°		0°		15°	
M_∞	7.5		5.0		5.0	
Re_D	0.13×10^6		2.32×10^6		2.32×10^6	
P_∞ (Pa)	210.		3748.		3748.	
ρ_∞ (kgm^{-3})	9.1×10^{-2}		5.1×10^{-3}		5.1×10^{-3}	
T_∞ (K)	138.9		138.9		138.9	
<i>Mesh</i>	128x96	64x48	128x96	64x48	128x96x32	

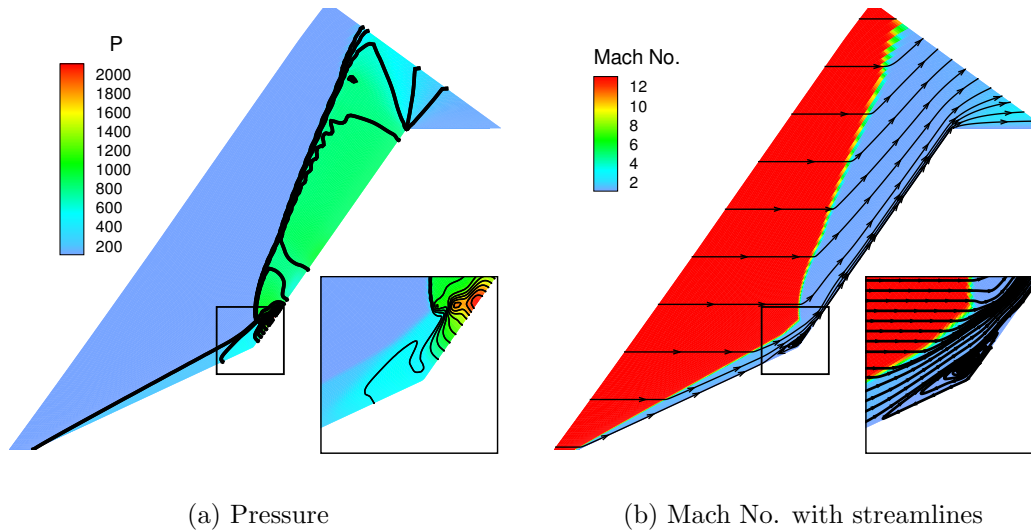


Figure 6.5: Run 2894 with WENO 9th in CNS3D.


The recirculation can also clearly be seen in the streamline visualisation. The pressure coefficient and Stanton number distributions for the runs with CNS3D and FLUENT are shown in figs. 6.13, 6.14, 6.15 & 6.16, 6.17, 6.18.

The first-order results obtained exhibit a remarkable similarity between the two codes. Increasing the order of accuracy above first-order offers a much greater improvement than doubling the resolution of the grid, and is undoubtedly cheaper. With CNS3D this improvement extends above second-order methods. Figure 6.6 shows the difference between the codes for run 2893 on the fine grid. It was previously mentioned that this flow could be characterised by the size of the separation zone and the values at the peak, and on this basis the results are clearly in favour of higher-order methods (Fig. 6.7). It is worth noting again at this point that any reconstruction can only be as accurate as the size of the local continuous stencil, a particular issue on coarser grids, and further work could be done in identifying whether a fifth or seventh order would be more appropriate in this case - previously ninth-order was found to be a suitable compromise between accuracy and computational cost.


Concerning the experimental data, it would appear that in the lower Reynolds number case grid convergence has been obtained, and a clear discrepancy between experiment and simulation exists for CNS3D (Fig. 6.8). Preliminary runs on finer grids have indicated that this is indeed the converged solution to the problem set in the simulations, although it is worth acknowledging that the measure of surface quantities, in particular the gradients for Stanton number, are dependent on the grid in a finite volume solver. It seems evident that one or more of the assumptions made in modelling this flow are insufficient. The matter has been addressed in various works in the literature, and it is difficult to hypothesise as to which assumptions are relevant - be it chemical equilibria or slip effects, unsteady flow or varying wall temperature and heating effects. Some of these would prove easier to assess than others, and there is also a potentially significant related issue of grid design. However, the high-order methods present a useful basis for further investigation.

6.3.2 HB-2

The results for HB-2 show similar trends to the double cone (Figs. 6.19), 6.20 & 6.21). Case 1 is the more complex of the axi-symmetric problems. First order of accuracy fails to capture any of the experimentally observed behaviour around the cylinder-flare transition. At second-order FLUENT seems not to be capturing any of the main separation, whereas Figures 6.9 & 6.10) clearly show that CNS3D captures the ‘double-step’ profile characteristic of a separation bubble. The reattachment point is well-matched to experiment, although the total separation length is somewhat smaller. Results have also been obtained using higher-order methods. Interestingly, the separation point moves forward to match experimental data but the reattachment point also moves. On observation of the development of the ninth-order simulation the reattachment point appears to oscillate. The average position may well be a



(a) Pressure coefficient



(b) Stanton No.

