Computational Ultrarelativistic Hydrodynamics

Ian Hawke
Queens' College, Cambridge

A dissertation submitted to the
University of Cambridge
for the degree of
Doctor of Philosophy

September 2001

The work in chapter 5 of this dissertation was done in collaboration with my supervisor, J. M. Stewart. In addition, parts of this work have made use of a computer code (the adaptive mesh refinement algorithm described in chapter 4) which is an improved version of a code originally written by my supervisor. All other work described in this dissertation is my own.

I was supported by an EPSRC scholarship. Some of this work was performed on the COSMOS supercomputer which is supported by SGI, HEFCE and PPARC.

Abstract

This thesis studies critical collapse in the early universe as modelled by a Friedman-Robertson-Walker (FRW) spacetime. This differs from most previous studies of critical collapse as the spacetime is not asymptotically flat or approximately stationary. This study allows us to look at the formation of black holes in the early universe.

In order to perform this study we consider in depth the numerical methods necessary to study the collapse of an ultrarelativistic fluid. The general solution of the Riemann problem with the ultrarelativistic equation of state is derived, along with previous results. The efficiency and accuracy of different exact and approximate Riemann solvers is compared.

The Riemann solvers are then used to implement and test a number of High Resolution Shock Capturing methods. A variety of high order methods are compared using a wide variety of tests. The stability, accuracy and efficiency of the methods are the focus of this part of the thesis.

These methods are combined with different gridding strategies. The Adaptive Mesh Refinement algorithm is described, together with simpler methods such as an uneven grid. The Adaptive Mesh algorithm is found to have some unexpected problems when the system of equations is not hyperbolic.

Finally we look at the behaviour of the collapse of a density perturbation in a FRW spacetime. It is shown that it does not conform to the standard critical collapse scenario. Various explanations for this are discussed, together with the implications for primordial black hole formation.

Acknowledgements

It is a great pleasure to acknowledge the many people who, either knowingly or unknowingly have helped and taught me the things that I have learnt over the past three years.

Firstly, I am very grateful to my supervisor, Dr. John Stewart, for all his help, advice and encouragement, even if the "small, simple problem to get you started" turned out to be just a little more difficult than expected.

Secondly, I am indebted to Dr. Nikolaos Nikiforakis for introducing me to the world of High Resolution methods, and for always having time for my questions.

Thirdly, I wish to thank Dr. Simon Hern for taking the time to listen to my ideas (and take them seriously!), for reading my work in such great detail, and for more encouragement than I expect he thought he gave.

I would also like to thank all those who make me welcome in my (all too brief) stay at the Universidad de Valencia, especially Professor J. M. Ibáñez and Dr. J. M. Martí for talking to me in such detail.

I am also grateful to the following people for all sorts of reasons: Dr. J. A. Pons, Daphne Sulston, Dr. Stuart Rankin, Dr. Paul Shellard, Marija Zamaklar, Selena Ng, Toby Wiseman, Dr. Caroline Low, Professor E. F. Toro, Blandine Walker, Owen Arthurs, Andrew Somers and all those in Cambridge University Bowmen.

My two final thanks go to those without whom I would not have even contemplated starting; to my family for all their support and, above all, to Leonie, for everything.

Contents

1 Introduction **1**

2 The Riemann initial value problem **6**

2.1 Formulation 7

2.2 Solution 9

2.2.1 Linear advection 9

2.2.2 Burger's equation 11

2.2.3 Linear system 14

2.3 Nonlinear systems 16

2.3.1 Conventions 16

2.3.2 Solution form 18

2.3.3 The Rezzolla–Zanotti approach 19

2.4 Ultrarelativistic hydrodynamics 24

2.4.1 Rarefactions 25

2.4.2 Shocks 27

2.5 Ideal gas relativistic hydrodynamics 29

2.5.1 Rarefactions 31

2.5.2 Shocks I 32

2.5.3 Shocks II 34

2.6 Ultrarelativistic hydrodynamics in more than one dimension . 35

2.6.1 Rarefactions 36

2.6.2 Shocks 39

2.7 Ideal gas relativistic hydrodynamics in more than one dimension 41

2.7.1 Rarefactions 42

2.7.2 Shocks 42

2.8 Approximate Riemann solvers 43

2.8.1 HLLE solver 43

2.8.2 Roe solver 45

2.8.3 Marquina solver 46

2.8.4 Accuracy of approximate solvers 48

2.9 Efficient numerical solutions 49

2.9.1 The iterative method 51

2.9.2 The starting value 52

2.9.3 Efficiency 53

2.9.4 Efficiency of the multidimensional solver 55

3 Numerical methods 57

3.1 Notation 58

3.1.1 Stability 61

3.2 Method of Lines 66

3.3 Ordinary differential equation solvers 67

3.3.1 Runge–Kutta methods 67

3.3.2 Gaussian quadrature 69

3.3.3 Adams–Bashforth methods 70
3.3.4 Adams–Moulton methods 70
3.3.5 Milne–Simpson methods 71
3.4 Reconstructions 71
3.4.1 First order reconstruction 72
3.4.2 Slope limited reconstruction 72
3.4.3 Essentially Non-Oscillatory reconstruction 74
3.5 The numerical flux 77
3.6 Source terms 80
3.7 Other numerical methods 81
3.7.1 SLIC 81
3.7.2 WAF 83
3.8 Boundary conditions 85
3.9 Hydrodynamical details 88
3.9.1 Finding the primitive variables 88
3.9.2 The spectral decomposition 92
3.10 Code validation tests 97
3.10.1 One dimensional tests 98
3.10.2 Relativistic blast wave interaction 114
3.10.3 Higher dimensional tests 119
3.10.4 Validation tests: conclusion 127

4 The Numerical grid **131**

4.1 The basic grid 131
4.2 Uneven grids 133

4.3 Adaptive uneven grids 135
4.4 Adaptive Mesh Refinement 137
4.4.1 The adaptive grids 138
4.4.2 Time stepping the grids 140
4.4.3 Adapting the grids 143
4.4.4 Error estimation and Richardson extrapolation 144
4.4.5 Boundary conditions and flux conservation 146
4.4.6 Interpolation onto child grids 148
4.4.7 Example 149

5 Primordial black hole formation 151

5.1 Physical background 152
5.1.1 General Relativity 152
5.1.2 Critical collapse 154
5.1.3 FRW spacetimes 158
5.1.4 Critical collapse in FRW spacetimes 159
5.2 Previous calculations 160
5.3 Coordinates, variables and equations 163
5.4 Initial data 167
5.5 Boundary conditions 169
5.6 The numerical grid 171
5.7 The evolution algorithm 176
5.8 Code tests 182
5.9 Results 189
5.9.1 General behaviour 189

5.9.2 Accretion 197

5.9.3 Critical behaviour 205

5.10 Conclusions 216

6 Conclusion 222

A Other work 226

A.1 Numerical methods for primitive variables 226

A.1.1 Primitive SLIC 226

A.1.2 Burgers equation in primitive form 231

A.2 Friedrich's equations 232

A.2.1 Equations 232

A.2.2 Evolution of density fluctuations 235

A.3 A null formulation 236

A.4 Convergence tests 239

List of Figures

2.1 The possible waves resulting from the Riemann problem . . . 9

2.2 The rarefaction shock is unstable under a small perturbation . 12

2.3 The solution of the Riemann problem 19

2.4 The algorithm to calculate the Marquina flux 47

2.5 Exact and approximate solutions to the Sod problem 49

3.1 The timestep is restricted by the CFL condition 60

3.2 The WAF method uses the Riemann problem solution 83

3.3 WAF for the linear advection test 99

3.4 ENO3 for the linear advection test 101

3.5 Second order methods applied to Burger's equation 102

3.6 Higher order methods applied to Burger's equation 104

3.7 Ultrarelativistic gas tests in one dimension: I 107

3.8 Ultrarelativistic gas tests in one dimension: II 109

3.9 The stencil of two TVD methods 110

3.10 Ideal gas tests in one dimension: I 111

3.11 Ideal gas tests in one dimension: II 113

3.12 Ideal gas tests in one dimension: III 115

3.13 Ideal gas tests in one dimension: IV 116
3.14 The Colella–Woodward problem - converged solution 118
3.15 Solutions to the Colella–Woodward problem: I 120
3.16 Solutions to the Colella–Woodward problem: II 121
3.17 Solutions to the Colella–Woodward problem: III 122
3.18 Solutions to the Colella–Woodward problem: IV 123
3.19 Ultrarelativistic gas tests with tangential velocities: I 124
3.20 Ultrarelativistic gas tests with tangential velocities: II 125
3.21 Ideal gas tests with tangential velocities: I 126
3.22 Ideal gas tests with tangential velocities: II 128

4.1 The basic computational grid 132
4.2 An uneven grid 134
4.3 Adaptive grids 139
4.4 Time stepping and AMR 141
4.5 Pseudo-code for time stepping the adaptive grids. 141
4.6 Pseudo-code for error checking the adaptive grids 144
4.7 AMR in action 149

5.1 Critical behaviour as a dynamical system 155
5.2 Possible slicings of FRW 164
5.3 AMR applied to critical collapse 174
5.4 Long time behaviour of the algorithm 184
5.5 Long time convergence behaviour of the algorithm 185
5.6 Convergence behaviour of the algorithm 187
5.7 Scaling behaviour for an asymptotically flat spacetime 188

5.8 The general behaviour of the pressure 190
5.9 The general behaviour of the velocity 191
5.10 Snapshots of the general behaviour: supercritical case 192
5.10 Snapshots of the general behaviour: supercritical case 193
5.11 The general behaviour of the pressure 195
5.12 The general behaviour of the velocity 196
5.13 Snapshots of the general behaviour: subcritical case 198
5.13 Snapshots of the general behaviour: subcritical case 199
5.14 Snapshots of the general behaviour: highly supercritical case . 200
5.14 Snapshots of the general behaviour: highly supercritical case . 201
5.15 Accretion behaviour is independent of black hole mass 202
5.16 Marginal behaviour in a collapse 204
5.17 Critical behaviour for the centred Gaussian 206
5.18 Critical behaviour for the centred Gaussian 207
5.19 Critical behaviour for the off centred Gaussian 208
5.20 Critical behaviour for the off centred Gaussian 209
5.21 The Evans–Coleman critical solution 211
5.22 Intermediate solutions at different times 212
5.23 Deviations from the Evans–Coleman critical solution 213
5.24 Time spent near the critical solution 214

A.1 Shocks and primitive variables 230

List of Tables

2.1 Initial data for the Sod problem 48

2.2 Tests for the relativistic ideal gas Riemann solver 54

2.3 Efficiency of the exact and approximate Riemann solvers . . . 55

2.4 Efficiency of exact Riemann solvers in more than one dimension 56

3.1 Convergence tests for the linear advection equations 100

3.2 Convergence tests for Burgers equation 105

3.3 Efficiency for the Colella–Woodward problem 119

5.1 Convergence behaviour for α 182

A.1 Convergence tests for the ultrarelativistic test 241

A.2 Convergence tests for the first ideal gas test 242

A.3 Convergence tests for second ideal gas test 243

A.4 Convergence tests for the ultrarelativistic test 244

A.5 Convergence tests for the ultrarelativistic test 245

Chapter 1

Introduction

This thesis is a numerical study of critical collapse in the early universe as modeled by a Friedman-Robertson-Walker (FRW) spacetime. Critical collapse is the study of an evolving matter distribution near the threshold of black hole formation - for a review see Gundlach (1999). This is interesting for two main reasons. The first is that although critical collapse has been studied in great depth both numerically and analytically in spherical symmetry (it has even been claimed that it is a "solved problem"), there are very few studies where the spacetimes considered are not asymptotically flat. This leads to some interesting problems. There are numerical problems due to the non-trivial outer boundaries. There are analytical problems as the spacetime is not asymptotically flat so there is no reason for the black hole to be stationary and so the black hole mass could be dynamic after formation. It is also unclear what effect the loss of asymptotic flatness may have on the critical behaviour. The second reason for interest is that under certain types of critical collapse, usually called type II collapse, infinitesimally small black holes may form. Recent advances in observational technology, particularly microlensing techniques such as Alcock et al. (1996), Renault et al. (1997), mean that very small black holes that could form under type II critical collapse might be observable. However, before a dedicated effort is made to find such small black holes, it would be sensible to find out how likely they are to exist.

Much of this thesis deals with the numerical methods necessary to study the collapse of a perfect fluid, particularly the *ultrarelativistic* fluid. This is the relativistic limit of a barotropic fluid, that is one with a constant sound speed, and in cosmological terms refers to the era when the fluid was kinetic energy dominated. Aside from a few studies (notably Smoller and Temple (1993), Neilsen and Choptuik (2000)) the ultrarelativistic fluid has been largely ignored. This is because the equation of state is too simple for the study of stellar formation and collapse where the ideal gas equation of state is the simplest realistic model, and for cosmological work the fluid is sufficiently homogenous so that many of the numerical issues addressed here, such as shock formation, are largely irrelevant. However, when studying large nonlinear perturbations in the early universe the ultrarelativistic equation of state is a reasonable approximation and, with the formation of shock waves, the use of the numerical techniques involving the Riemann problem is essential. It has also been argued that the ultrarelativistic equation of state is the relevant model for critical collapse. This is because self-similarity, which appears to be important for critical collapse, requires the ultrarelativistic equation of state. We shall look at the advantages of the ultrarelativistic equation of state over the ideal gas in some detail.

The reason for the particular presentation I have given can be expressed simply: every numerical evolution is an approximation, and for many interesting numerical evolutions an approximation where the errors are unknown. The pioneering work of Choptuik (1986, 1991) showed the importance of carefully testing the convergence of numerical codes to ensure that the results are believable. However convergence may not be enough, especially in the presence of shocks. Examples of convergent methods that give physically incorrect results can be seen in LeVeque (1998); a simple example and further discussion are given in appendix A. For this reason we concentrate on implementing many different methods, as a result will be more believable if it can be reproduced with different methods. The presentation here is meant to emphasize those parts of the numerical method that can be changed, what the approximations are and the effects these approximations have.

As with critical collapse, the numerical evolution of a system of hyperbolic partial differential equations has been said to be a "solved problem". This attitude is reasonable, given the effectiveness of the numerical methods outlined in chapters 2 and 3. Some of the material here is original; in particular, the solution of the Riemann problem for ultrarelativistic hydrodynamics in more than one dimension in section 2.6 and the spectral decomposition in section 3.9.2 have not appeared before in the literature. I am also not aware of any direct numerical comparisons between the SLIC and WAF methods of Toro and the widely used ENO codes based on the work of Harten et. al. that appears in section 3.10.

However the evolution of such a system of equations when coupled to other types of equations is not such a simple problem. Relativity demands that the coordinates in which a numerical calculation takes place are not fixed, but evolve dynamically with the matter. In the particular coordinate system we use to study black hole formation this requires the solution of an ordinary differential equation in space to determine the metric at each step. This can lead to problems with conventional numerical methods as outlined in chapter 5.

The outline of this thesis is as follows. In chapter 2 we shall look at the Riemann problem, an essential part of modern numerical methods for systems of hyperbolic partial differential equations. We shall look at the solution of the Riemann problem for various sets of equations. We shall also look at approximate solutions of the Riemann problem. Finally we shall look at how we solve the Riemann problem numerically, and discuss the efficiency of various methods. The original work in this chapter includes the extension of the Rezzolla–Zanotti approach to more general Riemann problems (section 2.3.3), the solution of the Riemann problem for the ultrarelativistic equation of state with tangential velocities (section 2.6) and the efficiency tests for exact Riemann solvers (section 2.9). In chapter 3 we will look at a number of different numerical methods. Although the focus will be on systems of hyperbolic equations, the work is presented so that the introduction of other equations coupled to the hyperbolic equations is as simple as possible. We will look at the performance of these methods on a number of test cases. The original work in this chapter includes various points useful in the construction of a numerical method, such as the spectral decomposition of the Jacobian matrix, in section 3.9, and the validation tests in section 3.10. Chapter 4 will outline possible ways of defining and adapting the numerical grid. The majority of the original work done on these topics did not relate to the problems explained here and so is explained only briefly. In chapter 5 we shall use the methods introduced in the rest of the thesis to study numerically the formation of primordial black holes. This chapter, with the exception of some of the validation tests, is original work. Finally, some unsuccessful approaches and numerical methods not related to the work in the main thesis are given in appendix A.

Chapter 2

The Riemann initial value problem

The Riemann problem is a central part of modern High Resolution Shock Capturing (HRSC) schemes. Firstly, the (possibly approximate) solution of the Riemann problem at cell boundaries is an essential part of such schemes. As this solution needs to be found many times in a realistic calculation, the computational efficiency of the solution is important. Secondly, the Riemann problem is used as the most basic test of such schemes.

The use of the Riemann problem as a basic test in Computational Fluid Dynamics (CFD) has been widespread since the start of the field. However, its use in a numerical method was introduced in the work of Godunov (1959). This numerical method was not widely used until revisited by van Leer (1979) as more efficient and accurate Riemann solvers became available and the central rôle of the Riemann problem became clear.

The use of Riemann solvers in numerical relativity started much more recently. The solution of the Riemann problem for the relativistic ideal gas was given by Martí and Müller (1994) in one dimension and by Pons et al. (2000) in more than one dimension. For the ultrarelativistic equation of state the solution in one dimension was given by Smoller and Temple (1993) and the solution in more than one dimension appears here for the first time.

A central paper for numerical relativity is Pons et al. (1998). This shows that by using the equivalence principle each Riemann problem can be viewed as taking place on a locally flat background. Hence any special relativistic Riemann solver can be used. For this reason we only consider special relativistic hydrodynamics in this chapter. We will see later how many previous formulations of problems in full general relativity took advantage of this without making it explicit.

This chapter will look at the formulation and solution of the Riemann problem both in the general case and also as applied to relativistic hydrodynamics. Exact and approximate solutions will be given where possible, with notes on the computational efficiency of the solution. This chapter is largely based on the accounts of Laney (1998), LeVeque (1998) and Toro (1999).

2.1 Formulation

The general Riemann initial value problem in one dimension consists of a hyperbolic system of partial differential equations in conservation law form,

$$\partial_t \mathbf{q} + \partial_x \mathbf{f}(\mathbf{q}) = 0, \tag{2.1}$$

where $\mathbf{q}$ is the vector of *conserved variables* and $\mathbf{f}$ the *flux*, with the piecewise constant initial data

$$\mathbf{q}(x,0) = \begin{cases} \mathbf{q}_L & \text{if } x < 0, \\ \mathbf{q}_R & \text{if } x > 0. \end{cases} \tag{2.2}$$

If the data is discontinuous then the derivatives in equation (2.1) may not exist. So instead we consider *weak* solutions of the equations. That is, solutions of the integral form

$$\frac{\mathrm{d}}{\mathrm{d}t} \int_{x_1}^{x_2} \mathbf{q} \, \mathrm{d}x = \mathbf{f}(\mathbf{q}(x_1,t)) - \mathbf{f}(\mathbf{q}(x_2,t)). \tag{2.3}$$

There may be many weak solutions of equation (2.3). We must use a physical principle to choose the correct solution. The (weak, physical) solution of the Riemann problem describing the evolution in time of the states, is self-similar in the similarity variable $\xi = x/t$.

If the system (2.1) is of size N, that is $\mathbf{q}$ contains N variables, then the solution is completely given by N *waves* separating $N+1$ constant *states*. If the wave has no spread then the solution will change discontinuously across it. This discontinuity is known as a *shock*, *contact discontinuity* or *tangential discontinuity* depending on the physical context. If the wave spreads then the solution changes continuously across it and the wave is known as a *rarefaction*. A hypothetical example is shown in figure 2.1.

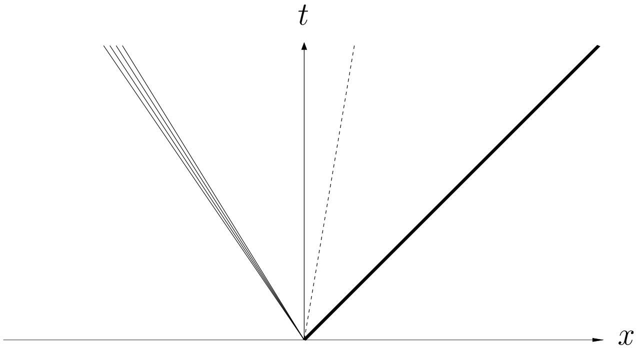

Figure 2.1: The possible waves resulting from the Riemann problem. From left to right we have a rarefaction, a contact and a shock.

2.2 Solution

The solution of the general Riemann problem cannot be given in a closed analytic form, even in one dimension. However, it is usually possible to find the answer numerically to any required accuracy. In that sense, the exact solution is said to be known. We shall illustrate how this solution is found by looking at some simple examples.

2.2.1 Linear advection

Suppose the system (2.1) is linear and of size one, that is there is only one conserved variable q. Equation (2.1) reduces to

$$\partial_t q + a \partial_x q = 0, \tag{2.4}$$

where a is a constant. This is the simplest useful form of the system (2.1). Generalizations of the linear advection equation are useful in many areas of fluid dynamics.

If we look for a self-similar solution in the similarity variable $\xi = x/t$ the equation becomes

$$(a - \xi)\frac{\mathrm{d}q}{\mathrm{d}\xi} = 0. \tag{2.5}$$

Hence the function $q(t, x)$ will be constant along the line $\xi = a$ as long as $\mathrm{d}q/\mathrm{d}\xi$ exists and is non-zero.

If $\mathrm{d}q/\mathrm{d}\xi = 0$ then there is no variation in the initial data and so the same solution can be said to hold. However, if the initial data contains a discontinuity then $\mathrm{d}q/\mathrm{d}\xi$ does not exist. Hence we must use the *Rankine–Hugoniot* conditions. These express conservation across the discontinuity. For a general hyperbolic conservation law in the form (2.1) these conditions have the form

$$V_s[\mathbf{q}] = [\mathbf{f}(\mathbf{q})], \tag{2.6}$$

where V_s is the shock speed which is a scalar. The notation $[\mathbf{q}]$ denotes the change of $\mathbf{q}$ across the discontinuity. In the case of the linear advection equation (2.4) this gives us

$$V_s[q] = a[q], \tag{2.7}$$

which gives the solution $V_s = a$ for a genuine discontinuity where $[q] \neq 0$.

This gives the solution of the Riemann problem for equation (2.2) as

$$q = \begin{cases} q_L & \text{if } x < at \\ q_R & \text{if } x > at. \end{cases} \tag{2.8}$$

For numerical work we often only need to consider the flux along the characteristic ray $\xi = 0$ which is given by

$$f(q) = \begin{cases} aq_L & \text{if } a < 0 \\ aq_R & \text{if } a > 0. \end{cases} \tag{2.9}$$

2.2.2 Burger's equation

Suppose instead the system (2.1) is of size one and nonlinear. As an example consider Burger's equation

$$\partial_t q + \partial_x(\tfrac{1}{2}q^2) = 0. \tag{2.10}$$

Burger's equation results from a particular approximation to the equations of Newtonian fluid dynamics; it is a simple model of the acoustic waves of the Euler equations. It is useful in illustrating the generic appearance of shocks in the Euler equations.

If the data is continuous then we can look for a solution in terms of the similarity variable ξ. We can also use the primitive form of equation (2.10) which gives

$$\partial_t q + q\partial_x q = 0. \tag{2.11}$$

Using the same approach as for the linear advection equation (2.4) we immediately have that the function $q(t, x)$ is constant along the line $\xi = q$.

If the data is discontinuous then the Rankine–Hugoniot conditions (2.6) give

$$V_s = \tfrac{1}{2}(q_L + q_R), \tag{2.12}$$

(a)

(b)

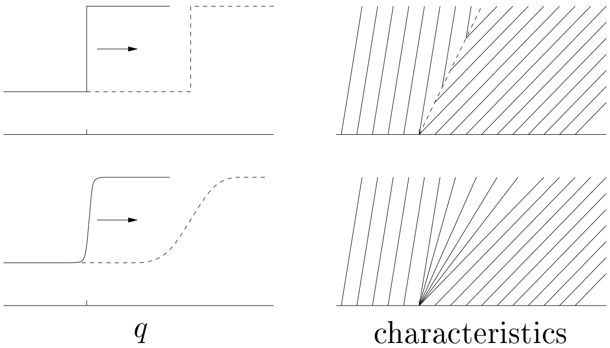

Figure 2.2: The evolution (left) and characteristics (right) for a rarefaction shock (a) and a small perturbation (b). For Burger's equation, the rarefaction shock predicted by the Rankine Hugoniot conditions is unstable under a small perturbation. This can be seen from the characteristics coming out of the shock (the dashed line). The rarefaction wave is the physical solution.

for a genuine discontinuity.

As a brief digression we note an example of why the conservation form is essential when we are dealing with shocks. Consider the hyperbolic conservation law

$$\partial_t \left(\tfrac{1}{2}q^2\right) + \partial_x \left(\tfrac{1}{3}q^3\right) = 0. \tag{2.13}$$

This has the same primitive form as Burger's equation (2.11), but the shock velocity is given by

$$V_s = \frac{2}{3}\frac{q_R^3 - q_L^3}{q_R^2 - q_L^2}. \tag{2.14}$$

In order for a shock to propagate at the correct velocity we must use the physically correct conservation form. As the conservation form of the equations is not uniquely determined by the primitive form of the equations, the correct shock speed cannot be deduced from the primitive form alone.

The similarity solution together with the shock relation do not give us the complete solution to the Riemann problem by themselves. This is because the predictions from the two methods above are not necessarily consistent. Consider the discontinuous initial data (2.2) with $q_L < q_R$. The Rankine–Hugoniot conditions (2.12) give the solution as a shock. However, under a small perturbation such that the initial profile is continuous, the similarity solution (2.11) gives a spreading rarefaction wave. These solutions are shown in figure 2.2. As the Rankine–Hugoniot solution is not stable it is said to be an unphysical *rarefaction shock*. This is one example of a hyperbolic conservation law having many weak solutions, and requiring a physical principle, here stability, to pick out the physical solution.

It follows that the solution of the Riemann problem is given by

$$q = \begin{cases} \left.\begin{array}{ll} q_L & \text{if } x < q_L t \\ x/t & \text{if } q_L t < x < q_R t \\ q_R & \text{if } x > q_R t \end{array}\right\} & \text{if } q_L < q_R \\ \\ \left.\begin{array}{ll} q_L & \text{if } x < \frac{1}{2}(q_L + q_R)t \\ q_R & \text{if } x > \frac{1}{2}(q_L + q_R)t \end{array}\right\} & \text{if } q_L > q_R. \end{cases} \tag{2.15}$$

From the solution (2.15) the flux along the characteristic ray $\xi = 0$ is given by

$$
f(q) = \begin{cases} \left.\begin{array}{rl} \frac{1}{2}q_L^2 & \text{if } q_L > 0 \\ 0 & \text{if } q_L q_R < 0 \\ \frac{1}{2}q_R^2 & \text{if } q_R < 0 \end{array}\right\} & \text{if } q_L < q_R \\ \\ \left.\begin{array}{rl} \frac{1}{2}q_L^2 & \text{if } \frac{1}{2}(q_L + q_R) > 0 \\ \frac{1}{2}q_R^2 & \text{if } \frac{1}{2}(q_L + q_R) < 0 \end{array}\right\} & \text{if } q_L > q_R. \end{cases} \tag{2.16}
$$

2.2.3 Linear system

Suppose the system (2.1) is of size N and linear. Equation (2.1) is written in the form

$$
\partial_t \mathbf{q} + A\partial_x \mathbf{q} = 0, \tag{2.17}
$$

where A is an $N \times N$ constant coefficient matrix.

Although the methods used in the previous sections will find the solution to the Riemann problem, it is more instructive to look for the solution in terms of *characteristic variables*. Let $\mathbf{l}^i$, $\mathbf{r}^i$ be the left and right eigenvectors of the matrix A with eigenvalues λ_i, normalized such that $\mathbf{l}^i \cdot \mathbf{r}^j = \delta_{ij}$. We shall assume that the eigenvectors span the space. The characteristic variables $\mathbf{w}^i$ are defined by

$$
\mathbf{w}^i = \mathbf{l}^i \cdot \mathbf{q}. \tag{2.18}
$$

Then equation (2.17) when written in terms of the characteristic variables becomes

$$
\partial_t \mathbf{w} + \Lambda\partial_x \mathbf{w} = 0, \tag{2.19}
$$

where Λ is the matrix containing the eigenvalues λ_i on the diagonals and zeros elsewhere. Hence each characteristic variable $\mathbf{w}^i$ obeys the linear advection equation (2.4) with velocity $a = \lambda_i$. This solves the Riemann problem in terms of characteristic variables.

In order to write the solution in terms of the original variables $\mathbf{q}$ we order the variables such that $\lambda_1 \leqslant \cdots \leqslant \lambda_N$. We also define the differences in the characteristic variables $\Delta\mathbf{w}^i = \mathbf{w}^i_L - \mathbf{w}^i_R$ across the i^{th} characteristic wave. These differences are scalars. We note that these differences can easily be found from the initial data using

$$\Delta\mathbf{w}^i = \mathbf{l}^i \cdot (\mathbf{q}_L - \mathbf{q}_R) \,. \tag{2.20}$$

As the change in the solution across each wave is precisely the difference in the associated characteristic variable, the solution of the Riemann problem in terms of characteristic variables can be written as either

$$\mathbf{w} = \mathbf{w}_L + \sum_{i=1}^{M} \Delta\mathbf{w}^i\mathbf{e}^i \quad \text{if } \lambda_M < \xi < \lambda_{M+1}, \tag{2.21}$$

or

$$\mathbf{w} = \mathbf{w}_R - \sum_{i=M+1}^{N} \Delta\mathbf{w}^i\mathbf{e}^i \quad \text{if } \lambda_M < \xi < \lambda_{M+1}, \tag{2.22}$$

or as the average

$$\mathbf{w} = \tfrac{1}{2}\left(\mathbf{w}_L + \mathbf{w}_R + \sum_{i=1}^{M} \Delta\mathbf{w}^i\mathbf{e}^i - \sum_{i=M+1}^{N} \Delta\mathbf{w}^i\mathbf{e}^i\right) \quad \text{if } \lambda_M < \xi < \lambda_{M+1}, \tag{2.23}$$

where $\mathbf{e}^i$ is the column vector $(\mathbf{e}^i)_j = \delta^i_j$.

Converting back to the original variables $\mathbf{q}$ we have the solution

$$\mathbf{q} = \tfrac{1}{2}\left(\mathbf{q}_L + \mathbf{q}_R + \sum_{i=1}^{M} \Delta\mathbf{w}^i\mathbf{r}^i - \sum_{i=M+1}^{N} \Delta\mathbf{w}^i\mathbf{r}^i\right) \quad \text{if } \lambda_M < \xi < \lambda_{M+1}. \tag{2.24}$$

In the case where we are only interested in the flux along the characteristic $\xi = 0$ we can write the solution in the simple form

$$\mathbf{f}(\mathbf{q}) = \tfrac{1}{2}\left(\mathbf{f}(\mathbf{q}_L) + \mathbf{f}(\mathbf{q}_R) - \sum_{i=1}^{N} |\lambda_i| \Delta\mathbf{w}^i\mathbf{r}^i\right). \tag{2.25}$$

2.3 Nonlinear systems

Although the results of the previous section do extend to the fully nonlinear Riemann problem, it is usually best to proceed on a case by case basis. However, there are a few general points that we shall address here.

2.3.1 Conventions

The four sections 2.4–2.7 give the relations across the waves necessary to solve the Riemann problem. The following conventions will be followed. A *wave* is a path in the (t, x) plane separating different states of the solution. For a linear problem, a wave is a *characteristic* of the solution, that is a path along which the associated *characteristic variable* is unchanged, but across which the associated characteristic variable may change. For a nonlinear problem the characteristic form and hence the definition of a wave still holds, but the will depend on the data at every point and so the definition only holds locally. If the values of $\mathbf{q}$ are only known on one side of the wave, the known data is the *pre*-wave data and the unknown the *post*-wave data. The pre-wave data has the subscript a and the post-wave the subscript b. When there is a choice of sign the upper sign will refer to the left wave and the lower to the right wave.

Investigating the stability of each wave to check whether the shock solution or rarefaction solution is physical is complex. Instead, for the hydrodynamical Riemann problems considered here, we can use the compressible nature of the gas. This ensures that the entropy, and hence the pressure, must increase across a shock (i.e., $p_b > p_a$).

For the rest of this section we will use the following notation. As outlined in section 2.1 the solution of the Riemann problem is given by the N wave speeds and the $N+1$ states between the waves. We shall denote these states by $\mathbf{q}_{*_i}$ where $i = 1, \ldots, N+1$. We note that $\mathbf{q}_{*_1} = \mathbf{q}_L$ and $\mathbf{q}_{*_{N+1}} = \mathbf{q}_R$. The other states are known as *intermediate* states. If there is only one intermediate state we shall denote it $\mathbf{q}_*$. If a variable q^A is unchanged across all the intermediate states, we shall denote denote it q^A_*.

In all the hydrodynamical cases the coordinate velocity v (or v^x in more than one dimension) and the pressure p are unchanged across all the intermediate states. This is because the intermediate waves are linear, either contact or tangential discontinuities. With certain choices of variables (where the linear waves have their characteristic form) the linear discontinuities do not couple to the rest of the system. This allows the linear and nonlinear parts of the system to be treated separately. For example, the propagation speed of the linear contact discontinuity in ideal gas hydrodynamics is given by the intermediate velocity. However the density ρ_0 in the intermediate state can be found from the relations between the velocity and the pressure and the equation of state; the jump in the density across the contact discontinuity does not need to be evaluated to find the solutions across the nonlinear waves.

The relations between the velocity and the pressure will be written $v_{L*}(p_*)$ or $v_{R*}(p_*)$. This denotes the velocity v_* in the intermediate state across the left or right wave given the known left or right data $\mathbf{q}_{L,R}$ and a value for the pressure p_* in the intermediate state. In section 2.3.3 we shall consider any system with two variables that are unchanged across all the intermediate states and adapt the notation accordingly.

2.3.2 Solution form

One general point is the form in which the solution is usually given. It is not always possible to give the complete solution in closed form, although in the cases we consider here it is possible to find the solution numerically to any desired accuracy. So the solution is usually given by the relations between the variables across each individual wave. If there is only one intermediate state or, as explained in the last paragraph of section 2.3.1, the other intermediate waves are linear and do not affect this particular subset of variables, then the solution is frequently represented geometrically as in figure 2.3. This shows the solution of a Riemann problem for relativistic hydrodynamics. The shock and rarefaction relations give smooth curves in the pressure–velocity plane. We construct the curves passing through $\mathbf{q}_L$ and $\mathbf{q}_R$. Their intersection defines the intermediate state $\mathbf{q}_*$. In this example there is a left going rarefaction wave connecting $\mathbf{q}_L$ to $\mathbf{q}_*$ and a right going shock wave connecting $\mathbf{q}_R$ to $\mathbf{q}_*$.

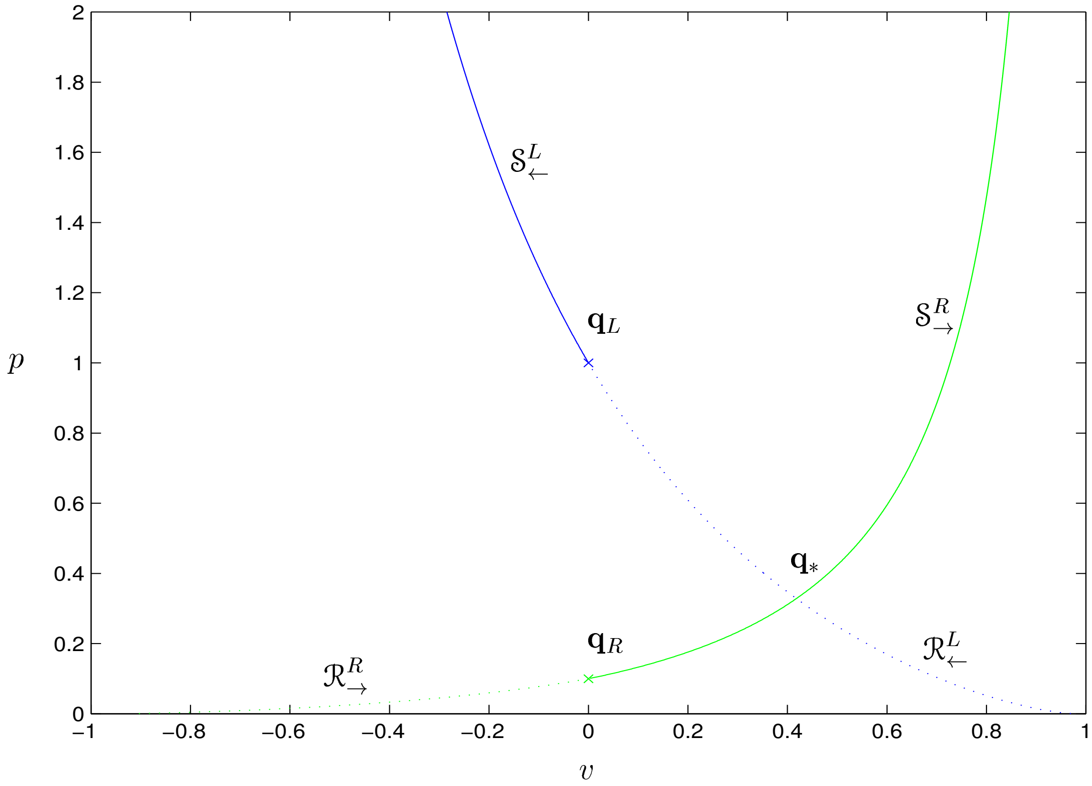

Figure 2.3: The solution of the Riemann problem shown geometrically in the pressure–velocity plane. $\mathcal{R}_{\leftrightarrows}$ are the left and right going rarefaction waves (dotted). $\mathcal{S}_{\leftrightarrows}$ are the left and right going shock waves (full). The intermediate state is labeled $\mathbf{q}_*$. The left and right initial states are $(p_L, \rho_L, v_L) = (1, 1, 0)$, $(p_R, \rho_R, v_R) = (0.1, 0.125, 0)$, and $\gamma = 5/3$. This is the relativistic version of the problem of Sod (1978).

2.3.3 The Rezzolla–Zanotti approach

Another general point is an extension of an approach of Rezzolla and Zanotti (2001). Rezzolla and Zanotti showed how the wave pattern of the solution of the Riemann problem could be found directly from the initial data for the relativistic ideal gas equations. This does not directly solve the Riemann problem itself. However it may improve the efficiency of the solution, as explained in section 2.9. In this section we extend this solution to *any* Riemann problem with only two nonlinear waves, that is, only two waves that are not contact or tangential discontinuities. This allows the Rezzolla–Zanotti approach to be extended directly to the ultrarelativistic equation of state given in section 2.4, or to the relativistic gas in more than one dimension as given in section 2.7, or other systems not considered here.

We note that each approach we have outlined above in sections 2.2.1–2.2.3 deals with only a single wave at a time. We have either found which sort of wave we expect and used the appropriate approach (as in the case of the linear system of section 2.2.3), or we have used the known data to predict what sort of wave would be physical (as in the case of Burger's equation in section 2.2.2). When we have more than one wave, we will not know the data on both sides of a given wave. Hence we will not be able to say immediately what the wave pattern is, and so we cannot find the unknown data.

The standard approach is iterative. By guessing some of the unknown data, we can find the wave pattern and hence find how inaccurate our guess is. We can then refine the guess until the solution is known to within some tolerance. However, in certain circumstances it is possible to determine the wave pattern in advance.

We make the following assumptions.

1. There are two truly nonlinear waves, i.e., either two shocks or two

rarefactions or one shock and one rarefaction.

2. There exist two independent variables which are constant across all the intermediate linear discontinuities and from which the other variables may be found, possibly through using the known initial data. We shall denote these variables by $q^{A,B}$. For hydrodynamics q^A will be the coordinate velocity v and q^B the pressure p.

3. We know the relations $q^A_{L*}(q^B_*)$, $q^A_{R*}(q^B_*)$, which give the variable q^A across the left or right wave respectively given the intermediate value q^B_* and the known initial data. Given a value of q^B in the intermediate state we can find the corresponding intermediate value of q^A. For hydrodynamics these relations are given in sections 2.4–2.7.

4. There exists a frame independent difference operator $\Delta(q^A_1, q^A_2)$ for the variable q^A. For non-relativistic hydrodynamics Δ is given by

$$\Delta(v_1, v_2) = v_1 - v_2 \tag{2.26}$$

and for special relativistic hydrodynamics by

$$\Delta(v_1, v_2) = \frac{v_1 - v_2}{1 - v_1 v_2}. \tag{2.27}$$

5. The known value $\Delta(q^A_L, q^A_R)$, when written as a function of the intermediate value q^B_* by

$$\Delta(q^A_L, q^A_R) = \Delta\left(\Delta\left(q^A_L, q^A_{L*}(q^B_*)\right), \Delta\left(q^A_R, q^A_{R*}(q^B_*)\right)\right), \tag{2.28}$$

is monotonic as a function of q_*^B.

With these assumptions we can determine the wave pattern. We do this by comparing the known value of $\Delta(q_L^A, q_R^A)$ given by the initial data with the value of the functional form (2.28) in the special cases given by setting q_*^B equal to either q_L^B or q_R^B. This is equivalent to looking at the special cases where one of the nonlinear waves is trivial. Due to the monotonicity assumption, these cases separate the parameter space of q^B into the separate wave patterns.

We illustrate this with the case of the relativistic ideal gas. By noting that $\Delta(q, q) = 0$ and $\Delta(q, 0) = -\Delta(0, q) = q$ we have, if $p_L > p_R$, the three possibilities

$$\Delta(v_L, v_R) > \Delta(v_L, v_{L*}(p_R)) \quad \Rightarrow \quad \mathcal{S}_{\leftarrow}\mathcal{C}\mathcal{S}_{\rightarrow} \quad (2.29a)$$

$$\Delta(v_L, v_{L*}(p_R)) > \Delta(v_L, v_R) > -\Delta(v_R, v_{R*}(p_L)) \quad \Rightarrow \quad \mathcal{R}_{\leftarrow}\mathcal{C}\mathcal{S}_{\rightarrow} \quad (2.29b)$$

$$-\Delta(v_R, v_{R*}(p_L)) > \Delta(v_L, v_R) \quad \Rightarrow \quad \mathcal{R}_{\leftarrow}\mathcal{C}\mathcal{R}_{\rightarrow}, \quad (2.29c)$$

where $\mathcal{S}_{\leftrightarrows}$ denotes a left or right going shock, $\mathcal{R}_{\leftrightarrows}$ a left or right going rarefaction and $\mathcal{C}$ a contact discontinuity. Symmetry gives the results for $p_L < p_R$. We have ignored the possibility of vacuum formation which is considered by Rezzolla and Zanotti (2001) because this would cause any Riemann problem based numerical method to break down.

We note that this method generalizes directly to higher dimensions for hydrodynamics. This is because the addition of more dimensions only introduces more linear waves. However it does not immediately generalize to any case with more nonlinear waves. This is because we would need to determine some of the intermediate data in order to find the correct value of q^B at which the special cases occur. Hence this is approach unlikely to be useful in more complex cases such as magnetohydrodynamics, considered by Ryu and Jones (1995).

Once the wave pattern is known, we have to find the intermediate data. To do this in closed form we have to note that $q^A_{L*} = q^A_* = q^A_{R*}$ and invert equation (2.28) for the frame invariant data difference Δ. This may not always be possible. For Newtonian and relativistic hydrodynamics it is not possible. For the two rarefaction wave pattern it is possible, with the explicit expressions given in Toro (1999) and Rezzolla and Zanotti (2001). However the other wave patterns cannot be found; for example even in the two shock case an iterative method is required as shown by Toro (1999) and Dai and Woodward (1997).

In practice the iterative approach is often more efficient even when the closed form solution is known. Given the relations between the variables across the waves we must then find the wave pattern and the intermediate states, from which we can calculate the wave speeds. In the case of hydrodynamics this is done by solving the implicit equation

$$v_{L_*}(p_*) = v_{R_*}(p_*), \tag{2.30}$$

where again $v_{L_*}(p_*)$ is the velocity in the intermediate state found by the relations across the left wave with the intermediate pressure p_*. In the case of ultrarelativistic hydrodynamics we typically write this as a function of the density ρ instead of pressure. This is equivalent to finding the root of the function f,

$$f(p_*) = v_{L_*}(p_*) - v_{R_*}(p_*). \tag{2.31}$$

Efficient numerical procedures for doing this are outlined in section 2.9.

2.4 Ultrarelativistic hydrodynamics

The equations for relativistic hydrodynamics in one dimension with an ultra-relativistic equation of state are derived from the conservation of the energy momentum tensor

$$T^{ab} = (p + \rho)u^a u^b + p\eta^{ab}, \tag{2.32}$$

where p is the pressure, ρ the mass-energy density, and η^{ab} the flat space metric. The system is closed with the equation of state

$$p = (\Gamma - 1)\rho \tag{2.33}$$

where $1 \leqslant \Gamma < 2$ is a constant.

We define the vectors of conserved variables and fluxes to be

$$\mathbf{q} = (\tau, S)^T, \tag{2.34a}$$

$$\mathbf{f} = (S, Sv + p)^T, \tag{2.34b}$$

where the definitions

$$\tau = (\rho + p)W^2 - p, \quad S = (\rho + p)W^2 v \tag{2.35}$$

have been made. The equations then have conservation form

$$\partial_t \mathbf{q} + \partial_x \mathbf{f}(\mathbf{q}) = 0. \tag{2.36}$$

Here $v = u^x/u^t$ is the coordinate velocity and W the Lorentz factor $W = (1 - v^2)^{-1/2}$. We define the sound speed $c_s^2 = \partial p / \partial \rho = \Gamma - 1$, and the useful auxiliary quantities

$$\alpha = (\Gamma - 1)/\Gamma, \qquad \mu = (\Gamma - 1)/\Gamma^2. \tag{2.37}$$

We are using units in which the speed of light is equal to 1.

The solution of the Riemann problem was first given by Smoller and Temple (1993). We outline it here in terms of the methods introduced above.

2.4.1 Rarefactions

Firstly, to look for a self-similar solution we write the hyperbolic conservation law (2.36) in the primitive form

$$\begin{aligned}
\left((1+c_s^2)W^2 - c_s^2\right)\partial_t\rho + \left(2\rho(1+c_s^2)W^4 v\right)\partial_t v + \\
\left((1+c_s^2)W^2 v\right)\partial_x\rho + \left(\rho(1+c_s^2)W^4(1+v^2)\right)\partial_x v &= 0, \qquad (2.38\text{a})\\
\left((1+c_s^2)W^2 v\right)\partial_t\rho + \left(\rho(1+c_s^2)W^4(1+v^2)\right)\partial_t v + \\
\left((1+c_s^2)W^2 v^2 + c_s^2\right)\partial_x\rho + \left(2\rho(1+c_s^2)W^4 v\right)\partial_x v &= 0. \qquad (2.38\text{b})
\end{aligned}$$

Simplifying the second equation by subtracting the first multiplied by v we have

$$
\begin{aligned}
\left((1+c_s^2)W^2 - c_s^2\right)\partial_t\rho + \left(2\rho(1+c_s^2)W^4 v\right)\partial_t v + \\
\left((1+c_s^2)W^2 v\right)\partial_x\rho + \left(\rho(1+c_s^2)W^4(1+v^2)\right)\partial_x v &= 0, \qquad (2.39a)\\
v c_s^2\partial_t\rho + \rho(1+c_s^2)W^2\partial_t v + \\
c_s^2\partial_x\rho + \rho(1+c_s^2)W^2 v\partial_x v &= 0. \qquad (2.39b)
\end{aligned}
$$

This leads to the equations in terms of the similarity variable ξ as

$$
\begin{aligned}
\left(\xi + (1+c_s^2)W^2 v(\xi v - 1)\right)\frac{\mathrm{d}\rho}{\mathrm{d}\xi} + \\
\left(\rho(1+c_s^2)W^4(2\xi v - 1 - v^2)\right)\frac{\mathrm{d}v}{\mathrm{d}\xi} &= 0, \qquad (2.40a)\\
(v\xi - 1)c_s^2\frac{\mathrm{d}\rho}{\mathrm{d}\xi} + \\
\rho(1+c_s^2)W^2(\xi - v)\frac{\mathrm{d}v}{\mathrm{d}\xi} &= 0. \qquad (2.40b)
\end{aligned}
$$

Regard these as simultaneous linear equations for $\mathrm{d}\rho/\mathrm{d}\xi$, $\mathrm{d}v/\mathrm{d}\xi$. For a non-trivial solution to exist the determinant of the arising matrix must vanish. Solving the equations gives the two solutions

$$
\xi = \frac{v \mp c_s}{1 \mp v c_s}. \qquad (2.41)
$$

Having found ξ we can substitute this in equation (2.40) to find how v

and ρ vary across the wave. This gives

$$\frac{\mp c_s}{\rho}\frac{\mathrm{d}\rho}{\mathrm{d}\xi} = \frac{1+c_s^2}{1-v^2}\frac{\mathrm{d}v}{\mathrm{d}\xi}, \tag{2.42}$$

which can be integrated to get the *Riemann invariants* $J_\pm$

$$J_\pm = \frac{1+v}{1-v}\rho^{\pm 2\alpha} \tag{2.43}$$

which are constant across the wave. Hence the post wave velocity is given by

$$v_b = \frac{A_\pm - 1}{A_\pm + 1}, \tag{2.44}$$

where

$$A_\pm = \frac{1+v_a}{1-v_a}\left(\frac{\rho_a}{\rho_b}\right)^{\pm 2\alpha}. \tag{2.45}$$

It will be useful later to note that

$$\frac{\mathrm{d}v_b}{\mathrm{d}\rho_b} = \mp\frac{4c_s A_\pm}{\rho_b(A_\pm + 1)^2}. \tag{2.46}$$

2.4.2 Shocks

If the initial data is not continuous, then we need to use the Rankine–Hugoniot conditions (2.6). Writing them as

$$V_s\left[\rho(1+c_s^2)W^2v^2 + \rho\right] = \left[\rho(1+c_s^2)W^2v\right], \tag{2.47a}$$

$$V_s\left[\rho(1+c_s^2)W^2v\right] = \left[\rho(1+c_s^2)W^2v^2 + \rho c_s^2\right], \tag{2.47b}$$

we solve this system by eliminating the shock speed and then solving for one variable. By eliminating W_b^2 we can write the shock speed as

$$V_s = \frac{rW_a^2 v_a(v_b - v_a) + \alpha(1-r)}{rW_a^2(v_b - v_a) + v_b\alpha(1-r)}, \tag{2.48}$$

where

$$r = (\rho_a/\rho_b)\,. \tag{2.49}$$

Hence the post shock velocity is given by

$$v_b = \frac{rv_a \mp (r-1)\,(1-v_a^2)\,\sqrt{\mu^2(r-1)^2 + \mu r}}{\mu(r-1)^2\,(1-v_a^2) + r}. \tag{2.50}$$

Again, it will be useful later to note that

$$\frac{\mathrm{d}v_b}{\mathrm{d}\rho_b} = \frac{dn' - nd'}{d^2} \tag{2.51}$$

where

$$\begin{aligned}
d &= \mu(r-1)^2\left(1-v_a^2\right) + r, && (2.52\text{a})\\
n &= rv_a \mp (r-1)\left(1-v_a^2\right) z, && (2.52\text{b})\\
z &= \sqrt{\mu^2(r-1)^2 + \mu r}, && (2.52\text{c})\\
d' &= r'(1 + 2\mu(r-1)(1-v_a^2)), && (2.52\text{d})\\
n' &= \rho_a r' \mp (1-v_a^2)r'\left(z + \tfrac{1}{2}\frac{(r-1)(2\mu^2(r-1)+\mu)}{z}\right). && (2.52\text{e})
\end{aligned}$$

2.5 Ideal gas relativistic hydrodynamics

The exact solution of the Riemann problem for relativistic hydrodynamics was given by Martí and Müller (1994). As the method used is largely the same as for the ultrarelativistic equation of state in section 2.4 we will only outline the results and the variations. Here the notation follows Neilsen (1999) instead of Martí and Müller (1994).

The equations for relativistic hydrodynamics in one dimension with the relativistic ideal gas equation of state are derived from the conservation of the energy momentum tensor

$$T^{ab} = \rho_0 h u^a u^b + p\eta^{ab}, \tag{2.53}$$

and the continuity equation

$$\left(\rho_0 u^a\right)_{,a} = 0, \tag{2.54}$$

where p is the pressure, ρ_0 the proper rest-energy density, η^{ab} the flat space metric and h the specific enthalpy defined by

$$h = 1 + \epsilon + \frac{p}{\rho_0}, \tag{2.55}$$

with ϵ being the specific internal energy. The system is closed with the equation of state

$$p = (\gamma - 1)\rho_0\epsilon, \tag{2.56}$$

where $1 \leqslant \gamma < 2$ is a constant, with the specific internal energy given by

$$\epsilon = \frac{p}{(\gamma - 1)\rho_0} \tag{2.57}$$

and the total energy by

$$\rho = \rho_0(1 + \epsilon). \tag{2.58}$$

Using the first law of thermodynamics for isentropic flow the equation of state (2.56) can be integrated to give

$$p = k(s)\rho_0^\gamma \tag{2.59}$$

where k depends only on the entropy. This form may only be used away from shocks where the entropy is conserved.