Figure 6.6: Run 2893 on 128x48 mesh

better fit to experiment and it may be the case that the axial position is not sufficiently well resolved in the grid at that point. It is possible the discrepancy is within the scope of experimental uncertainty and other test conditions could be modelled to investigate further.

Case 2 is comparatively straightforward. Viscous effects are not expected to be significant, and indeed the first- and second-order simulations all converged quite quickly. Figure 6.22 shows the normalised results against experimental data for the first-order simulations. There is little difference between either code or the experimental data. The slight discrepancies are essentially better






Figure 6.7: Comparison of separation zone with run 2894 on coarse grid (128x46) between CNS3D and FLUENT - Axial-velocity plotted with instantaneous streamlines

resolved in the higher-order simulations (Figs. 6.23 & 6.24), however, it is worth noting that FLUENT is less able to handle the coarser mesh at second-order.

Case 3 is more involved, requiring a three-dimensional solution. Figure 6.11 illustrates some of the flow features found as the object assumes an angle to the incoming flow. The case has been run in FLUENT at first order, and the results appear comparable to CNS3D also at first order, however, the ordering of data has made it impossible as yet to extract the information of



(a) Pressure coefficient



(b) Stanton No.

Figure 6.8: Run 2894 on 128x48 mesh

interest, and all results shown have been calculated with CNS3D. Figure 6.12 shows the pressure distribution against experiment at second-order for the three monitoring stations around the object. As expected, the attached flow running underneath the object is reasonably well captured but at the other points where the flow separates the solution does not resolve the details.

It is possible to integrate these pressure distributions over the body to calculate force and moment coefficients. Experiment has shown that other forces are negligible so this offers a reasonably reliable way of verifying the results. Table 6.3 shows the calculated forebody axial force coefficient at zero-lift with




Figure 6.9: Pressure distribution for HB-2 at Mach 5




Figure 6.10: Pressure distribution for HB-2 at Mach 5 at cylinder-flare transition

CNS3D for case 1. The comparison with experimental data is good. Similar calculations can be done for the 3D problem at inclined angle. It would be useful to investigate the trend over a range of angles of attack but the results in Table 6.4 do indicate that the simulation is capable of capturing these important parameters. The calculation does not include a correction for the base (which was not simulated) which may account for some of the discrepancy compared to experiment.

6.3.3 Efficiency and robustness




Figure 6.12: Pressure distribution for inclined HB-2 (Case 3)

In order to draw comparisons about the relative efficiency of the various numerical methods test runs were made on a single processor using the finer HB-2 grid to ascertain the cost per iteration (Table 6.5). This is an approximate measure intended only to show that the run times are of the same order of magnitude, and the relative effect of increasing the order of accuracy. It has been remarked that these flows are not necessarily steady, and do not necessarily converge, and this has been borne out by observation with an essentially steady result emerging rapidly followed by a long period of slight and reducing oscillation. In summary the results show: (i) that CNS3D is faster than FLUENT per iteration at second-order; (ii) the 9th-order method requires more CPU per iteration, however, higher-order methods offer exceedingly good value in terms of their efficacy on coarse grids, gaining a better level of accuracy for half the grid size in each direction, something that has been noted by within the group for a variety of compressible flows. FLUENT is more robust in so much as




Figure 6.11: Longitudinal velocity for inclined HB-2 (Case 3)

it has algorithms in place to force a stable solution, however the validity of such an approach is questionable. CNS3D will simply terminate should it be producing unphysical results, requiring user interaction to correct and advance the simulation.

6.4 Conclusions

The proposed geometries have been simulated using CNS3D with first-, second- and ninth-order methods and compared to experimental data and FLUENT results. CNS3D is both more accurate and faster than FLUENT at second-order, while it provides significantly more accurate results when used in conjunction with very high-order methods. This bears out the results seen in the previous chapters and demonstrates the wider applicability of the high-order high-resolution methods.

Table 6.3: Zero-lift Forebody Axial Force Coefficient


Grid	C_A
96x48	0.55
128x96	0.56
Experiment	0.57

Table 6.4: Normal Force Coefficient and pitching moment at 15°


	C_N	C_m
CNS3D	1.04	-0.83
Experiment	1.3	-0.9

Table 6.5: Computational cost normalised by CNS3D first-order.