In the limit where the average internal velocity of the fluid becomes relativistic the fluid energy is dominated by the internal energy, $\rho_0 \epsilon \gg \rho_0$, so the total energy of the fluid $\rho \approx \rho_0 \epsilon$. Dropping the equation for particle number conservation (2.54) we find the ultrarelativistic equations and equation of state of section 2.4.

We define the vectors of conserved variables and fluxes to be

$$\mathbf{q} = (D, S, \tau)^T, \tag{2.60a}$$

$$\mathbf{f} = (Dv, Sv + p, \tau v + pv)^T, \tag{2.60b}$$

where the definitions

$$D = \rho_0 W, \quad S = \rho_0 h W^2 v, \quad \tau = \rho_0 h W^2 - p - D \tag{2.61}$$

have been made. In the ultrarelativistic limit the two different definitions of τ may be treated equivalently and so to conform to the notation in the literature we do not give them different names. Then the equations have the conservation form

$$\partial_t \mathbf{q} + \partial_x \mathbf{f}(\mathbf{q}) = 0. \tag{2.62}$$

Here $v = u^x/u^t$ is the coordinate velocity and W the Lorentz factor $W = (1 - v^2)^{-1/2}$. We also define the proper energy density e by

$$e = \rho_0(1 + \epsilon) = \rho_0 + \frac{p}{\gamma - 1}, \tag{2.63}$$

and the sound speed c_s by

$$c_s^2 = \frac{K\gamma(\gamma - 1)\rho_0^{\gamma-1}}{(\gamma - 1) + K\gamma\rho_0^{\gamma-1}}. \tag{2.64}$$

2.5.1 Rarefactions

If the wave is a rarefaction ($p_b < p_a$) then the solution for the characteristic variable ξ gives

$$\xi = \frac{v \mp c_s}{1 \mp v c_s}, \tag{2.65}$$

which lead to the velocity across the wave as

$$v_b(p_b) = \frac{A_\pm(p_b) - 1}{A_\pm(p_b) + 1}, \tag{2.66}$$

where $A_\pm$ is given by

$$A_\pm(p_b) = \frac{1 + v_a}{1 - v_a} \left(\frac{(\gamma - 1)^{1/2} + c_s(p_a)}{(\gamma - 1)^{1/2} - c_s(p_b)} \frac{(\gamma - 1)^{1/2} - c_s(p_*)}{(\gamma - 1)^{1/2} + c_s(p_b)} \right)^{\pm 2(\gamma - 1)^{-1/2}}, \tag{2.67}$$

and where the speed of sound $c_s(p)$ is given by

$$c_s(p) = \sqrt{\frac{\gamma p}{\rho_0 h}}, \tag{2.68}$$

Again it is useful to note that

$$\frac{\mathrm{d}v_b}{\mathrm{d}p_b} = \pm \frac{c_s(1 - v_b^2)}{\gamma p_b}. \tag{2.69}$$

2.5.2 Shocks I

If the wave is a shock ($p_b > p_a$) then

$$v_b(p_b) = \frac{\frac{W_s}{j}(p_b - p_a) + h_a W_a v_a}{h_a W_a + (p_b - p_a)\left(\frac{W_s v_a}{j} + \frac{1}{W_a (\rho_0)_a} \right)}, \tag{2.70}$$

where j is the mass flux and W_s the Lorentz factor associated with the shock. To find these we first calculate the post shock enthalpy, which is the only positive root of the quadratic

$$h_b^2\left(1+\frac{\gamma-1}{p_b\gamma}(p_a-p_b)\right)-h_b\left(\frac{\gamma-1}{p_b\gamma}(p_a-p_b)\right)+\left(\frac{h_a}{(\rho_0)_a}(p_a-p_b)-h_a^2\right)=0. \tag{2.71}$$

The equation of state (2.56) then gives the density $(\rho_0)_b$ from which the mass flux is given by

$$j=\mp\sqrt{\frac{p_b-p_a}{\frac{h_a}{(\rho_0)_a}-\frac{h_b}{(\rho_0)_b}}}. \tag{2.72}$$

Finally, the velocity associated with the shock V_s from which we find directly the Lorentz factor W_s is given by

$$V_s=\frac{1}{j^2+((\rho_0)_aW_a)^2}\left(((\rho_0)_aW_a)^2v_a\mp j^2\sqrt{1+\left(\frac{(\rho_0)_a}{j}\right)^2}\right). \tag{2.73}$$

Again it is useful to note that

$$\begin{aligned}\frac{\mathrm{d}v_b}{\mathrm{d}p_b}=&\left(\left(\frac{W_s}{j}\right)'\frac{p_b-p_a}{W_a}\left(h_a+\frac{p_b-p_a}{(\rho_0)_a}\right)+\frac{h_aW_s}{W_aj}-h_a\frac{v_a}{(\rho_0)_a}\right)\cdot\\&\left(\left(h_aW_a+(p_b-p_a)\left(v_a\frac{W_s}{j}+\frac{1}{(\rho_0)_aW_a}\right)\right)^2\right)^{-1},\end{aligned} \tag{2.74}$$

where we have

$$
\begin{aligned}
(\tfrac{W_s}{j})' &= j^{-2}(W_s' j - j' W_s), && \text{(2.75a)}\\
W_s' &= V_s W_s^3 V_s', && \text{(2.75b)}\\
V_s' &= \left(\tfrac{V_s}{d}\right)^2 j j' \left(\frac{\pm d\,(2j^2 + (\rho_0)_a^2)}{\sqrt{j^4 + ((\rho_0)_a j)^2}} - \right.\\
&\quad \left. 2\left(v_a ((\rho_0)_a W_a)^2 \pm \sqrt{j^4 + ((\rho_0)_a j)^2} \right) \right), && \text{(2.75c)}\\
d &= j^2 + ((\rho_0)_a W_a)^2, && \text{(2.75d)}\\
j' &= \tfrac{1}{2} \frac{j^3}{(p_b - p_a)^2} \left(\tfrac{h_a}{(\rho_0)_a} - \tfrac{h}{(\rho_0)_b} + p_b - p_a \right) \frac{(\rho_0)_b h' - h(\rho_0)_b'}{(\rho_0)_b^2}, && \text{(2.75e)}\\
(\rho_0)_b' &= \tfrac{\gamma}{\gamma - 1} \frac{h - 1 - p_b h'}{(h-1)^2}, && \text{(2.75f)}\\
h' &= -\frac{h^2 x' + h y' z'}{2xh + y}, && \text{(2.75g)}\\
x' &= -\tfrac{\gamma - 1}{\gamma} \frac{p_a}{p_b^2}, && \text{(2.75h)}\\
y' &= -x', && \text{(2.75i)}\\
z' &= -\frac{h_a}{(\rho_0)_a}. && \text{(2.75j)}
\end{aligned}
$$

2.5.3 Shocks II

The relations across the shock wave given in the previous section have one notable problem. That is, equation (2.72) for the mass flux j is ill defined as $p_b \to p_a+$. This may cause numerical problems. An alternative is to look at the relative velocity difference across the shock. The expressions were given by Rezzolla and Zanotti (2001). However this is the first time it has been applied in the direct solution of the Riemann problem.

The relative velocity difference v_{ab} is given by

$$v_{ab} = \frac{v_a - v_b}{1 - v_a v_b}, \tag{2.76}$$

which is a relativistic invariant. In the frame comoving with the shock this can be given by

$$v_{ab} = \pm\sqrt{\frac{(p_b - p_a)(e_b - e_a)}{(p_b + e_a)(p_a + e_b)}}. \tag{2.77}$$

In order to find e_b we solve for the enthalpy h_b as in equation (2.71) and then calculate the density from the equation of state (2.56). Finally $v_a(p_b)$ is given by

$$v_a(p_b) = \frac{v_a - v_{ab}}{1 - v_a v_{ab}}. \tag{2.78}$$

2.6 Ultrarelativistic hydrodynamics in more than one dimension

To my knowledge the explicit solution of the Riemann problem for ultrarelativistic hydrodynamics in more than one dimension appears for the first time here.

We define the vectors of conserved variables and fluxes to be

$$\begin{aligned} \mathbf{q} &= (\tau, S^1, S^2, S^3)^T, & \text{(2.79a)} \\ \mathbf{f}^{(i)} &= (S^i, S^1 v^i + p\delta^{1i}, S^2 v^i + p\delta^{2i}, S^3 v^i + p\delta^{3i})^T, & \text{(2.79b)} \end{aligned}$$

where the definitions

$$\tau = (\rho + p)W^2 - p, \quad S^i = (\rho + p)W^2 v^i \tag{2.80}$$

have been made. Then the equations have then conservation form

$$\partial_t \mathbf{q} + \partial_{x_i} \mathbf{f}^{(i)}(\mathbf{q}) = 0. \tag{2.81}$$

We have also redefined the Lorentz factor to take into account the transverse velocities. The extra definitions are

$$v^t = \left((v^y)^2 + (v^z)^2\right)^{1/2}, \quad v = \left((v^x)^2 + (v^t)^2\right)^{1/2}, \quad W = \left(1 - v^2\right)^{-1/2}, \tag{2.82}$$

where v^t is the transverse speed. We are considering the solution in the x direction. Symmetry considerations give the solution in the other directions.

2.6.1 Rarefactions

We start by looking for a similarity solution. We expand the conservation form (2.81) into the primitive form, discarding the y and z derivatives and simplifying as in section 2.4.1. Rewriting in terms of the characteristic variable ξ gives

$$\begin{aligned}
\left\{(1+c_s^2)W^2(\xi - v^x) - \xi c_s^2\right\} \frac{\mathrm{d}\rho}{\mathrm{d}\xi} + \rho(1+c_s^2)W^4 \cdot \\
\left\{\left(2\xi v^x - 1 - (v^x)^2 + (v^t)^2\right) \frac{\mathrm{d}v^x}{\mathrm{d}\xi} + (\xi - v^x)\frac{\mathrm{d}(v^t)^2}{\mathrm{d}\xi}\right\} &= 0 \quad (2.83a) \\
(\xi v^x - 1)c_s^2 \frac{\mathrm{d}\rho}{\mathrm{d}\xi} + \rho(1+c_s^2)W^2(\xi - v^x)\frac{\mathrm{d}v^x}{\mathrm{d}\xi} &= 0 \quad (2.83b) \\
\xi v^y c_s^2 \frac{\mathrm{d}\rho}{\mathrm{d}\xi} + \rho(1+c_s^2)W^2(\xi - v^x)\frac{\mathrm{d}v^y}{\mathrm{d}\xi} &= 0 \quad (2.83c) \\
\xi v^z c_s^2 \frac{\mathrm{d}\rho}{\mathrm{d}\xi} + \rho(1+c_s^2)W^2(\xi - v^x)\frac{\mathrm{d}v^z}{\mathrm{d}\xi} &= 0. \quad (2.83d)
\end{aligned}$$

We note that this reduces to the one dimensional case (2.39) in the trivial case $v^y = v^z = 0$. Solving for ξ gives

$$\xi = \frac{v^x(1-c_s^2) \pm c_s\sqrt{(1-v^2)\left[1 - v^2 c_s^2 - (v^x)^2(1-c_s^2)\right]}}{1 - v^2 c_s^2} \tag{2.84}$$

or $\xi = v^x$, which corresponds to the tangential discontinuity. This appears as a discontinuity in the transverse velocities $v^{y,z}$ only.

After some manipulation of the similarity equations we find that

$$\begin{aligned}
(\xi v^x - 1)c_s^2 \frac{\mathrm{d}\rho}{\mathrm{d}\xi} + \rho(1+c_s^2)W^2(\xi - v^x)\frac{\mathrm{d}v^x}{\mathrm{d}\xi} &= 0 \quad (2.85a) \\
\rho^\alpha W v^y &= \text{constant} \quad (2.85b) \\
\rho^\alpha W v^z &= \text{constant}. \quad (2.85c)
\end{aligned}$$

The last two equations give that

$$\rho^\alpha W v^t = \text{constant}, \tag{2.86}$$

which gives the transverse velocity across the rarefaction in terms of v_*^x as

$$v_*^t = \rho_K^\alpha W_K v_K^t \left[\frac{1-(v_*^x)^2}{\rho_*^{2\alpha} + \rho_K^{2\alpha} W_K^2 (v_K^t)^2} \right]^{1/2}. \tag{2.87}$$

If we call the constant in equation (2.86) C then we have an ordinary differential equation for v^x and ρ of the form

$$\frac{\alpha}{c_s} \frac{\rho^\alpha}{C^2 + \rho^{2\alpha}} \sqrt{C^2(1-c_s^2) + \rho^{2\alpha}} \frac{\mathrm{d}\rho}{\rho} = \frac{\mathrm{d}v^x}{1-(v^x)^2}. \tag{2.88}$$

Defining $r = \rho^\alpha$ this gives

$$\frac{\sqrt{C^2(1-c_s^2) + r^2}}{c_s(C^2 + r^2)} \mathrm{d}r = \frac{\mathrm{d}v^x}{1-(v^x)^2}, \tag{2.89}$$

which is integrable. Hence the Riemann invariants can be written

$$\begin{aligned} J_\pm = \left(\frac{1+v^x}{1-v^x} \right) \left[r + \sqrt{C^2(1-c_s^2) + r^2} \right]^{\pm \frac{2}{c_s}} \cdot \\ \left[\left(\frac{r\sqrt{C^2(1-c_s^2)+r^2} + r^2 + C^2(1+c_s)}{r\sqrt{C^2(1-c_s^2)+r^2} + r^2 + C^2(1-c_s)} \right) \right]^{\pm 1}, \end{aligned} \tag{2.90}$$

and the post shock velocity v_b is given by

$$v_b = \frac{A_\pm - 1}{A_\pm + 1}, \tag{2.91}$$

where $A_{\pm}$ is given by

$$A_{\pm} = \frac{(J_{\pm})_a}{\left(\bar{J}_{\pm}\right)_b}, \tag{2.92a}$$

$$\begin{aligned}\bar{J}_{\pm} = {} & \left[r + \sqrt{C^2(1-c_s^2)+r^2}\right]^{\pm\frac{2}{c_s}} \cdot \\ & \left[\left(\frac{r\sqrt{C^2(1-c_s^2)+r^2}+r^2+C^2(1+c_s)}{r\sqrt{C^2(1-c_s^2)+r^2}+r^2+C^2(1-c_s)}\right)\right]^{\pm 1}.\end{aligned} \tag{2.92b}$$

2.6.2 Shocks

We write out the Rankine–Hugoniot conditions for the system (2.79) as

$$V_s\left[\rho\left((1+c_s^2)W^2-c_s^2\right)\right] = \left[\rho(1+c_s^2)W^2v^x\right] \tag{2.93a}$$

$$V_s\left[\rho(1+c_s^2)W^2v^x\right] = \left[\rho\left((1+c_s^2)W^2(v^x)^2+c_s^2\right)\right] \tag{2.93b}$$

$$V_s\left[\rho(1+c_s^2)W^2v^y\right] = \left[\rho(1+c_s^2)W^2v^xv^y\right] \tag{2.93c}$$

$$V_s\left[\rho(1+c_s^2)W^2v^z\right] = \left[\rho(1+c_s^2)W^2v^xv^z\right]. \tag{2.93d}$$

After some manipulation the last two equations give

$$\left[\frac{v^y}{v^z}\right] = 0, \tag{2.94}$$

so that (as in the rarefaction case) the orientation of the transverse velocity cannot change. We also have that

$$v_b^{y,z} = \frac{W_a^2 v_a^{y,z}(v_a^x - V_s)}{W_a^2(v_a^x - V_s) + V_s\alpha(1-r)}, \tag{2.95}$$

where $r = \rho_a/\rho_b$ as before. Hence the transverse velocity is given by

$$v_b^t = \frac{W_a^2 v_a^t (v_a^x - V_s)}{W_a^2 (v_a^x - V_s) + V_s \alpha (1 - r)}. \tag{2.96}$$

By a similar approach we have the shock velocity V_s given by

$$V_s = \frac{r W_a^2 v_a^x (v_b^x - v_a^x) + \alpha (1 - r)}{r W_a^2 (v_b^x - v_a^x) + v_b^x \alpha (1 - r)}. \tag{2.97}$$

By combining these we can write

$$W_b^{-2} = \frac{\left(\alpha(r-1) + W_a^2 (1 - (v_a^x)^2)\right)^2 (1 - (v_b^x)^2) - W_a^4 (v_a^t)^2 (1 - v_a^x v_b^x)^2}{\left[(r-1)\alpha + W_a^2 (1 - (v_a^x)^2)\right]^2}. \tag{2.98}$$

This gives the quadratic equation for v_b^x as

$$\begin{aligned} &\left\{(\mu(r-1)^2 + r W_a^2) - \Lambda(\Theta + W_a^2 (v_a^x)^2)\right\} (v_b^x)^2 \\ -2W_a^2 v_a^x (r - \Lambda) v_b^x - &\left\{(\mu(r-1)^2 - r W_a^2 (v_a^x)^2) + \Lambda(W_a^2 - \Theta)\right\} = 0, \end{aligned} \tag{2.99}$$

where we have defined the additional auxiliary quantities

$$\begin{aligned} \Theta &= \alpha(r-1) + W_a^2 (1 - (v_a^x)^2) \\ &= \alpha(r-1) + 1 + W_a^2 (v_a^t)^2 && \text{(2.100a)} \\ \Lambda &= \frac{\alpha(r-1) W_a^2 (v_a^t)^2}{\Theta}. && \text{(2.100b)} \end{aligned}$$

In the one dimensional limit, Λ vanishes and we recover the solution in section 2.4.2.

2.7 Ideal gas relativistic hydrodynamics in more than one dimension

The solution of the Riemann problem for the ideal gas equation of state in more than one dimension was given by Pons et al. (2000). There is only one major change from the one dimensional case outlined in section 2.5.

We define the vectors of conserved variables and fluxes to be

$$\mathbf{q} = (D, S^1, S^2, S^3, \tau)^T, \tag{2.101a}$$

$$\mathbf{f}^{(i)} = (Dv^i, S^1 v^i + p\delta^{1i}, S^2 v^i + p\delta^{2i}, S^3 v^i + p\delta^{3i}, S^i - Dv^i)^T, \tag{2.101b}$$

where the definitions

$$D = \rho_0 W, \quad S^i = \rho_0 h W^2 v^i, \quad \tau = \rho_0 h W^2 - p - D \tag{2.102}$$

have been made. Then the equations have the conservation form (2.81). We have also redefined the Lorentz factor to take into account the transverse velocities. The extra definitions are

$$v^t = \left((v^y)^2 + (v^z)^2\right)^{1/2}, \quad v = \left((v^x)^2 + (v^t)^2\right)^{1/2}, \quad W = \left(1 - v^2\right)^{-1/2}. \tag{2.103}$$

We are considering the solution in the x direction. Symmetry considerations give the solution in the other directions.

2.7.1 Rarefactions

The main difference between the multidimensional and one dimensional cases is that the solution across a rarefaction wave is given in terms of an ordinary differential equation. The condition on the similarity variable ξ is

$$\xi = \frac{v^x(1-c_s^2) \pm c_s\sqrt{(1-v^2)\left[1-v^2c_s^2-(v^x)^2(1-c_s^2)\right]}}{1-v^2c_s^2}. \tag{2.104}$$

This leads to the ordinary differential equation

$$\frac{\mathrm{d}v^x}{\mathrm{d}p} = \pm\frac{1}{\rho_0 h W^2 c_s}\frac{1}{\sqrt{1+g(\xi_\pm, v^x, v^t)}}, \tag{2.105}$$

where the function g is given by

$$g(\xi_\pm, v^x, v^t) = \frac{(v^t)^2(\xi_\pm^2-1)}{(1-\xi_\pm v^x)^2}. \tag{2.106}$$

With the added constraint

$$v_b^t = h_a W_a v_a^t\left[\frac{1-(v_b^x)^2}{h_b^2+(h_a W_a v_a^t)^2}\right]^{1/2}, \tag{2.107}$$

the ordinary differential equation (2.105) can be written as a function of p_b and v_b only, allowing a numerical solution to be found. The analytic solution involves elliptic integrals and so is not very useful for numerical work.

2.7.2 Shocks

If the wave is a shock then the only difference is that the known Lorentz factor W_a contains terms from the transverse velocity. The equation for the velocity remains

$$v_a(p_b) = \frac{\frac{W_s}{j}(p_b - p_a) + h_a W_a v_a^x}{h_a W_a + (p_b - p_a)\left(\frac{W_s v_a^x}{j} + \frac{1}{W_a (\rho_0)_a}\right)}, \tag{2.108}$$

where the mass flux j is given by equation (2.72) and the shock velocity from equation (2.73).

2.8 Approximate Riemann solvers

Even the fastest exact Riemann solver still has to solve an implicit equation. As the solution of Riemann problems takes a large portion of the time to run in a HRSC code, approximations that speed the calculation of the boundary flux are often essential. This is especially true in higher dimensions, where the solution of the ordinary differential equation (2.105) makes the use of an exact Riemann solver impractical.

Approximate Riemann solvers can have problems, as shown in depth by Quirk (1994). Hence it is best where possible to compare the results of as many different solvers as possible. Here we shall give the three approximate solvers used in this thesis, starting with the simplest.

2.8.1 HLLE solver

The Harten-Lax-van Leer-Einfeldt (HLLE) solver of Einfeldt (1988) is a simple two wave approximation. We assume that the maximum wave speeds $\xi_\pm$ are known. The solution of the Riemann problem is then given by requiring conservation to hold across the waves. The solution takes the form

$$\mathbf{q} = \begin{cases} \mathbf{q}_L & \text{if } \xi < \xi_- \\ \mathbf{q}_* & \text{if } \xi_- < \xi < \xi_+ \\ \mathbf{q}_R & \text{if } \xi > \xi_+, \end{cases} \tag{2.109}$$

and the intermediate state $\mathbf{q}_*$ is given by

$$\mathbf{q}_* = \frac{\xi_+ \mathbf{q}_R - \xi_- \mathbf{q}_L - \mathbf{f}(\mathbf{q}_R) + \mathbf{f}(\mathbf{q}_L)}{\xi_+ - \xi_-}. \tag{2.110}$$

If we just want the numerical flux along the boundary then this takes the form

$$\mathbf{f}(\mathbf{q}) = \frac{\widehat{\xi}_+ \mathbf{f}(\mathbf{q}_L) - \widehat{\xi}_- \mathbf{f}(\mathbf{q}_R) + \widehat{\xi}_+ \widehat{\xi}_- (\mathbf{q}_R - \mathbf{q}_L)}{\widehat{\xi}_+ - \widehat{\xi}_-}, \tag{2.111}$$

where

$$\widehat{\xi}_- = \min(0, \xi_-), \quad \widehat{\xi}_+ = \max(0, \xi_+). \tag{2.112}$$

Knowledge of the precise minimum and maximum characteristic velocities $\xi_\pm$ requires knowing the solution of the Riemann problem. Instead, the characteristic velocities are usually found from the eigenvalues of the Jacobian matrix $\partial \mathbf{f} / \partial \mathbf{q}$ evaluated at some intermediate state. To ensure that the maximum and minimum eigenvalues over the entire range between the left and right states are found we have evaluated the Jacobian in both the left and right states and taken the maximum and minimum over all eigenvalues. This ensures, for the systems of equations considered here, that the real maximum and minimum characteristic velocities are contained within $[\xi_-, \xi_+]$.

If we set $\alpha = \max(|\xi_-|, |\xi_+|)$ and replace the characteristic velocities $\xi_\pm$ with $\pm\alpha$, we find the Lax–Friedrichs flux

$$\mathbf{f}(\mathbf{q}) = \tfrac{1}{2}\left(\mathbf{f}(\mathbf{q}_L) + \mathbf{f}(\mathbf{q}_R) + \alpha(\mathbf{q}_L - \mathbf{q}_R)\right). \tag{2.113}$$

This is very diffusive, but also very stable.

2.8.2 Roe solver

The linearized solver of Roe (1981) is probably the most popular approximate Riemann solver. The simplest interpretation is that the Jacobian $\partial \mathbf{f}/\partial \mathbf{q}$ is linearized about some intermediate state. Then the conservation form reduces to the linear equation

$$\partial_t \mathbf{q} + A\partial_x \mathbf{q} = 0, \tag{2.114}$$

where A is a constant coefficient matrix. This is identical to equation (2.17) and so all the results of section 2.2.3 on linear systems hold. We reiterate that the standard form for the flux along the characteristic ray $\xi = 0$ is

$$\mathbf{f}(\mathbf{q}) = \tfrac{1}{2}\left(\mathbf{f}(\mathbf{q}_L) + \mathbf{f}(\mathbf{q}_R) - \sum_{i=1}^{N} |\lambda_i| \Delta \mathbf{w}^i \mathbf{r}^i\right). \tag{2.115}$$

There is a choice of which intermediate state the Jacobian should be evaluated at. Roe gives three criteria that ensure the consistency and stability of the numerical flux:

1. $A(\mathbf{q}_{\text{Roe}})\,(\mathbf{q}_R - \mathbf{q}_L) = \mathbf{f}(\mathbf{q}_R) - \mathbf{f}(\mathbf{q}_L)$,

2. $A(\mathbf{q}_{\text{Roe}})$ is diagonalizable with real eigenvalues,

3. $A(\mathbf{q}_{\text{Roe}}) \to \partial \mathbf{f}/\partial \mathbf{q}$ smoothly as $\mathbf{q}_{\text{Roe}} \to \mathbf{q}$.

A true Roe average for relativistic hydrodynamics, i.e., an intermediate state that satisfies all these conditions, has been constructed by Eulderink and Mellema (1995). However, frequently it is sufficient to use

$$\mathbf{q}_{\text{Roe}} = \tfrac{1}{2}\left(\mathbf{q}_R + \mathbf{q}_L\right), \tag{2.116}$$

which satisfies only the last two conditions. For simplicity we have implemented this arithmetic average.

2.8.3 Marquina solver

Unlike all the other Riemann solvers introduced so far, the Marquina solver as outlined in Donat and Marquina (1996) does not solve the Riemann problem completely. Instead only the flux along the characteristic ray $\xi = 0$ is given. It can be seen as a generalized Roe solver, as the results are the same except at sonic points. These points are where the fluid velocity is equal to the speed of sound of the fluid. In the context of Riemann problems sonic points are found when the ray $\xi = 0$ is within a rarefaction wave.

Firstly define the left $\mathbf{l}(\mathbf{q}_{L,R})$ and right $\mathbf{r}(\mathbf{q}_{L,R})$ eigenvectors and the eigenvalues $\lambda(\mathbf{q}_{L,R})$ of the Jacobian matrix $\partial \mathbf{f}/\partial \mathbf{q}$ evaluated at the left and right states. Next define left and right characteristic variables $\mathbf{w}_{L,R}$ and fluxes

$\boldsymbol{\phi}_{L,R}$ by

$$\mathbf{w}^i_{L,R} = \mathbf{l}^i(\mathbf{q}_{L,R}) \cdot \mathbf{q}_{L,R}, \quad \boldsymbol{\phi}^i_{L,R} = \mathbf{l}^i(\mathbf{q}_{L,R}) \cdot \mathbf{f}(\mathbf{q}_{L,R}). \tag{2.117}$$

Then the algorithm chooses the correct sided characteristic flux if the eigenvalues $\lambda_i(\mathbf{q}_L)$, $\lambda_i(\mathbf{q}_R)$ have the same sign, and uses a Lax–Friedrichs type flux if they change sign. In full, the algorithm is given in figure 2.4.

For $i = 1, \dots, N$ **do**
- **If** $\lambda_i(\mathbf{q}_L)\lambda_i(\mathbf{q}_R) > 0$ **then**
  - **If** $\lambda_i(\mathbf{q}_L) > 0$ **then**
    - $\boldsymbol{\phi}^i_+ = \boldsymbol{\phi}^i_L$
    - $\boldsymbol{\phi}^i_- = 0$
  - **else**
    - $\boldsymbol{\phi}^i_+ = 0$
    - $\boldsymbol{\phi}^i_- = \boldsymbol{\phi}^i_R$
  - **end if**
- **else**
  - $\alpha^i = \max(|\lambda_i(\mathbf{q}_L), \lambda_i(\mathbf{q}_R)|)$
  - $\boldsymbol{\phi}^i_+ = \frac{1}{2}\left(\boldsymbol{\phi}^i_L + \alpha^i \mathbf{w}^i_L\right)$
  - $\boldsymbol{\phi}^i_- = \frac{1}{2}\left(\boldsymbol{\phi}^i_R - \alpha^i \mathbf{w}^i_R\right)$
- **end if**

end do

Figure 2.4: The algorithm to calculate the Marquina flux.

Then the numerical flux is given by

$$\mathbf{f}(\mathbf{q}) = \sum_{i=1}^{N} \left[\boldsymbol{\phi}^i_+ \mathbf{r}^i(\mathbf{q}_L) + \boldsymbol{\phi}^i_- \mathbf{r}^i(\mathbf{q}_R)\right]. \tag{2.118}$$

p_L	$(\rho_0)_L$	v_L	p_L	$(\rho_0)_L$	v_L
1	1	0	0.1	0.125	0

Table 2.1: The initial data for the relativistic version of the problem of Sod (1978).

2.8.4 Accuracy of approximate solvers

How accurate are these approximate solvers? If we look at their performance on a single Riemann problem the answer is typically that they are very inaccurate. This is illustrated by figure 2.5, where the complete solution for the relativistic Sod problem is shown. The Sod problem (Sod (1978)), defined originally for the non-relativistic ideal gas in one dimension in slab symmetry, is a Riemann problem with the initial data given by table 2.1 and the fluid constant $\gamma = 5/3$. This is a very mild test, and yet both the Roe solver and the HLLE solver are obviously inaccurate in their calculation of the intermediate state. They are of course inaccurate inside the rarefaction fan.

When used within a finite difference or finite volume method to solve *local* Riemann problems over very small regions of space and time these inaccuracies are considerably less important. This will be explained in detail in chapter 3. However, the inaccuracies caused by approximate Riemann solvers may still have an effect. Some of these problems, such as the HLLE solver not detecting the contact surface, can be rectified (for the HLLE case this has been implemented by Toro et al. (1994) in the HLLC solver). In order to remove all of these problems one must use the exact solver, which may frequently be too computationally expensive, or carefully test the approximate

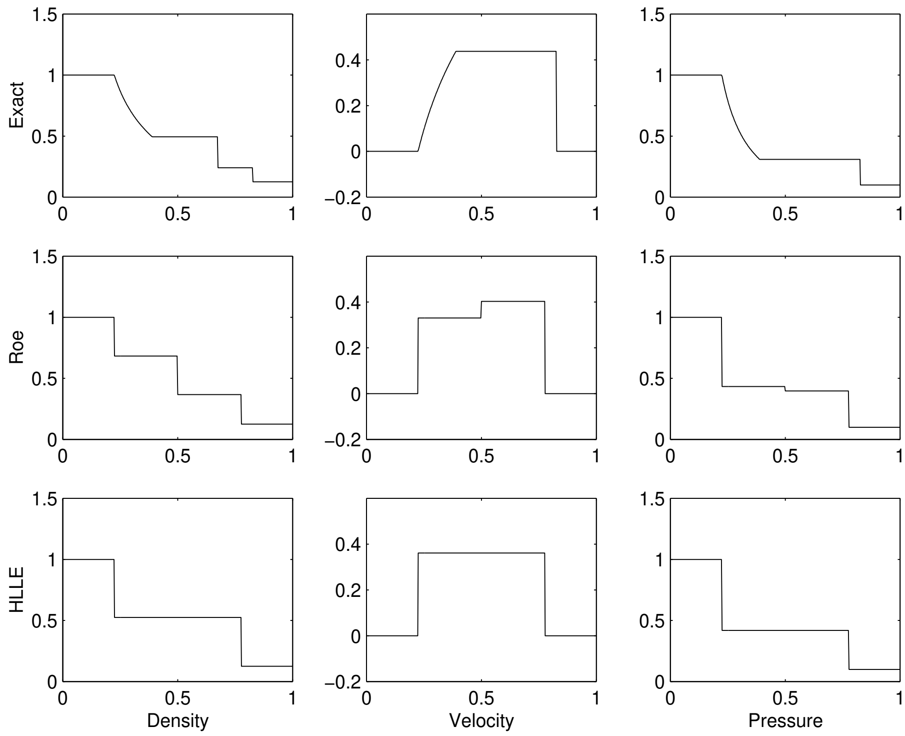

Figure 2.5: The solutions for the primitive variables for the relativistic Sod problem outlined in figure 2.3. The Roe solver linearizes about the arithmetic mean of the left and right states. The HLLE solver takes the maximum and minimum wave speeds from the eigenvalues of the Jacobian evaluated at the left and right states. The approximate solvers are inaccurate in many ways.

solver against the exact solver at all stages of the method.

2.9 Efficient numerical solutions

We shall concentrate on the solution of the Riemann problems for relativistic hydrodynamics in one dimension (although the techniques carry over directly to the multidimensional case). For these problems we are trying to find the two nonlinear wave speeds and the intermediate state. The speed of the contact discontinuity is given by the velocity in the intermediate state. As neither the velocity nor the pressure change across the contact surface we can construct all the unknown states from the intermediate pressure and the equation of state.

As mentioned in section 2.3.3 we are trying to solve the implicit equation

$$v_{L_*}(p_*) = v_{R_*}(p_*), \tag{2.119}$$

which is equivalent to finding a root of the function $f(p_*)$, with

$$f(p_*) = v_{L_*}(p_*) - v_{R_*}(p_*). \tag{2.120}$$

The states v_{L_*} and v_{R_*} are found by the appropriate relations across the left and right waves given an initial guess p_*. These relations are given in sections 2.4–2.7.

The Riemann problem is said to be solved by the solution of the implicit equation (2.120). This gives the value of the pressure in the intermediate state. All other quantities in the intermediate state can be derived from this. Then the wave speeds can be found.

Solving the implicit equation is an iterative process. This involves a number of steps. Here we split them into three main parts. The first, the equations, must be dealt with on a case by case basis. Whenever we give concrete examples these will be for the relativistic ideal gas of section 2.5. The other main steps are the numerical iterative method and the starting value for this method. One other issue, the criteria to stop the iteration, is dealt with alongside the iterative method.

2.9.1 The iterative method

To solve an implicit equation such as (2.120) various numerical approaches can be used. Broadly speaking these fall into two groups.

The first group is the interval bisection methods. Given an interval in which the root is known to lie, a point within the interval is chosen. The function f is calculated at this point and, by comparing the value with the values at the endpoints, a smaller interval within which the root lies is found. When either the interval or the value of f is sufficiently small, the root is said to be known. The method from this group that we shall use is the Brent–Djikaard solver which uses a combination of bisection and inverse quadratic and linear interpolation.

The second group calculates the function and its derivatives at individual points. Starting from an initial guess, the derivatives can be used to improve the estimate. When the change in p_* or the value of f is sufficiently small, the root is said to be known. The method from this group that we shall use is the Newton–Raphson method. Given a guess $p_*^{(n)}$, the function and its derivative are calculated, and a new guess $p_*^{(n+1)}$ is found from

$$p_*^{(n+1)} = p_*^{(n)} - \frac{f\left(p_*^{(n)}\right)}{\mathrm{d}f\left(p_*^{(n)}\right)/\mathrm{d}p_*}. \tag{2.121}$$

For relativistic hydrodynamics both these methods should work reasonably efficiently, as the function f has a unique root and is a monotonic function of p_*.

2.9.2 The starting value

Both of the numerical methods mentioned above need an initial guess. For a bisection method, this needs to be an interval in which the root is known to lie. For a Newton–Raphson type method any physical value will do.

In Newtonian hydrodynamics Toro (1999) studied the efficiency of various starting values. He finds that the exact value given by assuming the wave pattern contains two rarefactions is frequently the most efficient, even given the numerical expense to calculate it. However, this only requires the solution of a quadratic equation. In relativistic hydrodynamics the solution (as given by Rezzolla and Zanotti (2001)) requires the solution of a quartic equation, which is computationally expensive and prone to round-off errors. So, for our Newton–Raphson method the starting value is always the arithmetic average of the pressures of the initial states.

The Riemann solver of Martí and Müller (1999) found an initial interval by taking this arithmetic average and doubling or halving it until the root was known to lie in the interval. This method is simple, but can require many calculations of f in certain circumstances (e.g., strong shock heating).

A method of Rezzolla and Zanotti (2001) for determining the wave pattern in advance was outlined in section 2.3.3. The interval is then found from the initial data and the wave pattern. For example, if $p_L > p_R$ then we have

$$[p_{\min}, p_{\max}] = \begin{cases} [p_L, \infty] & \text{if the wave pattern is } \mathcal{S}_{\leftarrow}\mathcal{C}\mathcal{S}_{\rightarrow} \\ [p_R, p_L] & \text{if the wave pattern is } \mathcal{R}_{\leftarrow}\mathcal{C}\mathcal{S}_{\rightarrow} \\ [0, p_R] & \text{if the wave pattern is } \mathcal{R}_{\leftarrow}\mathcal{C}\mathcal{R}_{\rightarrow}. \end{cases} \tag{2.122}$$

In the actual numerical implementation, various details are slightly altered. Firstly, as noted in section 2.5.3, the Newton–Raphson solver may have problems in the limit $p_* \to p_a+$. To avoid problems with the shock calculation in the trivial case we reduce the starting value slightly,

$$p_*^{(0)} = (0.5 - \epsilon)(p_L + p_R), \tag{2.123}$$

where ϵ is a small number. This ensures that if the problem is trivial it will use the expression for a rarefaction which is regular as $p_* \to p_K$.

Secondly, if the interval is very small or trivial then numerical rounding may cause problems. Instead of using the tight bounds given by the Rezzolla–Zanotti procedure, we expand the interval slightly so that

$$[p_{\min}, p_{\max}] \to [(1 - \epsilon)p_{\min}, (1 + \epsilon)p_{\max}]\,. \tag{2.124}$$

2.9.3 Efficiency

For a large simulation in relativistic hydrodynamics the efficiency of the numerical method will be important. The more efficient the code, the higher the resolution for a given amount of computational resources. The efficiency of many popular approximate Riemann solvers in relativistic hydrodynamics was studied by Aloy et al. (1999). This is the first study of the efficiency of the exact solver in relativistic hydrodynamics.

We illustrate the efficiency of the different methods using five test problems for one dimensional ideal gas relativistic hydrodynamics. The tests are given in table 2.2. A further test, testing the efficiency of different Riemann solvers when used as part of an evolution method, will be described in section 3.10.2.

Test	γ	p_L	$(\rho_0)_L$	v_L	p_R	$(\rho_0)_R$	v_R
1	5/3	1	1	0	0.1	0.125	0
2	5/3	1000	1	0	0.01	1	0
3	5/3	10	1	-0.6	20	10	0.5
4	5/3	1	1	0.9	10	1	0
5	5/3	1	1	0	1.01	1.01	0.01

Table 2.2: Tests for the relativistic ideal gas Riemann solver.

The first test is the relativistic version of the test of Sod (1978). The next three are the tests given by Martí and Müller (1994). The wave patterns are rarefaction–shock, rarefaction–rarefaction and shock–shock respectively. The last is a near trivial problem typically solved by a HRSC method over most of the domain.

The tests were run on an SGI Origin R12000 with full optimization. The codes were as similar as possible. The times, given in seconds, are for the solution of 10^5 Riemann problems and are given in table 2.3. The Roe solver failed on the third test by calculating a negative pressure in the intermediate state.

As we can see, the Newton–Raphson and Rezzolla–Zanotti approaches

Test	MM	RZ	NR	Roe	HLLE
1	3.4	7.0	4.0	0.9	0.4
2	6.3	6.0	3.1	0.9	0.4
3	7.4	6.0	10.1	–	0.4
4	4.3	3.9	2.7	1.0	0.4
5	4.3	3.4	3.8	0.8	0.4

Table 2.3: Efficiency of the exact and approximate Riemann solvers.

are usually faster than the standard bisection approach used by Martí and Müller. The Newton–Raphson approach is only poor in the two rarefaction case. The Rezzolla–Zanotti approach is only poor for the Sod problem. Most importantly, both show a near 30% improvement in the near trivial case. However, the substantial advantages of the approximate solvers are obvious.

2.9.4 Efficiency of the multidimensional solver

Here we compare the efficiency of the multidimensional ultrarelativistic and ideal gas solvers. Although not solving the same problem, the difference in efficiency is illuminating. The initial data are

γ, Γ	p_L	$(\rho_0)_L, \rho_L$	v^x_L	p_R	$(\rho_0)_R, \rho_R$	v^x_R
5/3	1000	1500	0	1	1.5	0

This is a blast wave test similar to test 2 in section 2.9.3. The transverse velocities used, together with the times taken, are given in table 2.4. For the ultrarelativistic solver 10^5 Riemann problems were solved, whilst for the ideal gas solver just 100 Riemann problems were solved.

The ideal gas solver is between 2 and 3 orders of magnitude slower than the ultrarelativistic solver, and the ultrarelativistic solver is as fast (to within an order of magnitude) in higher dimensions as in one dimension.

v^t	Ultrarelativistic, 10^5 problems	Ideal gas 10^2 problems
0	7.0	1.6
0.5	7.9	3.5
0.9	11.7	8.6

Table 2.4: Efficiency of exact Riemann solvers in more than one dimension.

The Riemann solver codes tested here have not been optimized beyond machine optimization. This is true of the approximate solvers as well as the exact solvers. When used to solve local Riemann problems inside an evolution method as in chapter 3 it is possible that suitable "hand" optimizations may improve the efficiency of the exact Riemann solvers. As the approximate solvers are less complex it is likely that any efficiency improvements would have a smaller effect. Whether it is worth optimizing a Riemann solver instead of, for example, implementing a more accurate numerical method, will depend on the problem to be solved. For example, it might be practical to use the ultrarelativistic exact solver in a multidimensional HRSC calculation. It does not seem likely that it would be worth using the exact ideal gas solver, even after substantial optimizations.

Chapter 3

Numerical methods

As mentioned in the introduction to chapter 2, the basis of modern High Resolution Shock Capturing (HRSC) methods is the paper of Godunov (1959). He showed how the solution of the Riemann problem could be used to evolve a set of hyperbolic partial differential equations. He also showed that a linear method with order of accuracy greater than one could not be monotonicity preserving. That is to say it could introduce spurious numerical oscillations.

The extension to higher order methods was largely due to the work of van Leer, culminating in the introduction of useful second order accurate methods in van Leer (1979). The first useful extension beyond second order accuracy was given by the ENO schemes of Harten et al. (1987). To my knowledge all the methods that are applicable to general systems of hyperbolic conservation laws and are of higher than second order accuracy and do not introduce spurious oscillations, are based on the ENO methods. These other methods such as the WENO schemes of Jiang and Shu (1996) and the ADER scheme of Toro (2000) will not be considered here.

In this chapter we will introduce a number of numerical methods. Most will be based on the method of lines. This allows us to separate the effects of changing different parts of the method, such as the method of integrating in time of section 3.3, the method of reconstruction of section 3.4, and the method of calculating the numerical flux of section 3.5. We shall also consider some more compact methods, not based on the method of lines, where it is not possible to easily separate the algorithm in this fashion.

Finally we shall compare all the methods on a variety of different test cases. We perform these tests in order to find the best balance between accuracy and efficiency for the real problems that we wish to simulate in chapter 5. This is the largest comparative study of numerical methods applied to relativistic hydrodynamics that I am aware of. It is also the only study that I know of that directly compares the WAF and SLIC methods of Toro (1992), Toro and Billet (2000) with the ENO methods of Harten et al. (1987).

3.1 Notation

We are trying to solve numerically a hyperbolic system of partial differential equations in conservation law form,

$$\partial_t \mathbf{q} + \partial_x \mathbf{f}(\mathbf{q}) = \mathbf{s}(\mathbf{q}), \tag{3.1}$$

or more generally the integral form

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_{x_1}^{x_2} \mathbf{q}\,\mathrm{d}x = \mathbf{f}(\mathbf{q}(x_1,t)) - \mathbf{f}(\mathbf{q}(x_2,t)) + \int_{x_1}^{x_2} \mathbf{s}(\mathbf{q})\,\mathrm{d}x. \tag{3.2}$$

In order to solve this numerically we must approximate $\mathbf{q}$ with a discrete approximation $\mathbf{q}_i^n$. Here n and i represent the position of the approximation in time and space respectively. For the moment we will assume that the computational point $\mathbf{q}_i^n$ represents the integral average of the function $\mathbf{q}$ over the cell $I_i = [x_{i-1/2}, x_{i+1/2}]$, where $x_{i\pm1/2} = x_i \pm \frac{1}{2}\Delta x_i$, at time t^n. The grid will be considered in more detail in chapter 4.

Algorithms for solving (3.2) fall into two broad classes, *explicit* and *implicit* methods. As their names suggest, an explicit method obtains the data at time level $n+1$ in terms of known data at earlier time levels, whilst an implicit method relates data at time level $n+1$ to other unknown data at time level $n+1$, as well as the known data. Hence an implicit method typically requires a matrix inversion or a similarly computationally expensive operation, which usually makes implicit algorithms slower than explicit ones. This is compensated for by better stability properties and the possibility of taking larger time steps. Here we shall consider only explicit methods.

For explicit methods, there is a physical restriction on the length of time step that can be used. This is the Courant-Friedrichs-Lewy (CFL) condition. Let us consider one point of the numerical approximation $\mathbf{q}_i^n$. Expressed in the language of relativity, the CFL condition says that the numerical domain of dependence of $\mathbf{q}_i^n$ must include the physical domain of dependence of $\mathbf{q}_i^n$. Expressed in the language of Riemann problems, the CFL condition says that

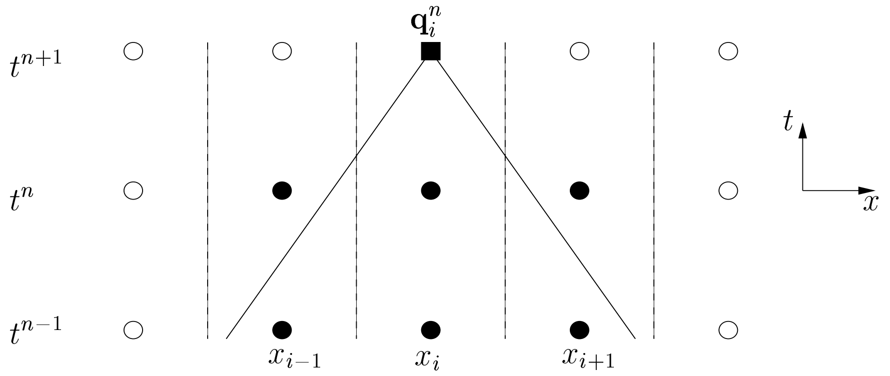

Figure 3.1: The CFL condition restricts the possible timestep $\Delta t = t^{n+1} - t^n$. The domain of dependence of $\mathbf{q}_i^n$ (the filled square) is traced back in time and space. The numerical domain of dependence that contains it is shown by the filled circles. The boundaries of the computational cells are given by the dashed lines, showing that some part of the data in all the cells in the numerical domain of dependence could affect the calculation of $\mathbf{q}_i^n$.

any computational cell that can be connected to $\mathbf{q}_i^n$ by a characteristic must be used by an algorithm calculating $\mathbf{q}_i^n$. This means that over any single timestep, a characteristic originating from one Riemann problem will not affect the boundary data for any other Riemann problem. These situations are shown in figure 3.1. The violation of the CFL condition typically leads to unphysical numerical oscillations. This restriction is typically written in terms of the grid spacing as

$$\Delta t \leq c \Delta x, \tag{3.3}$$

where c is the CFL number which depends on the algorithm and the data $\mathbf{q}^n$. When talking about an algorithm we typically give the CFL number that is required for the method to be stable. When talking about a numerical implementation the "CFL number" refers to the safety factor applied to the time step. This is in addition to any restrictions from the characteristic speeds of the data and the stability requirements of the algorithm.

Algorithms can also be classified by their *order of accuracy*. This describes how the error inherent in the numerical approximation changes as the grid spacing is changed. Formally, an algorithm is r^{th} order accurate if, given an exact solution at time level n, $\mathbf{q}(t^n)$, the error between the numerical approximation $\mathbf{q}^{n+1}$ and the exact solution $\mathbf{q}(t^{n+1})$ obeys

$$\mathbf{q}^{n+1} = \mathbf{q}(t^{n+1}) + \mathcal{O}\left((\Delta x)^{r+1}, (\Delta t)^{r+1}\right). \tag{3.4}$$

As the CFL condition means that it is necessary to take approximately $\mathcal{O}(\Delta t^{-1})$ steps in a typical calculation, this means that the global error (that is, the error over all the timesteps given exact initial data) obeys

$$\mathbf{q}^{N} = \mathbf{q}(t^{N}) + \mathcal{O}\left((\Delta x)^{r}, (\Delta t)^{r}\right). \tag{3.5}$$

Although the notion of order is applied to any numerical method solving equations like the hyperbolic system (3.2), we should mention that the analysis to prove that an algorithm actually obeys these rules is generally only performed for linear scalar equations with periodic boundaries and smooth data. There is no guarantee that these results hold for nonlinear equations, especially with complex boundary conditions.

3.1.1 Stability

To say that a numerical method is *stable* generically means that given specific initial data the evolution calculated by the numerical method will produce physical results. If the method is also *convergent*, that is the difference (in some sense) between the computed solution and the continuum solution reduces with increasing numerical resolution, then we might expect to calculate the correct answer.

Two sorts of error are of interest. The first is the *local truncation error*. Given an exact solution $\mathbf{q}$ to a differential equation, the local truncation error τ^n is given by

$$\sigma(t+\Delta t, x) = |\mathbf{q}(t+\Delta t, x) - \hat{\mathbf{q}}(t+\Delta t, x)|, \tag{3.6}$$

where $\hat{\mathbf{q}}(t+\Delta t, x)$ is calculated using the numerical method over a single timestep using the initial data $\mathbf{q}(t+\Delta t, x)$. The local truncation error can be estimated even if we do not know the exact solution (see section 4.4.4). A method is *consistent* if the local truncation error tends to zero as the grid spacing tends to zero. A linear method is said to have convergence order p (under certain restrictions - see LeVeque (1992)) if there exists a constant C_σ such that

$$|\sigma(t,x)| \leqslant C_\sigma (\Delta x)^p \tag{3.7}$$

for all grid spacings $\Delta x < (\Delta x)_0$.

The second error is the *global error*. Using the same notation, the global error $e(t)$ is given by

$$e(t) = |\mathbf{q}(t,x) - \hat{\mathbf{q}}(t,x)|, \tag{3.8}$$

where $\hat{\mathbf{q}}(t,x)$ is calculated using the numerical method over the range $t \in [t^0, t]$ using the initial data $\mathbf{q}(t^0, x)$. A method is said to be *convergent* if the global error tends to zero for any fixed t as the grid spacing is reduced to zero. However as the exact solution is frequently not known, a method is also said to be convergent if the error between numerically calculated solutions with different resolutions tends to zero in this limit. The Lax Equivalence theorem shows that for a linear method and linear equations, consistency and stability are necessary and sufficient for convergence (in the first sense). For nonlinear equations or methods, the Lax Wendroff theorem says only that if the solution converges then it will converge to a weak solution of the equations. This may not be the physical solution. For details, see LeVeque (1992).

The linear stability of a numerical method is usually checked using von Neumann analysis. The initial data is decomposed into Fourier modes and the evolution of each mode is checked individually for a single evolution step. The modes do not interact due to the linearity of the method. Then the evolution of the single wave

$$u_j = C_m \exp\left(\frac{2\pi I m j}{N}\right) \tag{3.9}$$

where m is the wave number, $I = \sqrt{-1}$, j the spatial position and N the number of grid points, can be written in the form

$$u_j^n = G^n e^{I\phi j} \tag{3.10}$$

where $\phi = 2\pi m/N$ and G is a function that depends only on ϕ and the method. here G^n means G to the power n, whereas u_j^n means the approximation to u and spatial position j and time level n. G is called the *amplification* *factor* and a method is linearly stable if $|G| < 1$ for all $\phi \in [-\pi, \pi]$. For many numerical methods this analysis gives the CFL condition.

It is often not possible to apply linear stability analysis to nonlinear methods. However, Godunov's theorem states that any linear method will not be monotonicity preserving if it is higher than first order accurate. Instead we must consider nonlinear stability. A number of nonlinear stability conditions are given by Laney (1998). Here we mention just three. A numerical method is *monotonicity preserving* if, for a scalar conservation law with monotone increasing (decreasing) initial data on an infinite spatial domain, the solution is monotone increasing (decreasing) for all time. Although this prevents spurious oscillations this condition may be too restrictive for many cases and does not cover non-monotone initial data.

The *total variation* of the numerical approximation $\mathbf{q}^n$ is given by

$$TV(\mathbf{q}^n) = \sup_i \sum \left| \mathbf{q}^n_{i+1} - \mathbf{q}^n_i \right| . \tag{3.11}$$

A numerical method is said to be *total variation diminishing* (TVD) if the total variation does not increase in any time step, i.e.,

$$TV(\mathbf{q}^{n+1}) \leq TV(\mathbf{q}^n). \tag{3.12}$$

A TVD method is monotonicity preserving.

A numerical method is said to be *Essentially Non-Oscillatory* (ENO) if the increase in total variation is proportional to some positive power r of the typical step length, i.e.,

$$TV(\mathbf{q}^{n+1}) \leq TV(\mathbf{q}^n) + \mathcal{O}\left((\Delta x)^r, (\Delta t)^r\right). \tag{3.13}$$

An ENO method need not be monotonicity preserving. However in the limit $\Delta x \to 0$ we recover the TVD condition which is monotonicity preserving.

We can also define the stability of a system of ordinary differential equations. Consider the linear system of ordinary differential equations

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{q} = A\mathbf{q} \tag{3.14}$$

where A is a constant $N \times N$ matrix. A numerical method will approximate this equation and the numerical solution $\mathbf{q}^n$ will satisfy

$$\mathbf{q}^{n+1} = g(\Delta t A)\mathbf{q}^n \tag{3.15}$$

where the function g will depend only on the method. Hence the method will be absolutely stable, that is any errors in the numerical solution will not grow exponentially, if $|g(z)| < 1$ for all complex z with negative real part. The region of stability of a numerical method is the set of points

$$\mathcal{S} = \{z \in \mathbb{C} : |g(z)| < 1\}. \tag{3.16}$$

Here we shall say that a method is stable if a neighbourhood of the origin is contained in the region of stability. This implies that there exists a timestep Δt such that errors will not grow exponentially. Although this will only be true for linear constant coefficient systems in practice it holds for reasonable nonlinear systems as well, where the Jacobian matrix is used in place of A. For more details, see Shampine and Gordon (1975), Stoer and Burlisch (1993).

3.2 Method of Lines

The method of lines reduces the solution of the partial differential equation (3.1) to the solution of an ordinary differential equation. This is done by assuming that the discretization of the function $\mathbf{q}$ is a *finite volume* one. That is,

$$\mathbf{q}_i = \int_{x_i - \frac{1}{2}\Delta x_i}^{x_i + \frac{1}{2}\Delta x_i} \mathbf{q} \, \mathrm{d}x. \tag{3.17}$$

The the weak form of the conservation law (3.2) gives the ordinary differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{q}_i = \mathbf{f}(\mathbf{q}(x_i - \tfrac{1}{2}\Delta x_i, t)) - \mathbf{f}(\mathbf{q}(x_i + \tfrac{1}{2}\Delta x_i, t)) + \mathbf{s}_i. \tag{3.18}$$

In order to solve this ordinary differential equation we require four things. The first is a pointwise ordinary differential equation solver such as the Runge–Kutta method. The second is a way of reconstructing the function $\mathbf{q}$ at the spatial points $x_i \pm \frac{1}{2}\Delta x_i$. The third is a method of calculating the numerical flux $\mathbf{f}$ at these points given the reconstruction. Finally we require some method for dealing with the source term $\mathbf{s}$. We will deal with each in turn.

3.3 Ordinary differential equation solvers

We are trying to solve the ordinary differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{q} = \mathbf{F}(\mathbf{q}). \tag{3.19}$$

We shall assume that $\mathbf{q}$ is known for $t < t^n$, and that we wish to know the solution at time t^{n+1}. We shall denote $t^{n+1} - t^n$ by h. Then there are a number of ways to solve equation (3.19) numerically.

3.3.1 Runge–Kutta methods

The simplest are the Runge–Kutta methods. These take a number of first order accurate steps which when combined give a higher order of accuracy. In full generality the N^{th} order Runge–Kutta method is written

$$\begin{aligned} \mathbf{q}^{(0)} &= \mathbf{q}^n, \\ \mathbf{q}^{(i)} &= \sum_{k=0}^{i-1} \left(\alpha_{ik}\mathbf{q}^{(k)} + \Delta t \beta_{ik}\mathbf{F}(\mathbf{q}^{(k)})\right), \qquad i = 1, \dots, N, \\ \mathbf{q}^{n+1} &= \mathbf{q}^{(N)}. \end{aligned} \tag{3.20}$$

If all the coefficients α_{ik}, β_{ik} are nonnegative then this is just a convex combination of explicit Euler forward difference operators with Δt replaced by $\beta_{ik}\Delta t/\alpha_{ik}$. The CFL constraint thus restricts the overall timestep by a factor c with

$$c = \min_{i,k} \frac{\alpha_{ik}}{\beta_{ik}}. \tag{3.21}$$

It has been shown by Shu and Osher (1988) that these methods are TVD (that is, TVD when applied to the method of lines with general operator F) with the CFL coefficient c provided that the coefficients α_{ik}, β_{ik} are nonnegative. Explicit TVD second and third order methods have been found. However it has been shown by Gottlieb and Shu (1998) that no TVD fourth order method exists. Higher order TVD methods do exist (a fifth order method was constructed by Gottlieb and Shu (1998)) but those constructed so far have restrictive CFL conditions and are very complex.

We will use only the optimal second, third and fourth order methods. The second order method can be written in the form

$$\begin{aligned} \mathbf{q}^{(1)} =& \mathbf{q}^n + \Delta t \mathbf{F}(\mathbf{q}^n), \\ \mathbf{q}^{n+1} =& \tfrac{1}{2}\left(\mathbf{q}^n + \mathbf{q}^{(1)} + \Delta t \mathbf{F}(\mathbf{q}^{(1)})\right). \end{aligned} \tag{3.22}$$

We note that this is equivalent to Heun's predictor corrector method. The CFL constraint is $c = 1$.

The third order method can be written in the form

$$\begin{aligned} \mathbf{q}^{(1)} =& \mathbf{q}^n + \Delta t \mathbf{F}(\mathbf{q}^n), \\ \mathbf{q}^{(2)} =& \tfrac{1}{4}\left(3\mathbf{q}^n + \mathbf{q}^{(1)} + \Delta t \mathbf{F}(\mathbf{q}^{(1)})\right), \\ \mathbf{q}^{n+1} =& \tfrac{1}{3}\left(\mathbf{q}^n + 2\mathbf{q}^{(2)} + 2\Delta t \mathbf{F}(\mathbf{q}^{(2)})\right). \end{aligned} \tag{3.23}$$

Again the CFL constraint is $c = 1$.

The fourth order method is not strictly TVD. However it seems stable for most numerical purposes, has a CFL constraint of $c = 1$ and is considerably simpler than other fourth order TVD methods. It can be written in the form

$$\begin{aligned}
\mathbf{q}^{(1)} =& \mathbf{q}^n + \tfrac{1}{2}\Delta t \mathbf{F}(\mathbf{q}^n),\\
\mathbf{q}^{(2)} =& \mathbf{q}^n + \tfrac{1}{2}\Delta t \mathbf{F}(\mathbf{q}^{(1)}),\\
\mathbf{q}^{(3)} =& \mathbf{q}^n + \Delta t \mathbf{F}(\mathbf{q}^{(2)}),\\
\mathbf{q}^{n+1} =& \tfrac{1}{6}\left(-2\mathbf{q}^n + 2\mathbf{q}^{(1)} + 4\mathbf{q}^{(2)} + 2\mathbf{q}^{(3)} + \Delta t \mathbf{F}(\mathbf{q}^{(3)})\right).
\end{aligned} \tag{3.24}$$

The Runge–Kutta methods given here are stable (that is there exists a stability region containing the origin, as explained in section 3.1), explicit and self-starting (that is only the data at $\mathbf{q}^n$ is required).

3.3.2 Gaussian quadrature

By far the most efficient method is that of Gaussian quadrature. However this requires knowing the form of $\mathbf{F}$ at points between times t^n and t^{n+1}. For example, the second order method is the midpoint rule

$$\mathbf{q}^{n+1} = \mathbf{q}^n + \Delta t \mathbf{F}(\mathbf{q}^{n+\frac{1}{2}}). \tag{3.25}$$

In certain circumstances it is possible to find the function $\mathbf{F}$ at these points using other information. For instance, when solving the hyperbolic system of conservation laws (3.1) we can use this system in order to give a time evolution equation for $\mathbf{F}$ in terms of spatial derivatives of $\mathbf{q}$, which we can approximate as we shall see in detail in section 3.4. Hence an approximation for the function $\mathbf{F}$ can be given inside the cell. However, for a large system of equations or for high order methods this approach rapidly becomes excessively complex as shown by Harten et al. (1987).

3.3.3 Adams–Bashforth methods

The Adams–Bashforth methods are explicit multistep methods. That is, they require data from time levels before t^n. Hence some other method (such as a Runge–Kutta method) is required to start. However once this data is given they are much more compact than a Runge–Kutta method for a given order of accuracy. The second order method can be written

$$\mathbf{q}^{n+1} = \mathbf{q}^n + \tfrac{1}{2}\Delta t \left(3\mathbf{F}(\mathbf{q}^n) - \mathbf{F}(\mathbf{q}^{n-1})\right), \tag{3.26}$$

and the third order method can be written

$$\mathbf{q}^{n+1} = \mathbf{q}^n + \tfrac{1}{12}\Delta t \left(23\mathbf{F}(\mathbf{q}^n) - 16\mathbf{F}(\mathbf{q}^{n-1}) + 5\mathbf{F}(\mathbf{q}^{n-2})\right). \tag{3.27}$$

This method is stable (see Shampine and Gordon (1975)) and imposes no extra restrictions on the timestep.

3.3.4 Adams–Moulton methods

The Adams–Moulton methods are implicit methods that sometimes require data from time levels before t^n. This makes them very stable but except in very special circumstances they are slow (as a matrix inversion must be performed). The second order method is written

$$\mathbf{q}^{n+1} = \mathbf{q}^n + \tfrac{1}{2}\Delta t \left(\mathbf{F}(\mathbf{q}^{n+1}) + \mathbf{F}(\mathbf{q}^n)\right), \tag{3.28}$$

and the third order method is written

$$\mathbf{q}^{n+1} = \mathbf{q}^n + \tfrac{1}{12}\Delta t\left(5\mathbf{F}(\mathbf{q}^{n+1}) + 8\mathbf{F}(\mathbf{q}^n) - \mathbf{F}(\mathbf{q}^{n-1})\right). \tag{3.29}$$

As this method is implicit there is no theoretical restriction to the timestep. It is also stable.

3.3.5 Milne–Simpson methods

The Milne–Simpson methods are implicit multistep methods and hence are slow except in very special circumstances. They are an application of the trapezium rule. The third order method is written

$$\mathbf{q}^{n+1} = \mathbf{q}^{n-1} + \tfrac{1}{3}\Delta t\left(\mathbf{F}(\mathbf{q}^{n+1}) + 4\mathbf{F}(\mathbf{q}^n) + \mathbf{F}(\mathbf{q}^{n-1})\right). \tag{3.30}$$

As this method is implicit there is no theoretical restriction to the timestep. It is also stable.

3.4 Reconstructions

We assume that we are given $\mathbf{q}_i$, where each of these data represents an integral average as in equation (4.1). We want to reconstruct the function $\mathbf{q}$ as a function of x over the entire domain in such a way that we minimize any numerical oscillations.

3.4.1 First order reconstruction

The simplest reconstruction possible is that the function is piecewise constant across each cell. That is

$$\mathbf{q}(x) = \mathbf{q}_i, \qquad x \in \left[x_i - \tfrac{1}{2}\Delta x_i, x_i + \tfrac{1}{2}\Delta x_i\right]. \tag{3.31}$$

This is first order accurate and TVD.

3.4.2 Slope limited reconstruction

The next simplest reconstruction is that the function is piecewise linear across each cell. That is

$$\mathbf{q}(x) = \mathbf{q}_i + \mathbf{S}_i(x - x_i), \qquad x \in \left[x_i - \tfrac{1}{2}\Delta x_i, x_i + \tfrac{1}{2}\Delta x_i\right], \tag{3.32}$$

where $\mathbf{S}_i$ is a "slope".

We need to choose the slope so that the reconstruction is as accurate as possible without introducing spurious oscillations. Three possible slopes are

$$\mathbf{S}_i^{\text{upw}} = \frac{\mathbf{q}_{i+1} - \mathbf{q}_i}{x_{i+1} - x_i}, \tag{3.33a}$$

$$\mathbf{S}_i^{\text{dow}} = \frac{\mathbf{q}_i - \mathbf{q}_{i-1}}{x_i - x_{i-1}}, \tag{3.33b}$$

$$\mathbf{S}_i^{\text{ave}} = \tfrac{1}{2}\left(\mathbf{S}_i^{\text{upw}} + \mathbf{S}_i^{\text{dow}}\right). \tag{3.33c}$$

These are the upwind, downwind and average slopes across the cell. All of these are second order accurate. However none of them will avoid introducing oscillations near a shock. Godunov (1959) has shown that any linear method which is higher than first order accurate will introduce oscillations. This motivates the technique of *slope limiting.*

The idea behind slope limiting is to monitor the shock position. Then in the neighbourhood of a shock the reconstruction is reduced to first order, whilst in the smooth regions the higher order reconstruction can be used. To do this we reconstruct the slope $\mathbf{S}_i$ in the form

$$\mathbf{S}_i = \phi(r)\mathbf{S}_i^{\mathrm{ave}}, \tag{3.34}$$

where r is the ratio of the upwind and downwind slopes. Whenever this ratio is near unity the function must be smooth, so $\phi(r)$ should be near 1. However whenever the ratio is far from unity there is probably a discontinuity, so $\phi(r)$ should be near 0. Some examples (from the most to the least diffusive) include a minmod type

$$\phi = \begin{cases} \min(1, \frac{4}{1+r}) & \text{if } r > 0 \\ 0 & \text{otherwise,} \end{cases} \tag{3.35}$$

the Van Leer monotized centred (MC) limiter

$$\phi = \begin{cases} \min(\frac{2r}{1+r}, \frac{2}{1+r}) & \text{if } r > 0 \\ 0 & \text{otherwise,} \end{cases} \tag{3.36}$$

and the Superbee limiter

$$\phi = \begin{cases} \min(2, r, \frac{2}{1+r}) & \text{if } r \geq 1 \\ 1 & \text{if } \frac{1}{2} \leq r < 1 \\ 2r & \text{if } 0 < r < \frac{1}{2} \\ 0 & \text{otherwise.} \end{cases} \tag{3.37}$$

These reconstructions are TVD and are second order accurate in smooth regions.

3.4.3 Essentially Non-Oscillatory reconstruction

Although TVD methods are second order accurate (except near discontinuities) they do not in general extend to higher orders of accuracy. This is because the construction of the limiter function typically relies on second order accuracy as explained in depth by Toro (1999). To construct higher order methods we must dispense with the TVD constraint (3.12) and instead use the Essentially Non-Oscillatory (ENO) constraint (3.13). Rewritten solely in terms of spatial reconstructions this is

$$TV(\mathbf{q}^{n+1}) \leq TV(\mathbf{q}^{n}) + \mathcal{O}\left((\Delta x)^k\right), \tag{3.38}$$

where k is the order of the reconstruction, a positive integer. This allows oscillations near extrema, but they should be small and bounded. An ENO reconstruction retains high order accuracy even near discontinuities and avoids the "clipping" of extrema found in TVD methods.

The first practical ENO reconstruction was given by Harten et al. (1987). Here we follow the method outlined in Shu (1997).

The essential idea is to choose a stencil including r cells in such a way as to minimize the possibility of including a discontinuity in the reconstruction. We also want the reconstruction to be conservative. Suppose we are reconstructing the scalar function q in cell i. We start with the cell I_i. We then add to the stencil cell I_j, where $j = i \pm 1$, where we choose j to minimize the Newton divided differences

$$\begin{aligned} q\left[x_{i-1}, x_i\right] =& \frac{q_i - q_{i-1}}{x_{i+\frac{1}{2}} - x_{i-1-\frac{1}{2}}} \\ q\left[x_i, x_{i+1}\right] =& \frac{q_{i+1} - q_i}{x_{i+1+\frac{1}{2}} - x_{i-\frac{1}{2}}}. \end{aligned} \tag{3.39}$$

We then recursively add more cells, minimizing the higher order Newton divided differences $q\left[x_{i-1}, \ldots, x_{i+j}\right]$ defined by

$$q\left[x_{i-1}, \ldots, x_{i+j}\right] = \frac{q\left[x_i, \ldots, x_{i+j}\right] - q\left[x_{i-1}, \ldots, x_{i+j-1}\right]}{x_{i+j+\frac{1}{2}} - x_{i-1-\frac{1}{2}}}. \tag{3.40}$$

The reconstruction at the cell boundary is given by a standard k^{th} order polynomial interpolation on the chosen stencil.

Shu (1997) has outlined as elegant way of calculating the cell boundary values solely in terms of the stencil and the known data. If the stencil is given by

$$S(i) = \left\{I_{i-r}, \ldots, I_{i+k-r-1}\right\}, \tag{3.41}$$

for some integer r, then there exist constants c_{rj} depending only on the grid

x_i such that the boundary values for cell I_i are given by

$$q_{i+1/2} = \sum_{j=0}^{k-1} c_{rj} q_{i-r+j}, \qquad q_{i-1/2} = \sum_{j=0}^{k-1} c_{r-1,j} q_{i-r+j}. \tag{3.42}$$

The constants c_{rj} are given by the rather complicated formula

$$c_{rj} = \left\{ \sum_{m=j+1}^{k} \frac{\sum_{l=0,l\neq m}^{k} \prod_{q=0,q\neq m,l}^{k} \left(x_{i+1/2} - x_{i-r+q-1/2} \right)}{\prod_{l=0,l\neq m}^{k} \left(x_{i-r+m-1/2} - x_{i-r+L-1/2} \right)} \right\} \Delta x_{i-r+j}. \tag{3.43}$$

This simplifies considerably if the grid is even. The coefficients for an even grid are given (up to seventh order) by Shu (1997).

If we are reconstructing a system $\mathbf{q}$ rather than a single variable we need to be more careful. We can apply ENO reconstruction *componentwise*, that is to each variable in turn. However, for a system any discontinuity will propagate with speed equal to one of the characteristic speeds, i.e., the discontinuity occurs in one (or more) of the characteristic variables, the others remaining smooth. This indicates that there would be a considerable gain in accuracy if we perform the reconstruction in the characteristic variables themselves, and convert back to the conserved variables.

This *characteristicwise* reconstruction is performed as follows. For a given cell I_i the characteristic variables $\mathbf{w}_j$ are defined by

$$\mathbf{w}_j^m = \mathbf{l}^m(\mathbf{q}_i) \cdot \mathbf{q}_j, \tag{3.44}$$

for $j = i-k-1, \ldots, i+k+1$, where the left eigenvectors $\mathbf{l}(\mathbf{q}_i)$ are calculated by the same procedure as was used for the linear system of section 2.2.3.

This provides enough data to reconstruct all the primitive variables at the cell boundaries, $\mathbf{w}_{i\pm1/2}$. We then convert back using the consistent right eigenvectors

$$\mathbf{q}_{i\pm1/2} = \mathbf{r}(\mathbf{q}_i) \cdot \mathbf{w}_{i\pm1/2}. \tag{3.45}$$

This last step seems inaccurate as the right eigenvectors are not defined at the same spatial position as the reconstruction. An alternative used by Shu (1997) is to use some average of the data in cells I_i and I_{i+1} to produce approximate left and right eigenvectors at the cell boundary $x_{i+1/2}$. The reconstructions at this boundary are then reconstructed using the relevant characteristic variables. In practice we have found no advantage to this approach and so have used the method outlined in equations (3.44)–(3.45). More complex approaches such as Fedkiw et al. (1998) which reconstruct the Jacobian itself at the cell boundary seem to be more successful in certain situations, but we have not found it necessary to implement them.

3.5 The numerical flux

The method we use to calculate the numerical flux $\mathbf{f}_{i+1/2}$ through a cell boundary is outlined as follows:

- Reconstruct the data on cells I_i and I_{i+1} to give two approximations $\mathbf{q}^{\pm}_{i+1/2}$ at the cell boundary.

- Solve the Riemann problem given initially piecewise constant states $\mathbf{q}^{\pm}_{i+1/2}$ using any appropriate exact or approximate Riemann solver from chapter 2 to get the state along the cell boundary.

- Calculate the flux from the cell boundary state.

We note that in the case of the Marquina solver of section 2.8.3 we calculate the boundary flux directly without calculating the boundary state. It is actually more efficient to calculate the flux directly for all the approximate solvers of section 2.8.

This step is only first order accurate, due to the piecewise constant approximation in the second step. There are a number of ways to ensure that this does not lower the order of accuracy of the entire method. Here we outline just two.

The first is the Cauchy–Kowalewski procedure (sometimes called the Lax–Wendroff method) which assumes that the function is piecewise continuous and gives the function along the cell boundary in terms of a high order expansion. For example, a third order expansion would be

$$\begin{aligned}\mathbf{q}(x,t) \approx \mathbf{q}(x_i,t^n) + &\frac{\partial \mathbf{q}}{\partial t}(x_i,t^n)(t-t^n) + \frac{\partial \mathbf{q}}{\partial x}(x_i,t^n)(x-x_i) + \\ &\frac{\partial^2 \mathbf{q}}{\partial t^2}(x_i,t^n)\frac{(t-t^n)^2}{2} + \frac{\partial^2 \mathbf{q}}{\partial t \partial x}(x_i,t^n)(t-t^n)(x-x_i) + \\ &\frac{\partial^2 \mathbf{q}}{\partial x^2}(x_i,t^n)\frac{(x-x_i)^2}{2}.\end{aligned} \tag{3.46}$$

The spatial derivatives can be evaluated directly from the reconstruction - this is especially simple if the reconstruction is written as a Taylor series. The time and mixed derivatives are evaluated using (derivatives of) the original equation and known data. For example,

$$\begin{aligned}\frac{\partial \mathbf{q}}{\partial t} &= -\frac{\partial \mathbf{f}(\mathbf{q})}{\partial x}, \\ \frac{\partial^2 \mathbf{q}}{\partial x \partial t} &= -\frac{\partial^2 \mathbf{f}(\mathbf{q})}{\partial x^2},\end{aligned} \tag{3.47}$$

and so on. By evaluating the expansion (3.46) at the correct points in time we can produce accurate initial data for an ordinary differential equation integrator.

This, combined with the Gaussian quadrature procedure of section 3.3.2, was the method outlined in the original ENO paper of Harten et al. (1987). It is compact and efficient but for complicated systems of equations the algebra required rapidly becomes unmanageable.

Although I implemented second order versions of the Cauchy–Kowalewski approach for scalar equations and the equations of ultrarelativistic hydrodynamics, results are not presented here. This is because the results were very similar, as expected, to those found using the Runge-Kutta approach outlined immediately below. The small advantages in computational expense were outweighed by the large increase in complexity when generalizing to other systems of equations or to higher order methods.

The second is to note that first order accuracy in time is sufficient if we are using Runge–Kutta time integration as outlined in section 3.3.1. Although each individual step is first order accurate in time, the combination of all the steps has a higher order of accuracy. This does require temporary storage space for a number of intermediate grids (for the fourth order method, three temporary grids are required), so it is not as efficient as the Cauchy–Kowalewski approach. However, the implementation of a Runge–Kutta method is particularly simple and so we have used it for all the results in this work.

3.6 Source terms

The time evolution of the source terms is given by the method of lines using equation (3.18), provided we can give a numerical estimate to the source $\mathbf{s}_i$ over the cell. The simplest method is just to use the source applied to the cell average value,

$$\mathbf{s}_i = \mathbf{s}(\mathbf{q}_i), \tag{3.48}$$

which is second order accurate.

In order to get a higher order of accuracy we have to evolve the reconstruction of the function $\mathbf{q}$ over the cell and then average the result over the cell. One possibility is to use some sort of numerical integration over the cell, such as Gaussian quadrature (see section 3.3.2). One third order method that was implemented was to use

$$\mathbf{s}_i = \tfrac{1}{6}\left(\mathbf{s}(\mathbf{q}(x_{i-1/2})) + 4\mathbf{s}(\mathbf{q}(x_i)) + \mathbf{s}(\mathbf{q}(x_{i+1/2}))\right), \tag{3.49}$$

where $\mathbf{q}(x_{i\pm1/2})$ and $\mathbf{q}(x_i)$ are found using any third order reconstruction (note that for a third order or higher reconstruction $\mathbf{q}(x_i) \neq \mathbf{q}_i$ in general). We found no advantages to using this higher order reconstruction. The results were nearly identical to the simple second order method of (3.48) and it was slightly less robust. All results use the simple second order method.

3.7 Other numerical methods

In this section we outline two numerical methods which are not based on the method of lines. These methods make a useful comparison against the numerical methods given in the rest of this chapter. In this section we will assume that there are no source terms and that the grid is even.

3.7.1 SLIC

The Slope Limited Centred (SLIC) method of Toro (1999), Toro and Billet (2000) is a second order TVD method which does not require the explicit solution of the Riemann problem. It gives similar results to a second order method using a Lax–Friedrichs flux.

The first step is a slope limited reconstruction, performed exactly as in section 3.4.2. For every cell I_i this gives the boundary values $\mathbf{q}_{i\pm1/2}$.

The second step is an evolution to an intermediate value. For every cell I_i the boundary values are evolved to intermediate boundary values $\bar{\mathbf{q}}_{i\pm1/2}$ using the update formula

$$\bar{\mathbf{q}}_{i\pm1/2} = \mathbf{q}_{i\pm1/2} + \tfrac{1}{2}\frac{\Delta t}{\Delta x}\left[\mathbf{f}(\mathbf{q}_{i-1/2}) - \mathbf{f}(\mathbf{q}_{i+1/2})\right]. \tag{3.50}$$

This step only involves the reconstructed data in cell I_i.

The final step is an evolution from the intermediate values using the First Order Centred (FORCE) flux. This is an average of the Richtmyer and Lax–Friedrichs fluxes. We first define $\mathbf{q}^L_{i+1/2}$ to be the intermediate value $\bar{\mathbf{q}}_{i+1/2}$ from cell I_i and $\mathbf{q}^R_{i+1/2}$ to be the intermediate value $\bar{\mathbf{q}}_{i+1/2}$ from cell I_{i+1}.

Then the Richtmyer "velocity" is given by

$$\mathbf{q}^{\mathrm{Ri}}_{i+1/2} = \tfrac{1}{2}\left[\left(\mathbf{q}^{L}_{i+1/2} + \mathbf{q}^{R}_{i+1/2}\right) + \frac{\Delta t}{\Delta x}\left(\mathbf{f}(\mathbf{q}^{L}_{i+1/2}) - \mathbf{f}(\mathbf{q}^{R}_{i+1/2})\right)\right], \tag{3.51}$$

which gives the Richtmyer flux $\mathbf{f}^{\mathrm{Ri}}$ as

$$\mathbf{f}^{\mathrm{Ri}}_{i+1/2} = \mathbf{f}\left(\mathbf{q}^{\mathrm{Ri}}_{i+1/2}\right). \tag{3.52}$$

When the Lax–Friedrichs flux was introduced in section 2.8.1 it depended on the maximum wave speed. However, this is constrained by the CFL condition so we can (at the cost of slightly increased dissipation) replace this by the ratio of the space and time steps, giving the Lax–Friedrichs flux $\mathbf{f}^{\mathrm{LF}}$ as

$$\mathbf{f}^{\mathrm{LF}}_{i+1/2} = \tfrac{1}{2}\left[\mathbf{f}(\mathbf{q}^{L}_{i+1/2}) + \mathbf{f}(\mathbf{q}^{R}_{i+1/2}) + \frac{\Delta x}{\Delta t}(\mathbf{q}^{L}_{i+1/2} - \mathbf{q}^{R}_{i+1/2})\right]. \tag{3.53}$$

Finally, the FORCE flux is given by

$$\mathbf{f}^{\mathrm{FORCE}}_{i+1/2} = \tfrac{1}{2}\left(\mathbf{f}^{\mathrm{Ri}}_{i+1/2} + \mathbf{f}^{\mathrm{LF}}_{i+1/2}\right). \tag{3.54}$$

The main advantage of the SLIC scheme is that no information apart from the form of the system of equations (3.1) is necessary. It is also efficient in comparison to an exact Riemann solver based scheme. However, it is quite diffusive when compared to an upwind scheme such as WAF.

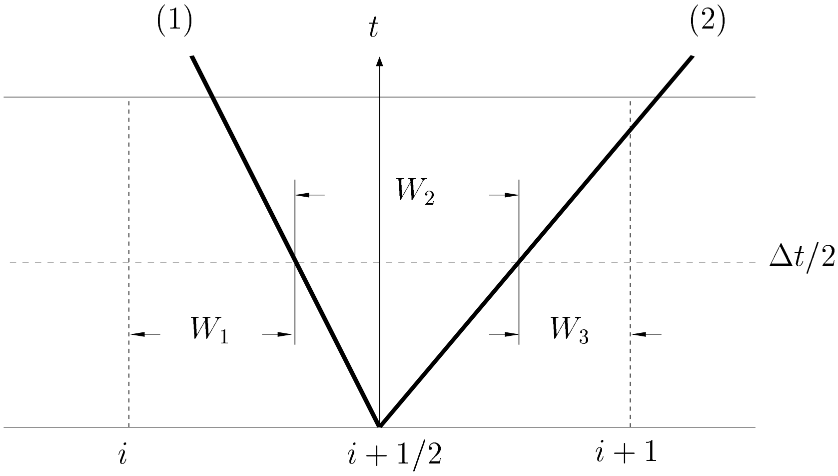

Figure 3.2: The numerical flux as a weighted average of the fluxes of the distinct states. The simplified structure of a hypothetical Riemann problem is shown, with the two waves and the three weights W_k.

3.7.2 WAF

The Weighted Average Flux (WAF) method of Toro (1992) is a second order TVD method that requires the complete solution of the Riemann problem, not just the solution along the computational boundary. It captures shocks very accurately, but can have problems at sonic points.

The WAF scheme takes a different approach to the weak form of the hyperbolic conservation law (3.2). Instead of the flux $\mathbf{f}_{i+1/2}$ being the flux through the boundary of the cell it is viewed as the integral average of the flux over the staggered cell $I_{i+1/2} = [x_i, x_{i+1}]$,

$$\mathbf{f}_{i+1/2} = \frac{1}{\Delta x} \int_{x_i}^{x_{i+1}} \mathbf{f}(\mathbf{q}(t, x)) \mathrm{d}x. \tag{3.55}$$

This is the same as the standard approach to second order.

The integral of the flux is calculated at time $t^{n+1/2}$ only. This is equivalent to approximating the time integral using the midpoint rule, which is the second order Gaussian quadrature (see section 3.3.2). By solving a Riemann problem we can find the data over the staggered cell $I_{i+1/2}$ at this time, and hence the flux.

To simplify the calculation of the flux integral (3.55) we assume that all the waves are represented by single waves through the "origin" $x_{i+1/2}$. Then the integral (3.55) becomes a summation,

$$\mathbf{f}_{i+1/2} = \sum_{k=1}^{N+1} W_k \mathbf{f}^{(k)}_{i+1/2}. \tag{3.56}$$

As in chapter 2, N represents the number of waves. W_k represents the geometric extent of the region k in which the flux takes the constant value $\mathbf{f}^{(k)}_{i+1/2}$. This is shown in figure 3.2.

W_k can be seen as a weight. If the wave speeds are denoted λ_k and the associated Courant number by $\nu_k = \Delta t \lambda_k / \Delta x$ then the weights are given by

$$W_k = \tfrac{1}{2}\left(\nu_k - \nu_{k-1}\right), \qquad \nu_0 = -1, \qquad \nu_{N+1} = 1. \tag{3.57}$$

The approximation of each wave as a single ray is inaccurate for a rarefaction wave. Toro found the most stable results from identifying the ray with the head of the rarefaction wave. This does not cover the case of sonic flow when the rarefaction includes the cell boundary. It is then necessary to replace the associated flux $\mathbf{f}^{(k)}$ with the flux evaluated at the cell boundary.

This method as outlined above is second order accurate. It can be shown that it is analogous, when applied to the linear advection equation, to the Richtmyer–Morton or two-step Lax–Wendroff scheme. It is also linear. By the theorem of Godunov (1959) it will be unstable. In order to prevent these instabilities from occurring *wave amplifiers* are used. These are similar to the limiters used in section 3.4.2, but may be larger than 1. These amplifiers are applied directly to the wave speeds λ_k, or equivalently to the Courant numbers ν_k. It can be shown that there is a direct relation between the amplifiers and the slope limiters of section 3.4.2. Denoting the limited Courant numbers $\Phi_k \nu_k$, we have

$$\Phi = \frac{1 - (1 - |\nu|)\,\phi}{|\nu|}. \tag{3.58}$$

3.8 Boundary conditions

All of the methods outlined above will not work directly at the boundaries of the numerical grid. This is because the methods require information not contained in the grid. If we denote the first computational cell I_1 and the last I_N then in order to calculate the intercell fluxes $\mathbf{f}_{1/2}$ and $\mathbf{f}_{N+1/2}$ we must know the solution in at least the cells I_0 and I_{N+1}.

The method we use is to attach *ghost* cells to the grid. These are extra computational cells outside the physical domain of the grid. They are labeled to correspond to the physical grid, that is I_0, I_{-1} on the left and I_{N+1}, I_{N+2} on the right. It is possible that there is no physical reality to these cells. For example, if the boundary is a solid wall then it makes little sense to talk of the density and velocity of the fluid there. However, the ghost cells are set so that the reconstructed boundary data is physically correct.

There are four simple cases where boundary conditions can be explicitly given.

- Exact solution.
- Periodic boundaries.
- Reflective boundaries.
- Outflow boundaries.

If the exact solution at the boundary is known then it can be imposed. This may not always be the boundary condition that minimizes the error. For example, if a small numerical error propagates to an exact boundary then some of this error may be reflected back into the grid.

Periodic boundaries occur when the cell boundaries $x_{1/2}$ and $x_{N+1/2}$ can be identified. The ghost cell data is then found by identifying cell I_0 with cell I_N and so on. Periodic boundaries greatly simplify the stability analysis of a numerical method. They can also be used where there is symmetry inherent in the problem.

Reflective boundaries occur where some of the functions are odd or anti-symmetric. This occurs for example at solid walls or at the origin in spherical symmetry, where the coordinate velocity in hydrodynamics becomes zero. If the cell boundary is identified with the physical boundary then in the finite volume formulation the boundary conditions

$$\mathbf{O}_{1-i} = -\mathbf{O}_i, \qquad \mathbf{E}_{1-i} = \mathbf{E}_i \tag{3.59}$$

are exact (for the left boundary), where $\mathbf{O}$ are the variables that are odd across the boundary, $\mathbf{E}$ are the variables that are even across the boundary, and i runs from 1 to the number of ghost cells.

Outflow boundary conditions occur when we truncate the physical domain artificially due to limited computational resources. At these boundaries we want any waves from inside the domain to propagate out without reflection. We also want the boundary condition not to introduce waves into the domain. Outflow boundaries are still very much a topic of ongoing research. For an overview of the results in hydrodynamics see Laney (1998), and in relativity see, for example, Rezzolla et al. (1999).

Outflow boundaries typically require some type of extrapolation. For the outflow boundaries used here we have either used zero order extrapolation

$$\mathbf{q}_{N+i} = \mathbf{q}_{N+1-i}, \tag{3.60}$$

or, as detailed in section 4.3, we bypass the problem altogether.

An important point about outflow boundary conditions is the number that should be specified in order to have a well-posed problem. As explained in detail in chapter 19 of Laney (1998), we are only free to specify the data on characteristics coming into the grid. Specifying data corresponding to outgoing characteristics is likely to be inconsistent with the data in the interior of the grid, even if the boundary condition is exact, as the interior is likely to be inaccurate due to numerical error. In practice this leads to small numerical errors propagating into the grid which will only be reduced by the dissipation of the numerical method, if at all. The method detailed in section 4.3 makes the outer boundary move so that the interior of the grid is (at the analytic level, for the hyperbolic equations) not dependent on the boundary.

3.9 Hydrodynamical details

Here we give the details essential for a numerical code solving relativistic hydrodynamics. The conversion from conservative to primitive variables for the relativistic ideal gas is given in full by Martí and Müller (1999). The conversion in the ultrarelativistic case was given (in terms of different variables) by Neilsen and Choptuik (2000) for the one dimensional case. The conversion in more than one dimension appears here for the first time. The spectral decomposition for the relativistic ideal gas was given by Martí et al. (1991) in one dimension and by Donat et al. (1998) in more than one dimension. The spectral decomposition for the ultrarelativistic case was given, in a considerably more complicated form, by Neilsen and Choptuik (2000) for the one dimensional case. The decomposition in more than one dimension appears here for the first time.

3.9.1 Finding the primitive variables

After every numerical step we will have new values of the conserved variables. From these we need to find the primitive variables. This will be necessary in order to construct the fluxes as well as for easy physical understanding.

Ultrarelativistic case in one dimension

We know the conserved variables

$$\tau = (p+\rho)W^2 - p = (p+\rho)W^2v^2 + \rho, \quad S = (p+\rho)W^2 v. \tag{3.61}$$

The equation of state gives the pressure $p = (\Gamma - 1)\rho$. Once the density ρ is known the velocity is given by

$$v = \frac{S}{\tau + p}. \tag{3.62}$$

The density is given by the solution of a quadratic equation. The physical solution is

$$\rho = \tfrac{1}{\Gamma - 1}\left(-2\beta\tau + \left[4\beta^2\tau^2 + (\Gamma - 1)(\tau^2 - S^2)\right]^{1/2}\right), \tag{3.63}$$

where $\beta = \frac{1}{4}(2 - \Gamma)$ is a non-negative constant.

When the fluid is highly relativistic, i.e., $|v| \to 1$, a numerical implementation of equation (3.62) may give velocities greater than the speed of light. To avoid this, define $\chi \equiv W^2 v$. In terms of the conserved variables

$$\chi = \frac{S}{\Gamma\rho}. \tag{3.64}$$

Then the velocity is given by

$$v = \frac{1}{2\chi}\left(\sqrt{1 + 4\chi^2} - 1\right). \tag{3.65}$$

Ultrarelativistic case in more than one dimension

We know the conserved variables

$$\tau = (\rho + p)W^2 - p, \quad S^i = (\rho + p)W^2 v^i. \tag{3.66}$$

The equation of state gives the pressure $p = (\Gamma - 1)\rho$. Once the density ρ is known the velocities are given by

$$v^i = \frac{S^i}{\tau + p}. \tag{3.67}$$

The density is given by the solution of a quadratic equation. The physical solution is

$$\rho = \tfrac{1}{\Gamma - 1}\left(-2\beta\tau + \left[4\beta^2\tau^2 + (\Gamma - 1)(\tau^2 - S^i S_i)\right]^{1/2}\right). \tag{3.68}$$

When the fluid is highly relativistic, i.e., $|v| \to 1$, a numerical implementation of equation (3.67) may give velocities greater than the speed of light. To avoid this, define $\chi^i \equiv W^2 v^i$. In terms of the conserved variables

$$\chi^i = \frac{S^i}{\Gamma\rho}. \tag{3.69}$$

Then the velocities are given by

$$v^i = \frac{\chi^i}{2\chi^j\chi_j}\left(\sqrt{1 + 4\chi^j\chi_j} - 1\right). \tag{3.70}$$

Ideal gas case in one dimension

It is necessary to solve an implicit equation to find the primitive variables in the ideal gas case. This is usually given as a function f of the pressure, of which the root is the desired value. The root is usually found by a Newton–

Raphson approach as outline in section 2.9.1. The function is given by

$$f(\bar{p}) = p(\bar{\rho}(\bar{p}), \bar{\epsilon}(\bar{p})) - \bar{p}, \tag{3.71}$$

where the estimated density $\bar{\rho}(\bar{p})$ and internal energy $\bar{\epsilon}(\bar{p})$ are given by

$$\bar{\rho}(\bar{p}) = \frac{D}{\bar{W}(\bar{p})} \tag{3.72}$$

and

$$\bar{\epsilon}(\bar{p}) = \frac{\tau - D(1 - \bar{W}(\bar{p})) + \bar{p}((1 - \bar{W}(\bar{p})^2)}{D\bar{W}(\bar{p})}, \tag{3.73}$$

where the estimated Lorentz factor $\bar{W}(\bar{p})$ is calculated from the velocity

$$\bar{v}(\bar{p}) = \frac{S}{\tau + D + \bar{p}}. \tag{3.74}$$

The derivative of the implicit function f can be approximated by

$$f' = \bar{v}(\bar{p})^2 c_s(\bar{p})^2 - 1. \tag{3.75}$$

Ideal gas case in more than one dimension

To recover the primitive variables in more than one dimension we use the same method. However we must calculate the velocity in each coordinate direction in order to calculate the Lorentz factor $\bar{W}(\bar{p})$. So the only change is that equation (3.74) becomes

$$\bar{v}^i(\bar{p}) = \frac{S^i}{\tau + D + \bar{p}}. \tag{3.76}$$

The derivative of the implicit function f can be approximated by

$$f' = \bar{v}^i(\bar{p})\bar{v}_i(\bar{p})c_s(\bar{p})^2 - 1. \tag{3.77}$$

3.9.2 The spectral decomposition

The spectral decomposition, that is the eigenvalues and eigenvectors of the Jacobian $\partial \mathbf{f}/\partial \mathbf{q}$, is required for every approximate Riemann solver and the ENO reconstruction. In the case of more than one dimension we give the decomposition for the three dimensional case in the x direction, from which the cases in lower dimensions and other directions can easily be found.

Ultrarelativistic case in one dimension

The eigenvalues are

$$\lambda_\pm = \frac{v \pm c_s}{1 \pm v c_s}. \tag{3.78}$$

The corresponding right eigenvectors are

$$\mathbf{r}_\pm = \begin{pmatrix} 1 \\ \lambda_\pm \end{pmatrix}, \tag{3.79}$$

and the corresponding left eigenvectors are

$$\mathbf{l}_\pm = \begin{pmatrix} \frac{\lambda_\mp}{\lambda_\mp - \lambda_\pm} \\ \frac{1}{\lambda_\pm - \lambda_\mp} \end{pmatrix}. \tag{3.80}$$

Ultrarelativistic case in more than one dimension

The eigenvalues are

$$
\begin{aligned}
\lambda_{\pm} &= \tfrac{1}{1-v^2c_s^2} & &\Big\{ v^x(1-c_s^2) \pm \\
& & & c_s\sqrt{(1-v^2)[1-(v^x)^2 - c_s^2(v^2-(v^x)^2]} \Big\}, \\
\lambda_0 &= v^x & &\text{(double).}
\end{aligned}
\tag{3.81}
$$

The corresponding right eigenvectors are

$$
\mathbf{r}_{\pm} = \begin{pmatrix}
(1-(v^x)^2)(v^x-\lambda_{\pm})(1-(\Gamma-1)v^2) \\
(1-(v^x)^2)\Big[(\lambda_{\pm}v^x - 1)(\Gamma-1)(1-v^2) + \\
\qquad v^x(2-\Gamma)(v^x-\lambda_{\pm})\Big] \\
v^y\Big[(\Gamma-1)\left(\lambda_{\pm}(1+(v^x)^2)v^2 - \lambda_{\pm}(v^x)^2 - 2v^xv^2\right) + \\
\qquad (2-\Gamma)(v^x)^2(\lambda_{\pm}-v^x) + \Gamma v^x - \lambda_{\pm}\Big] \\
v^z\Big[(\Gamma-1)\left(\lambda_{\pm}(1+(v^x)^2)v^2 - \lambda_{\pm}(v^x)^2 - 2v^xv^2\right) + \\
\qquad (2-\Gamma)(v^x)^2(\lambda_{\pm}-v^x) + \Gamma v^x - \lambda_{\pm}\Big]
\end{pmatrix},
\tag{3.82}
$$

and

$$
\mathbf{r}_{0,1} = \left(\tfrac{2v^y}{1+v^2-2(v^x)^2}, \quad \tfrac{2v^xv^y}{1+v^2-2(v^x)^2}, \quad 1, \quad 0 \right)^T \tag{3.83}
$$

$$
\mathbf{r}_{0,2} = \left(\tfrac{2v^z}{1+v^2-2(v^x)^2}, \quad \tfrac{2v^xv^z}{1+v^2-2(v^x)^2}, \quad 0, \quad 1 \right)^T, \tag{3.84}
$$

and the corresponding left eigenvectors are

$$\mathbf{l}_{\pm} = \pm\frac{1}{\Delta}\begin{pmatrix} (1-\Gamma)\lambda_{\mp}(v^x)^3v^2 + 3(\Gamma-1)v^x((v^y)^2+(v^z)^2)(v^x-\lambda_{\mp}) + \\ \Gamma(1-(v^x)^2) + \Gamma((v^y)^2+(v^z)^2) + \Gamma\lambda_{\mp}v^x(3(v^x)^2-2) + \\ \lambda_{\mp}v^x(3-4(v^x)^2) + (v^x)^4 - ((v^y)^2+(v^z)^2) - 1 \\ \Gamma((v^x)^2(1-(v^x)^2)(\lambda_{\mp}-v^x) + ((v^y)^2+(v^z)^2)(3\lambda_{\mp}-v^x)) + \\ ((v^y)^2+(v^z)^2)(\lambda_{\mp}-3v^x) + ((v^x)^2((v^x)^2-3((v^y)^2+(v^z)^2))) - \\ ((v^y)^2+(v^z)^2)-1)(\lambda_{\mp}-v^x) + 2v^x(1-(v^x)^2)((v^y)^2+(v^z)^2) \\ -2(\Gamma-1)v^y(1-(v^x)^2)(1-2\lambda_{\mp}v^x+(v^x)^2) \\ -2(\Gamma-1)v^z(1-(v^x)^2)(1-2\lambda_{\mp}v^x+(v^x)^2) \end{pmatrix}, \tag{3.85}$$

and

$$\mathbf{l}_{0,1} = \left(\frac{-v^y(1+v^2-2(v^x)^2)}{(1-(v^x)^2)(1-v^2)}, \quad \frac{v^xv^y(1+v^2-2(v^x)^2)}{(1-(v^x)^2)(1-v^2)}, \quad \frac{1-v^2+2(v^y)^2}{1-v^2}, \quad \frac{2v^yv^z}{1-v^2}\right), \tag{3.86}$$

$$\mathbf{l}_{0,2} = \left(\frac{-v^z(1+v^2-2(v^x)^2)}{(1-(v^x)^2)(1-v^2)}, \quad \frac{v^xv^z(1+v^2-2(v^x)^2)}{(1-(v^x)^2)(1-v^2)}, \quad \frac{2v^yv^z}{1-v^2}, \quad \frac{1-v^2+2(v^z)^2}{1-v^2}\right). \tag{3.87}$$

The auxiliary quantities are

$$\Delta = (1-(v^x)^2)^3(\Gamma-1)(\lambda_- - \lambda_+)(1-v^2)(1-(\Gamma-1)v^2). \tag{3.88}$$

Ideal gas case in one dimension

The eigenvalues are

$$\lambda_{\pm} = \frac{v \pm c_s}{1 \pm vc_s}, \qquad \lambda_0 = v. \tag{3.89}$$

The corresponding right eigenvectors are

$$
\begin{aligned}
\mathbf{r}_\pm &= \left(1,\quad hWA_\pm\lambda_\pm,\quad hWA_\pm - 1\right)^T, && (3.90)\\
\mathbf{r}_0 &= \left(W^{-1},\quad v,\quad 1 - W^{-1}\right)^T, && (3.91)
\end{aligned}
$$

and the corresponding left eigenvectors are

$$
\begin{aligned}
\mathbf{l}_\pm &= \mp\frac{h^2}{\Delta}\begin{pmatrix} hA_\mp\lambda_\mp(1-W) + v(hWA_\mp - 1) \\ 1 - hA_\mp \\ hA_\mp\lambda_\mp - v \end{pmatrix}^T, && (3.92)\\
\mathbf{l}_0 &= \frac{W}{h-1}\left(h - W,\quad Wv,\quad -W\right). && (3.93)
\end{aligned}
$$

The auxiliary quantity is

$$
A_\pm = \frac{1-v^2}{1-v\lambda_\pm}, \qquad \Delta = h^3W(h-1)(1-v^2)(\lambda_+A_+ - \lambda_-A_-). \tag{3.94}
$$

Ideal gas case in more than one dimension

The eigenvalues are

$$
\begin{aligned}
\lambda_\pm &= \tfrac{1}{1-v^2c_s^2} && \Big\{v^x(1-c_s^2)\pm \\
& && c_s\sqrt{(1-v^2)[1-(v^x)^2 - c_s^2(v^2-(v^x)^2]}\Big\}, && (3.95)\\
\lambda_0 &= v^x && \text{(triple).}
\end{aligned}
$$

The corresponding right eigenvectors are

$$
\mathbf{r}_\pm = \left(1,\quad hWA_\pm\lambda_\pm,\quad hWv^y,\quad hWv^z,\quad hWA_\pm - 1\right)^T, \tag{3.96}
$$

and

$$\mathbf{r}_{0,1} = \begin{pmatrix} W^{-1} \\ v^x \\ v^y \\ v^z \\ 1 - W^{-1} \end{pmatrix}, \quad \mathbf{r}_{0,2} = \begin{pmatrix} Wv^y \\ 2hW^2 v^x v^y \\ h(1 + 2W^2 (v^y)^2) \\ 2hW^2 v^y v^z \\ Wv^y (2hW - 1) \end{pmatrix}, \quad \mathbf{r}_{0,3} = \begin{pmatrix} Wv^z \\ 2hW^2 v^x v^z \\ 2hW^2 v^y v^z \\ h(1 + 2W^2 (v^z)^2) \\ Wv^z (2hW - 1) \end{pmatrix}, \tag{3.97}$$

and the corresponding left eigenvectors are

$$\mathbf{l}_{\pm} = \mp \frac{h^2}{\Delta} \begin{pmatrix} hWA_{\mp}(v^x - \lambda_{\mp}) - v^x - W^2(v^2 - (v^x)^2)\cdot \\ (2h-1)(v^x - A_{\mp}\lambda_{\mp}) + hA_{\mp}\lambda_{\mp} \\ 1 + W^2(v^2 - (v^x)^2)(2h-1)(1 - A_{\mp}) - hA_{\mp} \\ W^2 v^y (2h-1) A_{\mp}(v^x - \lambda_{\mp}) \\ W^2 v^x (2h-1) A_{\mp}(v^x - \lambda_{\mp}) \\ -v^x - W^2(v^2 - (v^x)^2)(2h-1)(v^x - A_{\mp}\lambda_{\mp}) + hA_{\mp}\lambda_{\mp} \end{pmatrix}^T \tag{3.98}$$

and

$$\mathbf{l}_{0,1} = \frac{W}{h-1}\left(h - W, \;\; Wv^x, \;\; Wv^y, \;\; Wv^z, \;\; -W\right), \tag{3.99}$$

$$\mathbf{l}_{0,2} = \frac{1}{h(1-(v^x)^2)}\left(-v^y, \;\; v^x v^y, \;\; 1 - (v^x)^2, \;\; 0, \;\; -v^y\right), \tag{3.100}$$

$$\mathbf{l}_{0,3} = \frac{1}{h(1-(v^x)^2)}\left(-v^z, \;\; v^x v^z, \;\; 0, \;\; 1 - (v^x)^2, \;\; -v^z\right). \tag{3.101}$$

The auxiliary quantities are

$$A_{\pm} = \frac{1-(v^x)^2}{1-v^x\lambda_{\pm}}, \qquad \Delta = h^3 W(h-1)(1-(v^x)^2)(\lambda_+ A_+ - \lambda_- A_-). \quad (3.102)$$

3.10 Code validation tests

Here we outline a number of simple tests to indicate the performance of the different numerical methods. We are trying to find the best balance between accuracy and efficiency. For the problem in chapter 5 we need accurate resolution of fine features including shocks over reasonably long evolutions whilst evolving complicated source terms. The best method would capture shocks and other discontinuities in very few points with good convergence properties and without introducing oscillations, whilst still being computationally efficient.

The tests here are from a variety of sources. The linear advection tests are standard. The test for Burger's equation can be found in Laney (1998). Various relativistic ideal gas tests have been performed, and a comprehensive list can be found in Martí and Müller (1999). However the only implementations of ENO schemes that I am aware of are those by Dolezal and Wong (1995), Yang et al. (1997), Font et al. (2000). The only previous tests of the ultrarelativistic gas are those in Neilsen and Choptuik (2000), Neilsen (1999) in one dimension. I know of no comparative study of any system of equations of the WAF and SLIC methods against the ENO type methods.

The schemes to be tested are the WAF and SLIC schemes and all the other schemes based on the method of lines. The method of lines schemes will be named by their reconstruction method, where “Slope” is a shorthand for slope limited reconstruction, and the Riemann solver used. The Runge–Kutta method will be used for all time integration, with the order of accuracy being the same as for the spatial reconstruction. For example, the method of lines scheme using the slope limiting TVD reconstruction of section 3.4.2, the Marquina approximate Riemann solver of section 2.8.3 and the second order Runge–Kutta ODE method will be labeled “Slope-Marquina”. We note that SLIC is also a slope limited scheme, but that we use the term “Slope” to apply only to those methods based on the method of lines. For all TVD schemes the results presented here will use a Van Leer type limiter, but the results are consistent with the other limiters given in section 3.4.2. In particular, where a TVD method is listed as failing to complete an evolution, this is true for all the limiters listed.