Solver	Order of reconstruction	CPU work units per iteration
CNS3D	1 st	1.0
CNS3D	2 nd	1.088
FLUENT	2 nd	1.109
CNS3D	9 th	1.893




(a) CNS3D 1st order: Pressure Coefficient



(b) FLUENT 1st order: Pressure Coefficient




(c) CNS3D 1st order: Stanton No.




(d) FLUENT 1st order: Stanton No.


Figure 6.13: Comparison of pressure coefficient and Stanton number for run 2893 for grids 128x48 and 256x96 at first order




(a) CNS3D 2nd order: Pressure Coefficient



(b) FLUENT 2nd order: Pressure Coefficient




(c) CNS3D 2nd order: Stanton No.




(d) FLUENT 2nd order: Stanton No.

Figure 6.14: Comparison of pressure coefficient and Stanton number for run 2893 for grids 128x48 and 256x96 at second order




(a) CNS3D 9th order: Pressure Coefficient




(b) CNS3D 9th order: Stanton No.


Figure 6.15: Comparison of pressure coefficient and Stanton number for run 2893 for grids 128x48 and 256x96 at ninth order




(a) CNS3D 1st order: Pressure Coefficient



(b) FLUENT 1st order: Pressure Coefficient




(c) CNS3D 1st order: Stanton No.




(d) FLUENT 1st order: Stanton No.


Figure 6.16: Comparison of pressure coefficient and Stanton number for run 2894 for grids 128x48 and 256x96 at first order




(a) CNS3D 2nd order: Pressure Coefficient



(b) FLUENT 2nd order: Pressure Coefficient




(c) CNS3D 2nd order: Stanton No.




(d) FLUENT 2nd order: Stanton No.

Figure 6.17: Comparison of pressure coefficient and Stanton number for run 2894 for grids 128x48 and 256x96 at second order




(a) CNS3D 9th order: Pressure Coefficient




(b) CNS3D 9th order: Stanton No.

Figure 6.18: Comparison of pressure coefficient and Stanton number for run 2894 for grids 128x48 and 256x96 at ninth order




(a) CNS3D




(b) FLUENT

Figure 6.19: Pressure distribution for HB-2 at Mach 5 at first order for CNS3D and FLUENT at both grid resolutions.



(a) CNS3D



(b) FLUENT

Figure 6.20: Pressure distribution for HB-2 at Mach 5 at second order for CNS3D and FLUENT at both grid resolutions.

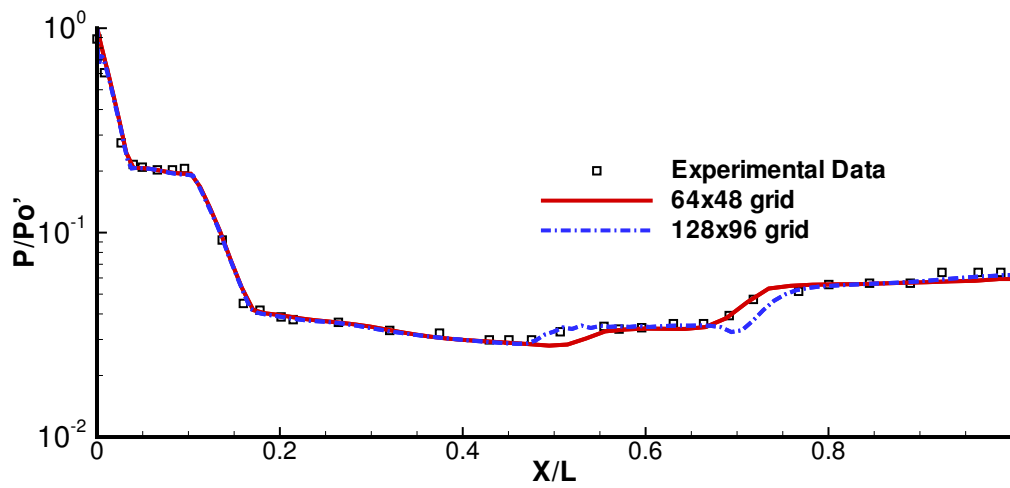




Figure 6.21: Pressure distribution for HB-2 at Mach 5 using WENO ninth-order in CNS3D.




(a) CNS3D




(b) FLUENT

Figure 6.22: Pressure distribution for HB-2 at Mach 7.5 at first order for CNS3D and FLUENT at both grid resolutions.