3.10.1 One dimensional tests

Linear advection

The solution of the Riemann problem for the linear advection equation

$$\partial_t q + a\partial_x q = 0 \tag{3.103}$$

was outlined in section 2.2.1. In fact the general solution is given from the characteristics. The initial profile $q(0, x)$ is propagated with speed a without otherwise changing. We set $q(0, x) = \sin 2\pi x$ with periodic boundaries and evolve up to $t = 1$. All methods used a CFL number of 0.7. We measure

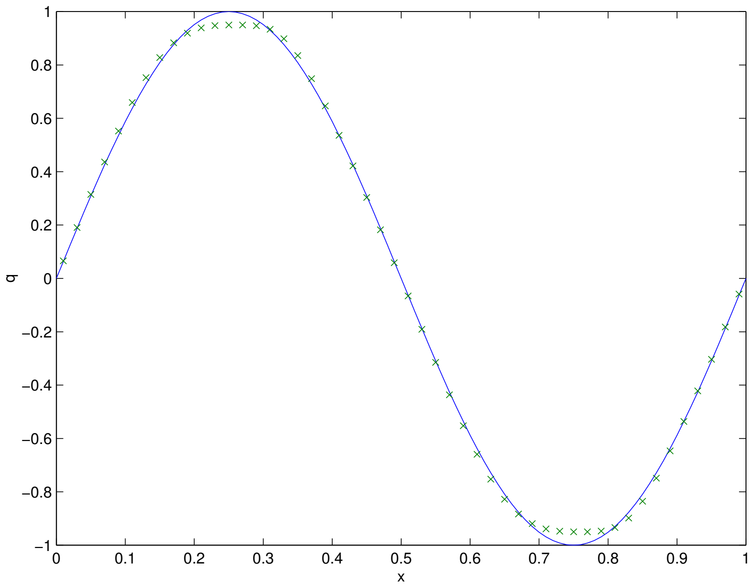

Figure 3.3: A sine wave is advected five times using periodic boundaries. The WAF method with 50 computational cells and a CFL number of 0.7 were used. The exact solution is given by the line. The extrema are clearly clipped.

the accuracy of each method and look at the convergence behaviour, listed in table 3.1. We note that any of the approximate Riemann solvers outlined in chapter 2 will be exact when applied to the linear advection equation.

Starting with the second order methods, we see that those based on the method of lines are less accurate than the single step WAF and SLIC schemes. Even the best second order scheme (WAF) "clips" the extrema to some extent, as shown in figure 3.3 (the figures are at time $t = 5$ to emphasize the behaviour). When we look at the higher order methods, the disadvantages of the method of lines approach are not apparent. The third order ENO method is a notable improvement over all the second order methods, with no clipping of the extrema, as shown in figure 3.4. The higher order methods

Method	Cells	L_1 error	L_1 order	L_2 error	L_2 order	L_∞ error	L_∞ order
SLIC	25	4.55e-2	–	1.10e-2	–	4.32e-3	–
	50	1.30e-2	1.81	2.43e-3	2.18	8.48e-4	2.35
	100	3.55e-3	1.87	5.35e-4	2.18	1.69e-4	2.33
	200	9.33e-4	1.93	1.18e-4	2.18	3.40e-5	2.31
WAF	25	1.72e-2	–	4.47e-3	–	2.06e-3	–
	50	4.31e-3	2.00	9.51e-4	2.23	4.00e-4	2.37
	100	1.04e-3	2.05	2.00e-4	2.25	7.23e-5	2.47
	200	2.40e-4	2.12	4.10e-5	2.29	1.37e-5	2.41
Slope	25	1.01e-1	–	2.33e-2	–	7.84e-3	–
	50	4.31e-2	1.23	6.98e-3	1.74	1.74e-2	2.17
	100	1.36e-2	1.66	1.71e-3	2.03	3.89e-4	2.16
	200	3.92e-3	1.79	3.94e-4	2.12	8.32e-5	2.25
ENO2	25	1.01e-1	–	2.34e-2	–	7.92e-3	–
	50	4.42e-2	1.19	7.16e-3	1.71	1.78e-3	2.15
	100	1.36e-2	1.70	1.72e-3	2.06	3.91e-4	2.19
	200	3.92e-3	1.79	3.94e-4	2.13	8.33e-5	2.24
ENO3	25	6.12e-3	–	1.36e-3	–	4.04e-4	–
	50	7.72e-4	2.99	1.21e-4	3.49	2.48e-5	4.03
	100	9.67e-5	3.00	1.07e-5	3.50	1.54e-6	4.00
	200	1.21e-5	3.00	9.51e-7	3.49	9.46e-8	4.00
ENO4	25	1.31e-3	–	3.28e-4	–	1.39e-4	–
	50	9.21e-5	3.83	1.92e-5	4.09	7.52e-6	4.21
	100	6.36e-6	3.85	1.13e-6	4.09	4.37e-7	4.11
	200	4.39e-7	3.86	6.62e-8	4.09	2.42e-8	4.18
ENO5	25	9.46e-5	–	2.09e-5	–	5.96e-6	–
	50	3.66e-6	4.70	5.76e-7	5.18	1.16e-7	5.68
	100	1.67e-7	4.45	1.86e-8	4.96	2.66e-9	5.45
	200	9.03e-9	4.21	7.28e-10	4.67	9.30e-11	4.84

Table 3.1: Convergence tests for the linear advection equations.

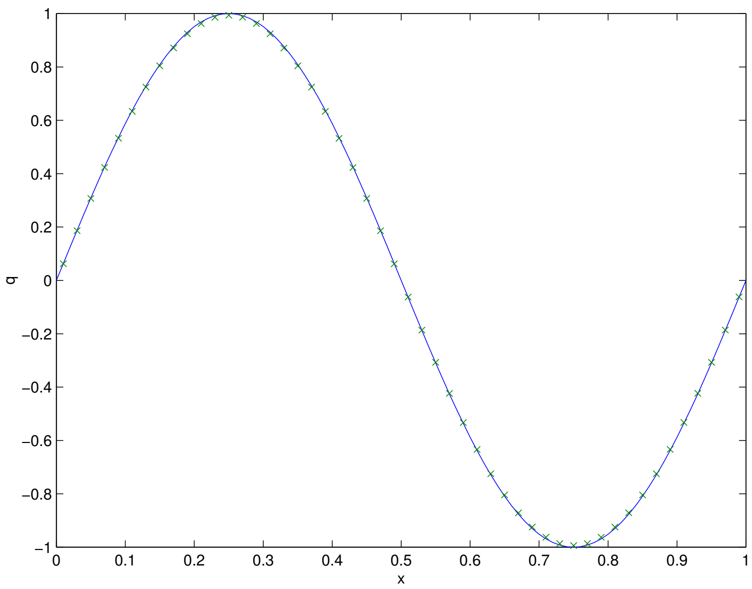

Figure 3.4: A sine wave is advected five times using periodic boundaries. The third order ENO method with 50 computational cells and a CFL number of 0.7 were used. The exact solution is given by the line. The solution is well captured with no phase error and very little clipping.

show even greater improvement. We note that the errors from ENO5 with 25 cells are already less than WAF, the best second order method, with 200 cells. We believe the convergence order for ENO5 tails off because the error is approaching machine zero.

Burger's equation

The solution of the Riemann problem for Burger's equation

$$\partial_t q + \partial_x (\tfrac{1}{2} q^2) = 0. \tag{3.104}$$

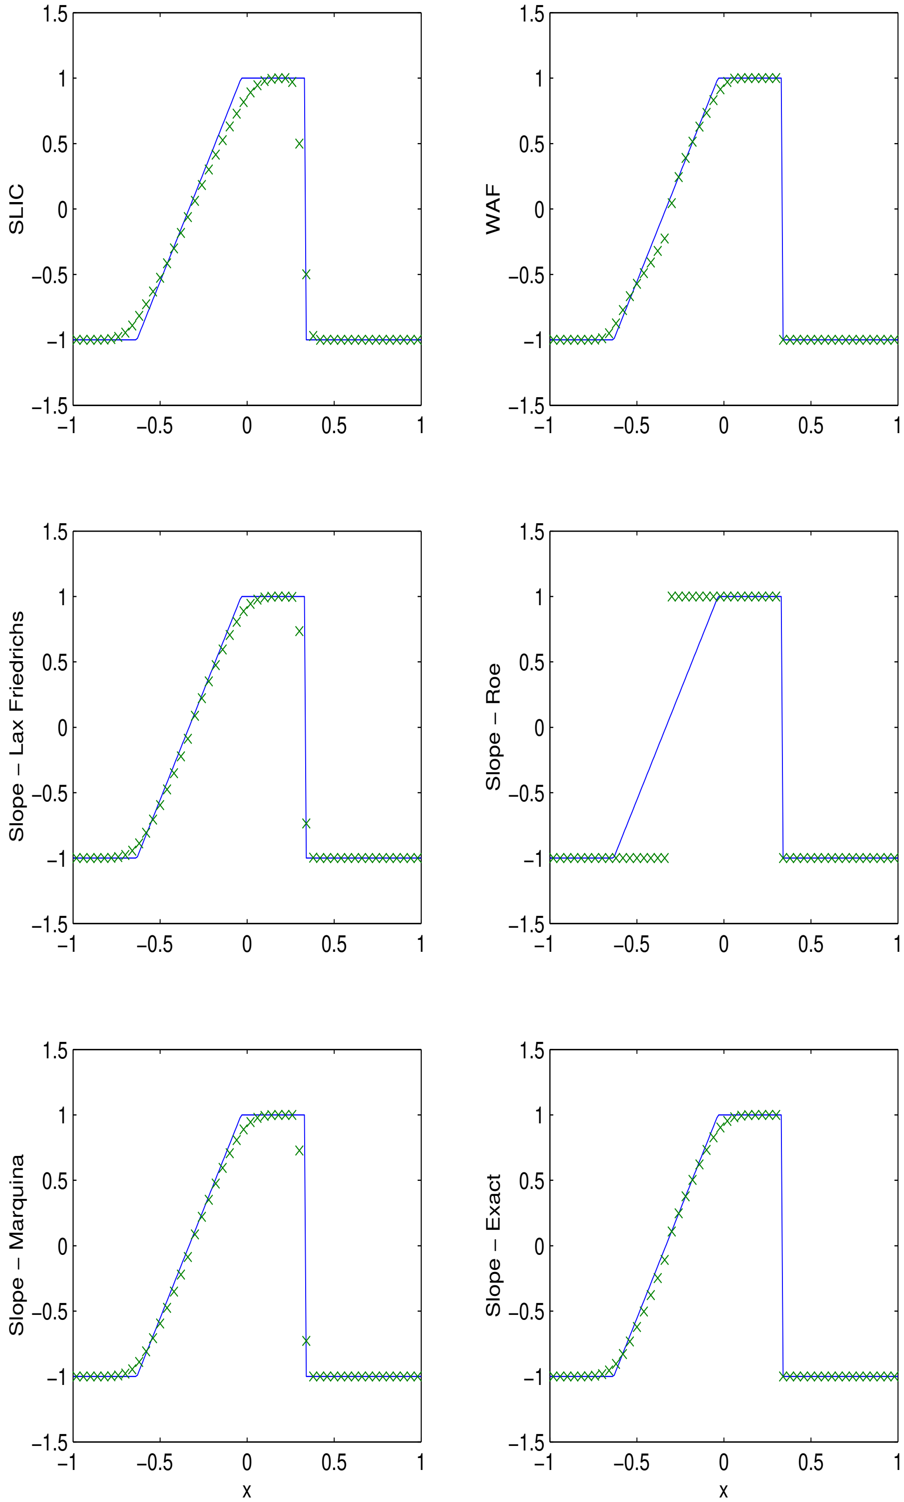

Figure 3.5: A shock tube type problem for Burger's equation. The exact solution is a sonic rarefaction on the left and a stationary shock on the right. All the methods are second order accurate. The failures of the WAF and Roe solver methods are clear. The exact solver captures the shock "perfectly".

was outlined in section 2.2.2. Here we look at a test from Laney (1998). The domain is $x \in [-1, 1]$, with the initial data

$$q = \begin{cases} 1 & |x| < \frac{1}{3} \\ -1 & \text{otherwise} \end{cases} . \tag{3.105}$$

The solution, which contains a rarefaction with a sonic point (i.e., a point where the characteristic velocity is the same as the fluid velocity) and a stationary shock, is shown at time 0.3. All the simulations use 50 points, outflow boundary conditions and a CFL of 0.7.

The solutions for the second order methods (with the exception of the ENO2 method, which gives results almost identical to the slope limited method) are shown in figure 3.5. The SLIC method does not capture the shock or the head and tail of the rarefaction particularly well. However the solution is perfectly physical. The WAF method captures the shock perfectly, and is good at the head and tail of the rarefaction. However, it has problems at the sonic point.

The slope limited methods show the difference between the Riemann solvers. The Lax–Friedrichs solver produces physical results, but the shock and rarefaction are not so well captured. The Roe solver captures the shock perfectly but produces an unphysical rarefaction shock in place of the rarefaction. It is possible to alter the Roe solver to avoid this problem, as shown by the Marquina solver which captures the rarefaction well. However, this adds a few points into the shock as the eigenvalues change sign here as well. The exact solver captures the shock perfectly and the rarefaction well, with

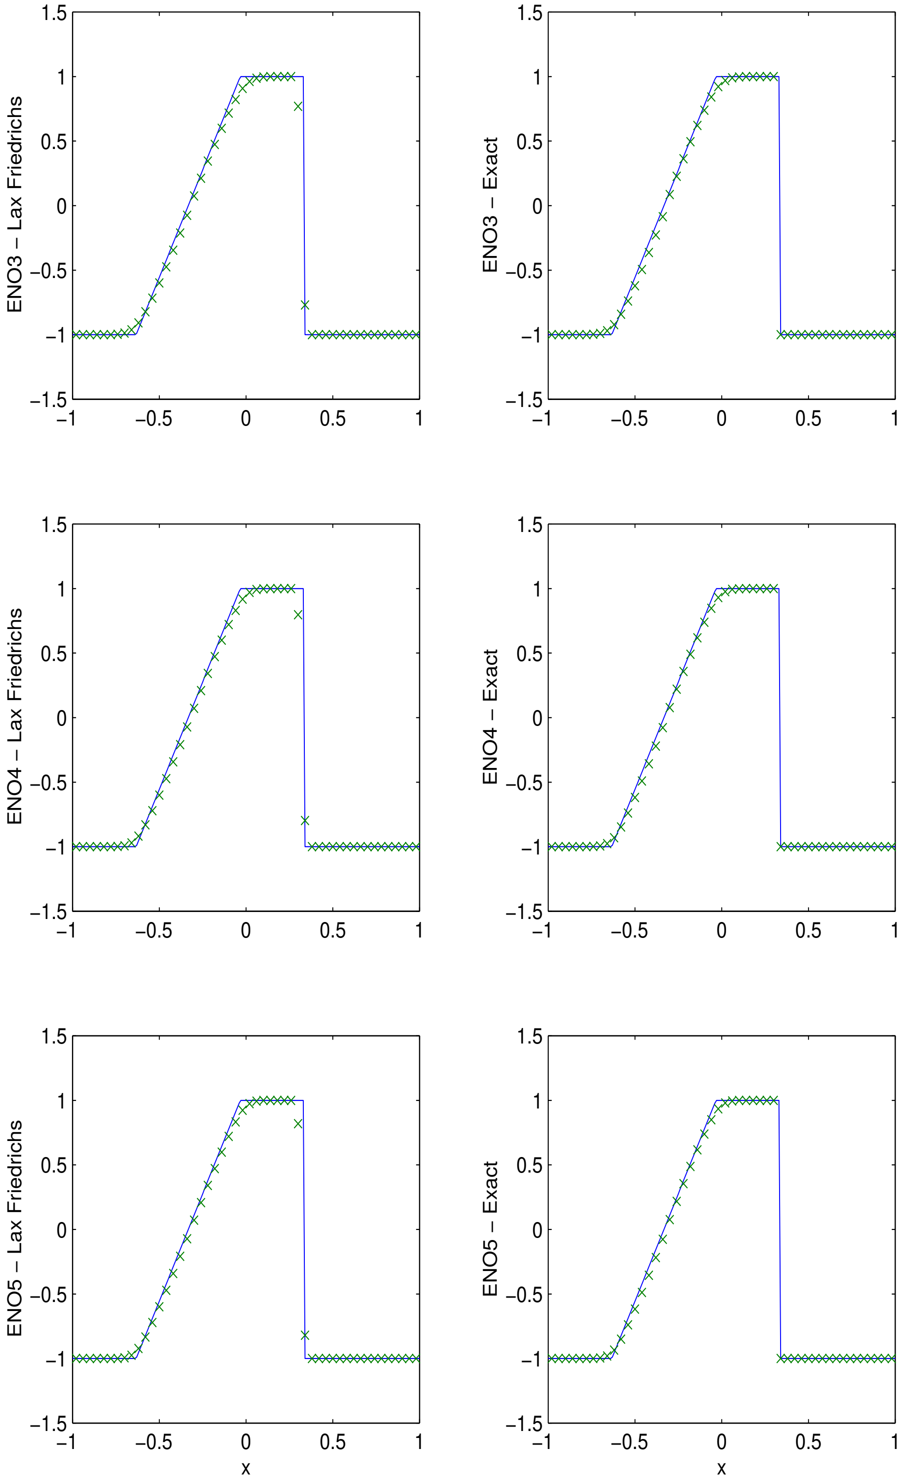

Figure 3.6: A shock tube type problem for Burger's equation. The exact solution is a sonic rarefaction on the left and a stationary shock on the right. All the methods are higher than second order accurate. All resolve the shock and rarefaction well. There is a very slight jump at the sonic point with the exact solver. However, the exact solver captures the shock better than the approximate solver.

Method	Cells	L_1 error	L_1 order	L_2 error	L_2 order	L_∞ error	L_∞ order
SLIC	50	5.54e-2	–	1.66e-2	–	5.00e-1	–
	100	2.41e-2	1.20	5.96e-3	1.48	3.57e-1	0.53
	200	1.34e-2	0.84	3.57e-3	0.74	4.71e-1	-0.44
WAF	50	2.26e-2	–	6.97e-3	–	2.04e-1	–
	100	1.30e-2	0.80	2.50e-3	1.48	8.76e-2	1.22
	200	5.81e-3	1.16	9.25e-4	1.43	4.51e-2	0.96
Slope-LF	50	3.11e-2	–	9.05e-3	–	2.65e-1	–
	100	1.59e-2	0.96	4.37e-3	1.05	2.75e-1	-0.06
	200	8.04e-3	0.99	2.14e-3	1.03	2.82e-1	-0.04
SLOPE-Roe	50	1.50e-1	–	4.48e-2	–	9.78e-1	–
	100	1.50e-1	0.00	3.16e-2	0.50	9.89e-1	-0.02
	200	1.50e-1	0.00	2.24e-2	0.50	9.95e-1	-0.01
SLOPE-Marquina	50	3.10e-2	–	9.16e-3	–	2.72e-1	–
	100	1.58e-2	0.97	4.42e-3	1.05	2.80e-1	-0.04
	200	7.95e-3	0.99	2.13e-3	1.05	2.82e-1	-0.01
SLOPE-Exact	50	1.84e-2	–	5.01e-3	–	9.60e-2	–
	100	9.53e-3	0.95	1.94e-3	1.37	7.44e-2	0.37
	200	4.81e-3	0.99	7.35e-4	1.40	4.97e-2	0.58
ENO3-LF	50	2.68e-2	–	7.84e-3	–	2.31e-1	–
	100	1.36e-2	0.99	3.67e-3	1.09	2.31e-1	0.00
	200	6.77e-3	1.00	1.75e-3	1.07	2.31e-1	0.00
ENO3-Exact	50	1.65e-2	–	4.23e-3	–	7.81e-2	–
	100	8.33e-3	0.98	1.61e-3	1.39	6.24e-2	0.32
	200	4.18e-3	1.00	5.99e-4	1.43	4.11e-2	0.60
ENO4-LF	50	2.45e-2	–	7.02e-3	–	2.03e-1	–
	100	1.23e-2	0.99	3.26e-3	1.11	2.03e-1	0.00
	200	6.16e-3	1.00	1.55e-3	1.08	2.03e-1	0.00
ENO4-Exact	50	1.56e-2	–	3.98e-3	–	7.05e-2	–
	100	7.90e-3	0.99	1.50e-3	1.41	5.66e-2	0.31
	200	3.93e-3	1.01	5.49e-4	1.44	3.67e-2	0.63
ENO5-LF	50	2.32e-2	–	6.47e-3	–	1.81e-1	–
	100	1.16e-2	0.99	2.97e-3	1.12	1.82e-1	-0.01
	200	5.83e-3	1.00	1.41e-3	1.07	1.84e-1	-0.02
ENO5-Exact	50	1.52e-2	–	3.85e-3	–	6.61e-2	–
	100	7.68e-3	0.99	1.44e-3	1.42	5.38e-2	0.30
	200	3.83e-3	1.01	5.26e-4	1.45	3.44e-2	0.64

Table 3.2: Convergence tests for Burgers equation. Convergence is not second order for any method because of the appearance of the shock and the nonlinearity of the problem.

only a small jump at the sonic point.

The solutions for the higher order methods are shown in figure 3.6. We only consider the Lax–Friedrichs and exact solvers, as the results from the Marquina solver will fall somewhere in between and those from the Roe solver will be unphysical. The results from all the methods are similar. The rarefaction is well captured, with only a very small jump at the sonic point with the exact solver. The shock is perfectly captured by the exact solver and well captured by the Lax–Friedrichs solver, with small improvements just visible with the increasing order of the reconstruction.

The convergence of the error when compared to the exact solution is given in table 3.2. For all methods the convergence is very poor, rarely being better than first order accurate. This is to be expected when a shock appears in the solution. The results here are similar to those found by Woodward and Colella (1984) where convergence of a numerical code for a blast wave in gas dynamics was considered. There a variety of codes not mentioned here were tested. The first order accurate Godunov method had a convergence rate between 0.1 and 0.4, and even the extremely accurate PPM method which is formally second order accurate had a convergence rate at best 1.1. With these results we should not expect convergence testing to give useful information in the neighbourhood of a shock. For the remaining tests the convergence tables can be found in section A.4.

Ultrarelativistic hydrodynamics tests

The equations for ultrarelativistic hydrodynamics are given, along with the solution to the Riemann problem, in section 2.4. Here we look at a simple

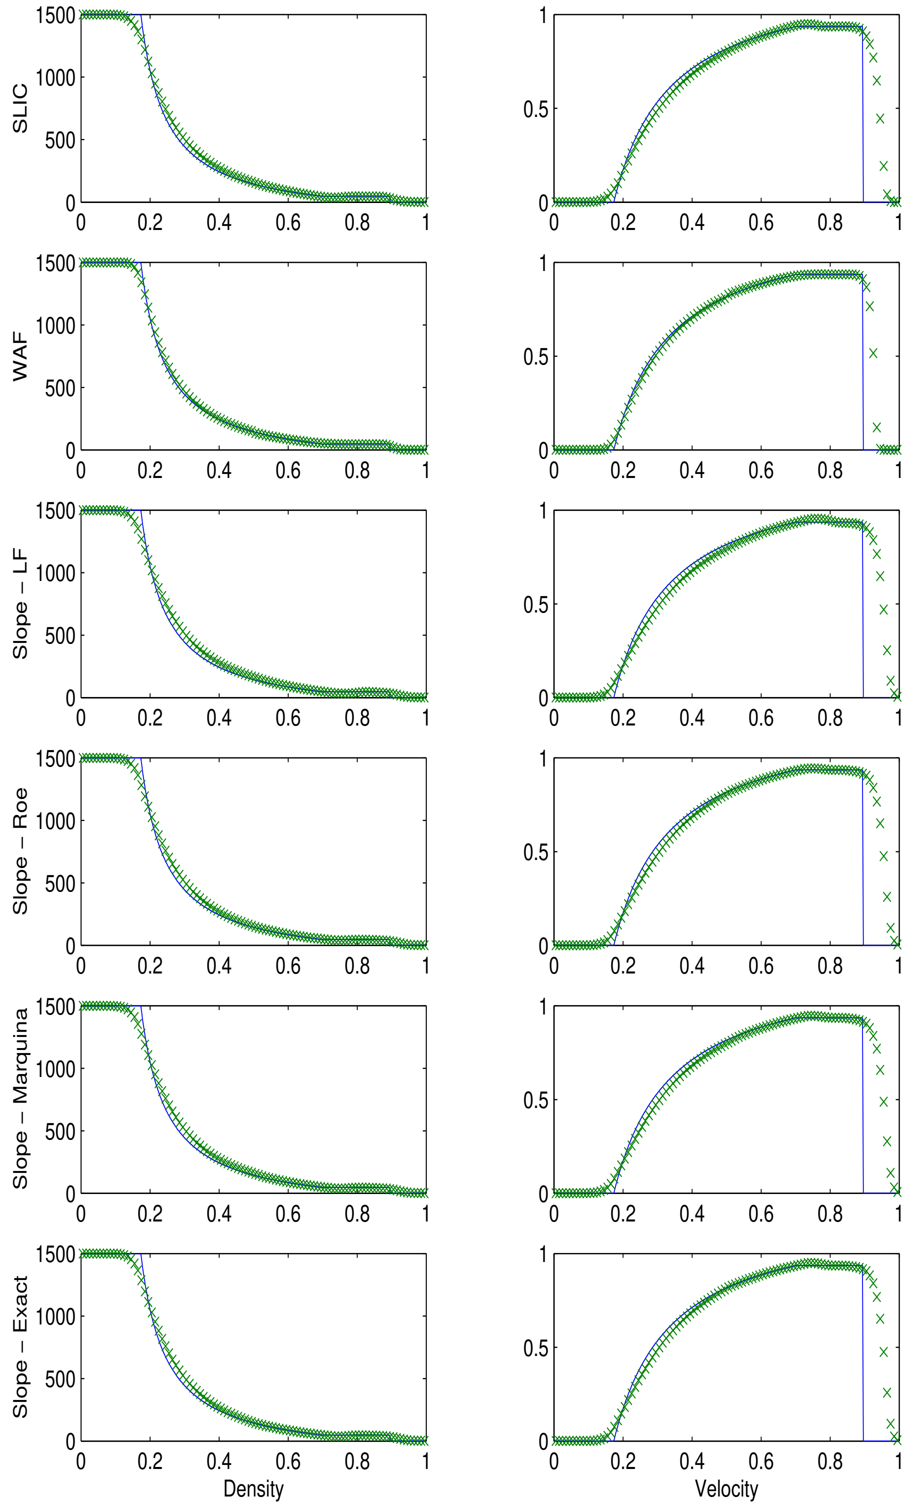

Figure 3.7: A shock tube problem for ultrarelativistic hydrodynamics. The exact solution is a sonic rarefaction on the left and a shock on the right. All the methods are second order accurate. The advantages of the SLIC and WAF codes are clear. There appear to be no problems at the sonic points.

shock tube problem. This domain is $x \in [0, 1]$, with the discontinuity at $x = 0.5$ and initial data

Γ	ρ_L	v_L	ρ_R	v_R
5/3	1500	0	1.5	0

This has a left going rarefaction with sonic point (the sonic point is at $v \approx 0.81$) and a right going shock. Each test uses 100 cells and a CFL number of 0.7. This test was used by Neilsen and Choptuik (2000), Neilsen (1999).

The solutions for the second order methods (with the exception of the ENO2 method, which gives results almost identical to the slope limited method) are shown in figure 3.7. The disadvantages of the codes based on the method of lines is clear. The WAF and SLIC methods both capture the shock far more accurately than the slope limited methods. The WAF method is also better on the rarefaction. It is also notable how little difference the Riemann solver makes in this case. None of these methods have a problem at the sonic point.

The solutions for the higher than second order methods are shown in figure 3.8. Again, only the most and least diffusive Riemann solvers are shown. The difference between the Riemann solvers is again very small. The capturing of the shock and the rarefaction is considerably improved over most of the second order methods, and clearly improves with the order of accuracy. However, we should note that the shock capturing of the WAF scheme is comparable to the fifth order ENO scheme.

The large differences in the shock capturing behaviour of the slope limited TVD methods which use the method of lines and, for example, SLIC, another slope limited TVD method can be explained in terms of the stencil used. The

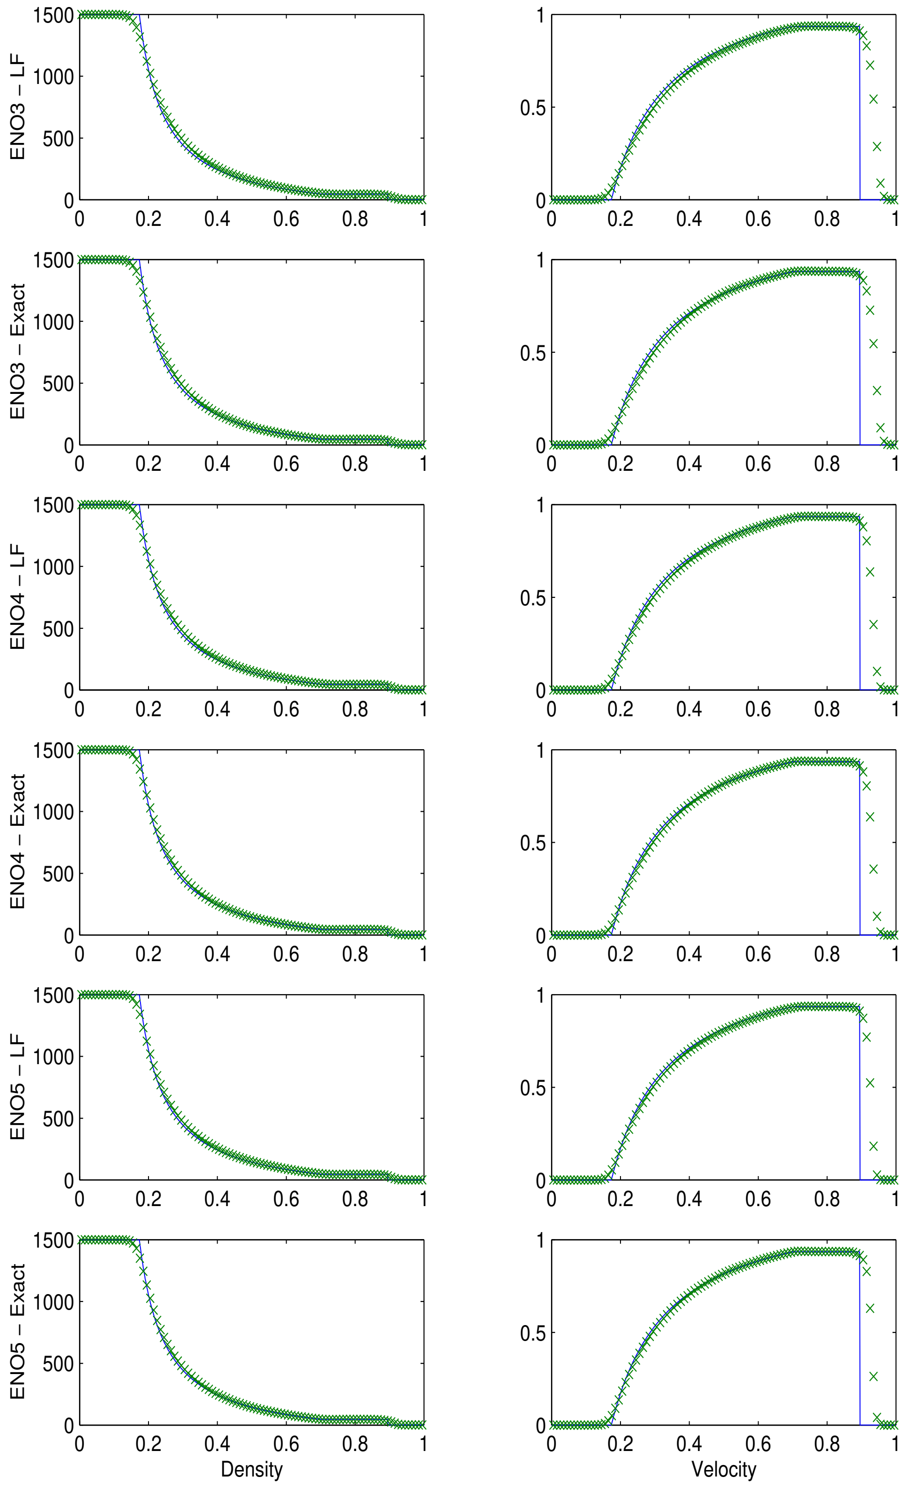

Figure 3.8: A shock tube problem for ultrarelativistic hydrodynamics. The exact solution is a sonic rarefaction on the left and a shock on the right. All the methods are higher than second order accurate. The advantages of the exact solver are small. However, the increasing order of accuracy does improve the shock capturing.

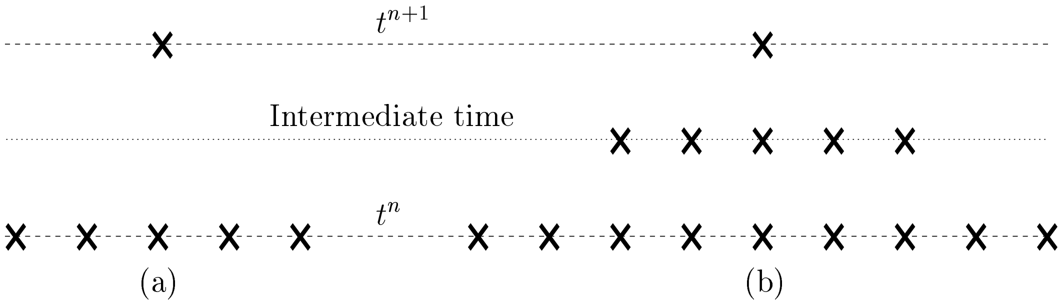

Figure 3.9: The stencil of the SLIC scheme (a) and the slope limited TVD method (b) based on the method of lines. The large increase in the stencil size probably leads to the poorer performance in the shock capturing of the method of lines based schemes.

stencil of a numerical method is the distribution and number of points used at time level n to construct a given point on time level $n+1$. The stencils of SLIC and the method of lines based TVD schemes are given in figure 3.9. The large increase in stencil width (from 5 to 9 points) for the method of lines approach is probably the cause of the extra smearing of the shock.

For the other methods given here the stencil is adaptive and depends on details of the flow. The WAF method will have a stencil size at most 5, but near a shock the upwind nature of the method will mean that the actual stencil used will probably be 3. Similarly, the ENO5-RK4 method using the method of lines will have a stencil size at most 41, but near a shock the actual stencil used may be as small as 6 points. The excellent shock capturing of WAF is due to this sort of upwinding and is not likely to be improved upon by higher orders of accuracy. However, the behaviour of the method in smooth regions will be improved.

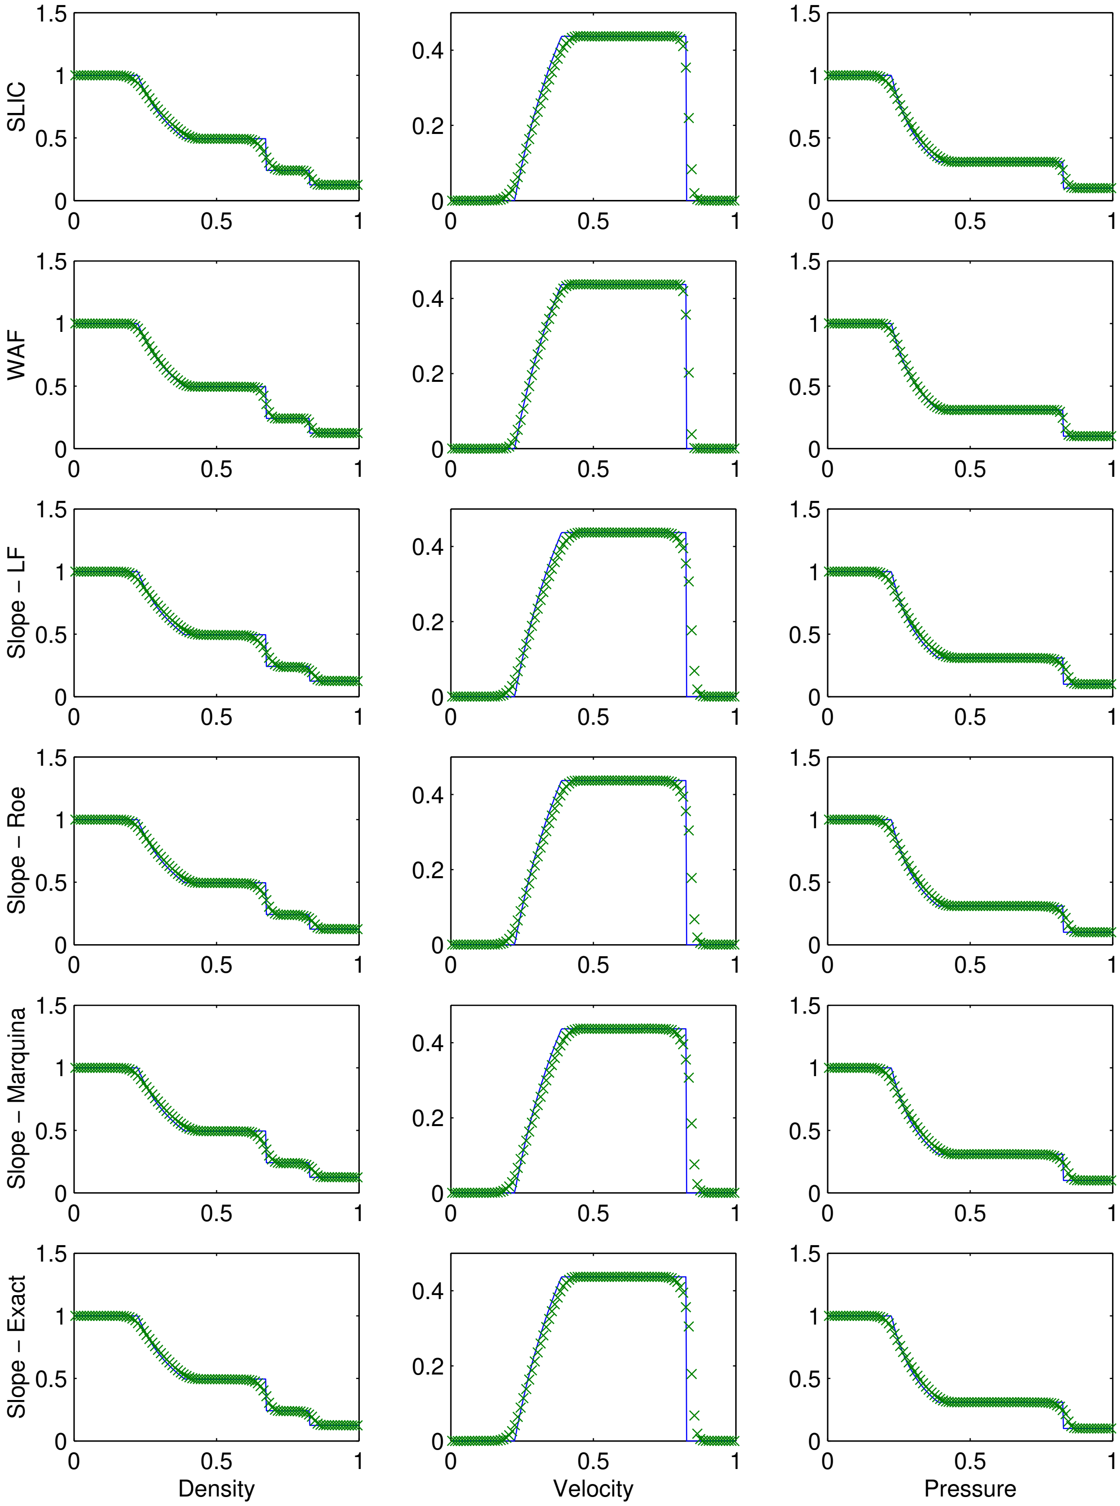

Figure 3.10: Sod's problem for the relativistic ideal gas. The exact solution is a rarefaction on the left and a shock on the right. All the methods are second order accurate. The advantages of the exact solver are small. The advantages of the WAF method are clear.

Relativistic ideal gas tests

The equations for the relativistic ideal gas are given, along with the solution to the Riemann problem, in section 2.5. Here we look two simple shock tube problems. The domain is $x \in [0, 1]$, with the discontinuity at $x = 0.5$ and initial data

γ	p_L	ρ_L	v_L	p_R	ρ_R	v_R
5/3	1	1	0	0.1	0.125	0
5/3	1000	1	0	0.01	1	0

Both have a left going rarefaction and a right going shock separated by a contact surface. The first test, Sod's problem, is fairly mild. The second, the blast wave from Martí and Müller (1994) is very hard to resolve as the contact and the shock are so close together. Both these tests were listed in section 2.9.3. For the Sod problem we use 100 cells and for the blast wave 200. Both tests use a CFL number of 0.7.

For the Sod problem the results from the second order methods are shown in figure 3.10. All the methods capture the shock reasonably well. The advantages of the WAF method are clear as it captures the shock and contact far better than any other second order method. The slope limited algorithms based on the method of lines show the advantages of the more accurate Riemann solvers, especially in the capturing of the contact surfaces.

The results from the higher order methods are shown in figure 3.11. The shock and contact are both extremely well captured in all cases. For these methods it is difficult to see the differences between the Riemann solvers, although there is a very slight improvement with the exact solver at the

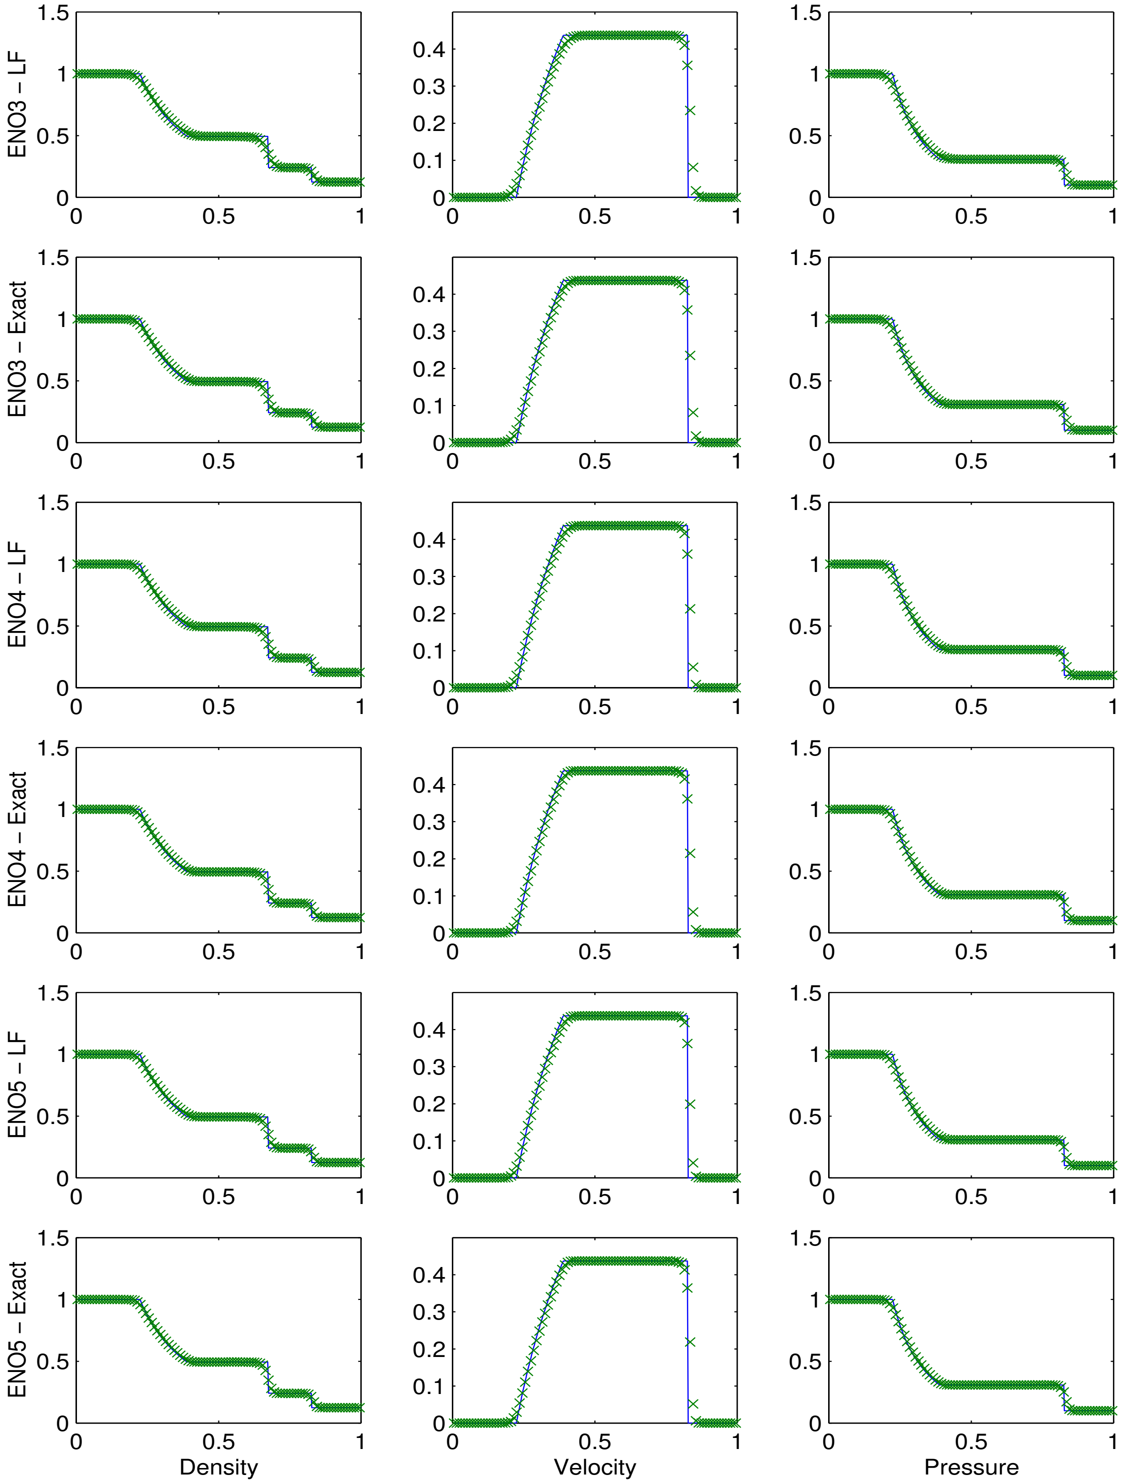

Figure 3.11: Sod's problem for the relativistic ideal gas. The exact solution is a rarefaction on the left and a shock on the right. All the methods are higher than second order accurate. The advantages of the exact solver are small. The advantages of increasing the order of accuracy are clear.

edges of the rarefaction. The improvements with the increasing order of the reconstruction are clearer to see.

For the blast wave the results for the second order methods are shown in figure 3.12. The WAF method failed to complete this test. All the other methods produce very similar results although, as in the ultrarelativistic case, the SLIC method captures the shock slightly better. Again, the Riemann solver appears to have very little effect.

The results for the higher than second order methods are shown in figure 3.13. The fifth order ENO method with exact solver failed to complete this test. However, it is clear that increasing the order of accuracy improves the capturing of the shock and the rarefaction, whilst the Riemann solver has very little effect.

3.10.2 Relativistic blast wave interaction

To test the efficiency of the various Riemann solvers for a real problem, we look at a relativistic version of the Colella–Woodward blast wave problem studied in Woodward and Colella (1984). This is a shock tube problem in one dimension with three initially piecewise constant states separated by two discontinuities. To my knowledge the full relativistic version of this test has only been performed by Yang et al. (1997). They used a relativistic kinetic beam scheme with a variety of different reconstruction methods (including high order ENO methods). Our results are broadly in agreement. However we note that our converged results, shown in figure 3.14, do not have the small density spikes at the shocks. To my knowledge this is the first study of

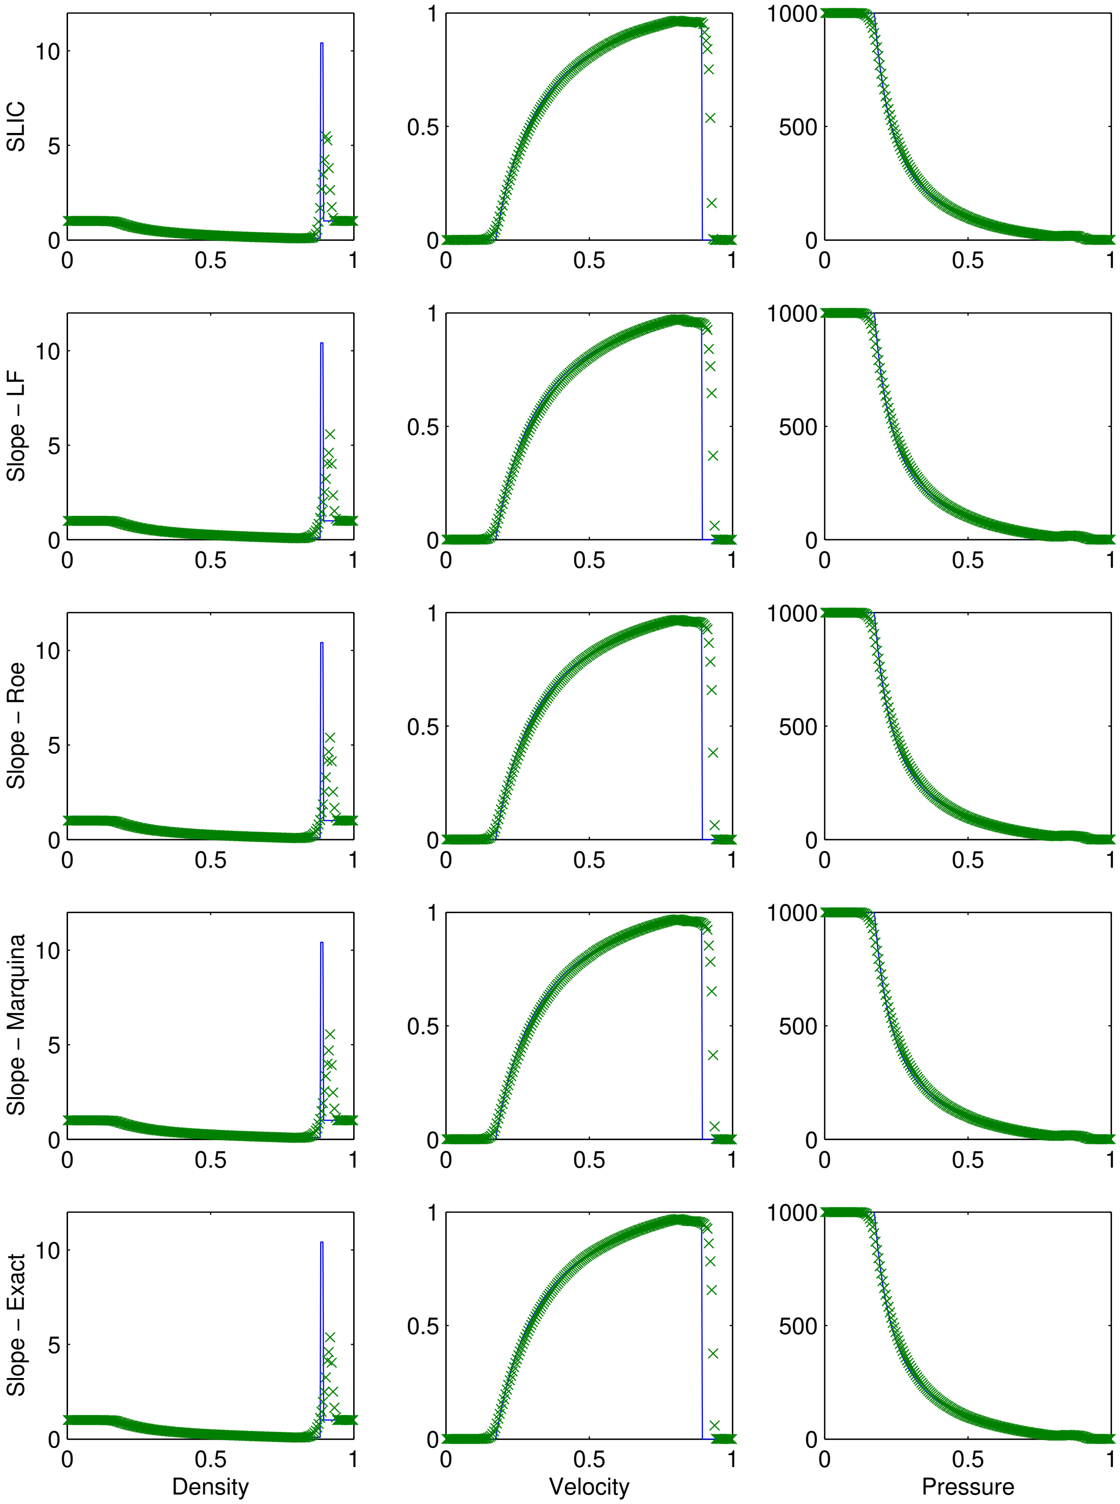

Figure 3.12: A shock tube problem for the relativistic ideal gas. The exact solution is a sonic rarefaction on the left and a shock on the right. All the methods are second order accurate. The advantages of the exact solver are small. There do not appear to be any sonic point problems.