(a) CNS3D



(b) FLUENT

Figure 6.23: Pressure distribution for HB-2 at Mach 7.5 at second order for CNS3D and FLUENT at both grid resolutions.




Figure 6.24: Pressure distribution for HB-2 at Mach 7.5 using WENO ninth-order in CNS3D.

I may not have gone where I intended to go, but I think I ended up where I intended to be.

Douglas Adams

7

Concluding Remarks

The course of this research has been long and varied. The aim was to develop and apply high-order methods to improve on current standards for accurate simulation of compressible turbulent flows. While there are many factors to consider the underlying principle of using higher-order reconstruction is overwhelmingly supported in the results shown.

The research began with development of a reliable and versatile code with which to carry out subsequent investigations. To this end the HLLC Riemann solver was implemented and a number of multi-component models have also been considered in an ongoing process. The early test cases in chapter 2 provided a basis for assessing various approaches to improving the accuracy of the code. From this the WENO schemes were taken forward and tested on more complex problems to simultaneously expand the range of problems for which they have been shown to be valid and gain physical insight from the greater accuracy thus provided. On the first count the high-order methods have been shown to be applicable to both complex turbulent mixing and high-speed shock-dominated flows with consistently good resolution of flow features at relatively low costs.

The improved resolution offered by the higher-order methods illustrated fundamental difficulties with the simulation of the single-mode Richtmyer-Meshkov instability (Ch. 3). They also allowed the simulations to run to very late times by capturing the secondary instabilities facilitating comparison with experiments and growth models. Chapters 4 and 5 utilised the excellent preservation of small-scales in the higher-order methods to more rapidly achieve grid convergence in problems of turbulent mixing. The results also compared

favourably to those provided by external codes, extending the validation of these methods. Finally the methods were applied to a wholly new area with the investigations of the hypersonic problems (Ch. 6). The same properties seen in the previous test cases were found to exist with higher-order methods giving physical results without the need for prohibitively large numbers of grid points.

Specifically it has been demonstrated that:

- the HLLC Riemann solver is an accurate and robust basis for developing methods to solve complex compressible turbulent problems all the way into the hypersonic regime
- high-order WENO methods achieve greater resolution of small scales and rapidly varying flows than early MUSCL methods without being undermined by stability issues around discontinuities;
- on a basic cost-analysis improving the reconstruction method is notably better value than increasing the number of grid points;
- the various improvements to the WENO weightings proposed in literature do indeed carry through to our results;
- treatment to correct the low-mach dependency of the solution method can offer significant improvement in the resolution of certain flows;
- the single-mode Richtmyer-Meshkov Instability is a difficult problem to initialise accurately, however there is strong correlation with experiment on the large-scale features and growth rates to late time
- the growth of the single-mode RMI can be considered to be driven by the vortex pair, and development of the vortex model may offer good agreement to very late time with simulations
- the results repeatedly show evidence of true turbulent properties in accordance with theory
- more complex multiple mode instabilities can be well resolved and compare well to the experimental data available
- the high-order methods continue to offer improved accuracy on coarser grids in challenging hypersonic geometries.

7.1 Future Work

The research naturally leads in two directions from this point. The first of these is further development of the code. Although the methods tested have been found to be highly effective there are other areas alluded to in chapter 2 which need to be addressed to improve the flexibility and robustness of the code for new problems. Specifically, work can be done on the time integration, not only for solving steady-state problems but also making use of the higher-order methods as suggested in the ADER scheme. The multi-component model still leaves room for improvement; along with extension to multi-phase problems and other systems of equations there is much to be done integrating the high-order methods into simulations for an ever greater range of problems. There is also the high-order methods themselves to be considered. Already improvements have been found in calculating the weights in WENO schemes, and for certain cases more advanced schemes altogether are being applied. Having ascertained the value of high-order accuracy reconstruction, there need be no limit on how that is extended.

The alternative direction is using the high-order methods to conduct more in-depth analysis of the physics of the problems presented here or others. The single-mode Richtmyer-Meshkov instability is an apparently simple problem which could be better understood by extending the simulations begun here, in particular in terms of modelling the growth rate and late-time behaviour. The next step beyond this is to simulate the equivalent three-dimensional problem. The results of which could prove very useful in extending the understanding of the multi-mode RMI problem. Work can also be done looking at the mixing parameters in the turbulent regions generated by these shock-induced instabilities. In particular, work continues on the effect of different large-scale perturbations to the resulting flow. There is already progress being made in other areas on the simulation of the hypersonic problems however it is symbolic of different and perhaps unexpected areas of fluid dynamics where the high-order methods could be used to gain an extra level of accuracy in a cost-effective manner.

References

- [1] G. Allaire, S. Clerc, and S. Kokh. A five-equation model for the simulation of interfaces between compressible fluids. *J. Comput. Phys.*, 181:577–616, 2002.
- [2] D. S. Balsara and C.-W. Shu. Monotonicity preserving weighted essentially non-oscillatory schemes with increasingly high order of accuracy. *Journal of Computational Physics*, 160:405–452, 2000.
- [3] G. N. Barakos. *Study of unsteady aerodynamics phenomena using advanced turbulence closures*. 1999.
- [4] K. R. Bates, N. Nikiforakis, and D. Holder. Richtmyer-Meshkov instability induced by the interaction of a shock wave with a rectangular block of SF6. *Physics of Fluids*, 19, 2007.
- [5] P. Batten, M. A. Leschziner, and U. C. Goldberg. Average-state jacobians and implicit methods for compressible viscous and turbulent flows. *Journal of Computational Physics*, 137:38–78, 1997.
- [6] S. J. Billet and E. F. Toro. On WAF-type schemes for multidimensional hyperbolic conservation laws. *Journal of Computational Physics*, 130:1–24, 1997.
- [7] M. Brouillette and B. Sturtevant. Experiments on the Richtmyer-Meshkov instability: single-scale perturbations on a continuous interface. *Journal of Fluid Mechanics*, 263:271–292, 1994.
- [8] J. Casper and H. L. Atkins. A finite-volume high-order ENO scheme for two-dimensional hyperbolic systems. *Journal of Computational Physics*, 106(1):62–76, 1993.
- [9] P. R. Chapman and J. W. Jacobs. Experiments on the three-dimensional incompressible Richtmyer-Meshkov instability. *Physics of Fluids*, 18, 2006.
- [10] J. J. Coblish, M. S. Smith, T. Hand, G. V. Candler, and I. Nompelis. Double-cone experiment and numerical analysis at AEDC hypervelocity wind tunnel no.9. 2005-0902, AIAA, 2005.

- [11] R. H. Cohen. Three-dimensional simulation of a Richtmyer-Meshkov instability with a two-scale initial perturbation. *Physics of Fluids*, 14(10), 2002.
- [12] I. Cooper, 2008. www.metavu.com/index.php?date=20080222.
- [13] R. Courant, E. Isaacson, and M. Rees. On the solution of nonlinear hyperbolic differential equations by finite differences. *Comm. Pure. Appl. Math.*, 5:243–255, 1952.
- [14] Crab Nebula. Hubble Space Telescope. www.spacetelescope.org/images/html/heic0515a.html.
- [15] L. da Vinci. speechbubble.wordpress.com.
- [16] P. Davey, 2008. www.metavu.com/index.php?date=20080613.
- [17] X. Deng and H. Zhang. Developing high-order weighted compact nonlinear schemes. *Journal of Computational Physics*, 165:22–44, 2000.
- [18] D. Drikakis. Advances in turbulent flow computations using high-resolution methods. *Progress in Aerospace Sciences*, 39:405–424, 2003.
- [19] D. Drikakis and A. Bagabir. Numerical experiments using high-resolution schemes for unsteady, inviscid, compressible flows. *Computer methods in applied mechanics and engineering*, 193:4675–4705, 2004.
- [20] D. Drikakis and W. Rider. *High-Resolution Methods for Incompressible and Low-Speed Flows*. Springer, 2005.
- [21] M.-C. Druguet, G. V. Candler, and I. Nompelis. Effect of numerics on navier-stokes computations of hypersonic double-cone flows. *AIAA Journal*, 43(3), 2005.
- [22] M. Dumbser, M. Kaser, V. A. Titarev, and E. F. Toro. Quadrature-free non-oscillatory finite volume schemes on unstructured meshes for nonlinear hyperbolic systems. *Journal of Computational Physics*, 226:204–243, 2007.
- [23] A. Eberle. Characteristic flux averaging approach to the solution of Euler’s equations. von Karman Institute for Fluid Dynamics; Lecture Series 1987-04, 1987.
- [24] M. G. Edwards. The dominant wave-capturing flux: A finite-volume scheme without decomposition for systems of hyperbolic conservation laws. *Journal of Computational Physics*, 2006. Article in Press.