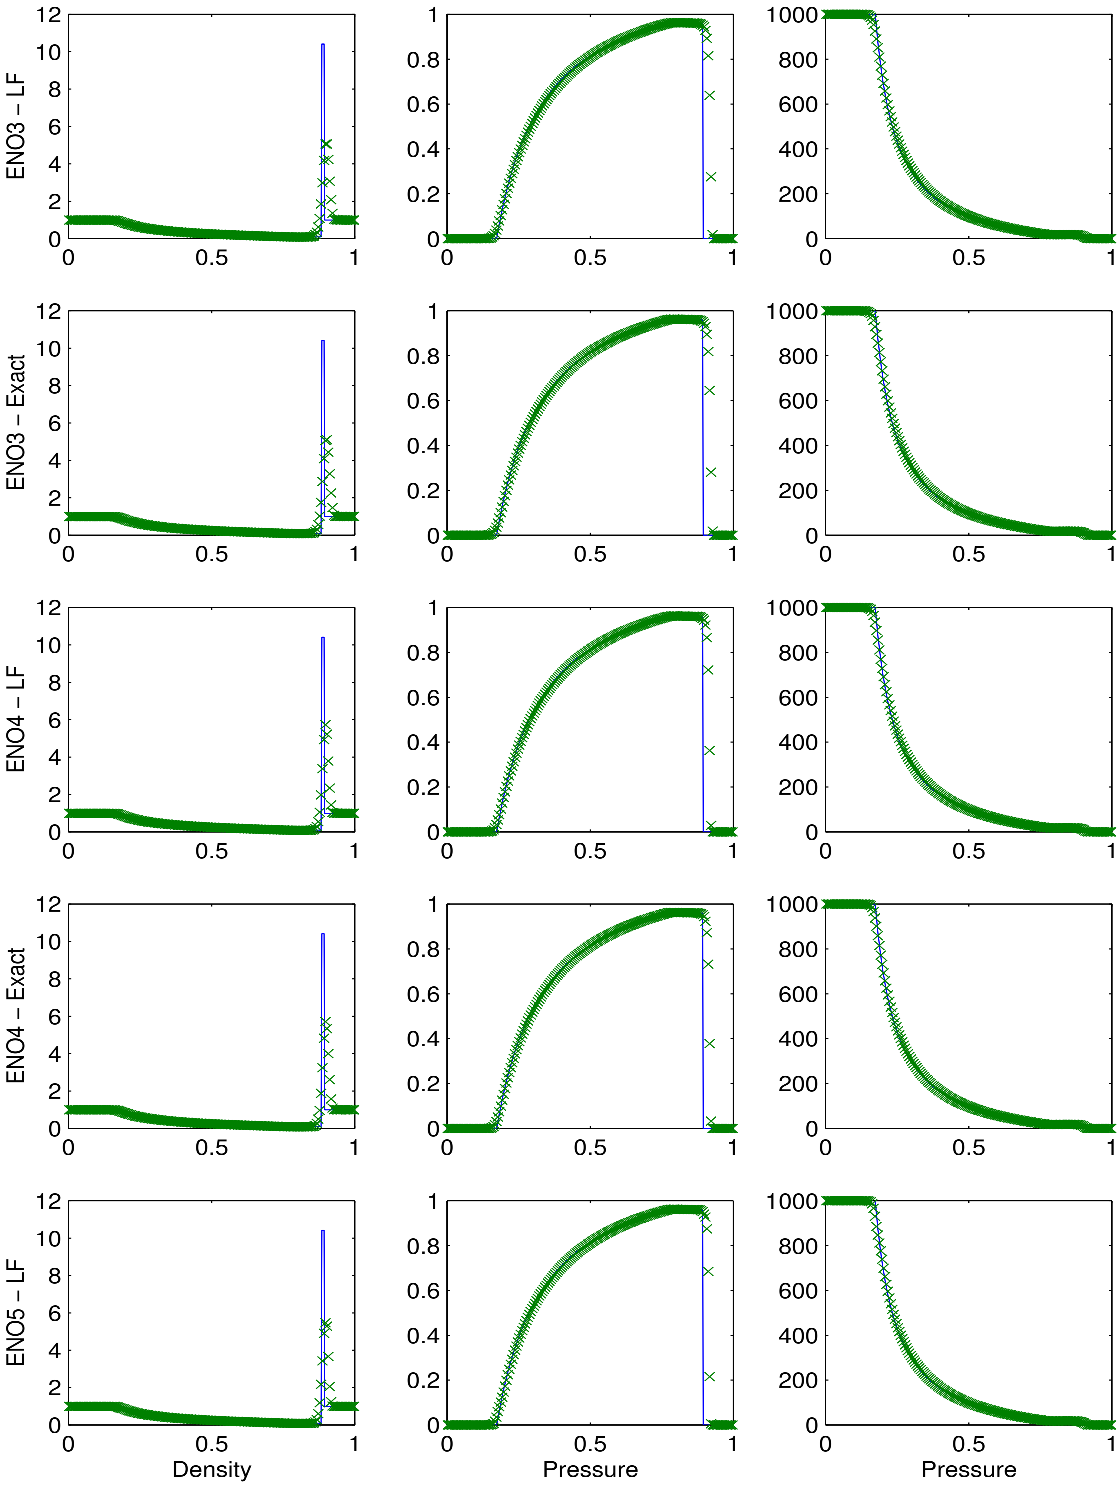

Figure 3.13: A shock tube problem for the relativistic ideal gas. The exact solution is a sonic rarefaction on the left and a shock on the right. All the methods are higher than second order accurate. The advantages of the exact solver are small. However, the increasing order of accuracy does improve the shock capturing.

the efficiency of relativistic HRSC methods using exact Riemann solvers on problems involving shock interactions, which should be more realistic than the simple Riemann problem tests found in the previous sections.

The computational domain is one unit long, with the initial discontinuities at $x_L = 0.1$ and $x_R = 0.9$. Reflective boundary conditions are enforced at either end of the domain. The initial data is

p_L	$(\rho_0)_L$	v_L	p_M	$(\rho_0)_M$	v_M	p_R	$(\rho_0)_R$	v_R
1000	1	0	0.01	1	0	100	1	0

The final time is $t_{\text{end}} = 0.85$; this gives results similar to the original Colella–Woodward problem at time 0.038.

As in the original Colella–Woodward problem there is no known exact solution. However we consider the results given in figure 3.14, calculated with 2000 points using a fourth order ENO scheme with a Marquina solver to be sufficiently close to the exact solution for our purposes.

With 400 computational cells and a CFL number of 0.7, this requires around 480 steps. For a second order method this requires the solution of around 385,000 Riemann problems and for a third order method around 575,000 Riemann problems. Despite the complicated shock wave structure, we expect the majority of the Riemann problems to be (nearly) trivial. Using an Essentially Non Oscillatory (ENO) reconstruction, the times (in seconds) for second and third order methods are given in table 3.3.

These results indicate the advantages of the Newton–Raphson method, with an improvement of 40-50% over the simplest Martí–Müller approach. The improvement is less when using the higher order method as the better reconstruction should result in easier Riemann problems to solve. The

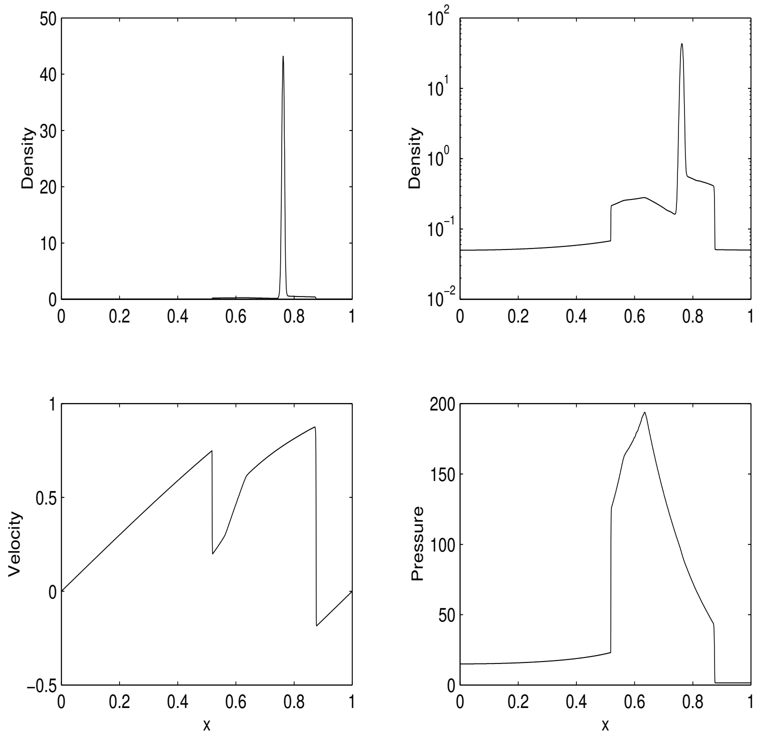

Figure 3.14: The relativistic Colella–Woodward problem at time 0.85. A fourth order ENO method with a Marquina solver, 2000 computational cells and a CFL of 0.5 were used. The second density plot has a logarithmic scale to show all the interesting features.

Rezzolla–Zanotti method with improvements of around 20% is less spectacular, but will probably be more stable as there is no possibility of producing an unphysical pressure.

Given the speed advantage of the Marquina solver over the most efficient exact solver, it is questionable whether it is worth using an exact solver at all. This is emphasized by the results shown in figures 3.15–3.18 which seem to depend far more on the order of the reconstruction than the Riemann solver used. The very slight improvement in the capturing of the shocks by

Solver	Second order	Third order
Marquina	8.3	12.2
Exact (MM)	23.3	31.6
Exact (RZ)	19.2	25.1
Exact (NR)	11.6	19.5

Table 3.3: Times in seconds for the solution of the Colella–Woodward problem.

the exact solver is compensated by the slight oscillations when used with the higher order methods. This is most noticeable in the pressure plot of figure 3.18 and seems to be due to slight problems at the sonic points being exaggerated after the shocks reflect from the boundaries.

However, we still believe that the exact solver can be useful in many numerical computations. The main reasons are

- All numerical Riemann solvers have problems somewhere. The more solvers that can be used, the more likely these problems are to be avoided.
- In higher dimensions the situation is complicated by the coupling of terms in the extra dimensions (such as the transverse velocity). It is likely that an exact solver will be necessary to capture these effects accurately.

3.10.3 Higher dimensional tests

Ultrarelativistic hydrodynamics tests

For the ultrarelativistic gas test we use a variant of the shock tube problem considered in the one dimensional tests. We keep the same initial data but

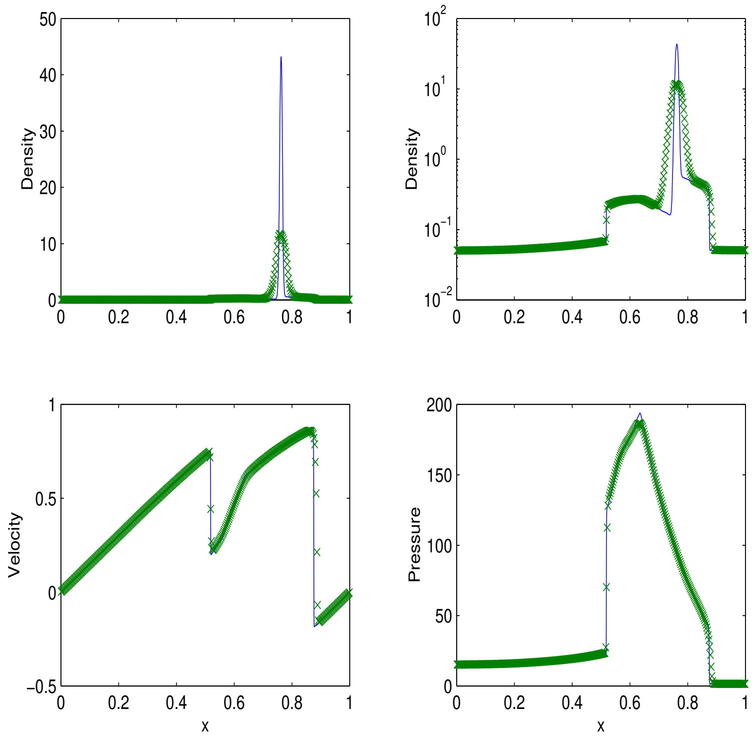

Figure 3.15: The relativistic Colella–Woodward problem at time 0.85. A second order ENO method with a Marquina solver, 400 computational cells and a CFL of 0.7 were used. The converged solution is given by the solid line.

set $v_L^t = 0.5$. The results were calculated using 200 cells and a CFL number of 0.7.

The results for the second order accurate methods are shown in figure 3.19. The SLIC method failed to complete this test. The slope limited methods all introduced considerable oscillations before the tangential discontinuity and the slope limited method with the exact solver failed to complete the test. Instead we include the second order ENO method with the exact solver. The reason for these problems is probably the reconstruction

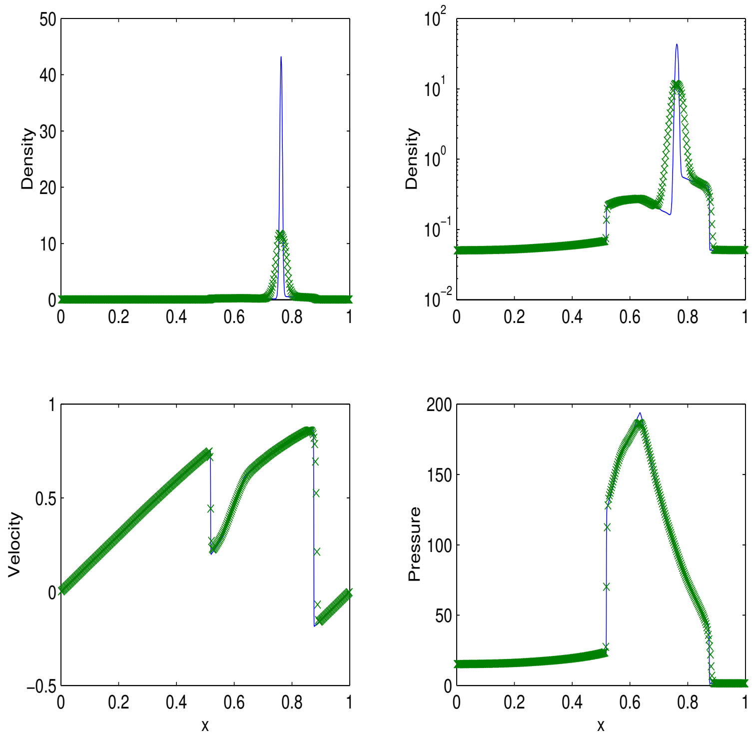

Figure 3.16: The relativistic Colella–Woodward problem at time 0.85. A second order ENO method with an exact solver, 400 computational cells and a CFL of 0.7 were used. The converged solution is given by the solid line.

which is performed on the conservative rather than characteristic variables in these cases. This is shown by comparing the slope limited method using the Lax–Friedrichs solver and the second order ENO method using the Lax–Friedrichs solver. The main difference is that the ENO methods reconstruct on the characteristic variables, avoiding the oscillations. The WAF method also avoids these problems as it only performs a first order reconstruction of the data. The advantages of using an exact solver are clear from the results of the WAF and ENO2 methods.

The results for the higher order methods are shown in figure 3.20. The

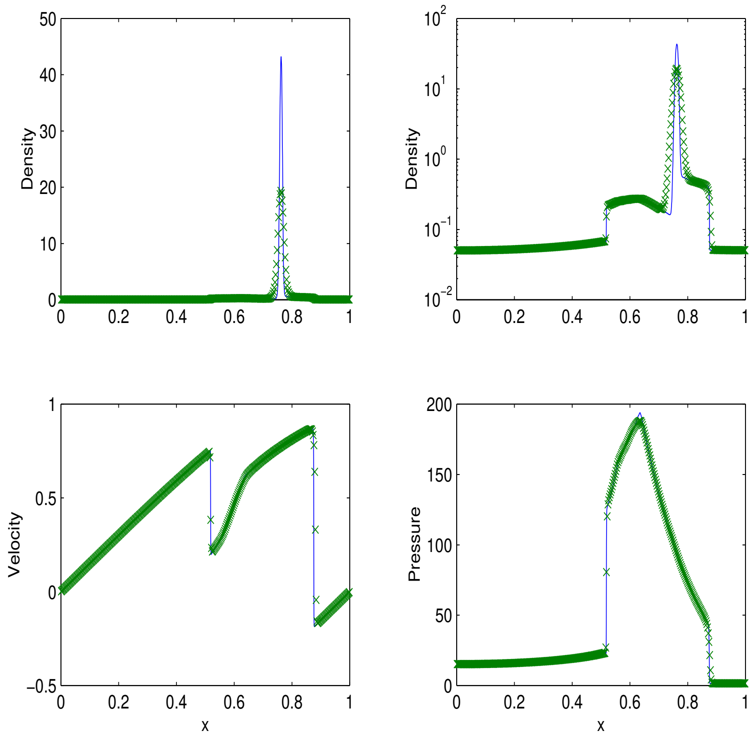

Figure 3.17: The relativistic Colella–Woodward problem at time 0.85. A third order ENO method with a Marquina solver, 400 computational cells and a CFL of 0.7 were used. The converged solution is given by the solid line.

fourth order ENO method with the exact solver failed to complete this test. When the reconstruction is checked to ensure it is physical across the cell, the implementation used is only fully correct for a third order method, probably causing the failure in this case. The advantages of the higher order methods and the more accurate Riemann solvers are clear.

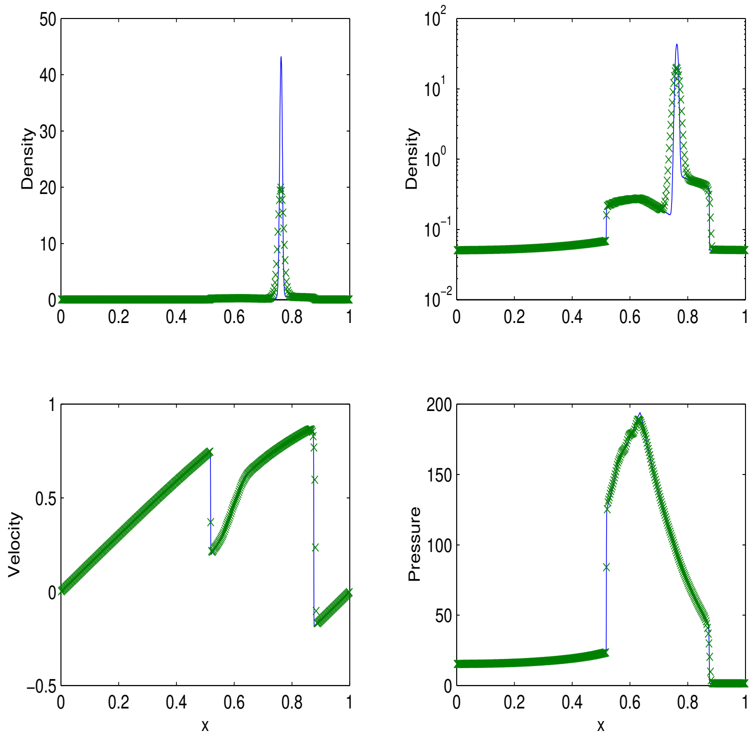

Figure 3.18: The relativistic Colella–Woodward problem at time 0.85. A third order ENO method with an exact solver, 400 computational cells and a CFL of 0.7 were used. The converged solution is given by the solid line.

Relativistic ideal gas tests

For the relativistic ideal gas we look at a variant of the Sod problem considered in the one dimensional tests. We keep the same initial data but set $v_L^t = 0.5$. The results were calculated using 100 cells and a CFL number of 0.7.

The results for the second order accurate methods are shown in figure 3.21. We did not implement the exact solver as it is so inefficient and hence we did not implement the WAF method either. Every method over-

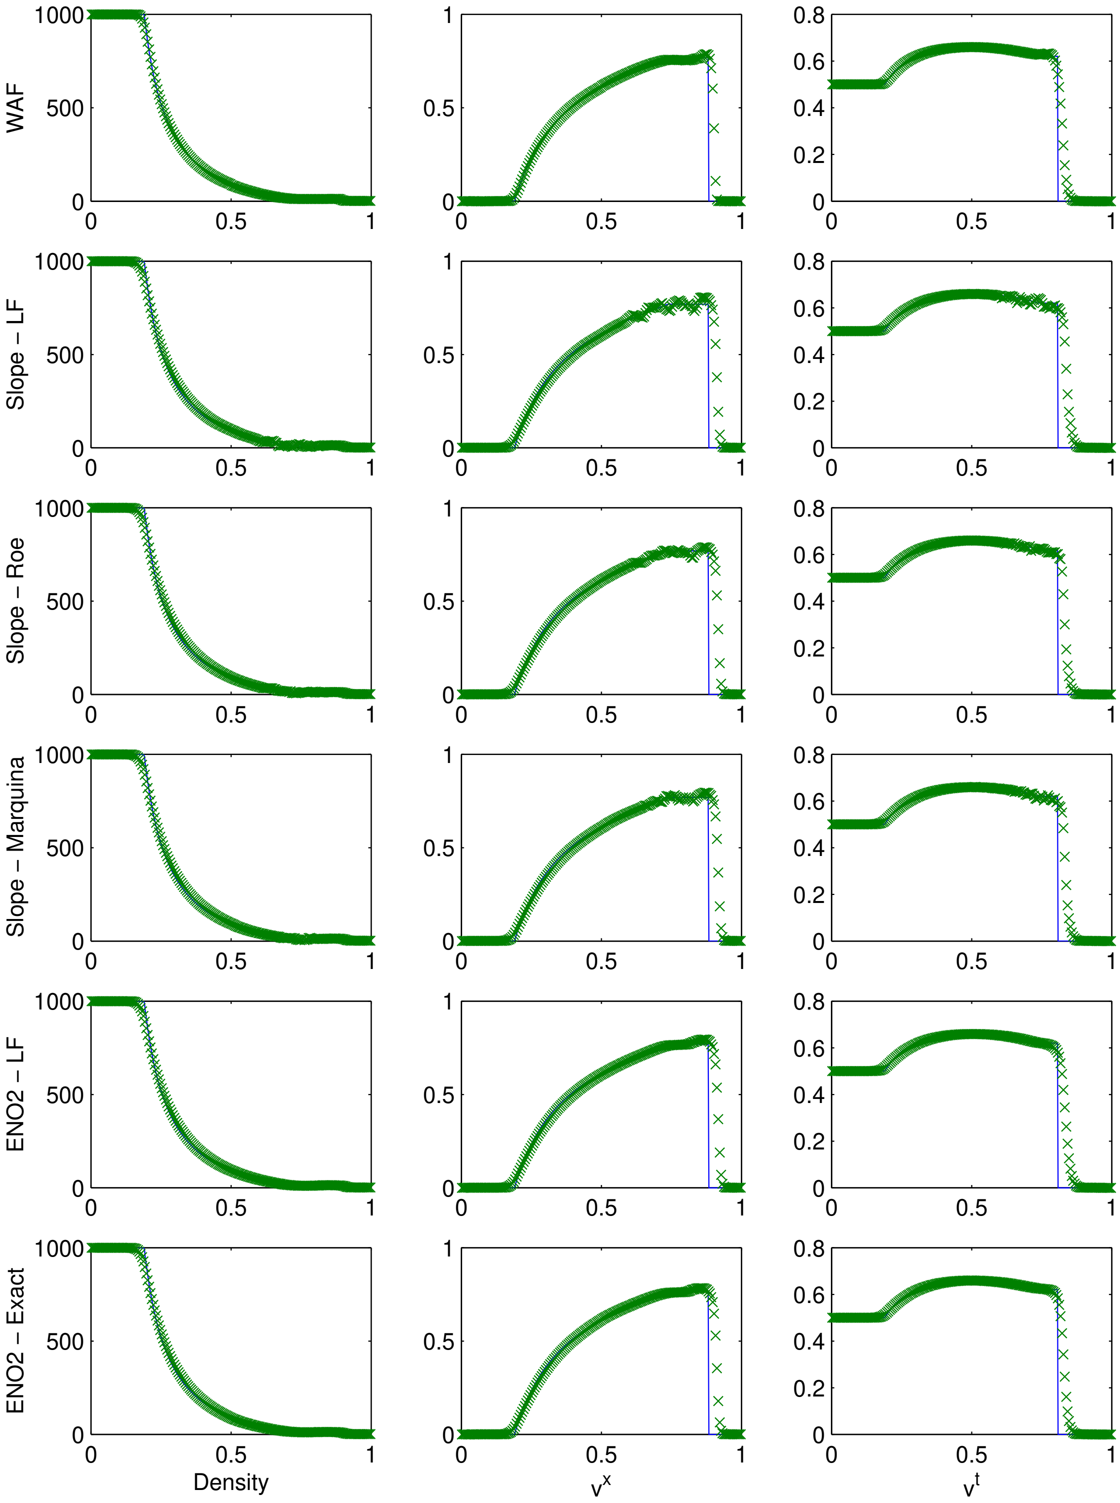

Figure 3.19: A shock tube problem for the ultrarelativistic gas. The exact solution is a rarefaction on the left and a shock on the right separated by a tangential discontinuity. All the methods are second order accurate. Each method overshoots before the tangential discontinuity. The less diffusive Riemann solvers capture the tangential discontinuity considerably better. The slope limited methods are very poor at this test.

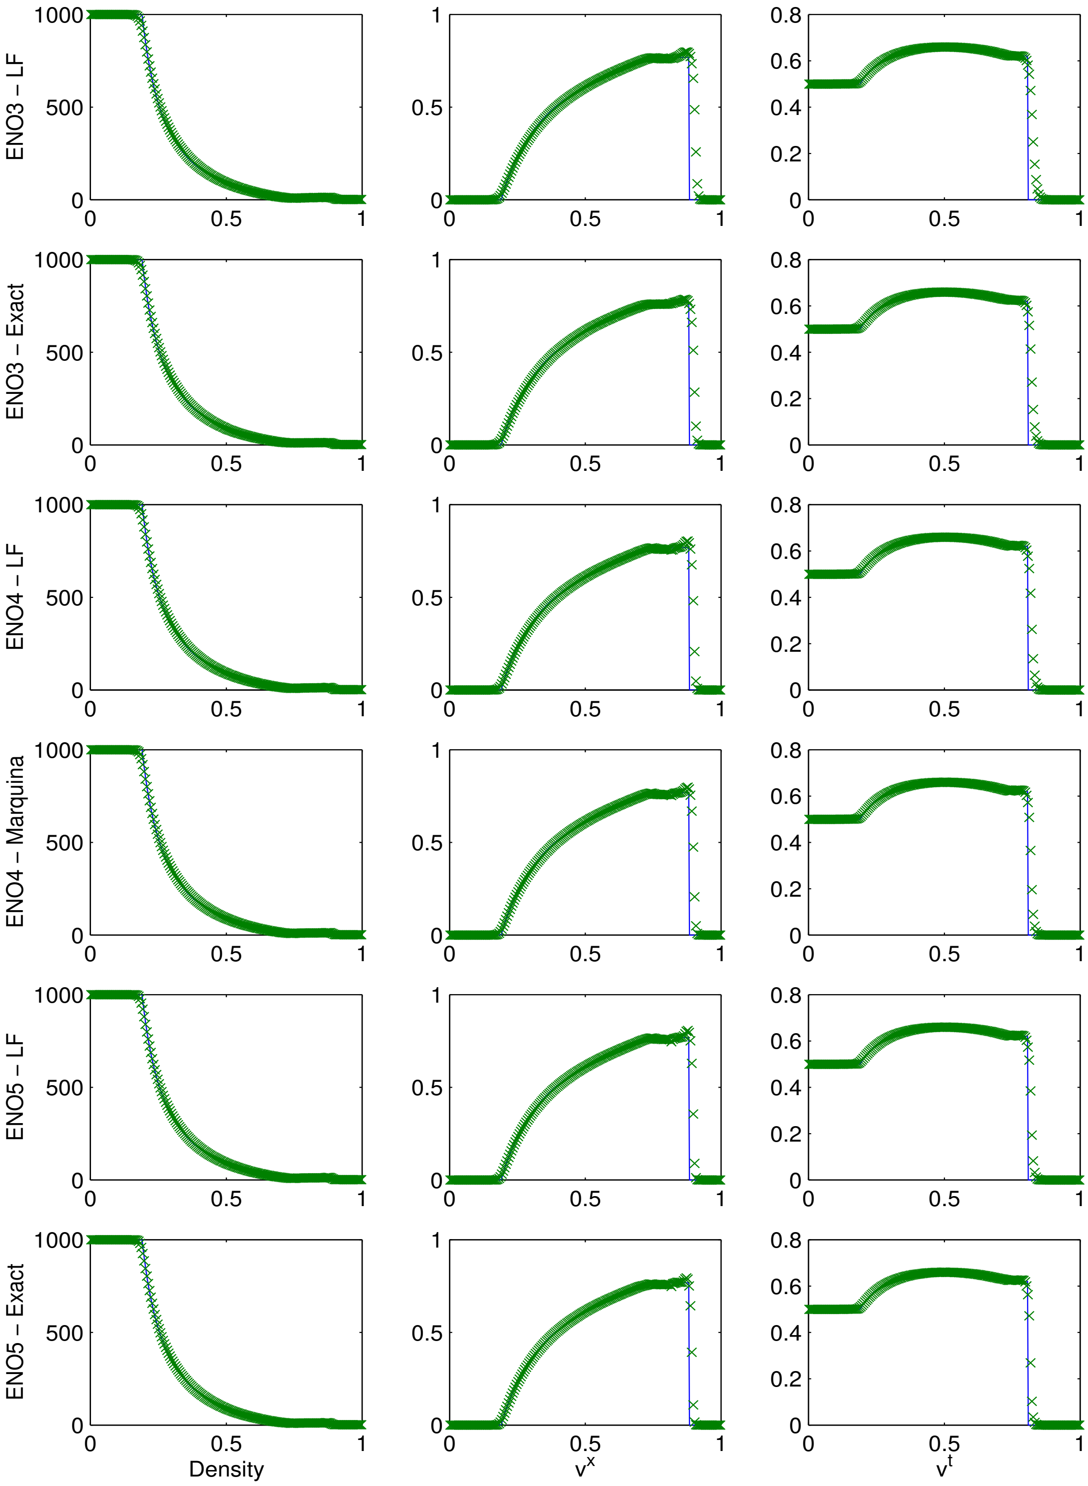

Figure 3.20: A shock tube problem for the ultrarelativistic gas. The exact solution is a rarefaction on the left and a shock on the right separated by a tangential discontinuity. All the methods are higher than second order accurate. Each method overshoots before the tangential discontinuity. The less diffusive Riemann solvers capture the tangential discontinuity considerably better. The increase in the order of accuracy considerably improves the shock capturing.

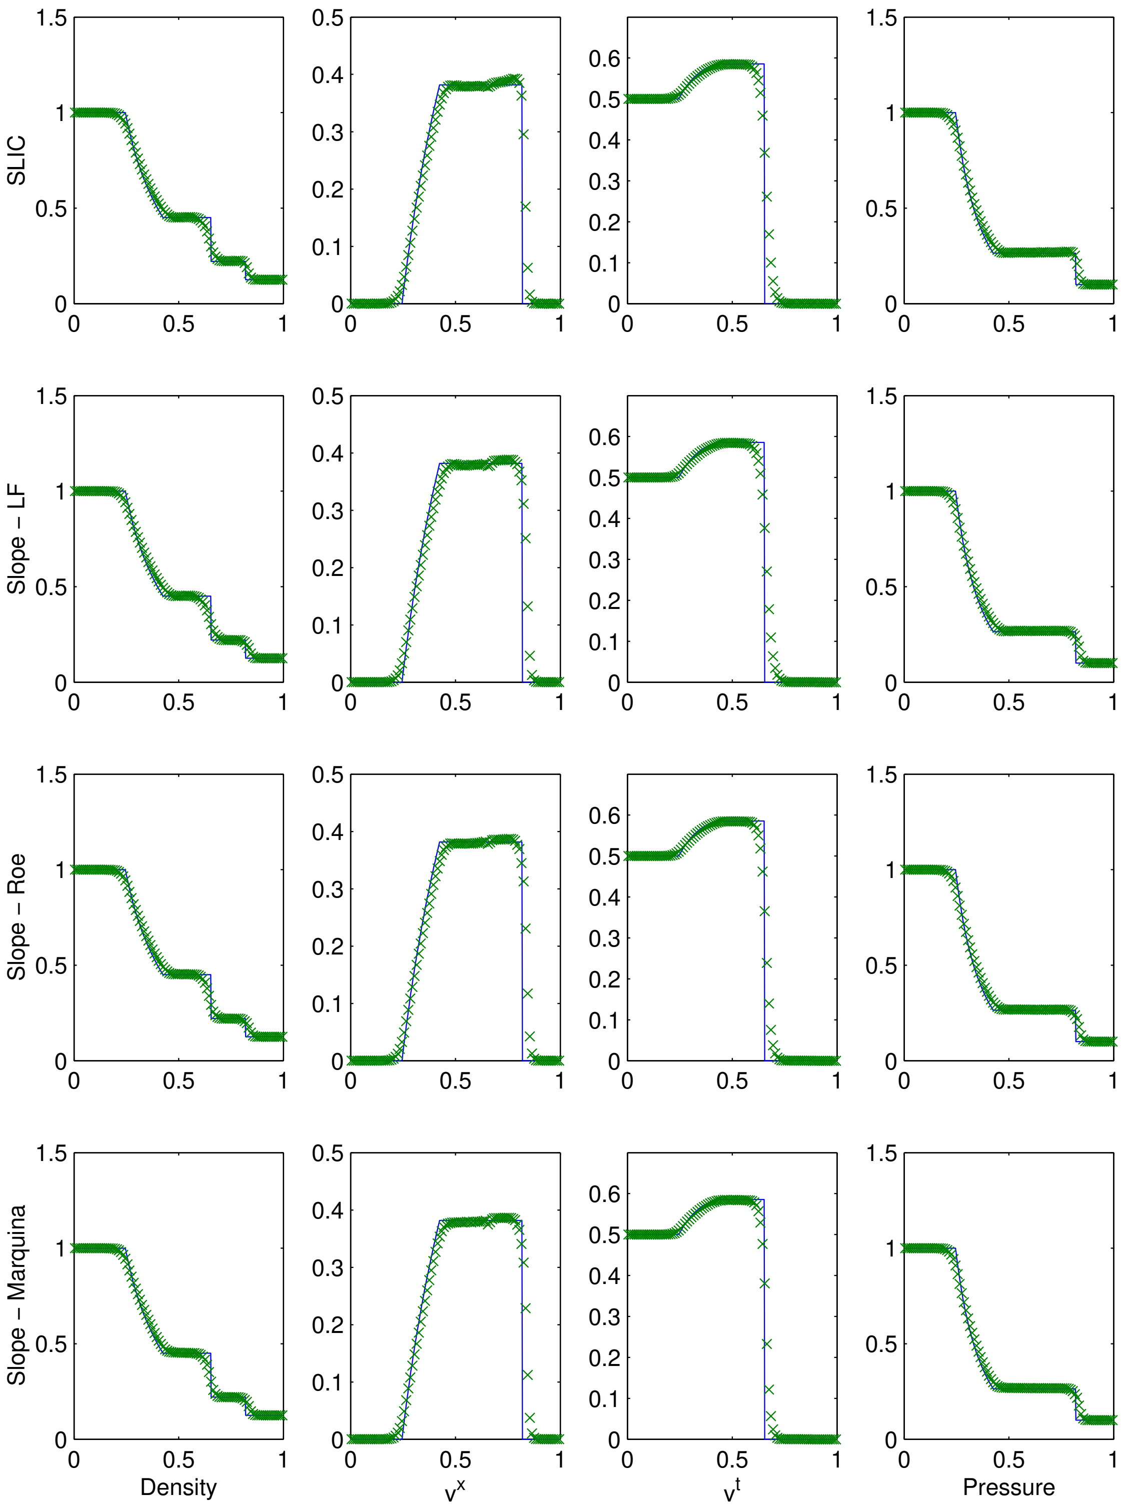

Figure 3.21: A shock tube problem for the relativistic ideal gas. The exact solution is a rarefaction on the left and a shock on the right separated by a contact with a tangential discontinuity. All the methods are second order accurate. Each method overshoots before the contact. The less diffusive Riemann solvers capture the contact considerably better.

shoots in the density to the left of the contact. It is clear how the capturing of the contact surface and tangential discontinuity is improved by the more accurate Roe and Marquina solvers.

The results for the higher order methods are shown in figure 3.22. The advantages of the higher order methods are obvious, especially in the capturing of the tangential discontinuity. Even with the fifth order method it is now possible to see improvements using the more accurate Marquina solver. However the difference between the solvers is much less than between different orders of the reconstruction algorithm.

3.10.4 Validation tests: conclusion

In the introduction to the tests we stated our requirements for the numerical scheme. It should capture shocks and other discontinuities in very few points with good convergence properties and without introducing oscillations, whilst still being computationally efficient. Given the results of the tests we believe that these requirements are best fulfilled using a third order ENO method with the Marquina flux.

From the point of view of accuracy, it is clear that for a reasonably long evolution the higher than second order methods are better. This is shown in the linear advection tests and most vividly shown in figures 3.3 and 3.4. For nonlinear systems it is also illustrated by the results of the relativistic Colella–Woodward test of section 3.10.2. Although the compact second order methods, especially WAF, are sometimes able to capture shocks and rarefactions as well as the higher order methods, they can also increase any

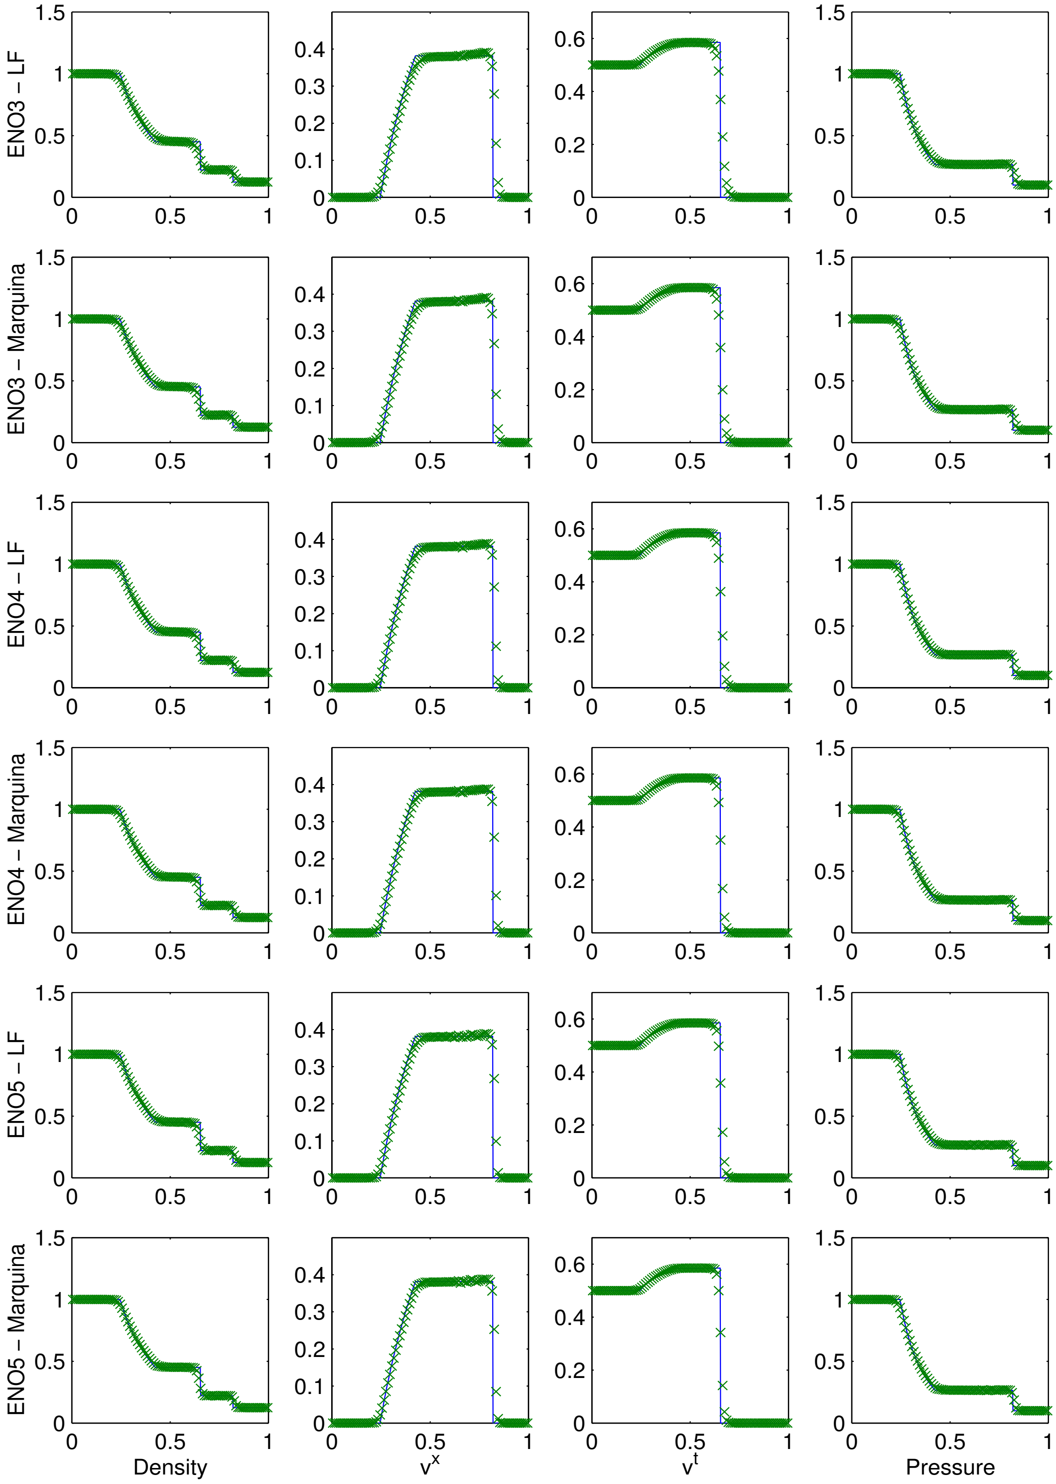

Figure 3.22: A shock tube problem for the relativistic ideal gas. The exact solution is a rarefaction on the left and a shock on the right separated by a contact with a tangential discontinuity. All the methods are higher than second order accurate. The advantages of the Marquina solver are visible. However, the increasing order of accuracy provides greater improvements.

overshoots near discontinuities (for example, compare the results for SLIC in figure 3.21 with the higher order results in figure 3.22) or have problems (usually small) near sonic points (for example, compare the results for WAF in figure 3.5 with the higher order results in figure 3.6). Given that the shock capturing behaviour of the method seems to be strongly affected by the stencil size as explained in section 3.10.1 it is perhaps surprising that the methods based on the Cauchy–Kowalewski approach, which have considerably smaller stencils, did not show improvements over the method of lines based schemes. This is probably because only second order methods were implemented. The rewards of implementing higher order methods using this approach seemed unlikely to be worth the computational complexity of implementing them.

From the point of view of stability, meaning a combination of monotonicity preservation and consistent physical evolution (that is, not computing an unphysical state), the third order methods are better than those of higher order. This is because it is easily possible, given the reconstructed boundary data, to check that a third order reconstruction is physical across the entire computational cell. With the higher order reconstructions this is more difficult and so was not fully implemented. With the more compact WAF and SLIC methods it is not always possible to check that a reconstructed state or flux is physical, which can cause the method to fail. The third order method also does not introduce the very slight oscillations that can be seen near the shock in figure 3.22 which are probably due to the non-TVD nature of the fourth order Runge–Kutta solver.

Finally, from the point of view of efficiency, we note that the very slight advantages of the exact solver do not seem to justify using it in a realistic calculation (although it may be essential in testing the validity of results). The improvements in accuracy from using the Marquina solver over the more diffusive Lax–Friedrichs solver, most notable in figure 3.22, are sufficient to justify the use of the more complex approximation in a realistic calculation. The use of higher than third order methods is difficult to justify on efficiency grounds unless the simulation is for a very long period of time and very small features need to be resolved. Long term gravitational wave simulations may benefit from such methods, but the hydrodynamical simulations of chapter 5 will probably not.

Chapter 4

The Numerical grid

In attempting to solve any differential equation numerically we must make some approximations. One of the central approximations is the numerical grid that underlies how the data is stored on the computer. The grid provides constraints on the algorithms we can use, such as the CFL condition outlined in section 3.1. How we view the grid also changes how we view the algorithms that we use. We shall describe the different viewpoints in the opening section. The remainder of the chapter will be devoted to the uneven and Adaptive Mesh grids that allow greater accuracy for a given size of grid.

4.1 The basic grid

For our purposes we shall consider a grid to be a collection of points distributed in space at a given time. They will be denoted $\mathbf{q}_i$, where i runs over $1, \ldots, N$. The points need not be evenly distributed in space. We shall denote the physical position of each point by x_i. The numerical algorithms

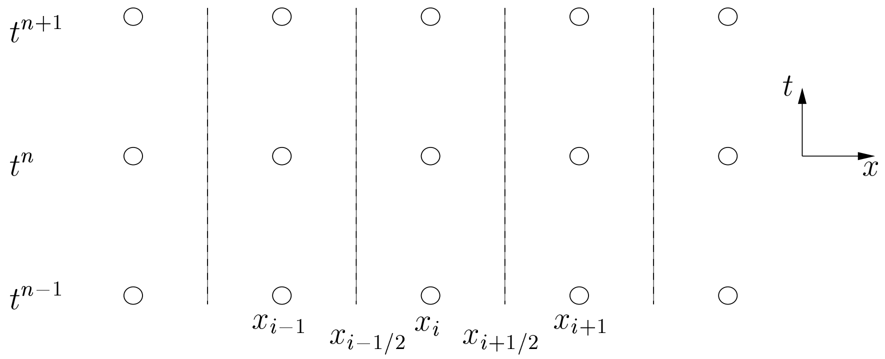

Figure 4.1: Three computational grids. The circles represent the points in computational space. The dashed lines represent the boundaries of the computational cells.

described in chapter 3 will take the data from a grid at time t^n, denoted $\mathbf{q}_i^n$, and use this (and possibly other conditions) to calculate the data on the grid at time t^{n+1}, denoted $\mathbf{q}_i^{n+1}$.

Each point $\mathbf{q}_i^n$ in computational space can represent many different things in physical space. The simplest representation is that the computational point $\mathbf{q}_i^n$ represents the value of the function $\mathbf{q}$ at the spatial position x_i and time t^n. This *finite difference* representation is useful for simple equations. However, this viewpoint has problems dealing with discontinuous functions.

The viewpoint generally used for HRSC methods is that the computational point represents an average of the function $\mathbf{q}$ over a *cell* in physical space. Specifically, we define the *cell-width* Δx_i so that the cell boundaries $x_{i\pm1/2} = x_i \pm \frac{1}{2}\Delta x_i$ coincide and that x_i is at the centre of the cell, as illustrated in figure 4.1. Then the computational point $\mathbf{q}_i^n$ represents the integral average of the function $\mathbf{q}$ over the cell $\left[x_{i-1/2}, x_{i+1/2}\right]$,

$$\mathbf{q}_i^n = \frac{1}{\Delta x_i}\int_{x_{i-1/2}}^{x_{i+1/2}} \mathbf{q}(t^n, x)\mathrm{d}x. \tag{4.1}$$

This *finite volume* viewpoint is natural when dealing with the weak form of a partial differential equation such as (2.3), and so is natural when dealing with discontinuities. It is however more difficult to deal with source terms, because we apply the source function to $\mathbf{q}$ and not the integral average. Higher order methods are also less efficient in higher dimensions, because simply performing reconstructions in one dimension at a time only gives the value of $\mathbf{q}$ as an integral average over the appropriate face of the computational cell. This is a second order accurate approximation to the value of $\mathbf{q}$ at the central point of the face. If we require a higher order of accuracy we have to perform another reconstruction on these face averages to find the values of $\mathbf{q}$ at individual points to solve the Riemann problem, greatly increasing the cost of the reconstruction procedure.

For the rest of this chapter we shall take the finite volume viewpoint. For more details on the advantages and disadvantages of finite difference viewpoints and methods we refer to Toro (1999), Laney (1998) and Shu (1997).

4.2 Uneven grids

The best way to improve the accuracy of a numerical simulation is typically to add more computational points. However, the CFL condition outlined in section 3.1 then restricts the time step Δt. In a generic simulation there will be large parts of the domain which are "uninteresting" and which do not need the finest grids. Rather than using a uniform grid over the entire domain there are many strategies for cutting down the computational expense

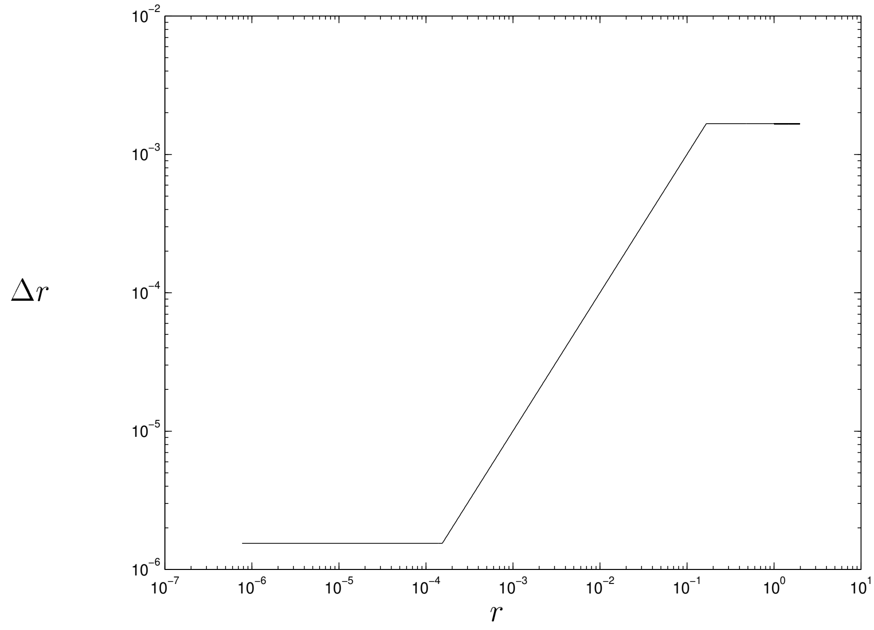

Figure 4.2: An example of an unevenly spaced grid. The cell width Δr is plotted against r. Even spacing is used near the origin $r = 0$ and at the outer boundary. The logarithmic spacing in between allows a wide range of cell widths to be covered.

without harming the accuracy.

The simplest is to use an uneven grid. If the "interesting" parts of the domain are known then the points can be increased in this area of the grid. For example, simulations in spherically symmetric critical collapse such as Neilsen and Choptuik (2000), Garfinkle (1995) frequently use grids that are a compromise between even and logarithmic spacing. A version of this which we use later is shown in figure 4.2. Less than 2000 points were required for this grid, whilst an even grid would have required over 1.2 million points for the same resolution near the origin $r = 0$.

The formulation of a numerical method is identical on an even or uneven grid. However, as noted in chapter 3, the implementation of a numerical method is frequently considerably simpler on an even grid. Another problem is that the CFL condition imposes a very severe restriction on an uneven grid of the type shown in figure 4.2. Typically the timestep is determined by the size of the smallest cells. So the timestep required for stability of the entire grid is orders of magnitude smaller than the timestep required by the coarser cells. This will cause large amounts of extra dissipation of any steep gradients resolved on the coarse regions of the grid. This then means that in order to accurately capture steep gradients and shocks in the coarse part of the grid, higher order numerical methods are required which increases the computational expense. This can be avoided using a more complicated method.

4.3 Adaptive uneven grids

The restriction of the timestep for an uneven grid such as is shown in figure 4.2 is most simply avoided by only using a fine grid where and when such resolution is required. The most complete method that implements this idea is that of Adaptive Mesh Refinement, outlined in section 4.4. A simpler method is to use one fixed grid that may vary in time. This will be the main approach used in chapter 5 for reasons explained there.

For the evolution step the grid is held to a fixed size and resolution. After each evolution step an indicator is checked to see whether more resolution is required. For the problem in chapter 5 we are trying to find small black holes in spherical symmetry. For this we add cells near the origin if the apparent horizon of the black hole is resolved in too few cells. The indicator itself is given explicitly in chapter 5.

If more resolution is required then the extra cells are added to the grid and filled using interpolation. The interpolation method used is that appropriate to the evolution method. That is, when the evolution method uses third order ENO reconstruction, the interpolation is performed using a third order ENO method. The spatial position of the extra cells is set so that they vary linearly in $\log(r)$ as in figure 4.2. Some care is required to ensure that the cell boundaries match both at the origin and at the end of the varying region.

To ensure this matching we set the variation of the space step from the number of extra points we wish to add. If we wish to add N cells near the origin, then we enforce that the new grid will have a constant space step Δr_i for the first N cells, but that space step for the next N cells will vary as

$$\Delta r_i = \Delta r_{i-1} * \left(\frac{2N+1}{2N} \right)^2 . \tag{4.2}$$

This ensures that the cell boundaries match and also that when more resolution is added the interpolation will be using a grid that has a constant space step. This may reduce interpolation error.

A slightly different algorithm was implemented by John Stewart. Instead of using a grid that smoothly varies in space, the spatial step is reduced by a constant factor (here always 2) on the refined grid. This considerably simplifies the interpolation algorithm, minimizing interpolation error. However the discontinuous change in the space step increases the chances of error at the edges of the refined grid. In practice both methods produced near identical results with no evidence of errors from the grid adaption.

The algorithm to be used in chapter 5 also varies at the outer boundary. We were unable to find a stable outer boundary condition. Instead the grid was expanded at this outer boundary to ensure that it remained as close as possible to a known exact solution. Then this known solution could be used at the outer boundary. In practice, this just required the addition of extra cells at the outer edge of the grid, with the space step the same as at the edge of the grid, and the values of the functions set by the known solution.