-
- [25] P. T. Fink and W. K. Soh. A new approach to roll-up calculations of vortex sheets. *Proc. R. Soc. Lond. A.*, 362:195–209, 1978.
- [26] E. Garnier, M. Mossi, P. Sagaut, P. Comte, and M. Deville. On the use of shock-capturing schemes for large-eddy simulation. *J. Comput. Phys.*, 153:273–311, 1999.
- [27] A. R. Gilchrist and M. J. Williams. Pressure distributions and forces on AGARD models HB-1 and HB-2 at $m=7.5$. Aerodynamics note 346, ARL, 1974.
- [28] P. A. Gnoffo. CFD validation studies for hypersonic flow prediction. 2001-1025, AIAA, 2001.
- [29] J. A. Greenough and W. J. Rider. A quantitative comparison of numerical methods for the compressible Euler equations: fifth-order WENO and piecewise-linear Godunov. *Journal of Computational Physics*, 196:259–281, 2004.
- [30] I. M. G. Grotowsky and J. Ballmann. Efficient time integration of Navier-stokes equations. *Computers & Fluids*, 28:243–263, 1999.
- [31] M. Hahn. Large-eddy simulation of compressible turbulent mixing for large-scale initial perturbations. In *IWPCTM 11*, 2008.
- [32] C. Harle, G. F. Carey, and P. L. Varghese. Analysis of high speed non-equilibrium chemically reacting gas flows. part ii. a finite volume/finite element model and numerical studies. *International Journal for Numerical Methods in Fluids*, 32:691–709, 2000.
- [33] C. Harle, G. F. Carey, and P. L. Varghese. Analysis of high speed non-equilibrium chemically reacting gas flows. part ii. a finite volume/finite element model and numerical studies. *International Journal for Numerical Methods in Fluids*, 32(6):691–709, 2000.
- [34] A. Harten, B. Engquist, S. Osher, and S. R. Chakravarthy. Uniformly high order accurate essentially non-oscillatory schemes, iii. *Journal of Computational Physics*, 71(2):231–303, 1987.
- [35] A. Harten, P. D. Lax, and B. van Leer. On upstream differencing and godunov-type schemes for hyperbolic conservation laws. *SIAM review*, 25(1):35–61, 1983.
- [36] J. K. Harvey, M. S. Holden, and T. P. Wadhams. Code validation study of laminar shock/boundary layer and shock/shock interactions in hypersonic flow part B: Comparison with Navier-Stokes and dsmc colutions. 2001-1031, AIAA, 2001.

-
- [37] A. K. Henrick, T. D. Aslam, and J. M. Powers. Mapped weighted essentially non-oscillatory schemes: Achieving optimal order near critical points. *Journal of Computational Physics*, 207:542–567, 2005.
- [38] D. J. Hill, C. Pantano, and D. I. Pullin. Large-eddy simulation and multi-scale modelling of a Richtmyer-Meshkov instability with reshock. *Journal of Fluid Mechanics*, 557:29–61, 2006.
- [39] M. S. Holden and T. P. Wadhams. Code validation study of laminar shock/boundary layer and shock/shock interactions in hypersonic flow part A: Experimental measurements. 2001-1031, AIAA, 2001.
- [40] D. A. Holder, A. V. Smith, C. J. Barton, and D. L. Youngs. Shock-tube experiments on Richtmyer-Meshkov instability growth using an enlarged double-bump perturbation. *Laser and Particle Beams*, 21:411–418, 2003.
- [41] J. W. Jacobs and V. V. Krivets. Experiments on the late-time development of single-mode Richtmyer-Meshkov instability. *Physics of Fluids*, 17, 2005.
- [42] J. W. Jacobs and J. M. Sheeley. Experimental study of incompressible Richtmyer-Meshkov instability. *Physics of Fluids*, 8, 1995.
- [43] A. Jameson. Numerical solution of the euler equations by finite volume methods using Runge-Kutta time-stepping schemes. Aiaa 1981-1259, AIAA Journal, 1981.
- [44] A. Jameson. Computational algorithms for aerodynamic analysis and design. *Applied Numerical Mathematics*, 13:383–422, 1993.
- [45] G.-S. Jiang and C.-W. Shu. Efficient implementation of Weighted ENO schemes. *Journal of Computational Physics*, 126:202–228, 1996.
- [46] E. Johnsen and T. Colonius. Implementation of WENO schemes in compressible multicomponent flow problems. *Journal of Computational Physics*, 219:715–732, 2006.
- [47] J. H. Jones. Pressure tests on the standard hypervelocity ballistic model HB-2 at mach 1.5 to 5. Aedc-tdr-64-246, AEDC, 1964.
- [48] M. K. Kadalbajoo and R. Kumar. A high resolution total variation diminishing scheme for hyperbolic conservation law and related problems. *Applied Mathematics and computation*, 2005. Article in Press.
- [49] A. Kolmogorov. The local structure of turbulence in an incompressible fluid at very high reynolds numbers. *Dokl. Akad. Nauk. SSSR*, 30:299, 1941.