Although this method reduces the computational expense of using an uneven grid it does not eliminate it, and it also does not help with the problem of the dissipation of steep gradients and shocks caused by the small timestep. In order to totally avoid these problems a considerably more complicated method is required.

4.4 Adaptive Mesh Refinement

The Adaptive Mesh Refinement (AMR) algorithm first outlined by Berger and Oliger (1984) can give high accuracy using a small number of computational points without the problems of the CFL condition associated with uneven grids. It has been used frequently in many forms since. An overview of existing work is given in Plewa (2001). Some of these codes are intended to be general purpose allowing for the simple implementation of different sets of equations, such as MacNeice et al. (2000), Fryxell et al. (2000), Cactus (2001) and Brown et al. (1997).

AMR was introduced to numerical relativity by Choptuik (1992) in his study of critical collapse. Since then AMR has been used in critical collapse simulations such as Choptuik et al. (1996), Hamadé and Stewart (1996), Hamadé et al. (1996), as well as the ongoing development of various codes to enable the simulation of fully three dimensional black hole collisions such as Cactus (2001), Papadopoulos et al. (1998), Brügmann (1996). In what follows we shall use an AMR code originally written by John Stewart. Previous versions of this code were used in Hamadé and Stewart (1996), Hamadé et al. (1996), Hern and Stewart (1998), Hern (1999). The most comprehensive review of the code is given in Hern (1999) which this presentation largely follows. The main improvements to the code have largely been to the time stepping and interpolation routines and are appearing here for the first time.

4.4.1 The adaptive grids

The underlying idea behind AMR is that extra computational points should only be added where and when they are required, without affecting those regions where they are not. To achieve this, at any given time there may be an arbitrary number of grids, each having a different distribution of cells. For simplicity we assume that each grid has evenly spaced cells. Then the grids are grouped into *levels* L^k each containing an arbitrary number of grids G_j^k where any grid on level k has cell width Δx_k. The base level L^1 covers the entire domain (typically with a single grid) using a coarse cell spacing. An example is shown in figure 4.3. The finer levels are added only where necessary. Typically the cell width reduces using the rule

$$\Delta x_{k+1} = \frac{\Delta x_k}{N_{\text{refine}}}. \tag{4.3}$$

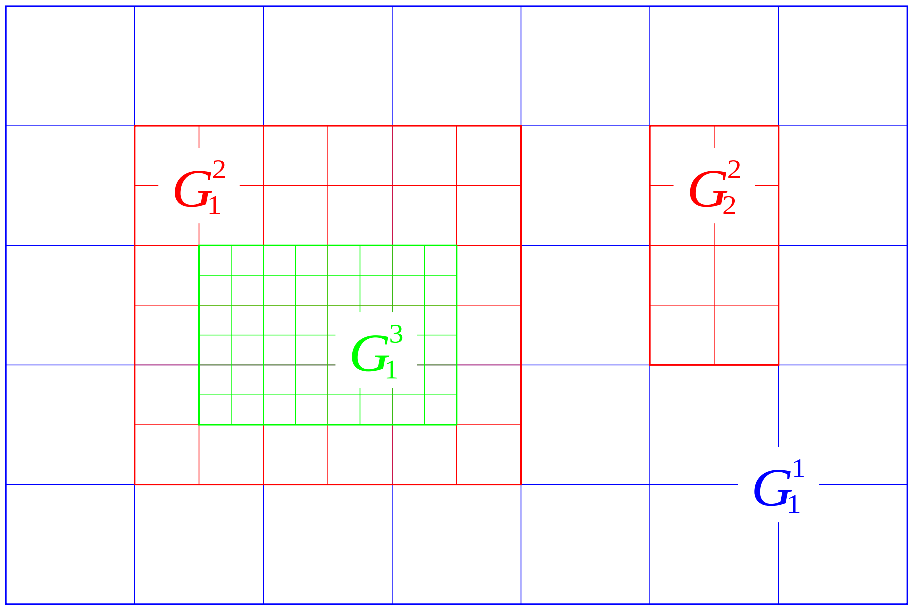

Figure 4.3: Three levels of adaptive grids. The finer grids can be evolved almost entirely independently of the coarse grids. This allows for increased accuracy at a small computational expense. This figure follows Hern (1999).

Here we are implicitly assuming that each grid is even. This is not necessary, but simplifies the implementation.

As each grid is evolved independently (except for finding the boundary conditions on fine grids - see section 4.4.5), the majority of the domain can be evolved using large timesteps. The finer grids will require considerably shorter timesteps, but this should affect only a small number of computational points.

In what follows we use the convention that the computational point at time t^n and spatial point x_i on level L^k is denoted $(I_i)^k$ with data $(\mathbf{q}_i^n)^k$, and that the computational points on the finer levels are numbered as if they were part of an even grid covering the entire domain. This means that the computational cell covered by $(I_i)^k$ is also covered by the union of $\left(I_{(i-1)N_{\text{refine}}+j}\right)^{k+1}$ where $j = 1, \ldots, N_{\text{refine}}$.

The AMR algorithm requires a lot of management to deal with the issues of multiple grids. Each is dealt with in turn.

4.4.2 Time stepping the grids

In order to make AMR as effective as possible each level should be evolved with its own timestep. This typically means that the finer levels will not be evolved as fast as the coarser. In order that boundary conditions can be supplied for the grids on the finer levels we use a recursive algorithm, illustrated in figure 4.4, and written as pseudo-code for level k in figure 4.5.

The order in which the levels are advanced ensures that there are always grids on level L^k from which boundary data for grids on L^{k+1} can be found.

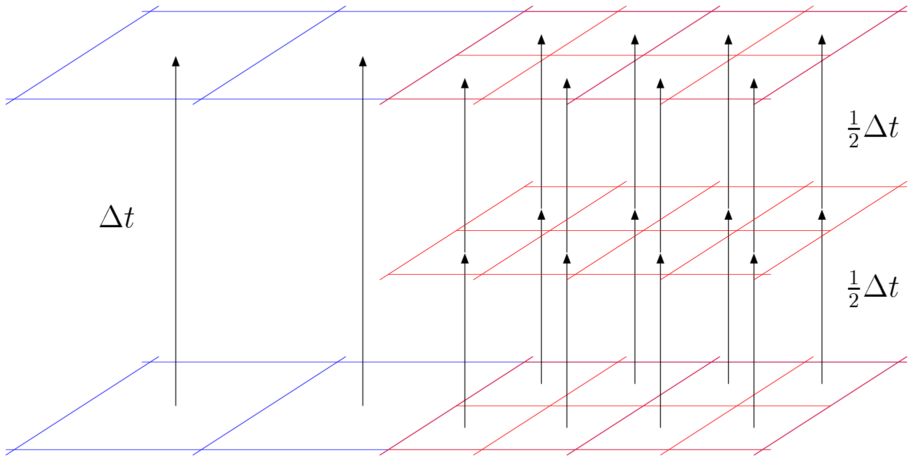

Figure 4.4: The finer grids are evolved using their own timestep, bypassing the restrictions of the CFL condition on the majority of the domain. This figure follows Hern (1999).

`timestep`(L^k) :
 Evolve level L^k forward by Δt^k
 If There are finer levels **then**
  `timestep`(L^{k+1})
 end if
 If Level L^k and level L^{k-1} are at the same time **then**
  Inject the data from level L^k into level L^{k-1}
  Finish evolving this level
 end if

Figure 4.5: Pseudo-code for time stepping the adaptive grids.

This requires interpolation in time, which means that the AMR algorithm must store more data than a typical method. However, this comes in useful in other ways as shown in section 4.4.4. If the grids on the finer levels lie on the physical boundary we should be sure to use the physical boundary conditions used by the coarse grid L^1.

When two levels reach the same time we want them to be consistent. As we assume that the finer grids are more accurate, we simply inject the data from the finer level to the coarser. In order to retain numerical conservation we use

$$(\mathbf{q}_i^n)^k = \tfrac{1}{N_{\text{refine}}} \sum_{j=1}^{N_{\text{refine}}} \left(\mathbf{q}^n_{(i-1)N_{\text{refine}}+j}\right)^{k+1}, \tag{4.4}$$

where we are using the conventions of section 4.4.1.

It is usual to refine the time step Δt in the same way as the space step Δx is refined in equation (4.3). This simplifies the timestepping, as each fine level requires N_{refine} steps to synchronize with the coarser level. It also simplifies the interpolation of any boundary data, shown in section 4.4.5. However, if the feature being refined is very poorly resolved on the coarsest grid the time step produced may not be small enough for stability. In particular for a very small black hole the apparent horizon may be very poorly resolved at the top level with the velocity of the fluid relatively small ($\approx -0.6c$). The characteristic speeds, given by equation (3.78) are therefore over 10% too small on the finer grids where $v \to -c$. For this reason we have implemented variable time steps over all levels. This does not seem to affect either the accuracy or efficiency of the code, but does increase the stability.

4.4.3 Adapting the grids

The central point of the AMR algorithm is the finer levels with their finer grids. So the central question must be their creation, adaption and destruction, as well as how they interact with the other grids.

Finer grids (or levels) are created when some error measure E_i becomes greater than some tolerance $\texttt{tol}_1$ on the cell $(I_i)^k$. The finer grid created covers the cells where the error measure is greater than the tolerance and some buffer zone around this area to ensure that all the "interesting" features are refined. The data on the fine level L^{k+1} is created by interpolation from the grids on the coarse level L^k, which will be covered in section 4.4.6. If fine grids abut or overlap they are merged.

Finer grids (or levels) are destroyed when the error measure E_i drops below some tolerance $\texttt{tol}_2$ on the cell $(I_i)^k$. Typically we take $\texttt{tol}_2 = \texttt{tol}_1$. As we wish to ensure *proper nesting*, that is that any grid on level $k+1$ is completely contained in a unique grid on level k, we must ensure that we do not destroy a grid without first ensuring that all of its child grids should also be destroyed. This procedure is again recursive over all levels. The procedure is outlined as pseudo-code for level k in figure 4.6.

In one dimension all that is required is some error measure - this is discussed in section 4.4.4. In more than one dimension matters are more complicated. The cells marked in error will typically not form a rectangular grid. A simple rectangular grid surrounding the marked cells will typically be inefficient as it will contain too many unmarked cells. Algorithms to "cluster" the new child grids in the most efficient manner are quite complex - for the

```
regrid(L^k) :
    If There are finer levels then
        regrid(L^{k+1})
    end if
    Inherit the errors from finer levels
    If E_i > tol_1 then
        Mark I_i to be refined
    end if
    Create the new child grids, copying data from previous children if possible
    Destroy the old children and, if necessary, the finer level
```

Figure 4.6: Pseudo-code for error checking the adaptive grids.

details we refer to Bell et al. (1994).

4.4.4 Error estimation and Richardson extrapolation

In order to decide where we should refine we must estimate the error arising from the numerical approximation on each grid. That is, we want to find the *local truncation error* σ, the error the numerical algorithm makes by taking a single step from given data:

$$(\sigma)_i = \left| \mathbf{q}(t^{n+1}, x_i) - \mathbf{q}_i^{n+1} \right|, \tag{4.5}$$

where $\mathbf{q}(t^{n+1}, x_i)$ is the exact solution to our evolution equation assuming initial data $\mathbf{q}_i^n$ at time t^n.

As the definition of the local truncation error requires the knowledge of the exact solution at time t^{n+1} we cannot measure it directly. However we can get a good approximation by using *Richardson extrapolation.* Assume we are using a k^{th} order accurate numerical method (see section 3.1) so that, starting from the exact solution at time t^{n-1} the numerical solution is

$$\mathbf{q}_i^n = \mathbf{q}(t^n, x_i) + ((\sigma_1)_i h)^k + \mathcal{O}(h^{k+1}), \tag{4.6}$$

where σ_1 is some constant and h a step length. Assume that the grids at time t^{n-1}, t^n and t^{n+1} are evenly spaced in time. Then we can find two numerical solutions at time t^{n+1}. The first, taking the two small steps, has

$$\left(\mathbf{q}_i^{n+1}\right)_h = \mathbf{q}(t^{n+1}, x_i) + ((\sigma_h)_i h)^k + \mathcal{O}(h^{k+1}), \tag{4.7}$$

whilst the second, going directly from time t^{n-1} to t^{n+1} has

$$\left(\mathbf{q}_i^{n+1}\right)_{2h} = \mathbf{q}(t^{n+1}, x_i) + ((\sigma_{2h})_i 2h)^k + \mathcal{O}\left((2h)^{k+1}\right). \tag{4.8}$$

The difference between these two approximations is

$$\left(\mathbf{q}_i^{n+1}\right)_{2h} - \left(\mathbf{q}_i^{n+1}\right)_h = ((\sigma_{\text{diff}})_i h)^k + \mathcal{O}\left((h^{k+1}\right), \tag{4.9}$$

which is proportional to the local truncation error.

We note that this analysis relies on the Taylor expansion of the solution and so is not valid in the region of a discontinuity. However in such a region the local truncation error will typically be $\mathcal{O}(1)$. So using this method may not give a good approximation to the local truncation error, but should give an error estimate sufficiently large to ensure that regridding takes place.

Richardson extrapolation may give an accurate error estimate, but it requires evolving an extra level of data which may be very slow. An alternative is the refinement criteria used by the FLASH code of Fryxell et al. (2000) which is adapted from finite element ideas of Löhner (1987). The error estimate depends on the local gradients of the data,

$$\mathbf{E}_i = \frac{|\mathbf{q}_{i+1} - 2\mathbf{q}_i + \mathbf{q}_{i-1}|}{|\mathbf{q}_{i+1} - \mathbf{q}_i| + |\mathbf{q}_i - \mathbf{q}_{i-1}| + \epsilon\left[|\mathbf{q}_{i+1}| - 2\,|\mathbf{q}_i| + |\mathbf{q}_{i-1}|\right]}, \tag{4.10}$$

where ϵ is a small constant (typically 10^{-4}) which allows the last term to act as a filter to avoid the refinement of small ripples. This error estimator refines steep gradients and extrema without needing an extra evolution step.

In practice we have found that Richardson extrapolation is the most efficient method. The error estimate (4.10) based on the local gradients tended to refine too much, considerably slowing the calculations. All results shown later are using Richardson extrapolation.

4.4.5 Boundary conditions and flux conservation

Conservative numerical methods work because the flux through a cell boundary $\mathbf{f}_{i+1/2}$ is used to update both adjoining cells I_i and I_{i+1}. At the boundary between a fine and a coarse grid this will no longer be the case. Because the fine grid takes more steps to update to the same time, boundary conditions must be interpolated in time and the boundary fluxes (averaged over the coarse time step) may not be the same.

In order to restore conservation, Berger and Oliger (1984) suggest *refluxing* the coarse grid at every step. All coarse cells which abut a fine grid are recalculated using a flux that is equivalent to the flux leaving the fine grid.

This retains conservation and is assumed to be more accurate (as the flux is calculated on the finer grid). Because of the local nature of hyperbolic equations only these boundary cells need updating. The refluxing idea has been extended by, for example, Almgren et al. (1998) to cover non-local elliptic and parabolic equations.

We have not implemented refluxing in our AMR code. This is for the following reasons:

- Conservation violations are at the order of the numerical method. We use high order methods, so any violations are small compared to errors elsewhere.

- As we are only working in one dimension we can easily increase the buffer zone for our child grids, so that the boundaries are in smooth regions. In these regions any errors will be very small.

- The refluxing algorithm is not simple to implement in our numerical code, and we have not found any errors that appear to be due to failure of conservation.

- The equations that we evolve typically have large source terms, so conservation will not hold anyway.

- Very high resolution runs performed without AMR show that the results of our problem, given in chapter 5, do not require the refluxing.

In higher dimensions where the buffer zone will have to be smaller, or with different forms of the equations these may not be sufficient reasons to avoid implementing refluxing.

4.4.6 Interpolation onto child grids

To create the data on a child grid we will usually have to interpolate from the parent. We attempt to minimize interpolation by copying data from old fine grids. This is to reduce the amount of high frequency "noise" introduced into the calculation. Here we discuss how to interpolate data without introducing spurious oscillations and retaining conservation.

In essence interpolation is identical to the reconstruction problem of section 3.4. There we are reconstructing the function $\mathbf{q}$ over a computational cell so that we can find the value of the function at a single point, usually the boundary. Here we wish to know the integral average of the function $\mathbf{q}$ over some part of the cell.

To retain conservation we are actually reconstructing the *primitive function* $\mathbf{Q}$ where

$$\frac{\partial}{\partial x}\mathbf{Q} = \mathbf{q}. \tag{4.11}$$

It follows that the integral average of $\mathbf{q}$ on the cell $I_{AB} = [x_A, x_B]$ is given by

$$\mathbf{q}_{AB} = \frac{\mathbf{Q}_B - \mathbf{Q}_B}{x_B - x_A}. \tag{4.12}$$

From the definition (4.11) we can easily construct the primitive function at the cell boundaries by

$$\mathbf{Q}_{i+1/2} = \sum_{j=0}^{i} \Delta x_j \mathbf{q}_j, \tag{4.13}$$

with the initial condition $\mathbf{Q}_{1/2} = 0$. From here we can use any interpolation procedure to calculate values of the primitive function at the boundaries of

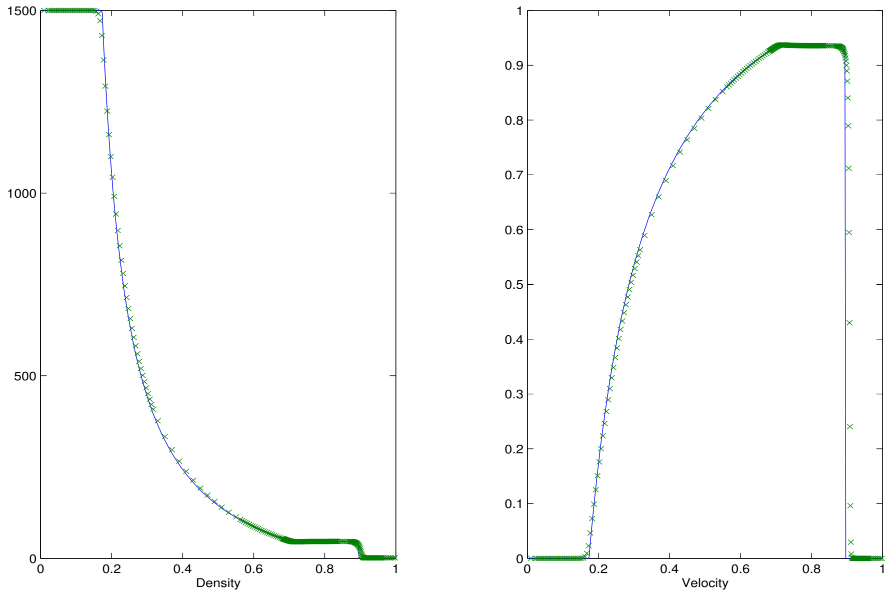

Figure 4.7: The ultrarelativistic blast wave solved using the slope limited Lax–Friedrichs method with 4 levels of Adaptive Mesh Refinement. The capturing of the shock position and the rarefaction fan is comparable to the high resolution methods seen in figures 3.7 and 3.8. However the number of points in the shock on the finer levels is still quite large as expected.

the fine cells and then use equation (4.12) to produce the child grid values. In practice we have been consistent with the numerical method implemented. That is we either use an slope limited TVD method as in section 3.4.2 or an ENO method as in section 3.4.3 depending on the required accuracy.

4.4.7 Example

As a short example we use the one dimensional ultrarelativistic hydrodynamics test used in section 3.10. This coarsest grid has 50 cells and 3 levels of refinement are used. We use the most diffusive method, the slope limited reconstruction with the Lax–Friedrichs solver. The results are given in figure 4.7. These should be compared with the results in figures 3.7 and 3.8. It is clear that the capturing of the shock position is better than in any of the methods, even the fifth order ENO codes, and that the capturing of the rarefaction fan is comparable to the high order codes. Note that the number of points in the shock is comparable to the number in the slope limited method without refinement, as we should expect.

Chapter 5

Critical collapse and primordial black hole formation

This chapter will look at the formation of primordial black holes and the relevance for critical collapse. The phenomenon of critical collapse, where the initial data is "tuned" to produce the smallest possible black hole, could shed light on the possible masses of black holes formed in the early universe. This in turn may be relevant to the study of dark matter.

At the start of this chapter I outline the physical background. This includes the current understanding of critical collapse and a brief outline of the Friedman-Robertson-Walker (FRW) spacetime that we use as our background. The numerical method and grid will be briefly outlined and validated, as much of this material has been covered in chapters 3 and 4. We then outline the general behaviour of a collapse evolution before looking at the behaviour near the black hole threshold.

5.1 Physical background

5.1.1 General Relativity

The outline of relativistic hydrodynamics given in chapters 2 and 3 ignored gravity. To include gravity we must consider general relativity. The essential principle of general relativity is the *equivalence principle*. As given by Will (2001), it states

1. The Weak Equivalence Principle (that the path of a freely falling object is independent of its internal structure and composition) is valid,
2. The outcome of any local non-gravitational experiment is independent of the velocity of the freely falling reference frame in which it is performed,
3. The outcome of any local non-gravitational experiment is independent of where and when in the universe it is performed.

To incorporate the effects of gravity spacetime is assumed to be a curved manifold. Any point in spacetime can locally be given coordinates, and any physical quantity must be independent of these coordinates. A quantity that is independent of the coordinates is called *covariant*. The equivalence principle then ensures that

1. The spacetime manifold has a symmetric metric,
2. The paths of freely falling objects are geodesics of the metric,
3. In locally freely falling frames special relativity holds.

We can use the equivalence principle to simply extend our numerical methods for special relativistic hydrodynamics to general relativity. This was explained in detail by Pons et al. (1998). Briefly, the equivalence principle is used to find a freely falling frame which is local to the cell boundary under consideration. The Riemann problem is solved using the equations of special relativistic hydrodynamics to find the flux through the cell boundary. It should be noted that this cell boundary may not be at rest in the freely falling frame. It is however simple to extend the Riemann solvers of chapter 2 by assuming that the cell boundary moves at a fixed velocity. The inverse coordinate transformation is then used to give the flux in the original coordinate system.

This should not be expected to work when the coordinate transformation is singular as it will be near the origin in spherical symmetry or near an apparent horizon in Schwarzschild coordinates. We can also no longer expect strict conservation to hold, as there will be source terms due to the geometry and also gravity. The conserved quantities only have a physical meaning in local freely falling frames, which is where we solve the Riemann problems.

Any numerical calculation will also fail when a genuine singularity forms. To avoid a singularity forming on the numerical grid the freedom to choose the coordinates is used so that “time” runs more slowly in those parts of the domain where there is large amounts of matter. This typically leads to coordinate singularities when black holes form but allows the calculation to continue. However numerical problems frequently arise due to “slice stretching” where the time coordinate is evolving at such different rates on different parts of the domain. For more details see Lehner (2001).

5.1.2 Critical collapse

The phenomenon of critical collapse was discovered by Choptuik (1992). Following analytic investigations by Christodoulou (1987, 1991, 1993), Choptuik was looking at the collapse of a massless scalar field in otherwise empty space. Describing the initial scalar field as a one-parameter family of asymptotically flat smooth data, he investigated the critical region where black holes are just formed. Writing the parameter as p he discovered the now famous scaling relation

$$M_{BH} \approx C(p - p_*)^\gamma, \tag{5.1}$$

where M_{BH} is the mass of the black hole, p_* the critical value of the parameter p, C a constant depending on the form of the initial data, and γ a constant *independent* of the initial data. He also discovered that the near the critical parameter the solution followed a complicated scale-periodic behaviour with a logarithmic scale period Δ again independent of the initial data. Finally he found that near the critical parameter, the solution at intermediate times appeared to be nearly independent of the initial data. This is referred to as *universal* behaviour. Small deviations from this *critical* solution led to either a black hole or dispersal.

The current understanding of critical collapse is based in part on a dynamical systems picture and is reviewed in detail in Gundlach (1999). In the phase space of pairs of three metrics and extrinsic curvatures, together with any matter variables, there are three fixed points. The black hole solution (typically Schwarzschild as we are restricting to spherical symmetry) and the background solution (typically flat space) are stable. The critical solution is

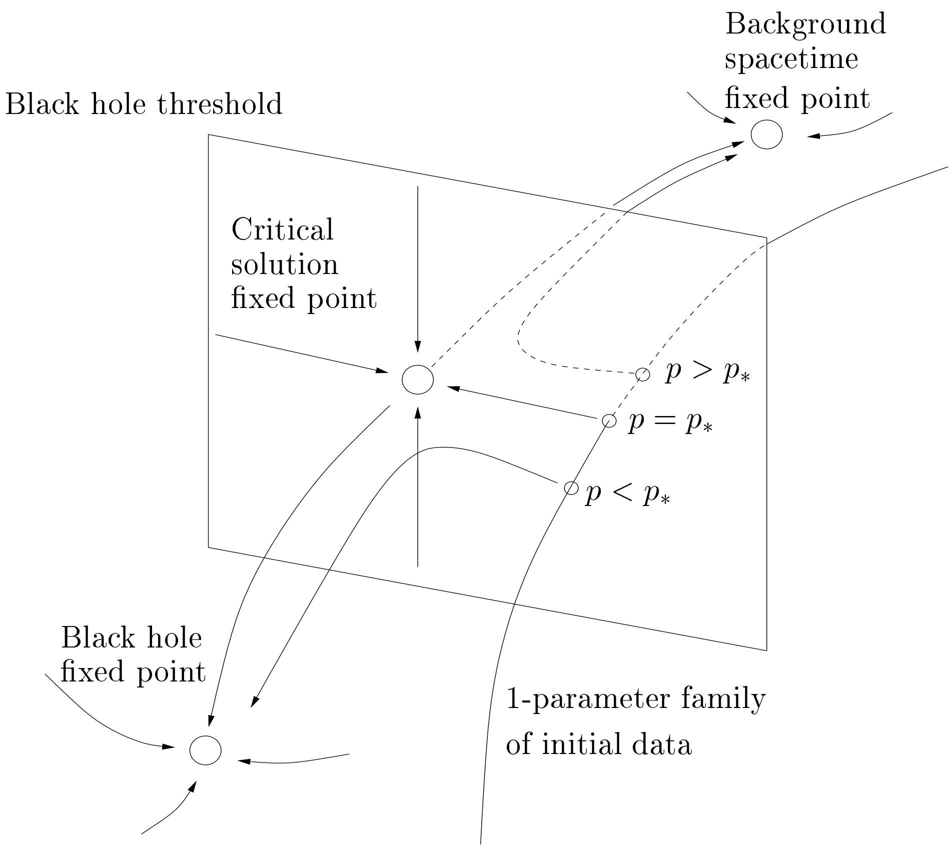

Figure 5.1: The phase space picture for a continuously self-similar solution. The arrow lines are time evolutions, corresponding to spacetimes. The line without an arrow is a 1-parameter family of initial data that crosses the black hole threshold at $p = p_*$. Solutions that start near the threshold spend time "near" the critical solution before either dispersing or forming a black hole. This diagram follows Gundlach (1999).

an unstable fixed point. This is illustrated by figure 5.1. The precise nature of the critical behaviour depends on the critical solution.

The nature of the critical solution determines the type of critical behaviour. If the critical solution is self-similar then we may expect to find the Choptuik mass scaling relation (5.1), which has been labeled type II behaviour. There may also exist scaling behaviour for other physical variables, as discussed later. If the critical solution is static or periodic then the black hole threshold solution may have a finite mass. This is labeled type I behaviour. Here the scaling is only apparent in the length of time an evolution stays "near" the critical solution, t_p, which obeys

$$t_p = -\gamma \log(p - p_*) + \text{constant}. \tag{5.2}$$

Further fine structure including the "echoing" behaviour is due to the type of self-similarity of the critical solution.

For type II collapse, the critical exponent γ can be found analytically (with certain assumptions). This is done by constructing the critical solution $Z_*(x)$, which is assumed to be self-similar in the similarity variable x which is some combination of the time and space coordinates t and r, and looking at the linear perturbations. Writing this perturbation in terms of a time coordinate τ we have

$$\delta Z(x, \tau) = \sum_i C_i e^{\lambda_i \tau} Z_i(x), \tag{5.3}$$

because the unperturbed solution is independent of τ. By definition only one of the complex exponents λ_i has positive real part, and is in fact strictly real. The critical exponent is the reciprocal of this.

The vast majority of investigations of critical phenomena have been restricted to asymptotically flat spherically symmetric models. The numerical exceptions are a study by Abrahams and Evans (1993) of the collapse of axisymmetric gravitational waves in vacuum, work by Pretorius and Choptuik (2000), Husain and Olivier (2001) on the collapse of a massless scalar field in anti de Sitter space, and the studies of collapse in an expanding FRW universe outlined below. Analytically, it has been shown by Gundlach (1998), Martín-García and Gundlach (1999) that small perturbations away from spherical symmetry have no effect on the critical solution or the exponent.

A useful point for determining the critical exponent from collapse simulations is an observation of Garfinkle and Duncan (1998). This argues that on dimensional grounds for type II behaviour any physical curvature invariant such as the maximum of the Ricci scalar $\mathcal{R}$ over the spacetime or on the axis (in those coordinates where this has a physical meaning) should obey

$$|\mathcal{R}| \approx C(p_* - p)^{-2\gamma} \tag{5.4}$$

for subcritical evolutions. It is simpler to measure this accurately than the black hole mass itself because of the numerical problems associated with singularity formation.

5.1.3 FRW spacetimes

The Friedman-Robertson-Walker (FRW) spacetimes are the basis for the standard model of cosmology. They are found from the Einstein equations by assuming that the spacetime is *isotropic*. This requires that it is impossible to construct a geometrically preferred tangent vector orthogonal to u^a, the tangent vector field to the congruence of timelike curves of the spacetime. For a full derivation of the FRW spacetimes we refer to Wald (1984).

The different possible FRW universes can be characterized by a number K describing the geometry of the spaces of constant curvature. If K is positive then the spaces are 3-spheres. If K is zero then the geometry is flat 3-space. If K is negative then the surfaces are three dimensional hyperboloids. We shall consider only the spatially flat $K = 0$ models.

In the $K = 0$ model the line element is usually expressed in the form

$$\mathrm{d}s^2 = -\mathrm{d}\tau^2 + A^2(\tau)\left(\mathrm{d}R^2 + R^2\left(\mathrm{d}\theta^2 + \sin^2\theta\mathrm{d}\phi^2\right)\right). \tag{5.5}$$

Assuming that the matter is a radiation dominated perfect fluid, that is an ultrarelativistic ideal gas with $\Gamma = 4/3$, the metric function a and density ρ have the form

$$A \propto \tau^{1/2}, \qquad \rho \propto \tau^{-2}. \tag{5.6}$$

Another useful coordinate system is given by rescaling the time coordinate so that the line element is

$$\mathrm{d}s^2 = A^2(T)\left(-\mathrm{d}T^2 + \mathrm{d}R^2 + R^2\left(\mathrm{d}\theta^2 + \sin^2\theta\mathrm{d}\phi^2\right)\right). \tag{5.7}$$

One notable physical point is the horizon problem. The assumption of isotropy is reasonable if the spacetime is able to reach thermal equilibrium, that is if the past horizon of any observer includes all comoving observers. This is not the case. In this context the *particle horizon* is the boundary of the set of points that can be causally connected to the point $(\tau = 0, R = 0)$ in the coordinates of the line element (5.5). In the isotropic coordinate system the horizon is simply given by the line $T = R$. All points within the horizon can be assumed to be in thermal equilibrium.

5.1.4 Critical collapse in FRW spacetimes

What type of scaling behaviour should we expect in a FRW spacetime? Previous results, listed in section 5.2 suggest that type II behaviour should be expected. However, those calculations either had problems with shocks or failed to investigate sufficiently small black holes.

If the critical behaviour were to be of type II then we might expect that the critical solution would correspond to that constructed by Evans and Coleman (1994). This would require the critical solution to be matched across a discontinuity to the solution at infinity - this sort of matching has been performed analytically by Cahill and Taub (1971). However, a number of other possible self similar solutions have been constructed by Carr et al. (2000), Carr and Coley (2000), with different asymptotic behaviour. So it is not clear what self similar solution (if any) we should expect.

If the critical behaviour were to be of type I then we might expect the critical solution to be one of the unstable static Tolman-Oppenheimer-Volkoff solutions, as suggested by Neilsen (1999). Again these would require some sort of matching to the solution at infinity. Although the previous results, and the self similar solutions of Evans and Coleman (1994) would suggest type II behaviour, there is an intrinsic length scale (the Hubble horizon) in an FRW spacetime. This may suggest, in analogy with the results in Brady et al. (1997), Choptuik et al. (1996), Olabarrieta and Choptuik (2001) that the critical behaviour should be type I.

5.2 Previous calculations

The formation of primordial black holes due to density fluctuations in a radiation dominated FRW universe has been studied both analytically and numerically. The early analytical calculations of Carr (1975), Carr and Hawking (1974) predicted that black holes could be formed with a range of masses and that these masses would not increase appreciably by accretion. The density perturbations were parametrized by δ, the extra gravitational mass within the Hubble horizon, and the critical value for black hole formation was given as lying approximately in the region $1/3 \leqslant \delta_c \leqslant 1$. Although the work considered general fluid equations of state, the numerical values for δ_c hold only for the radiation dominated ultrarelativistic perfect fluid, i.e., $p = \frac{1}{3}\rho$. We also note that the upper bound on the parameter, $\delta \leq 1$ is due to an approximation by Harrison (1970) that in the perturbed spacetime all spatial derivatives can be neglected. This allows the perturbations to be treated as sections of a Friedman model. As we shall see later this is an unreasonable assumption.

The parameter δ allows for the simple calculation of the black hole mass spectrum in approximate analytical calculations. It can be made covariant following Ellis and Bruni (1989). The fractional density gradient χ_a is defined as

$$\chi_a = \frac{h_b^a}{\rho} \partial_b \rho \tag{5.8}$$

where $h_{ab} = g_{ab} + u_a u_b$ projects into the spaces orthogonal to the fluid velocity. Then the density contrast is given by integrating over the boundary of the region V containing the perturbation,

$$\delta = \int_{\Sigma} |\chi_a| \mathrm{d}V^a. \tag{5.9}$$

In the FRW spacetime the natural surface over which to integrate is the horizon.

The earliest numerical calculations of which we are aware are Bicknell and Henriksen (1979), Nadezhin et al. (1978). The numerical methods used broke down when shocks occurred. More recently the problem has been used to test forms of the equations intended for use in higher dimensions, as in Shibata and Sasaki (1999). The studies of accretion seem to agree that it is largely unimportant for primordial black holes. However few details were given of the numerical method, or of the simulations near the critical parameter.

Interest in the subject was reawakened after the paper of Niemeyer and Jedamzik (1998). This suggested that if the critical behaviour observed by Choptuik was found in an FRW spacetime, then black holes of very small mass could be found naturally. This is due to the exponential decline of the probability function for primordial density fluctuations, allowing a natural "tuning" to the critical parameter. As outlined in a later paper by Niemeyer and Jedamzik (1999a) such small black holes may now be observable due to advances in observational techniques. The existence of such black holes would give important information about the early universe, as it is well known that black holes formed by stellar collapse have a relatively large minimum mass.

This idea was tested numerically for the simple radiation dominated equation of state by Niemeyer and Jedamzik (1999a) and also for a more complex matter model in Niemeyer and Jedamzik (1999b). They used a numerical code of Baumgarte et al. (1995) which uses a coordinate system with one null coordinate and one radial comoving coordinate. The code is second order accurate (except for the lapse equation) and uses artificial viscosity. The power law (5.1) was confirmed with similar critical exponent as $\gamma \approx 0.36$ found for the calculations in flat space (see Evans and Coleman (1994), Maison (1996), Koike et al. (1995)). They also found that the critical parameter $\delta_c \approx 0.7$ to be independent of the initial data. However the black hole masses found vary over only less than one order of magnitude and the smallest mass found was greater than 10% of M_h, the mass within the Hubble horizon of the unperturbed spacetime.

We attempted to repeat these calculations using a comoving form of the Einstein equations given by Friedrich (1998). This reduces the Einstein equations to a set of symmetric hyperbolic partial differential equations. However the variables evolved include the spatial derivatives of various metric functions which, in comoving coordinates, is equivalent to evolving the spatial derivatives of the density. These are not defined, even in the weak sense, when shock waves are present. When such discontinuities appear they are coordinate singularities, so problems occurred when they reflected from the origin, itself a coordinate singularity. We can avoid this problem using the same methods as Niemeyer and Jedamzik (1999a), which is the introduction of artificial viscosity. However, this is untrustworthy in relativistic calculations as shown by Norman and Winkler (1986). We found that it was possible to construct a black hole where none had existed before simply by increasing the artificial viscosity.

Other attempts were based on non-comoving coordinates. We found that attempts using an ingoing null coordinate (or some adaptive version that is null only near the axis) ran into numerical problems as soon as a perturbation was introduced. We believe that this can be explained from the equivalence principle applied to the Riemann problem following Pons et al. (1998). Where the slice is null, the boundary of the computational cell is moving at the speed of light. This means that any method is effectively first order accurate as the flux at the boundary is always evaluated at the same side. It also means that material will not be able to reflect from the origin, as the boundary conditions cannot be taken into account. With these problems we found a purely spacelike slicing, described below, was the best option.

5.3 Coordinates, variables and equations

After some experiments we settled on the “Schwarzschild-like” chart used by Neilsen and Choptuik (2000), with line element

$$ds^2 = -\alpha(t,r)^2\,dt^2 + a(t,r)^2\,dr^2 + r^2(d\theta^2 + \sin^2\theta\,d\phi^2). \tag{5.10}$$

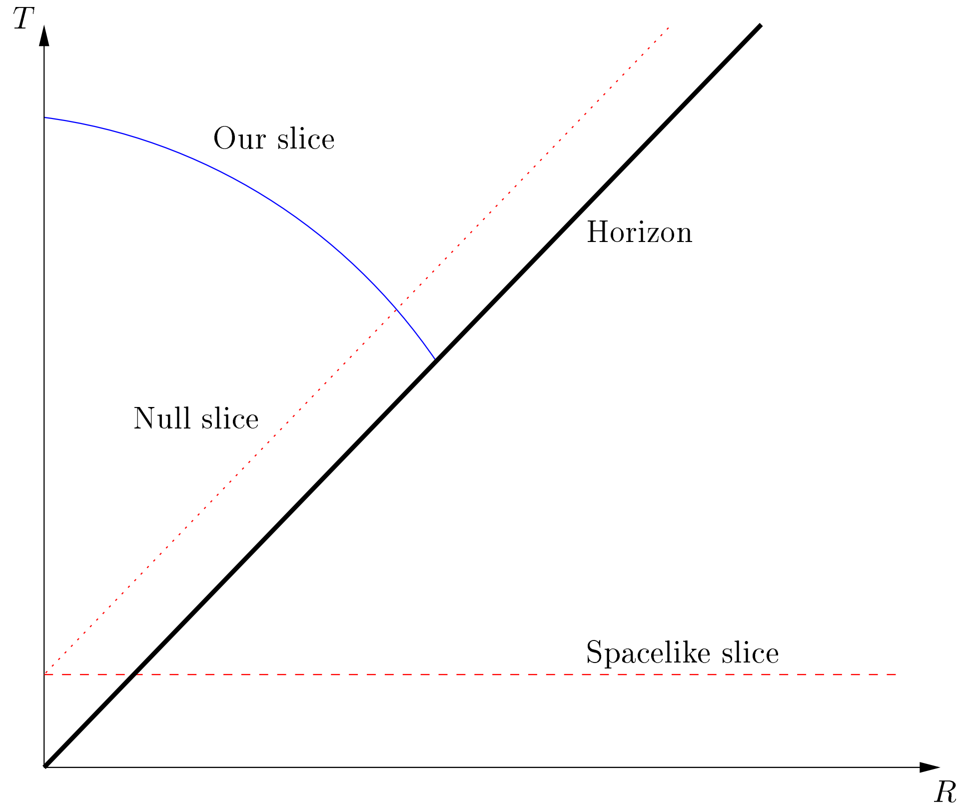

Figure 5.2: Different charts lead to different slicings of the FRW spacetime. The axes assume isotropic coordinates (5.7). From these the standard comoving form is just a rescaling of the time and so the slicing is the same, but the horizon (thick line) behaves differently. The two slicings used by Niemeyer and Jedamzik (1999a) are given by the red dashed line, for the spacelike slice used to construct the initial data, and the red dotted line for the null slice on which the evolution was performed. Our slice is given by the blue line. This becomes null at the horizon causing various metric quantities to diverge.

It should be emphasized that t and r are not the usual Friedman-Robertson-Walker coordinates introduced in section 5.1.3. This is clarified by looking at the slicings in the different coordinates, shown in figure 5.2. If we take $A(\tau) = \tau^{1/2}$ (see (5.6)) then the coordinates are related by

$$\tau = \tfrac{1}{2}(t + \sqrt{t^2 - r^2}), \qquad R = \sqrt{2\left(t - \sqrt{t^2 - r^2}\right)}. \tag{5.11}$$

The restriction $r \leqslant t$ means that our chart only covers the region within the horizon, but that is sufficient for our purposes.

We are considering a radiation fluid with energy momentum tensor

$$T^{ab} = (\rho + p)u^a u^b + p g^{ab}, \tag{5.12}$$

and equation of state $p = \frac{1}{3}\rho$. Here ρ is the energy density, p the pressure and u^a is the fluid 4-velocity normalized via $u_a u^a = 1$. The fluid coordinate velocity and associated Lorentz gamma factor W are

$$v(t, r) = \frac{a u^r}{\alpha u^t}, \qquad W(t, r) = \alpha u^t. \tag{5.13}$$

With these definitions the equations of motion of the fluid take a form similar to the ultrarelativistic fluid of section 2.4. Using the conserved variables of equation (2.35) and the fluxes of equation (2.36) the equations of motion of the fluid are

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{1}{r^2}\frac{\partial}{\partial r}\left(\frac{\alpha}{a} r^2 \mathbf{f}\right) = \mathbf{s}. \tag{5.14}$$

As before we shall refer to the conserved variables $\mathbf{q}$, the flux $\mathbf{f}$ and the source terms $\mathbf{s}$. We can convert from the conservative to primitive variables using the methods of section 3.9.1.

Following Neilsen and Choptuik (2000) we now redefine the conservative variables as

$$\mathbf{q} = \begin{pmatrix} \Pi \\ \Phi \end{pmatrix} = \tfrac{1}{2} \begin{pmatrix} \tau + S \\ \tau - S \end{pmatrix}. \tag{5.15}$$

To retain causality these conservative variables must both be strictly positive. Enforcing this condition appears to have no physical effect on the evolution and ensures that physical primitive variables can always be recovered. In near vacuum regions the original conservative variables (τ, S) become sufficiently close that numerical error may produce an unphysical state.

In terms of these new conserved variables the equations of motion (5.14) become

$$\frac{\partial}{\partial t} \begin{pmatrix} \Pi \\ \Phi \end{pmatrix} + \frac{1}{r^2}\frac{\partial}{\partial r}\left\{ \frac{\alpha}{a} r^2 \begin{pmatrix} \frac{1}{2}(\Pi - \Phi)(1+v) + p \\ \frac{1}{2}(\Pi - \Phi)(1-v) - p \end{pmatrix} \right\} = \begin{pmatrix} \Sigma \\ -\Sigma \end{pmatrix}, \tag{5.16}$$

where

$$\Sigma = -\tfrac{1}{2}\Big[\Pi(1-v) + \Phi(1+v)\Big]\alpha a\left(8\pi r p + \frac{m}{r^2}\right) + \alpha a p \frac{m}{r^2} + \frac{2\alpha p}{a r}, \tag{5.17}$$

and

$$m(t, r) = \frac{r}{2}\left(1 - \frac{1}{a^2}\right) \tag{5.18}$$

is the *mass aspect* function. Equations governing the metric variables were given in Neilsen and Choptuik (2000). The polar slicing condition, the condition that the line element retains the form (5.10) for all t implies

$$\frac{\partial}{\partial r}\log\alpha = a^2\left(2\pi r\Big[(\Pi-\Phi)v+2p\Big]+\frac{m}{r^2}\right). \tag{5.19}$$

The momentum constraint produces an evolution equation for a,

$$\frac{\partial a}{\partial t} = -2\pi r\alpha a^2(\Pi-\Phi), \tag{5.20}$$

which we shall use in the evolution, and the energy constraint produces another equation for a

$$\frac{\partial a}{\partial r} = a^3\left(2\pi r(\Pi+\Phi)-\frac{m}{r^2}\right), \tag{5.21}$$

which we use initially, at the outer boundary and to check the consistency and accuracy of the numerical results.

5.4 Initial data

Next we need to consider the initial conditions. Equations (5.16) permit arbitrary initial values for the conservative variables or equivalently for the primitive ones. If we define $d=\sqrt{t^2-r^2}$ then for a FRW spacetime

$$\rho_{FRW} = \frac{3}{8\pi G(t+d)^2}, \quad v_{FRW} = \frac{r}{t+d}, \quad a_{FRW} = \alpha_{FRW} = \sqrt{\tfrac{1}{2}(1+t/d)}. \tag{5.22}$$

We choose initially $v(t_0, r) = v_{FRW}(t_0, r)$ for $0 \leqslant r \leqslant r_0$ where $r_0 < t_0$ is the initial right hand end of the grid. We also choose

$$\rho(t, r) = \rho_{FRW}(t, r)\Big[1 + Cf(r)\Big]. \tag{5.23}$$

Here $f(r)$ is a *density profile* function. It has a single maximum and $f(r) \to 0$ exponentially fast as $r \to r_0$. The parameter C is freely specifiable and governs the strength of the density perturbation. In our coordinates it is not simple to relate the parameter C to the density contrast indicator δ used in Niemeyer and Jedamzik (1999a). This is addressed in more detail in section 5.9.3.

Three different profiles were examined. The first two were based on a Gaussian form for f,

$$f(r) = e^{-w(r-r_c)^2}, \tag{5.24}$$

where w is a constant (here chosen to be 100) and r_c the centre of the Gaussian. The two profiles used were a centred ($r_c = 0$) and off-centre ($r_c = 0.25$) Gaussian. The third profile was a Mexican Hat type,

$$f(r) = (1 - \tfrac{1}{2}wr^2)e^{-wr^2}, \tag{5.25}$$

where w is a constant. The total gravitational mass on a surface of constant t is unchanged using the Mexican Hat profile.

The off-centre Gaussian is not analytic at the origin as the radial derivative of the pressure profile does not vanish. However numerically this does not matter. As explained by Laney (1998) (chapter 19) the use of ghost cells with the boundary at the origin combined with the assumption that the derivative vanishes will reconstruct a function that is analytic at the boundary and a second order accurate approximation to the true function. In other words, for the off-centre Gaussian profile the initial data specified for the density is no longer exact, but only second order accurate in a way that is consistent with the boundary conditions to be imposed. This would not the case if a computational cell (interior or ghost) were centred on the origin.

5.5 Boundary conditions

If the spacetime is to be non-singular at $r = 0$ then $a(t, 0) = 1$. This furnishes the single boundary condition required by equation (5.21) and which can be integrated to furnish $a(t, r)$ for $0 \leqslant r \leqslant r_0$. We choose $\alpha(t, r_0)$ to take its FRW value and this provides the single boundary condition needed to integrate equation (5.19).

We note that although the profiles of our initial data are similar to those of Niemeyer and Jedamzik (1999a), we may not necessarily get the same results. This is because the radial extent of our initial data is completely contained within the Hubble horizon because of the limitations of our coordinate chart. In the comoving coordinates used in Niemeyer and Jedamzik (1999a) the behaviour is regular at the horizon and so the radial extent of the initial perturbation is larger than the horizon. In analogy with the results of Brady et al. (1997), where different critical behaviour was seen when the "size" of the perturbation was larger or smaller than the natural length scale of the spacetime, we may see different critical behaviour in the different cases.

Next consider the boundary conditions required at $r = 0$. On physical grounds we expect ρ to be an even function of r and v to be an odd function of r. Near $r = 0$ this implies

$$\Pi(t, -r) = \Phi(t, r), \qquad \Phi(t, -r) = \Pi(t, r), \tag{5.26}$$

as outlined in section 3.8. On physical grounds α and a should be even functions of r. Near $r = 0$ this implies

$$a(t, -r) = a(t, r), \qquad \alpha(t, -r) = \alpha(t, r). \tag{5.27}$$

Equations (5.26) and (5.27) will be used to generate the boundary values needed at $r = 0$ for the evolution.

The outer boundary is more problematic. Physically we expect a possible outflow of matter into an essentially FRW region, but no inflow except that due to the FRW spacetime itself. The exact FRW background cannot be used because eventually some additional matter from the perturbation will reach the outer boundary. Attempts to set up outflow boundary conditions to model this invariably led to numerical instabilities. Many different extrapolations on either primitive, conservative or characteristic variables were attempted. The best results were obtained with rational polynomial extrapolation which is naturally suggested by the form of the background solution given by equation (5.22), but even for this evolutions of the pure

FRW spacetime developed large errors too soon to be of any use in practice.

We sidestepped this problem in the following manner. After each step the ghost cells at the outer boundary were computed using the solution for the exact FRW spacetime. Equation (5.21) was then used on the boundary cells to ensure that that metric variable a was correct. Then if the coordinate time t had advanced sufficiently that we were no longer covering the same proportion of the slice (i.e., $r_0/t_0 < r_{\max}/t$) then an extra computational cell was added to the outer boundary. This grid expansion produces few errors and those that appear are typically advected outwards due to the underlying expansion of the spacetime. These errors might be coupled to the interior by changing the lapse α which must be solved at every step using the slicing condition (5.19). This should just renormalize the time coordinate and have no effect on the physics. However, this may stop the slicing avoiding the singularity. Alternatively, very large errors near the original grid could cause large changes in the gradients of α, possibly changing the characteristic speeds in such a way that the interior would be causally connected with the expanding part of the grid. In practice the expansion is sufficiently accurate that this does not happen, as seen in section 5.8.

5.6 The numerical grid

As mentioned in section 5.5 we bypass the problems with the outer boundary and the particle horizon by placing the outer boundary of the numerical grid inside the horizon, but expanding the grid so that for all coordinate time it covers the same amount of the space. In order to avoid having to add a large number of points unnecessarily we therefore want the grid spacing to be as large as possible at the outer boundary.

However, in order to find the mass of a black hole we need to know the position of the apparent horizon, r_{AH}. As we wish to be able to resolve very small black holes, we need the grid spacing near the origin to be similarly small. As the apparent horizon is given by the metric function a diverging to infinity, we found it necessary that there be a large number of points (at least 100) between the origin and the apparent horizon for stability.

This would seem like an ideal situation to use Adaptive Mesh Refinement, outlined in section 4.4. AMR is a good method to use for hyperbolic systems because any errors produced locally at a point will be physically propagated at a finite speed. This means that by increasing the resolution in some neighbourhood of this point, the error will definitely decrease. The extension of AMR to systems that are not in hyperbolic form is not an easy task, as seen in the implementation of Almgren et al. (1998). When elliptic equations are coupled to the hyperbolic system the effect of any error is immediately propagated over the entire grid. This will affect the boundaries of all child grids, and so the addition of resolution is not guaranteed to decrease the error.

In the studies of critical collapse that used AMR such as Choptuik (1992), Hamadé and Stewart (1996), Hamadé et al. (1996), the system of equations were not in hyperbolic form. Despite this, AMR worked perfectly. This might be because the coupling between the hyperbolic equations and (in these cases) the ordinary differential equations was sufficiently weak that the effects of the errors on the boundaries of child grids was negligible. Other studies of critical collapse such as Neilsen and Choptuik (2000) did not use AMR.

We attempted to use AMR to study the critical collapse of a perfect fluid on an FRW background. It worked very well for black holes of mass greater than about $10^{-2}M_h$. However for the very small black holes that we need to study in critical collapse the code failed. We believe that this is due to the coupling to the slicing condition (5.19) combined with the finite volume formulation.

Consider a small black hole where the apparent horizon is resolved in very few cells (around 10 is normal for the cases where AMR fails) on the coarsest grid. As the plots in section 5.9.1 show we expect the black hole to be indicated by a "spike" in the velocity that approaches $-c$, a "spike" in $2m/r$ that approaches 1, and a compact region of very high density. When this is poorly resolved the finite volume formulation ensures that the minimum cell averaged value of the velocity will be truncated away from $-c$, with the spike in $2m/r$ similarly truncated. However as the density is an even function this will still have a high value at the origin. This is shown in figure 5.3. On the coarse grid the pulse disperses because neither self-gravity nor hydrodynamical forces can maintain the compressed state of this region.