- [50] A. Kolmogorov. A refinement of previous hypotheses concerning the local structure of turbulence in a viscous incompressible fluid at high reynolds number. *J. Fluid Mech.*, 13:82–85, 1962.
- [51] Kyu Hong Kim and Chongam Kim. Accurate, efficient and monotonic numerical methods for multi-dimensional compressible flows. part i: Spatial discretization. *Journal of Computational Physics*, 208:527–569, 2005.
- [52] Kyu Hong Kim and Chongam Kim. Accurate, efficient and monotonic numerical methods for multi-dimensional compressible flows. part ii: Multi-dimensional limiting process. *Journal of Computational Physics*, 208:570–615, 2005.
- [53] M. Latini, O. Schilling, and W. S. Don. Effects of WENO flux reconstruction order and spatial resolution on reshocked two-dimensional Richtmyer-Meshkov instability. *Journal of Computational Physics*, 221:805–836, 2007.
- [54] O. A. Likhachev and J. W. Jacobs. A vortex model for Richtmyer-Meshkov instability accounting for finite atwood number. *Physics of Fluids*, 17, 2005.
- [55] R. Liska and B. Wendroff. Comparison of several difference schemes on 1d and 2d test problems for the euler equations. Available at <http://www-troja.fjfi.cvut.cz/~liska/CompareEuler/compare8/>, 2002.
- [56] X.-D. Liu, S. Osher, and T. Chan. Weighted essentially non-oscillatory schemes. *Journal of Computational Physics*, 115:200–212, 1994.
- [57] M. MacLean and M. Holden. Validation and comparison of WIND and DPLR result for hypersonic, laminar problems. 2004-0529, AIAA, 2004.
- [58] L. G. Margolin and W. J. Rider. A rationale for implicit turbulence modelling. *International Journal for numerical methods in fluids*, 39:821–841, 2002.
- [59] M. K. Mawlood, S. Basri, A. S. Mokhtar, M. M. H. M. Ahmad, W. Asrar, and A. A. Omar. Flux limiting with high-order compact schemes. *AIAA*, (607), 2005. 43rd AIAA Aerospace Sciences Meeting and Exhibit.
- [60] E. Meshkov. Instability of the interface of two gases accelerated by a shock wave. *Fluid Dyn.*, 43(5):101–104, 1969.
- [61] Mt. St. Helen's. photo-mond.t2i.info/?s=volcano.
- [62] C. E. Niederhaus and J. W. Jacobs. Experimental study of the Richtmyer-Meshkov instability of incompressible fluids. *Journal of Fluid Mechanics*, 485, 2003.

- [63] J. Olejniczak, M. J. Wright, and G. V. Candler. Numerical study of inviscid shock interactions double-wedge geometry. *Journal of Fluid Mechanics*, 352:1–25, 1997.
- [64] S. A. Orzag and G. S. Patterson. Numerical simulation of three-dimensional homogeneous isotropic turbulence. *Phys. Rev. Lett.*, 28:76–79, 1972.
- [65] S. Piperno and S. Depeyre. Criteria for the design of limiters yielding efficient high resolution TVD schemes. *Computers & Fluids*, 27(2):183–197, 1998.
- [66] S. Pirozzoli. Conservative hybrid compact-WENO schemes for shock-turbulence interaction. *Journal of Computational Physics*, 178:81–117, 2002.
- [67] J. Qiu and C.-W. Shu. On the construction, comparison and local characteristic decomposition for high-order central weno schemes. *Journal of Computational Physics*, 183:187–209, 2002.
- [68] R. Richtmyer. Taylor instability in shock acceleration of compressible fluids. *Comm. Pure Appl. Math.*, 13:297–319, 1960.
- [69] W. J. Rider, J. A. Greenough, and J. R. Kamm. Accurate monotonicity- and extrema-preserving methods through adaptive nonlinear hybridizations. *Journal of Computational Physics*, 225:1827–1848, 2007.
- [70] P. Roe. Approximate Riemann solvers, parameter vectors and difference schemes. *J. Comput. Phys.*, 43:357–372, 1981.
- [71] R. D. Sandberg. Governing equations for a new compressible navier-stokes solver in general cylindrical coordinates. Afm-07/07, University of Southampton, 2007.
- [72] C.-W. Shu. High-order finite difference and finite volume WENO schemes and discontinuous galerkin methods for CFD. *International Journal of Computational Fluid Dynamics*, 17(2):107–118, 2003.
- [73] G. A. Sod. A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. *Journal of Computational Physics*, 27:1–31, 1977.
- [74] A. Suresh and H. T. Huynh. Accurate monotonicity-preserving schemes with Runge-Kutta time stepping. *Journal of Computational Physics*, 136:83–99, 1997.