If the system were purely hyperbolic this would not matter as the error would be corrected on the fine grids. However the lapse changes by orders of magnitude when the fluid disperses. This causes the form of the lapse to be inaccurate over the entire grid. These inaccuracies are fed back into the fine grids through the boundary conditions which are interpolated from the parent. These errors will accumulated eventually causing the evolution to

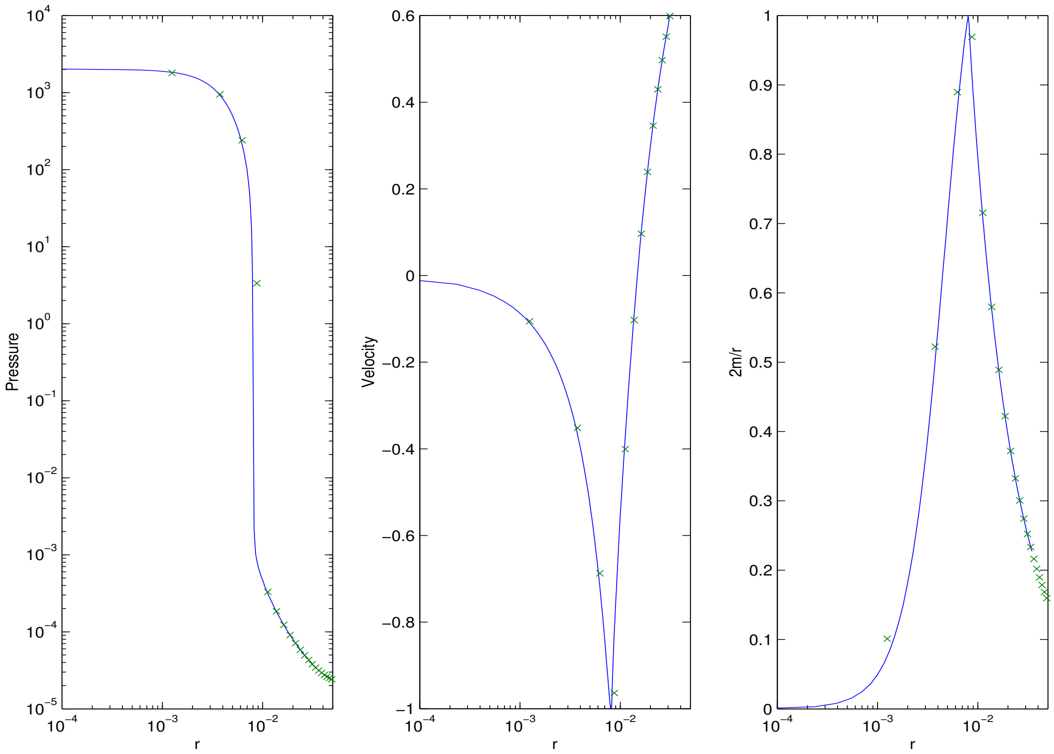

Figure 5.3: The pressure, velocity and $2m/r$ for a black hole evolution using AMR. The functions given by the data on fine grids are shown as the line, and on the coarsest grid by the points. The black hole is very poorly resolved on the coarse grid, with only 3-4 points between the origin and the horizon, leading to the truncation of the spikes in the velocity and $2m/r$. This is consistent with the finite volume formulation. However, the pressure is not truncated. On the coarse grid this pulse will disperse. This leads to large errors in the evaluation of the lapse which affect the boundaries of the fine grids. These errors accumulate and cause the run to fail.

fail, usually when the slicing on the finer grid fails to be singularity avoiding because of the incorrect normalization at the outer boundary.

To find the masses of very small black holes we have used a method similar to the work of Neilsen and Choptuik (2000). The grid is unevenly spaced in r, with a very fine spacing near the origin and coarse spacing at the outer boundary. Two different types of grid spacing were implemented. One is similar to that used by Neilsen and Choptuik (2000) and is shown in figure 4.2. The grid spacing is constant near the origin and at the outer boundary, and in between varies smoothly in $\log(r)$. The alternative, implemented by John Stewart, is the grid spacing being piecewise constant. The former implementation avoids any problems due to sudden changes in grid spacing. In the latter implementation it is simpler to interpolate the data onto the new grid accurately. Very similar results were found with both methods.

We add the grid spacing whenever the number of grid cells between the origin and any possible apparent horizon falls below a preassigned value (typically 300). We monitor the position of any possible apparent horizon by looking for a sharp positive peak in a. The number of cells added is typically of the order of 100. The new grid is filled using conservative interpolation. We have to balance the problems of adding too many points, which makes the code run too slowly, and too few, which forces it to do more interpolation which can introduce noise.

5.7 The evolution algorithm

The outline plan of the evolution algorithm that we use closely follows Neilsen and Choptuik (2000). The equations for the fluid variables are discretized using the method of lines. The metric factor $r^2\alpha/a$ in the flux term is ignored in the calculation of the numerical flux, and only introduced in the update step. This may seem inaccurate, but it is precisely what is required if we follow the approach of Pons et al. (1998) briefly outlined in section 5.1.1. This approach was also used by Romero et al. (1996). This uses the equivalence principle to transform the coordinates at a cell boundary so that they are freely falling and in slab symmetry. This allows any of the Riemann solvers given in chapter 2 to be used. This will be inaccurate at the origin in spherical symmetry as the coordinate transformation is singular there. However, with careful finite differencing the errors are of the order of the truncation error of the method.

Using the method of lines we can also evolve a. It is also possible to constraint a using equation (5.21). We preferred to evolve a, and all the results shown here use the evolution. This is because it is simpler to check the accuracy of the evolution using the constraint (5.21), it allows for longer runs and is marginally more computationally efficient. The increase in the runs is caused by the constrained evolutions failing, that is producing nonsensical values, when the gradient of a is very steep as happens at an apparent horizon. Such evolutions ran for sufficient time to determine the location of the apparent horizon but not long enough to study accretion. The difference between a constrained and a free evolution before this point is negligible.

All that is then required is to determine the lapse α, for which we must used the slicing equation. This is normalized to the exact FRW value at the outer boundary.

Before giving the full details of the evolution code, we note a small point mentioned in Neilsen and Choptuik (2000). Analytically, the final term in the fluxes in the system (5.16) should partially cancel one of the source terms. However, as this involves a factor of r^{-1} this is unlikely to happen numerically. Hence it is better to explicitly write the finite differencing in such a way that this part of the source term vanishes. We do this by splitting the flux into two parts,

$$\mathbf{f}^{(1)} = \tfrac{1}{2}\begin{pmatrix}(\Pi-\Phi)(1+v)\\ (\Pi-\Phi)(1-v)\end{pmatrix}, \qquad \mathbf{f}^{(2)} = \begin{pmatrix}p\\ -p\end{pmatrix}, \tag{5.28}$$

in terms of which the system (5.16) becomes

$$\frac{\partial}{\partial t}\begin{pmatrix}\Pi\\ \Phi\end{pmatrix} + \frac{1}{r^2}\frac{\partial}{\partial r}\left(\frac{\alpha}{a}r^2\mathbf{f}^{(1)}\right) + \frac{\partial}{\partial r}\left(\frac{\alpha}{a}\mathbf{f}^{(2)}\right) = \begin{pmatrix}\bar{\Sigma}\\ -\bar{\Sigma}\end{pmatrix}, \tag{5.29}$$

where the new source term $\bar{\Sigma}$ is given by

$$\bar{\Sigma} = -\tfrac{1}{2}\Big[\Pi(1-v)+\Phi(1+v)\Big]\alpha a\left(8\pi r p + \frac{m}{r^2}\right) + \alpha a p\frac{m}{r^2}. \tag{5.30}$$

However, we must make sure that the eigenvalues and eigenvectors are calculated using the original fluxes. These are simply found from the results for ultrarelativistic hydrodynamics in section 3.9.2. In full, the eigenvalues remain

$$\lambda_\pm = \frac{v \pm c_s}{1 \pm v c_s}. \tag{5.31}$$

The corresponding right eigenvectors are

$$\mathbf{r}_\pm = \begin{pmatrix} 1 + \lambda_\pm \\ 1 - \lambda_\pm \end{pmatrix} \tag{5.32}$$

and the corresponding left eigenvectors are

$$\mathbf{l}_\pm = \tfrac{1}{2} \begin{pmatrix} \frac{1-\lambda_\mp}{\lambda_\pm - \lambda_\mp} \\ \frac{1+\lambda_\mp}{\lambda_\mp - \lambda_\pm} \end{pmatrix}. \tag{5.33}$$

The full evolution algorithm takes the fluid variables $\mathbf{q}_i^n$ and metric variables a_i^n and α_i^n, together with a stable timestep Δt. For each step of a third order Runge-Kutta evolution the algorithm implemented is as follows.

- Reconstruct the fluid variables using a third order ENO method to find $\mathbf{q}^\pm_{i+1/2}$.
- Solve the Riemann problem to find the intercell fluxes $\mathbf{f}^{(1)}_{i+1/2}$ and $\mathbf{f}^{(2)}_{i+1/2}$.
- Evolve the fluid variables $\mathbf{q}$ using the finite differencing

$$\mathbf{q}_i^* = \mathbf{q}_i + \Delta t \left\{ \frac{3\left[\left(r^2 X \mathbf{f}^{(1)}\right)_{i-1/2} - \left(r^2 X \mathbf{f}^{(1)}\right)_{i+1/2}\right]}{r^3_{i+1/2} - r^3_{i-1/2}} + \frac{\left[\left(X \mathbf{f}^{(2)}\right)_{i-1/2} - \left(X \mathbf{f}^{(2)}\right)_{i+1/2}\right]}{r_{i+1/2} - r_{i-1/2}} + \mathbf{\Sigma}_i \right\}. \tag{5.34}$$

The metric quantity is evaluated using a simple arithmetic average,

$X_{i+1/2} = (\alpha_{i+1} + \alpha_i)/(a_{i+1} + a_i)$.

- Use the evolution equation to find a_i^*. This is incorporated in to the method of lines step.

- Use the slicing condition to find α_i^*. A third order Adams–Moulton method is used.

- Enforce the boundary conditions.

The complex finite differencing is due to considering the evolution equations (5.16) in terms of the radial coordinate r^3. For an even grid there is no difference between this approach and the simpler approach used by Romero et al. (1996). However on an uneven grid this finite differencing is considerably more stable near the origin.

This algorithm is only second order accurate. This is due to the way the source terms are evaluated, and the simple arithmetic average used to evaluate the metric quantity X at the cell boundary. It is possible to construct a scheme that is fully third order accurate. One possible algorithm that was implemented is as follows.

- Assume that we are given the fluid variables $\mathbf{q}_i$ and the geometric variables defined at the cell boundaries $\alpha_{i+1/2}$ and $a_{i+1/2}$.

- Reconstruct the fluid variables using a third order ENO method to find $\mathbf{q}^{\pm}_{i+1/2}$ and the point value $\mathbf{q}_i$.

- Solve the Riemann problem to find the cell boundary state $\mathbf{q}_{i+1/2}$.

- Using the cell boundary states, evaluate the numerical fluxes $\mathbf{f}^{(1)}_{i+1/2}$ and $\mathbf{f}^{(2)}_{i+1/2}$.

- Using the cell boundary states and the point value $\mathbf{q}_i$, evaluate the source terms at three points for every cell, $\mathbf{\Sigma}_{i\pm1/2}$, $\mathbf{\Sigma}_i$. Use numerical quadrature to construct the average over the cell $\bar{\mathbf{\Sigma}}_i$.

- Evolve the fluid variables $\mathbf{q}$ using the finite differencing

$$\mathbf{q}^*_i = \mathbf{q}_i + \Delta t \left\{ \frac{3\left[\left(r^2 X \mathbf{f}^{(1)}\right)_{i-1/2} - \left(r^2 X \mathbf{f}^{(1)}\right)_{i+1/2}\right]}{r^3_{i+1/2} - r^3_{i-1/2}} + \frac{\left[\left(X \mathbf{f}^{(2)}\right)_{i-1/2} - \left(X \mathbf{f}^{(2)}\right)_{i+1/2}\right]}{r_{i+1/2} - r_{i-1/2}} + \bar{\mathbf{\Sigma}}_i \right\}. \tag{5.35}$$

  The metric quantity is exact, $X_{i+1/2} = \alpha_{i+1/2}/a_{i+1/2}$.

- Use the cell boundary values to find the source terms to evolve $a_{i+1/2}$. This is incorporated in to the method of lines step.

- Use the cell boundary values and the slicing condition to find $\alpha^*_{i+1/2}$. A third order Adams–Moulton method is used.

- Enforce the boundary conditions.

We note that it is necessary to reconstruct the geometric terms at the cell centres in order to find the source terms for the fluid variables. An average of the possible third order ENO reconstructions was used which should be third order accurate in itself.

In practice the first algorithm was used rather than the second. The reasons for this were as follows.

- Computational expense.

- No noticeable differences between the results.

- The third order method was more likely to produce "spikes" at discontinuities. These are individual cells where the conserved variables have to be reset to be positive. Although this allows the evolution to continue, it appears that when this condition is frequently imposed in a small region the evolution may produce unphysical results near the shock or even fail.

- The most stable results with the first algorithm were found using the Marquina flux. This flux cannot be used with the third order algorithm as it does not construct the cell boundary state.

Of the methods listed in chapter 3 all except SLIC were implemented. The choice of the third order ENO reconstruction with the Marquina solver is largely in line with the reasons given in the conclusion to that chapter, section 3.10.4. For this specific problem some points are worth noting. Firstly the use of the Lax-Friedrichs flux produced oscillations near the origin even in the pure FRW case, making it useless. Secondly, oscillations would occasionally arise if a sonic point in the flow was insufficiently well resolved. A profile where the fluid velocity changed from zero to near the speed of light in about 15 cells would be an example. Higher order methods were less likely to suffer this flaw, as was the Marquina solver. The only true cure is additional resolution, however for most of the runs we used the "best" method. Selected runs were compared against the Roe and exact solvers and those performed

Method	Cells	L_1 error	L_1 order	L_2 error	L_2 order	L_∞ error	L_∞ order
AM3	50	9.17e-6	–	1.50e-6	–	6.22e-7	–
	100	1.48e-6	2.63	1.74e-7	3.11	6.97e-8	3.16
	200	2.15e-7	2.78	1.87e-8	3.22	7.95e-9	3.13
	400	2.91e-8	2.89	1.98e-9	3.24	9.33e-10	3.09
AB3	50	1.08e-4	–	1.59e-5	–	4.14e-6	–
	100	1.69e-5	2.68	1.76e-6	3.18	4.07e-7	3.35
	200	2.40e-6	2.82	1.77e-7	3.31	4.09e-8	3.31
	400	3.21e-7	2.90	1.71e-8	3.37	4.38e-9	3.22
MS2	50	5.46e-5	–	9.47e-6	–	2.44e-6	–
	100	1.40e-5	1.96	1.80e-6	2.40	3.17e-7	2.94
	200	3.54e-6	1.98	3.36e-7	2.42	4.03e-8	2.98
	400	8.91e-7	1.99	6.12e-8	2.46	5.08e-9	2.99

Table 5.1: Convergence behaviour for different methods of integrating the slicing equation (5.19). This illustrates why we use Adams–Moulton methods in our algorithm.

using second reconstructions. Where enough resolution was used to avoid these oscillations the difference was always well below the convergence error and the plots given below would be visually indistinguishable.

5.8 Code tests

The simplest tests are to check the convergence of the slicing algorithms for the geometric variables. We do this by looking at the error in α on the unperturbed FRW spacetime. We look at α as equation (5.19) requires knowledge of a as well. The initial data is given by equation (5.22) for the density and velocity, with the metric variables a and α found by integrating the ordinary differential equations (5.21) and (5.19). We compare the three different integration methods for α in table 5.1.

The best results are from the third order Adams–Moulton method. Both the Adams–Moulton and Adams–Bashforth methods display third order convergence, with the Adams–Bashforth method actually displaying slightly better convergence. However the error is considerably lower with the Adams–Moulton method. The Milne–Simpson method although formally third order accurate only displays second order convergence except in the L_∞ norm. This seems to be due to oscillations caused at the initial step of the slicing equation which are propagated across the entire grid.

For long time stability we look at the convergence behaviour up to coordinate time 21.2. Snapshots of the behaviour of the method applied to the FRW spacetime are shown in figure 5.4. It is clear that the agreement is very good even to late times. More quantitavely we look at the convergence of the solution against itself, the convergence of the energy constraint equation (5.21) and the convergence against the exact solution (5.22). By convergence of the solution against itself we mean that the difference between two solutions calculated at different resolutions also scales as expected. This difference is a sum over all evolved variables (Π, Φ, α, a). The results are given in figure 5.5. Calculations with 200, 400 and 800 cells were used. The results of the final error measure, against the exact solution, are considerably less useful than the first two. This is because the error is so small that it is strongly affected by numerical round off. The results from the first two error measures are much cleaner.

Initially the convergence of the first two measures is slightly less than second order. This is probably due to initial small transients from the boundaries. This behaviour from the boundary is clear from the large "hump" near

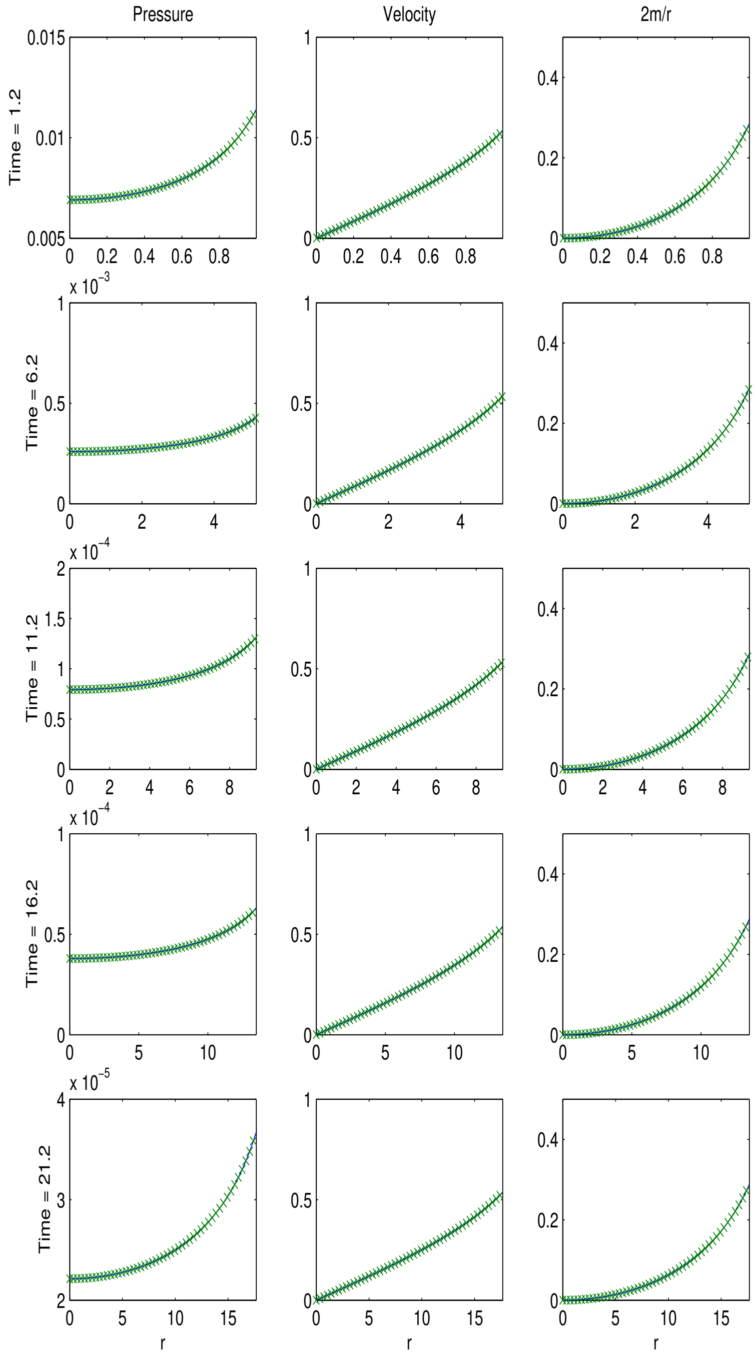

Figure 5.4: Long time behaviour of the algorithm applied to the exact FRW spacetime. The pressure, velocity and $2m/r$ are shown at a selection of coordinate times. The original grid had 200 cells. By the final coordinate time the grid contained about 14000 points. For clarity only 50 points are shown in each plot. The agreement between the FRW solution and the evolution even to this late time is extremely good.

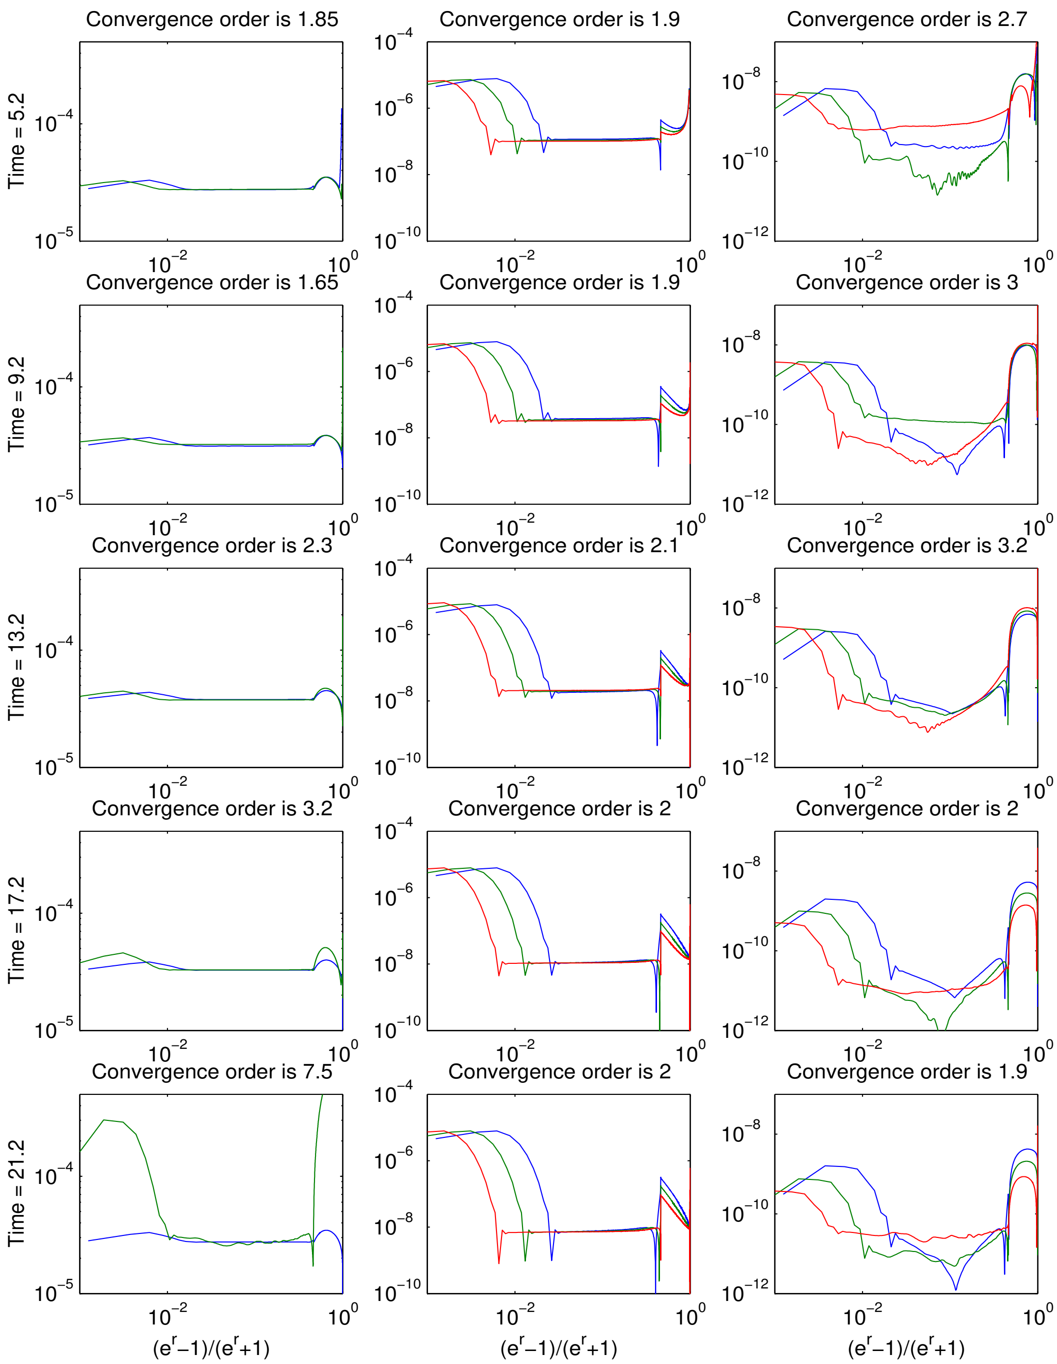

Figure 5.5: Long time stability of the algorithm is tested by the convergence of the solutions (left hand plots), of the error as measured by the energy constraint, equation (5.21) and by convergence against the exact solution, equation (5.22). The radial axis is scaled to concentrate on the original grid, but the whole grid is displayed. The order of convergence is indicated for each time. Apart from the transient behaviour near the origin the convergence is almost exactly second order in each of the first two measures. The late time problems with the solution convergence are due to the fluid quantities becoming very small, so that this error measure is no longer meaningful for low resolutions. The error against the exact solution is uniformly small so this test does not give such good results.

the origin. It has a different radial scale as it is purely numerical - the error is spread over the same number of grid points at each resolution. At intermediate times we see nearly exactly second order convergence. At late times the convergence still exists for the energy constraint, but the fluid variables are now so small (e.g., $\Pi \sim \Phi \sim 10^{-4}$) that the convergence on the solution itself is now affected by the rounding error for the low resolution runs. This convergence test is no longer very meaningful. However at all times for both of the first two error measures the absolute value of the error is small and does not grow appreciably. The error is also smallest in the original unexpanded part of the grid. This demonstrates the long time stability and convergence of the code.

Another test is to look at the convergence behaviour on a run that includes a perturbation. Here we have to look at the violation of the energy constraint, equation (5.21). We use the off-centre Gaussian profile of equation (5.24) which is near critical and evaluate equation (5.21) using second order finite differences. The results are shown in figure 5.6 at coordinate time 2.2. These results show precisely the same convergence behaviour as in the background. That is, second order convergence in the original part of the grid and slightly worse behaviour in the expanded part, with the errors propagating outwards.

A final test is to use the code for a spacetime with asymptotically flat boundary conditions as done by Evans and Coleman (1994), Neilsen and Choptuik (2000). Here the lapse is normalized such that $\alpha a = 1$ at the outer boundary of the grid. The scaling of the black hole mass and the maximum of the density over the spacetime given in figure 5.7, clearly show type II behaviour, with the expected critical exponent $\gamma \approx 0.36$. This shows that

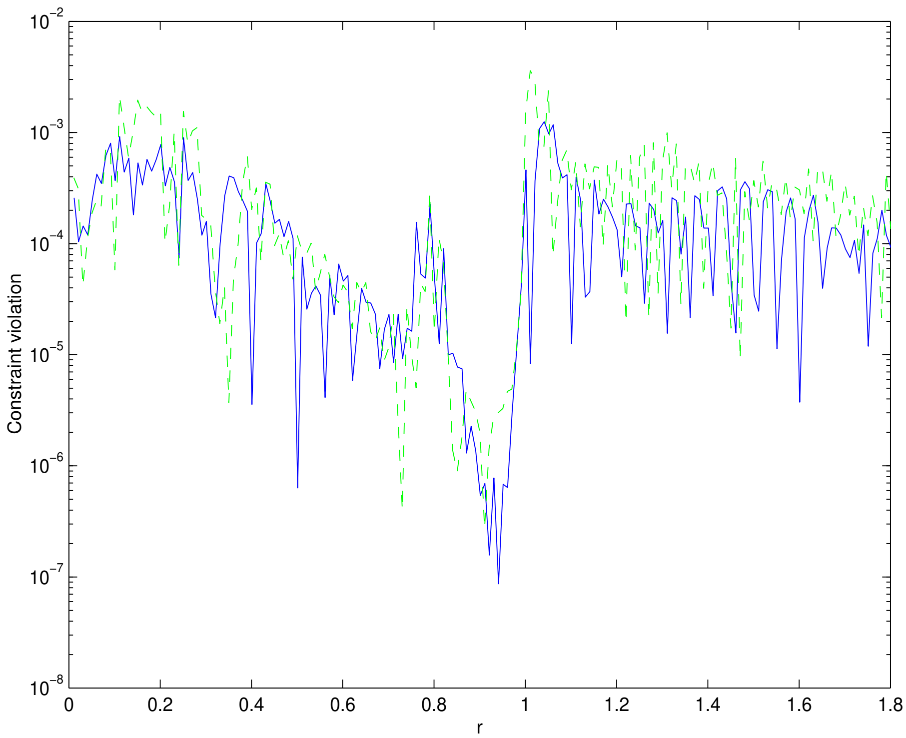

Figure 5.6: The violation of the energy constraint, equation (5.21) is measured by second order finite differencing. The results for 400 cells (solid blue line) are compared with the results for 800 cells (dashed green line) multiplied by 4. The initial data is FRW with an off-centred Gaussian perturbation. This shows that, at least in the original part of the grid ($r < 1$) that we have second order convergence. The convergence appears to be slightly worse for $r > 1$ but these problems are not propagating inwards.

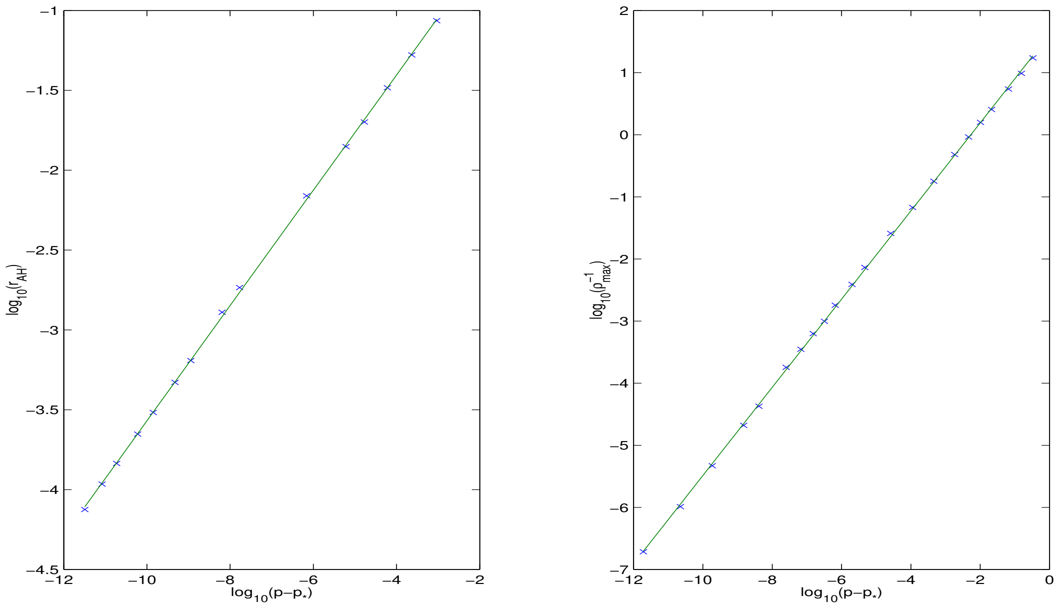

Figure 5.7: Plots of the mass of the black hole and the maximum of the density against the parameter C near the critical parameter for a critical search where the spacetime is asymptotically flat. Type II scaling is clearly found by the code with the minimum mass black hole being less than 10^{-4} and the critical exponent $\gamma \approx 0.36$ as expected.

our code is capable of resolving sufficiently small black holes to show type II behaviour, should it exist.

5.9 Results

5.9.1 General behaviour

For the three density profiles given in section 5.4 we were able to form black holes with both of the Gaussian forms, but with the Mexican Hat profile an apparent horizon was either present in the initial slice or did not form during the evolution. We are therefore unable to measure the mass of black holes formed with a Mexican Hat profile. The distributions used in Niemeyer and Jedamzik (1999a) are different in that the perturbation is larger than the horizon. We are unable to produce that type of perturbation in our coordinates.

We first consider the supercritical case. The initial profile splits into two parts. The part closest to the origin collapses under its own self gravity. The pressure rapidly grows by 4 or more orders of magnitude. The velocity profile is uniformly negative here. Eventually an apparent horizon will form.

The second part disperses from the origin. Most of the matter is at the "back" of the pulse (the part closest to the origin), due to the shape of the initial profile. Initially this pulse is stable and compact. As it travels into the unperturbed region a shock forms due to the higher velocity of the pulse than the matter in the unperturbed region.

In the region between the pulses a void forms. Most of the matter in this

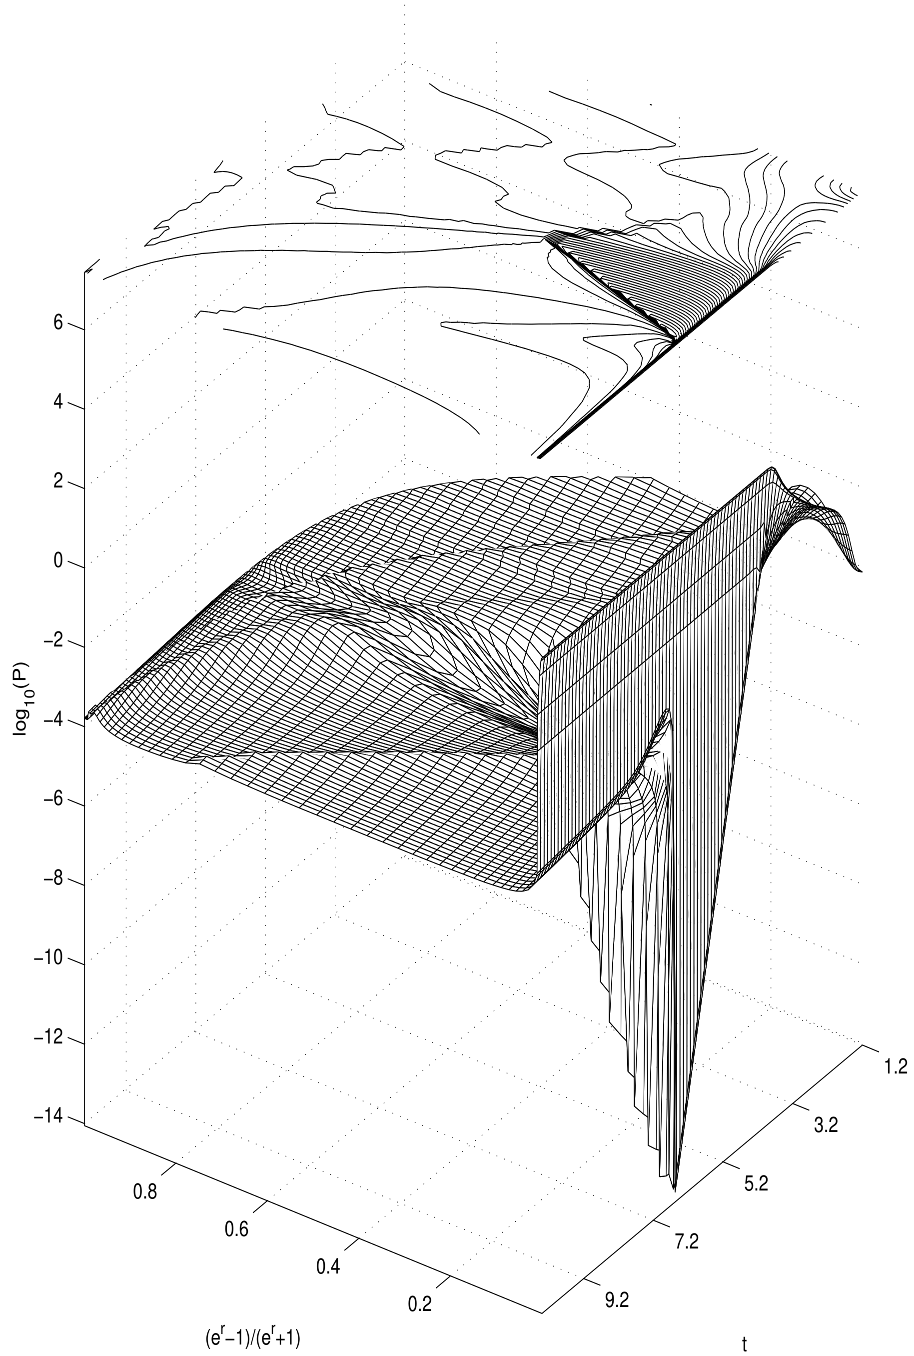

Figure 5.8: The pressure for an initially off centred supercritical Gaussian perturbation. The initial pulse splits. The pulse at the origin is grows under its own self-gravity. The pulse leaving the origin leaves a void in between. When the pressure gradient is sufficiently large the hydrodynamical forces cause some of the outgoing matter to fall in. This fills the void and accretes onto the black hole. The radial coordinate r goes from right to left and the time coordinate t from back to front.

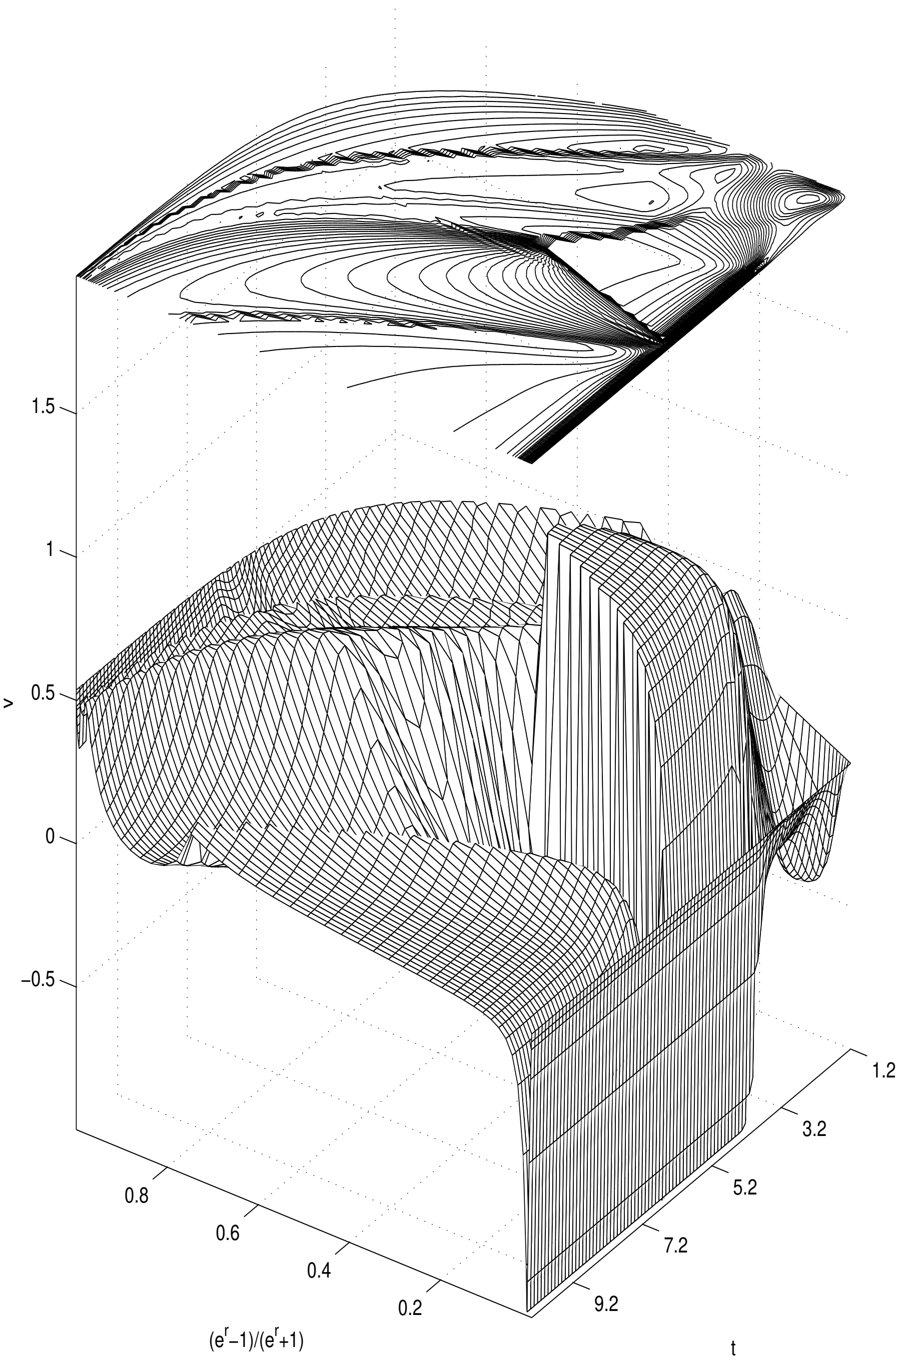

Figure 5.9: The velocity for an initially off centred supercritical Gaussian perturbation. The apparent horizon is indicated by the large negative velocity. The outgoing pulse is the large feature in the centre of the plot. This leads to the infalling shock, accreting onto the apparent horizon. The radial coordinate r goes from right to left and the time coordinate t from back to front.

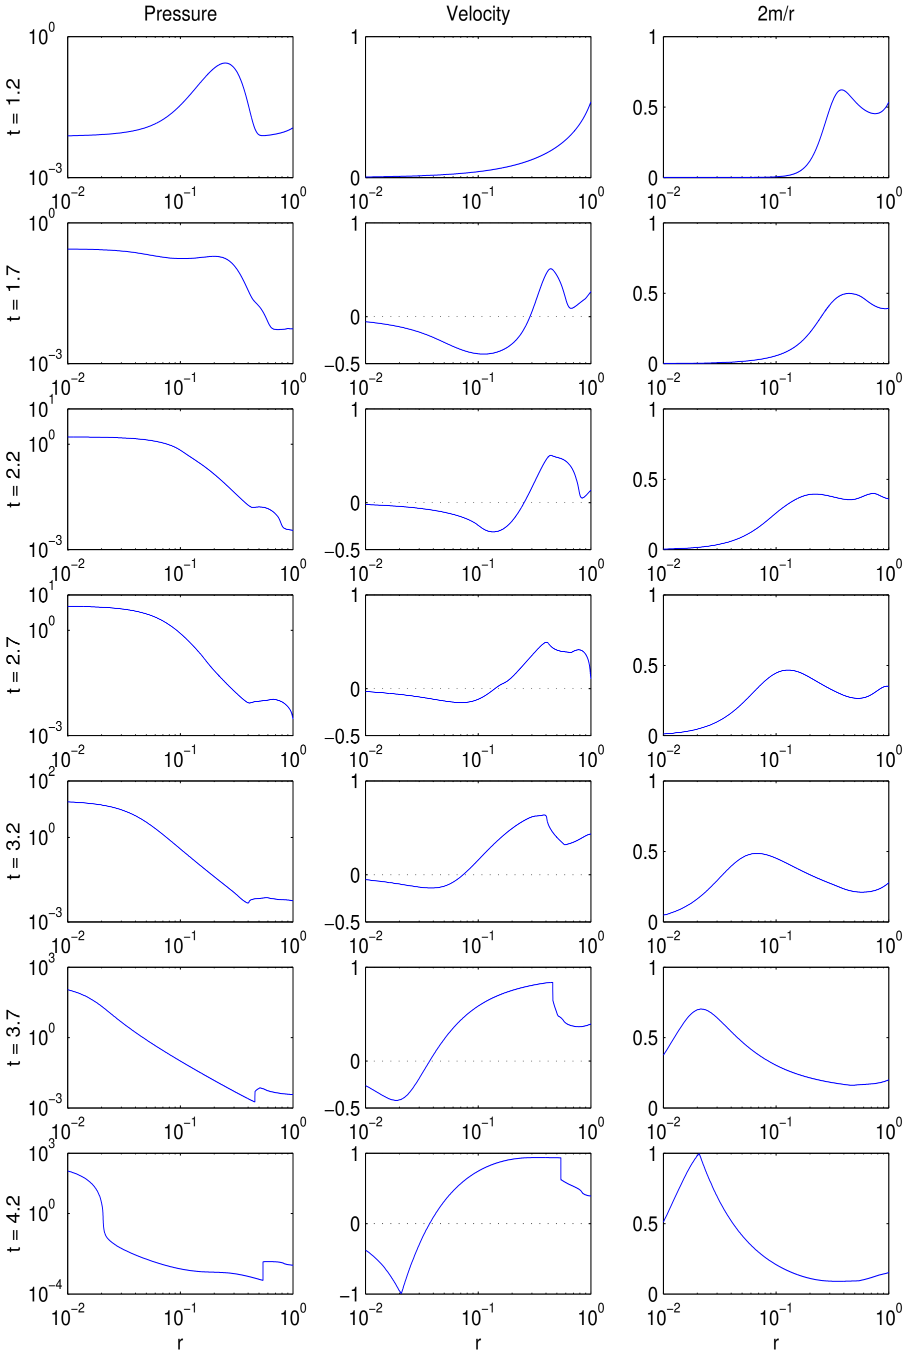

Figure 5.10: The pressure, velocity and $2m/r$ at selected coordinate times for an initially off centred supercritical Gaussian perturbation. The apparent horizon is indicated by $2m/r \to 1$; the velocity also tends the -1. The formation of the void is clear from the pressure plots.

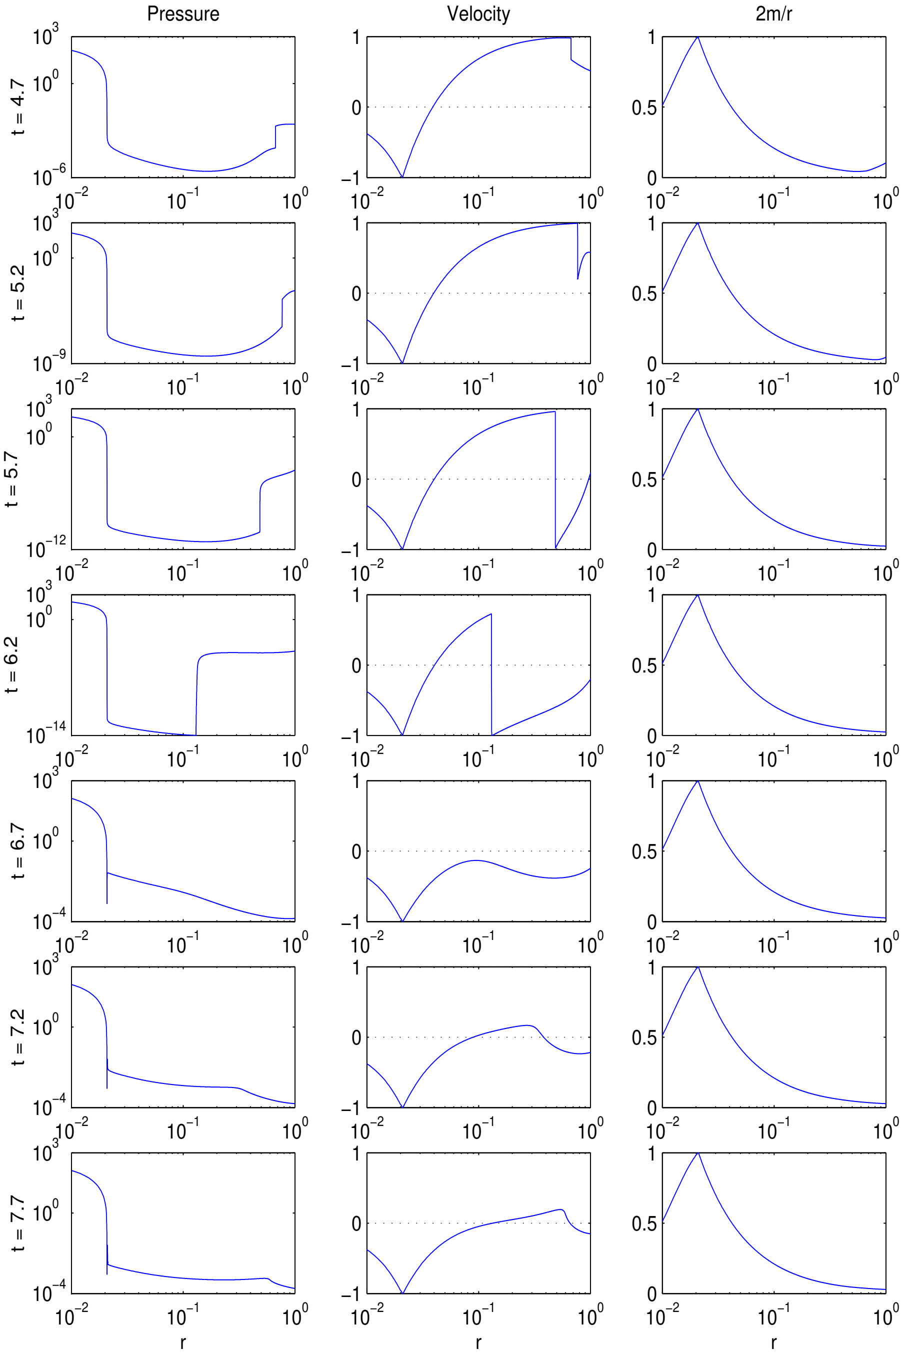

Figure 5.10: The pressure, velocity and $2m/r$ at selected coordinate times for an initially off centred supercritical Gaussian perturbation. At the apparent horizon before the void is filled the pressure changes by nearly 20 order of magnitude in a very small space. The "spike" in the pressure in the last plot is not a numerical instability. The apparent horizon has expanded due to the accretion of the infalling shock. In this region the lapse α has collapsed and "frozen" the evolution before the void is completely filled.

area is either accreting onto the region near the origin, or being washed out by the underlying expansion of the universe. If an apparent horizon forms, the pressure in this region may be between 15 and 20 orders of magnitude less than that inside the horizon. It is also substantially less than the pressure of the outgoing pulse. In this case, the pressure gradient and the gravitational field of the black hole force the outgoing pulse to split again. The majority of the outgoing pulse continues to escape, but some of the matter falls into the black hole. As this matter passes through the void a very strong shock forms across which the velocity may change from near c to near $-c$. After the accretion of this matter, the black hole very slowly accretes matter from the background, but this process appears to be negligible.

Mesh plots of the pressure for the supercritical case are shown in figure 5.8 and of the coordinate velocity in figure 5.9. This data is for a black hole with mass about $10^{-2}M_h$. The general behaviour outlined above can be seen. However, in order to make it clearer we refer to the snapshots from coordinate time 1.2 to 7.7, shown in figure 5.10. These show the pulse splitting, with the outgoing pulse rapidly leaving the domain of the plot. The apparent horizon, indicated by the steady spike with velocity $-c$ forms quite rapidly in this case. The void causes a shock to form which falls inwards. Little changes once this matter has accreted onto the black hole.

For the subcritical case, the pulse near the origin is initially held by its own self-gravity. However the hydrodynamic forces cause it to disperse, leaving a void near the origin. This travels outwards with high velocity. Once the pressure gradient between the void and the outgoing pulse is sufficient, an infalling shock forms. This is similar to the supercritical case except here

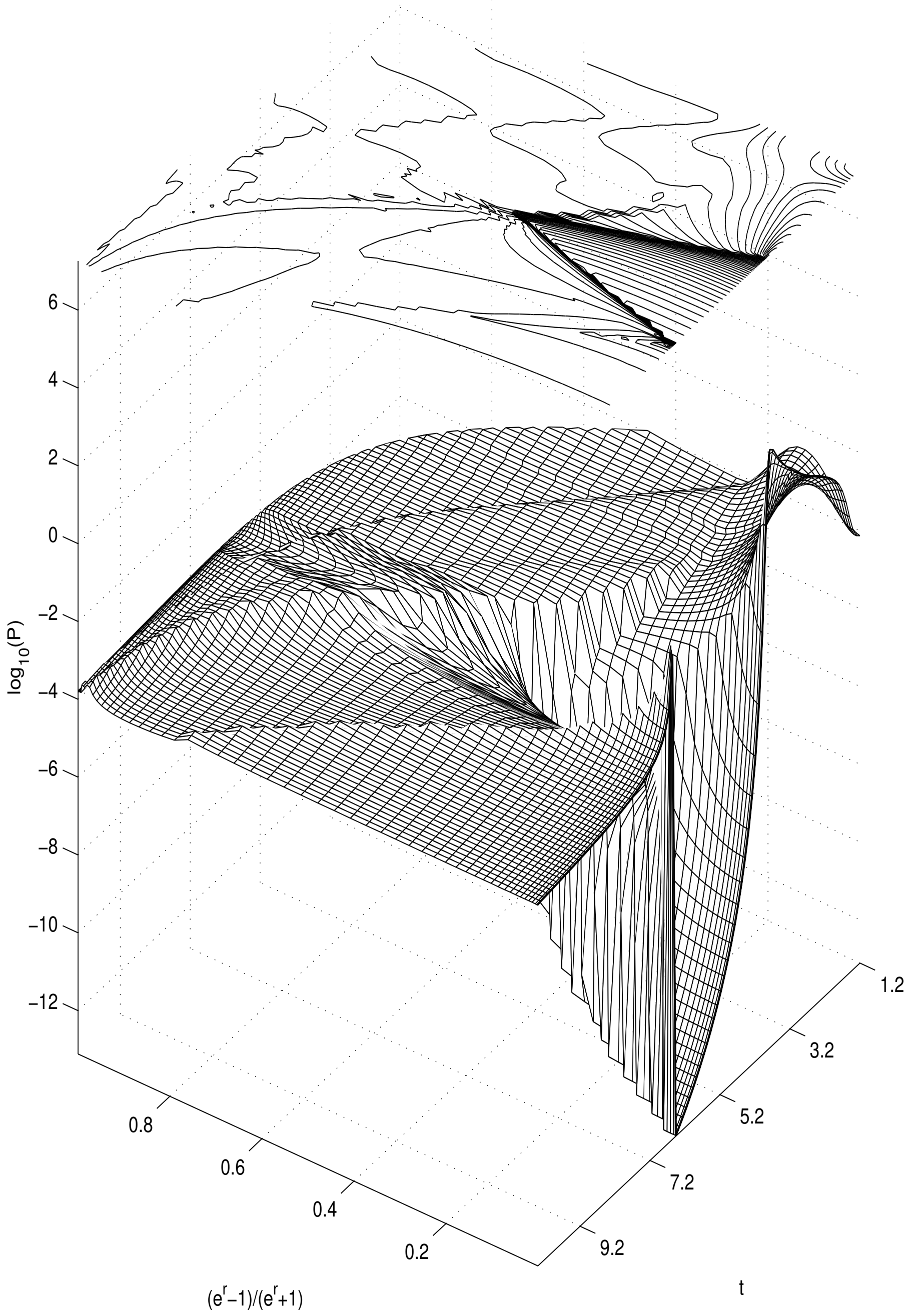

Figure 5.11: The pressure for an initially off centred subcritical Gaussian perturbation. The initial pulse splits. The pulse at the origin is does not have enough self-gravity and eventually disperses from the origin leaving a void. When the pressure gradient is sufficiently large the hydrodynamical forces cause some of the outgoing matter to fall in. This leaves the FRW background near the origin. The radial coordinate r goes from right to left and the time coordinate t from back to front.