- [75] P. K. Sweby. High resolution schemes using flux limiters for hyperbolic conservation laws. *SIAM Journal of numerical analysis*, 21(5), October 1984.
- [76] L. Tang and J. D. Baeder. Improving godunov-type reconstructions for simulation of vortex-dominated flows. *Journal of Computational Physics*, 213:659–675, 2006.
- [77] E. M. Taylor, M. Wu, and M. P. Martin. Optimization of nonlinear error for weighted essentially non-oscillatory methods in direct numerical simulations of compressible turbulence. *Journal of Computational Physics*, 223:384–397, 2007.
- [78] B. Thornber. *Implicit Large Eddy Simulation for unsteady multi-component compressible turbulent flows*. 2007.
- [79] B. Thornber, A. Mosedale, and D. Drikakis. Large eddy simulation of compressible turbulent mixing. ...submitted for conference, 2007.
- [80] B. Thornber, A. Mosedale, and D. Drikakis. On the implicit large eddy simulation of homogenous decaying turbulence. *Journal of Computational Physics*, 226:1902–1929, 2007.
- [81] B. Thornber, A. Mosedale, and D. Drikakis. An improved reconstruction method for compressible flows with low mach number features. *Journal of Computational Physics*, 227:4873–4894, 2008.
- [82] V. A. Titarev and E. F. Toro. ADER scheme for the three-dimensional non-linear hyperbolic systems.
- [83] V. A. Titarev and E. F. Toro. Finite-volume WENO schemes for three-dimensional conservation laws. *Journal of Computational Physics*, 201:238–260, 2004.
- [84] V. A. Titarev and E. F. Toro. WENO schemes based on upwind and centred TVD fluxes. *Computers & Fluids*, 34:705–720, 2005.
- [85] E. F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer, second edition, 1999.
- [86] E. F. Toro and V. A. Titarev. MUSTA fluxes for systems of conservation laws. *Journal of Computational Physics*, 216:705–720, 2006.
- [87] B. van Leer. Towards the ultimate conservative difference scheme. v. *Journal of Computational Physics*, 32:101–136, 1979.
- [88] M. Vetter and B. Sturtevant. Experiments on the Richtmyer-Meshkov instability of an air/SF6 interface. *Shock Waves*, 4:247–252, 1995.

- [89] S. P. Wang, M. H. Anderson, J. G. Oakley, and M. L. C. and R. Bonazza. A thermodynamically consistent and fully conservative treatment of contact discontinuities for compressible multi-component flows. *Journal of Computational Physics*, 195:528–559, 2004.
- [90] M. J. Wright, K. Sinha, J. Olejniczak, and G. V. Candler. Numerical and experimental investigations of double-cone shock interactions. *AIAA Journal*, 38(12), 2000.
- [91] K. Xu, M. Mao, and L. Tang. A multi-dimensional gas-kinetic BGK scheme for hypersonic viscous flow. *Journal of Computational Physics*, 203:405–421, 2005.
- [92] Z. Xu and C.-W. Shu. Anti-diffusive flux corrections for high-order finite difference WENO schemes. *Journal of Computational Physics*, 205:458–485, 2005.
- [93] D. Youngs. Three-dimensional numerical simulation of turbulent mixing by Rayleigh-Taylor instability. *Phys. Fluids A*, 3(5):1312–1320, 1991.
- [94] D. Youngs. Application of MILES to Rayleigh-Taylor and Richtmyer-Meshkov mixing. *AIAA-2003-4102*, 2003.
- [95] D. Youngs. Effect of initial conditions on self-similar turbulent mixing. In *IWPCTM 9*, 2004.
- [96] H. Yu and Y.-P. Liu. A second-order accurate, component-wise TVD scheme for nonlinear, hyperbolic conservation laws. *Journal of Computational Physics*, 173:1–16, 2001.
- [97] Q. Zhang and S.-I. Sohn. An analytical nonlinear theory of Richtmyer-Meshkov instability. *Physics Letters A*, 212:149–155, 1996.
- [98] Q. Zhang and S.-I. Sohn. Pade approximation to an interfacial fluid mixing problem. *Appl. Math. Lett.*, 10(5):121–127, 1997.