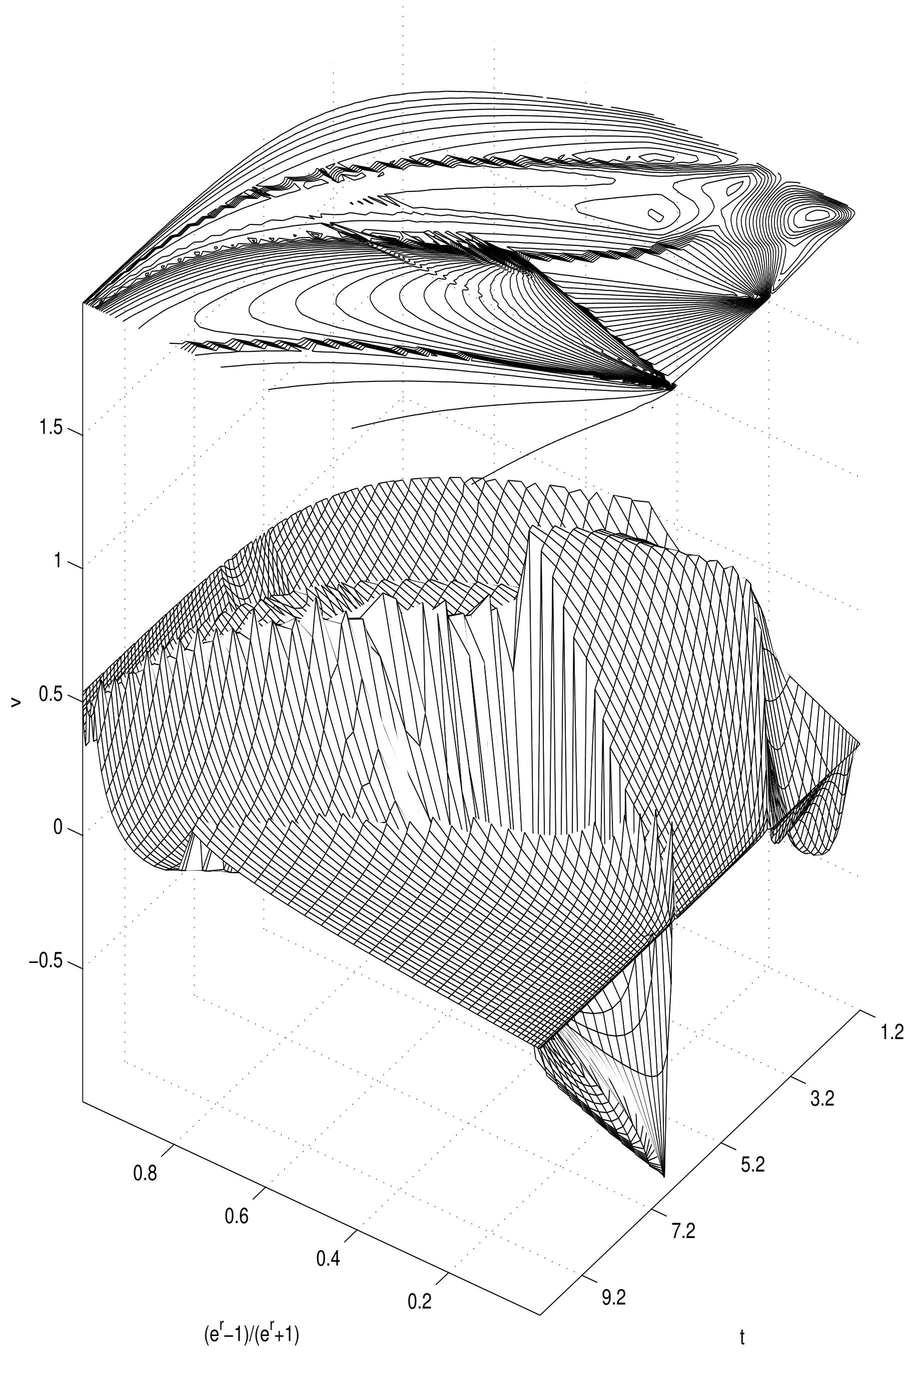

Figure 5.12: The velocity for an initially off centred subcritical Gaussian perturbation. The dispersal of the initial pulse and reflection of the infalling shock are on the right of the plot. The outgoing pulse leading to the infalling shock is in the centre. The radial coordinate r goes from right to left and the time coordinate t from back to front.

it cannot accrete onto a black hole. Instead the shock reflects from the origin leaving the FRW background in the interior. At late times the spacetime is approximately FRW with large perturbations propagating outwards.

Mesh plots of the pressure for the subcritical case are shown in figure 5.11 and of the coordinate velocity in figure 5.12. The general behaviour outlined above can be seen. However, in order to make it clearer we refer to the snapshots from coordinate time 1.2 to 7.6, shown in figure 5.13. These show the pulse splitting, with the outgoing pulse rapidly leaving the domain of the plot. The reflecting pulse goes out with velocity near c forming a void. The hydrodynamical forces then cause a very strong shock to form which falls in. After this reflects from the origin the FRW background is left in the interior.

When the perturbation is highly supercritical the shock does not form. This seems to be due to the self gravity of the fluid that is forming a horizon. The outgoing pulse is drawn back in, accreting quickly enough that no shock formation is visible. This is shown in figure 5.14.

5.9.2 Accretion

Once the black hole has formed it may increase in mass by accretion. Unlike in flat space there is no reason for the final spacetime to be stationary, so this accretion may continue indefinitely. However, we have found that the vast majority of the accretion occurs rapidly in an infalling shock. This agrees with the results of Niemeyer and Jedamzik (1999a). Once this material has accreted, the change in the black hole mass is negligible.

However, the most important point from the astrophysical sense is that

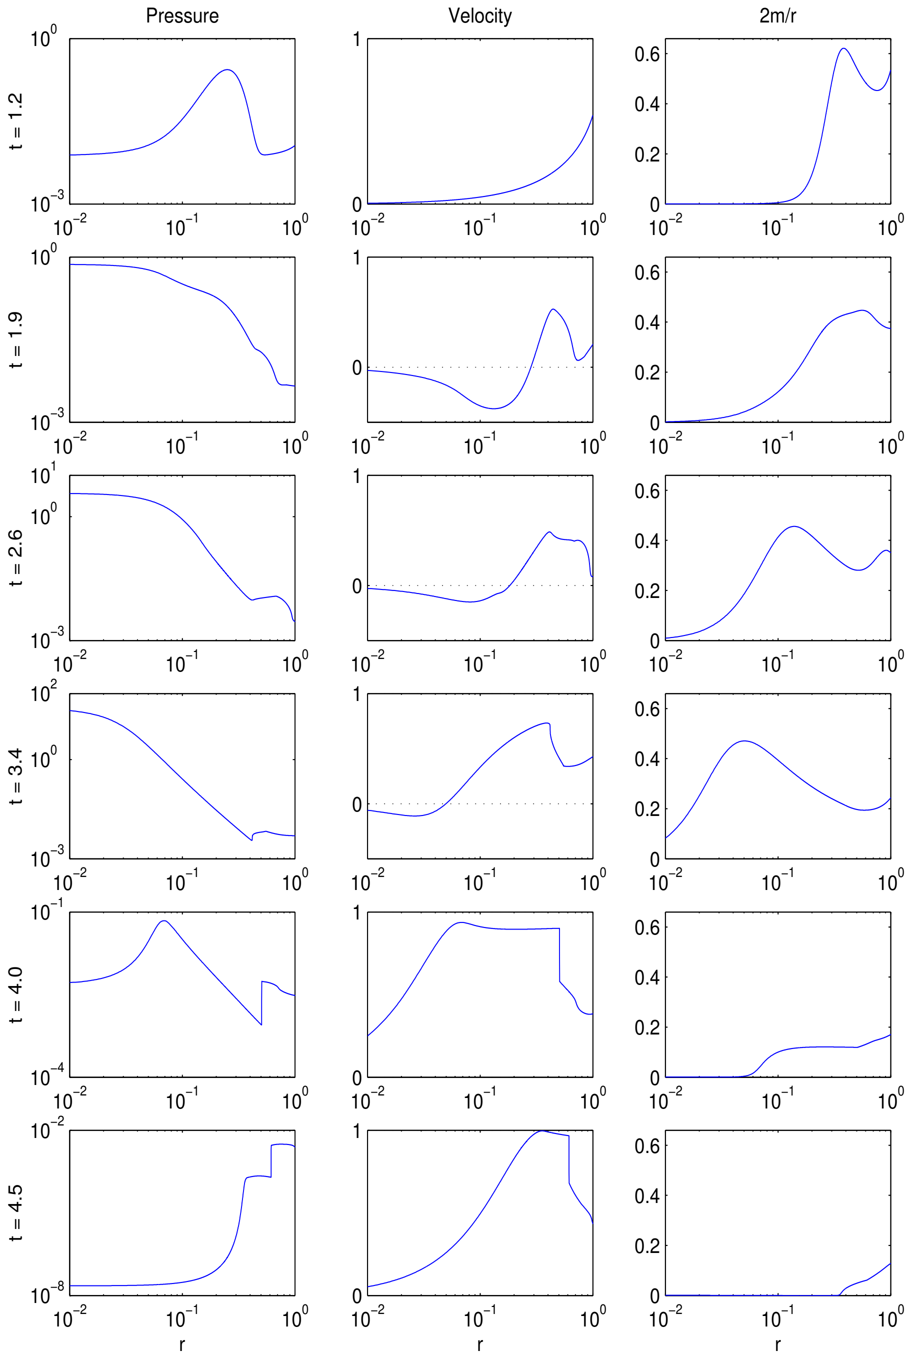

Figure 5.13: The pressure, velocity and $2m/r$ at selected coordinate times for an initially off centred subcritical Gaussian perturbation. The dispersal of the fluid near the origin leads to the very strong shock across which the velocity changes from near c to near $-c$.

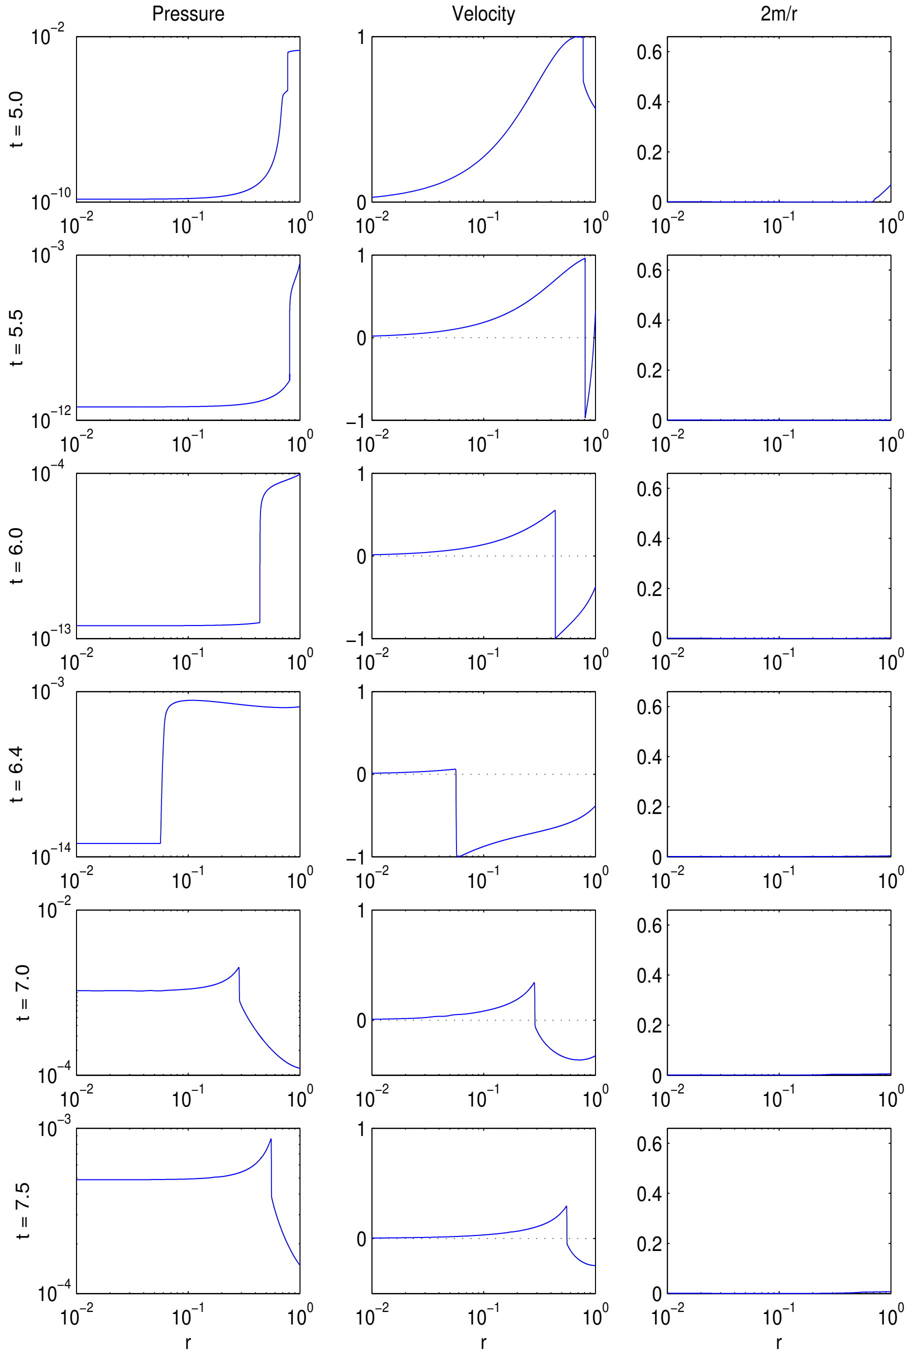

Figure 5.13: The pressure, velocity and $2m/r$ at selected coordinate times for an initially off centred subcritical Gaussian perturbation. After the shock has reflected from the origin we are left with the FRW background in the interior.

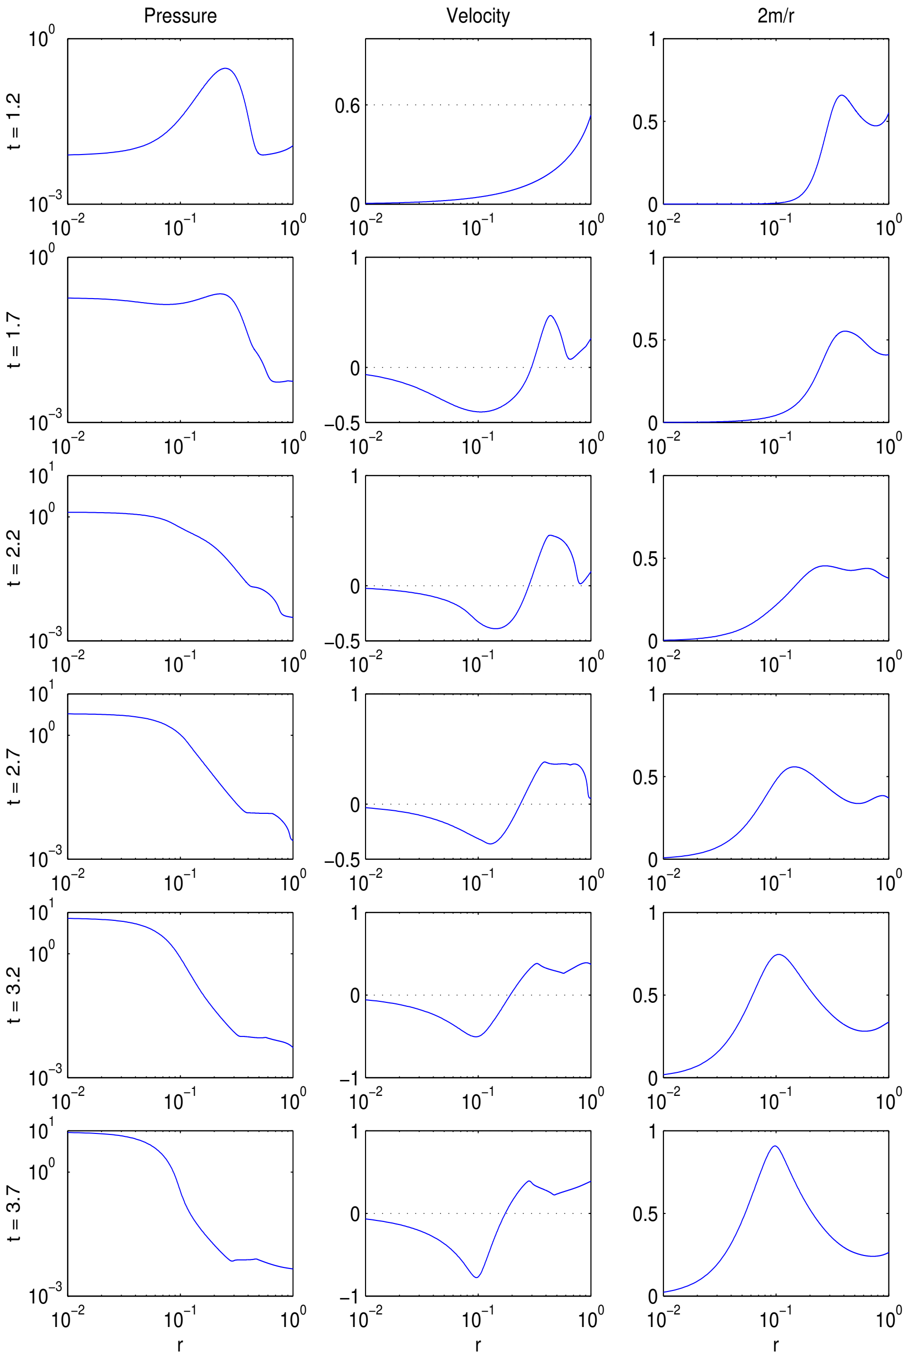

Figure 5.14: The pressure, velocity and $2m/r$ at selected coordinate times for an initially off centred highly supercritical Gaussian perturbation.

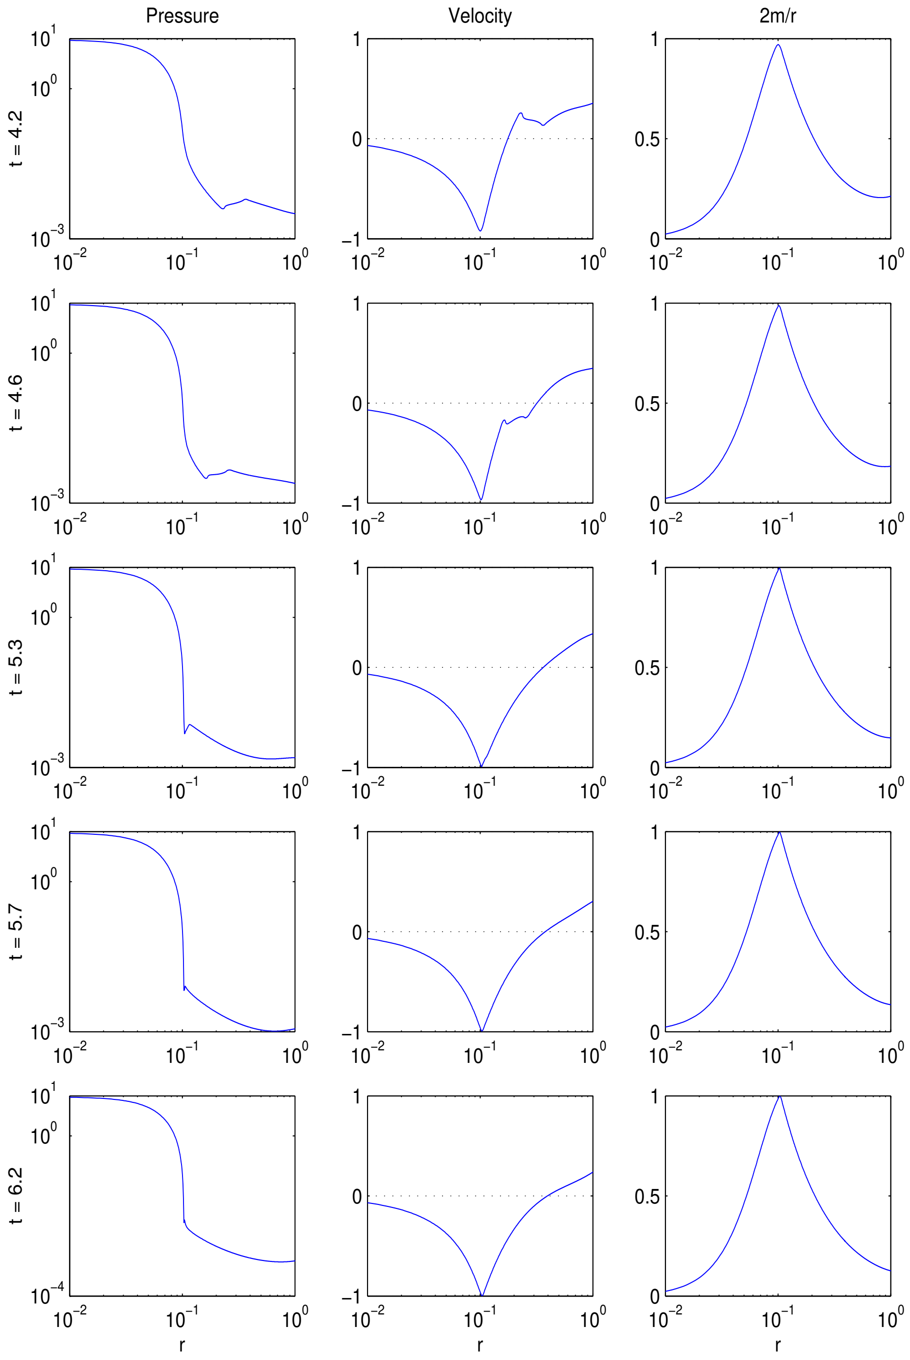

Figure 5.14: The pressure, velocity and $2m/r$ at selected coordinate times for an initially off centred very supercritical Gaussian perturbation.

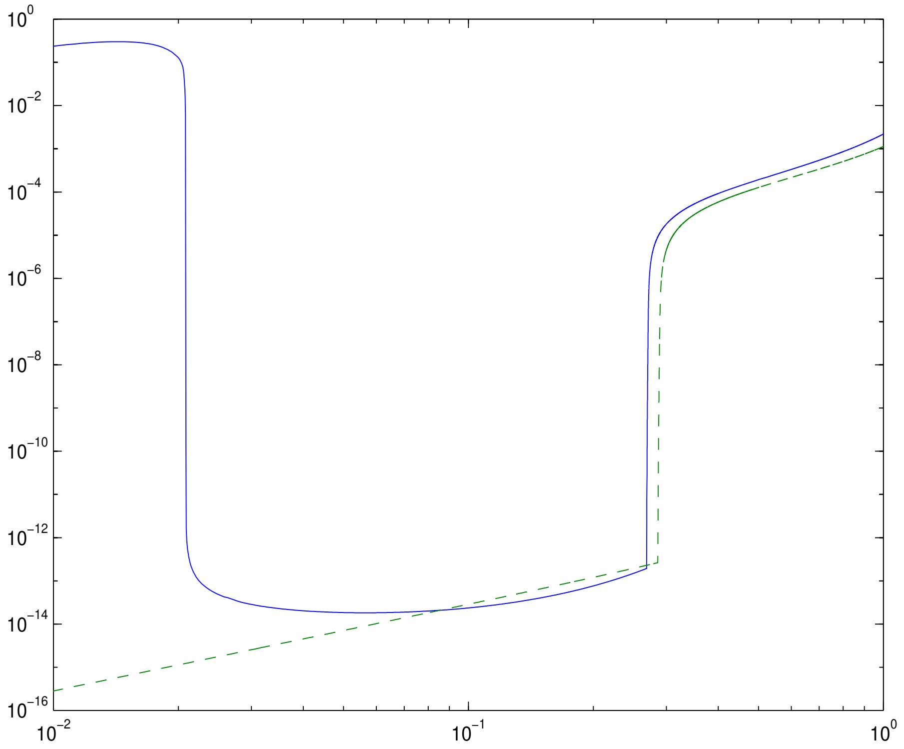

Figure 5.15: Two snapshots of the variable $\omega = 4\pi r^2 a\rho$ at the infalling shock. The blue solid line gives a supercritical evolution, and the green dashed line a subcritical evolution. It is clear that the amount of mass in the infalling shock is largely independent of the existence of a black hole. This shows that the accretion of the material in the infalling shock, which will add $\sim 10^{-4} M_h$ to the black hole mass, puts a lower bound on the mass of a black hole formed from realistic data. This is independent of the type of critical behaviour.

the amount of mass accreted is largely independent of the mass of the black hole. In fact in some senses it can be said to be independent of the existence of the black hole. The infalling shock is caused by the geometry and the pressure gradient caused by the formation of a void. The amount of mass in the infalling shock that accretes onto the black hole (if it exists) depends on the gradient between the void and the outgoing pulse. Gravitational effects only play a part for very large (that is $p - p_* \sim \mathcal{O}(1)$) black holes.

This is illustrated by figure 5.15. This shows two snapshots with the infalling shock at similar stages for a supercritical and a subcritical evolution. The gravitational mass of the material in the shock, given by the difference in the fluid variable $\omega = 4\pi r^2 a\rho$ across the shock, is clearly very similar in both cases. With this initial data the accretion of this material will add about $10^{-4}M_h$ to the black hole mass. The agreement for subcritical and supercritical evolutions is emphasized by the snapshots shown in figure 5.16. This shows margninally subcritical and supercritical evolutions until the point where the subcritical evolution is clearly dispersing. The shock position and strength is very similar.

Although it may be argued that different initial data would lead to different amounts of accretion, it is intuitively clear that the set of initial data that has no accretion cannot be generic. For a realistic black hole formed in this way we should expect a minimum mass from accretion alone, independent of the type of critical behaviour found.

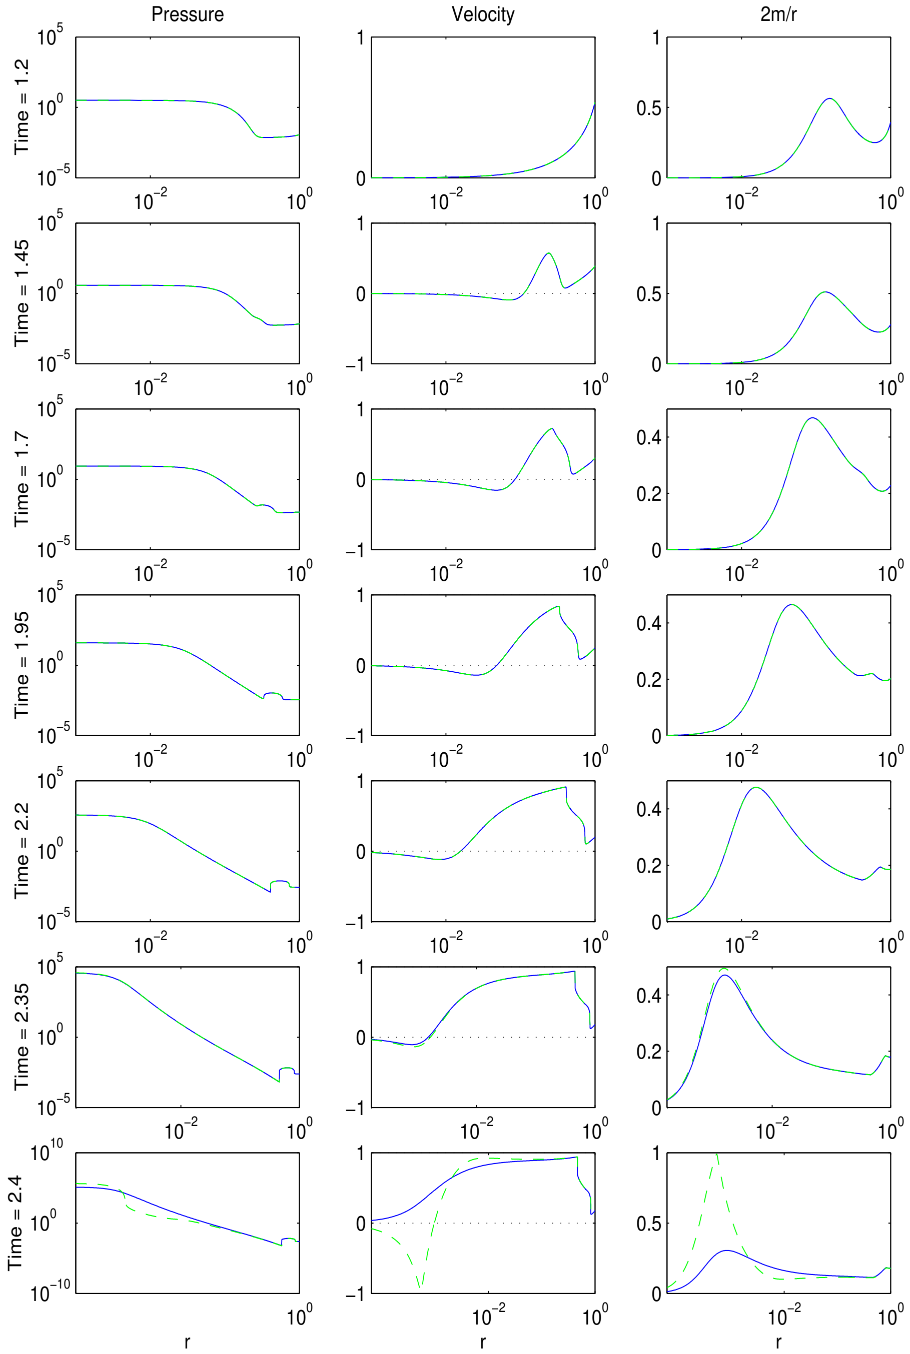

Figure 5.16: Snapshots of the pressure, velocity and $2m/r$ for a marginally subcritical (green dashed line) and supercritical (blue solid line) collapse. The centred profile was used. The collapse is shown until it is clear that the subcritical pulse is dispersing. The shock forms in both cases and is very similar in strength and position.

5.9.3 Critical behaviour

For the profiles that formed a black hole we have been able to construct black holes with masses that vary over around 3 orders of magnitude. This corresponds to determining the critical parameter with a relative error of less than 10^{-10}. For parameters greater than those shown here, we were unable to determine the black hole mass as the apparent horizon was present in the initial slice. Over some of this range of parameter space, typically between $1\frac{1}{2}$ and 2 orders of magnitude in the black hole mass, the black hole masses follow a power law scaling of the form of equation (5.1) with critical exponent $\gamma \approx 0.36$. This is at least twice the range given by Niemeyer and Jedamzik (1999a), which shows scaling over about half an order of magnitude. However in the range of parameter space shown by Niemeyer and Jedamzik (1999a) we do not see perfect scaling. This may be due to our initial data.

However we find a minimum mass for the black holes of about $2 \times 10^{-4} M_h$ for the centred Gaussian and about $3 \times 10^{-4} M_h$ for the off centred Gaussian where M_h is the mass inside the Hubble horizon on the initial slice. This might indicate type I behaviour. However there is evidence that the critical solution is the self-similar Evans–Coleman solution. We discuss the precise nature of the critical behaviour in section 5.10.

Confirmation of this behaviour can be found by looking at the scaling of the maximum of the density ρ for the subcritical evolutions. This is a dimensionful quantity and so should scale as

$$\rho \propto |p_* - p|^{-2\gamma}. \tag{5.36}$$

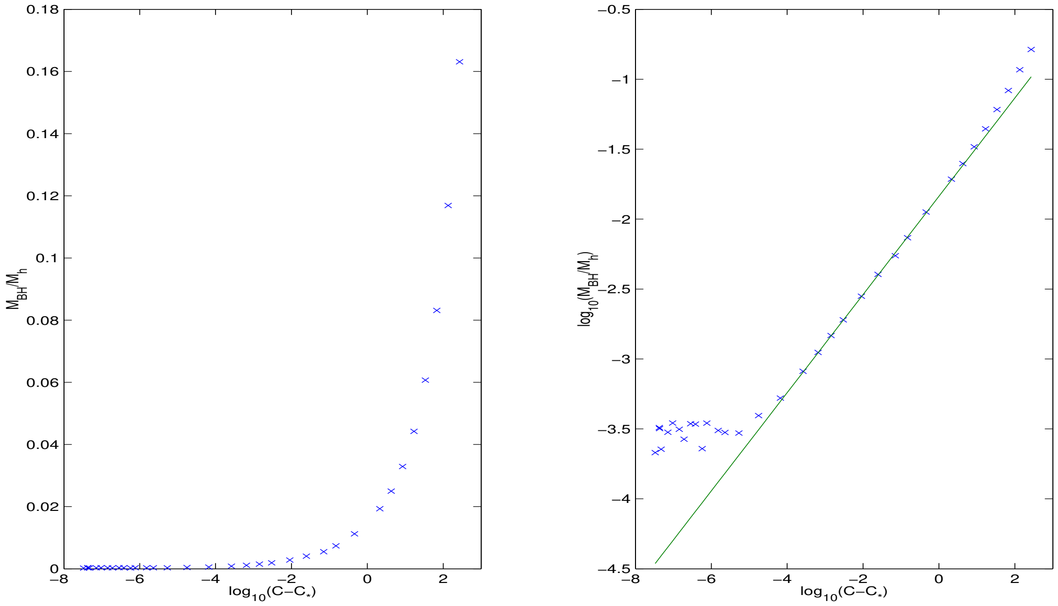

Figure 5.17: Plots of the mass of the black hole against the parameter C near the critical parameter for the centred Gaussian initial data. The minimum mass for the black hole is evident from the first plot. Away from the critical parameter we see power law scaling over about 2 orders of magnitude in the black hole mass with $\gamma \approx 0.35$. For very large values of $C - C_*$ the scaling again fails.

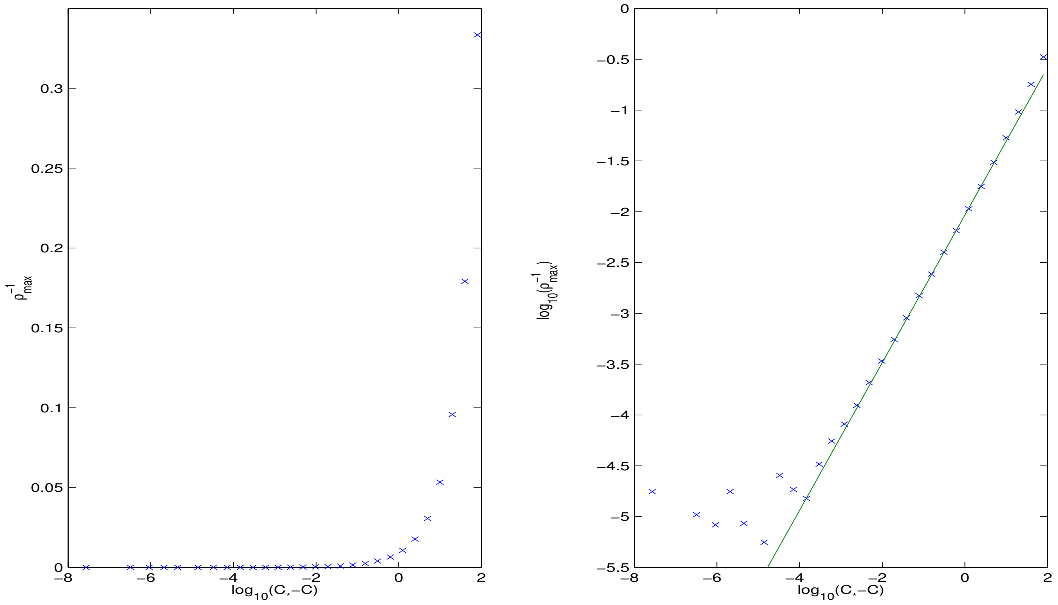

Figure 5.18: A plot of the inverse of the maximum of the density ρ against the parameter C near the critical parameter for the centred Gaussian initial data. Away from the threshold there is power law scaling with $\gamma \approx 0.36$. Near the threshold the scaling fails, as for the mass.

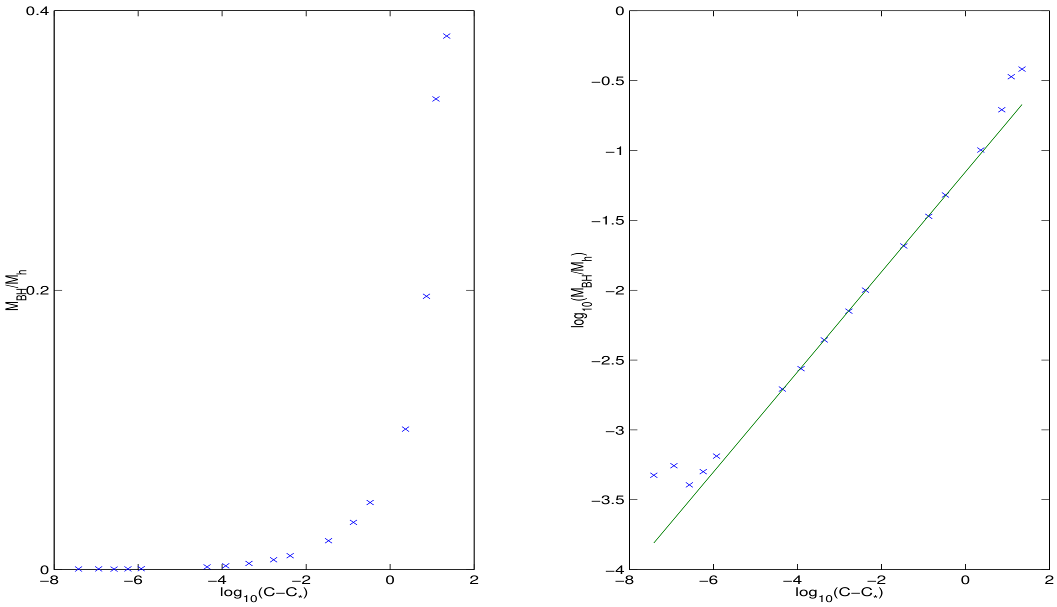

Figure 5.19: Plots of the mass of the black hole against the parameter C near the critical parameter for the off centred Gaussian initial data. The minimum mass for the black hole is evident from the first plot. Away from the critical parameter we see power law scaling over less than 2 orders of magnitude of the black hole mass with $\gamma \approx 0.36$.

The original argument of Garfinkle and Duncan (1998) only checked the Ricci scalar on axis. However for our equation of state the Ricci scalar is proportional to the trace of the energy momentum tensor due to Einstein's equations, and as for a perfect fluid $T \propto \rho - 3p$ this vanishes for the radiation dominated equation of state we use here.

The results for the black hole masses for the centred Gaussian are shown in figure 5.17, and for the off-centre Gaussian in figure 5.18. The scaling for the density is shown for the centred Gaussian in figure 5.19, and for the off centred Gaussian in figure 5.20. The scaling behaviour for the density holds over around 3 orders of magnitude, approximately twice that of the black

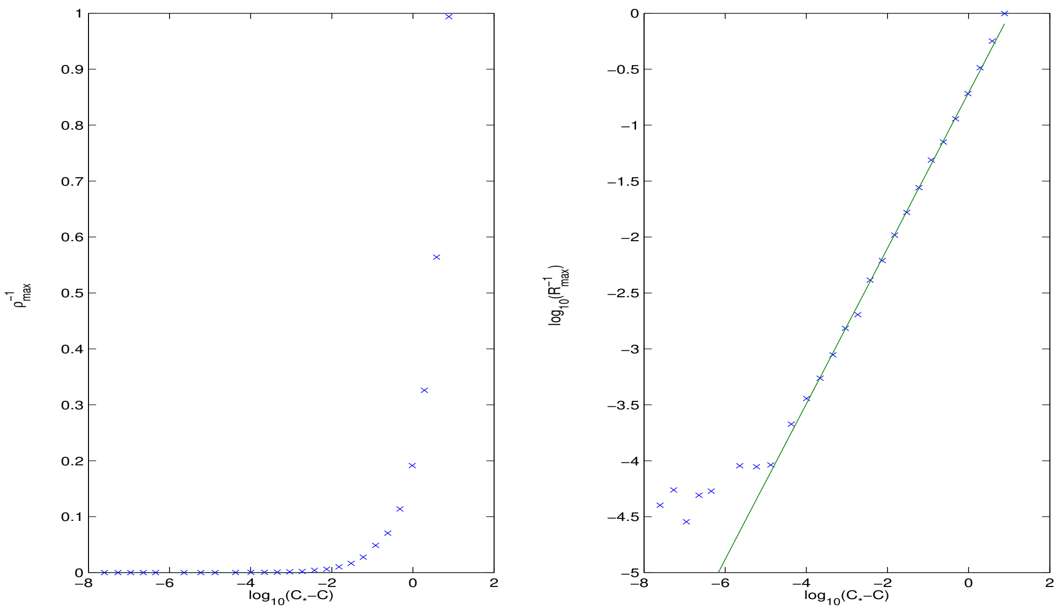

Figure 5.20: A plot of the inverse of the maximum of the density ρ against the parameter C near the critical parameter for the off centre Gaussian initial data. Away from the threshold there is power law scaling with $\gamma \approx 0.35$. Near the threshold the scaling fails, as for the mass.

hole mass. This is to be expected if the scaling holds over the range in the parameter C. The factor 2 is precisely that required by equation (5.36).

The approximate critical solutions generated by the collapse simulations appear to approach the solution of Evans and Coleman (1994). Evidence of this is shown in figure 5.21, and for a single solution over a range of times in figure 5.22, with the deviations from the exact self-simlar solution shown in figure 5.23. The self-similar solution was constructed using the methods outlined by Koike et al. (1999), Neilsen (1999). However we are unable to match the collapse simulations with the self-similar solution using the methods outlined by Neilsen (1999) as we do not observe type II behaviour. Instead we set the logarithmic radial coordinate in such a way that the peak in the metric function a agrees. The clear agreement near the origin indicates that this is the universal critical solution. However the matching to the outer boundary occurs at a shock, and near this shock the numerical solution differs from the Evans–Coleman solution.

Although the left hand plots of figure 5.21 may appear to indicate that even the shock wave is universal this in an artifact of the logarithmic radial scale. A close up of this area, shown in the right hand plots of figure 5.21, shows that although the profiles are similar the shock positions are quite different. As the coordinate freedom has already been used to match to the Evans–Coleman solution near the origin, it is clear that universality fails near the shock.

Finally we look for evidence of type I critical behaviour. This should be apparent in the lifetime of the critical solution, t_p, which should follow the scaling law (5.2). This is measured by the time the evolution spends near

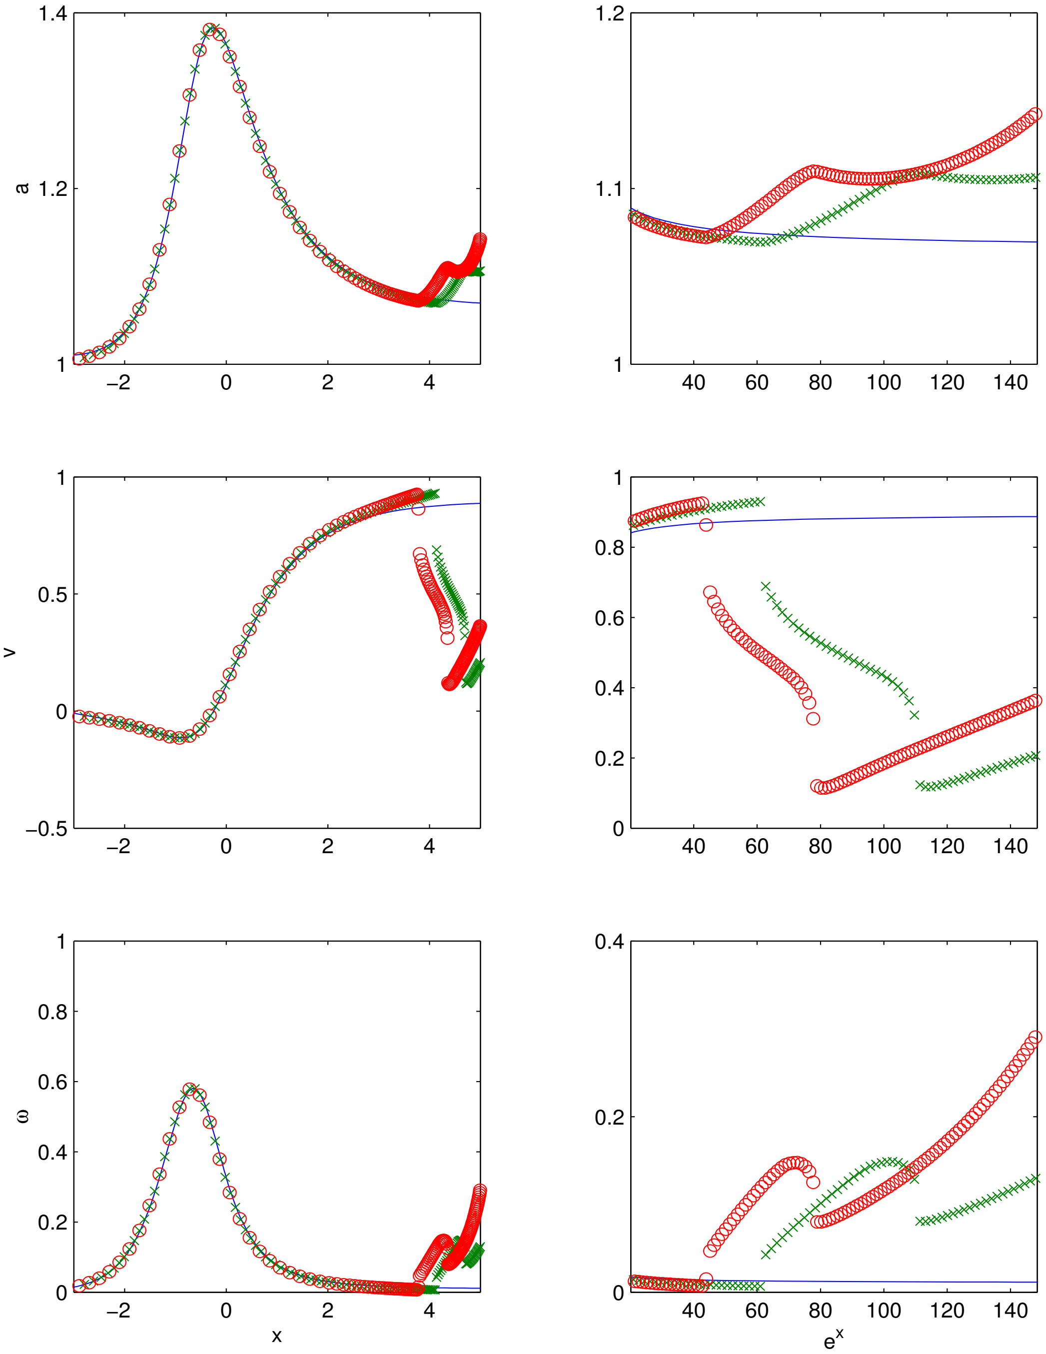

Figure 5.21: Approximate critical solutions compared to the Evans–Coleman self-similar solution. The metric function a, fluid velocity v and combined variable $\omega = 4\pi r^2 a^2 \rho$ are shown. The radial coordinate x is logarithmic with $x = 0$ being the sonic point. The Evans–Coleman solution is the line. The solution found from the centred Gaussian is marked with circles, and the off centred Gaussian with crosses. Only one in twenty points is shown. The Evans–Coleman solution is the critical solution near the origin. The is matched to the FRW solution at infinity across two shocks, which breaks the self-similarity. The solution is no longer universal at this point, as is shown by the right hand plots.

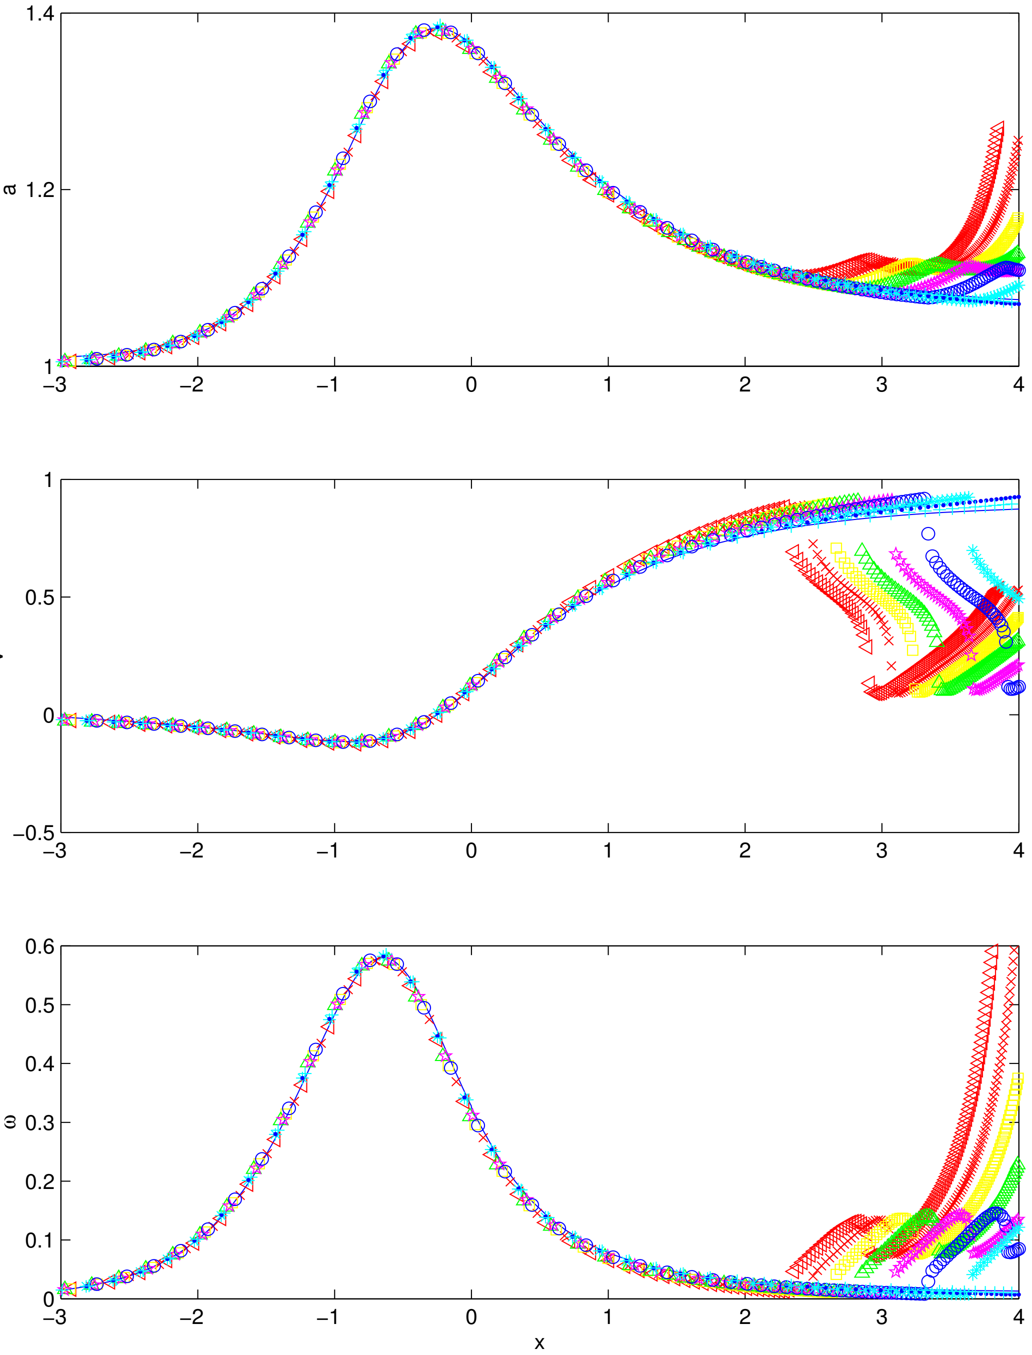

Figure 5.22: An intermediate solution is shown in the range $t_* - t \in [0.012, 0.12]$. The Evan–Coleman self-similar solution is also plotted as the solid line. The agreement is very good over this range.

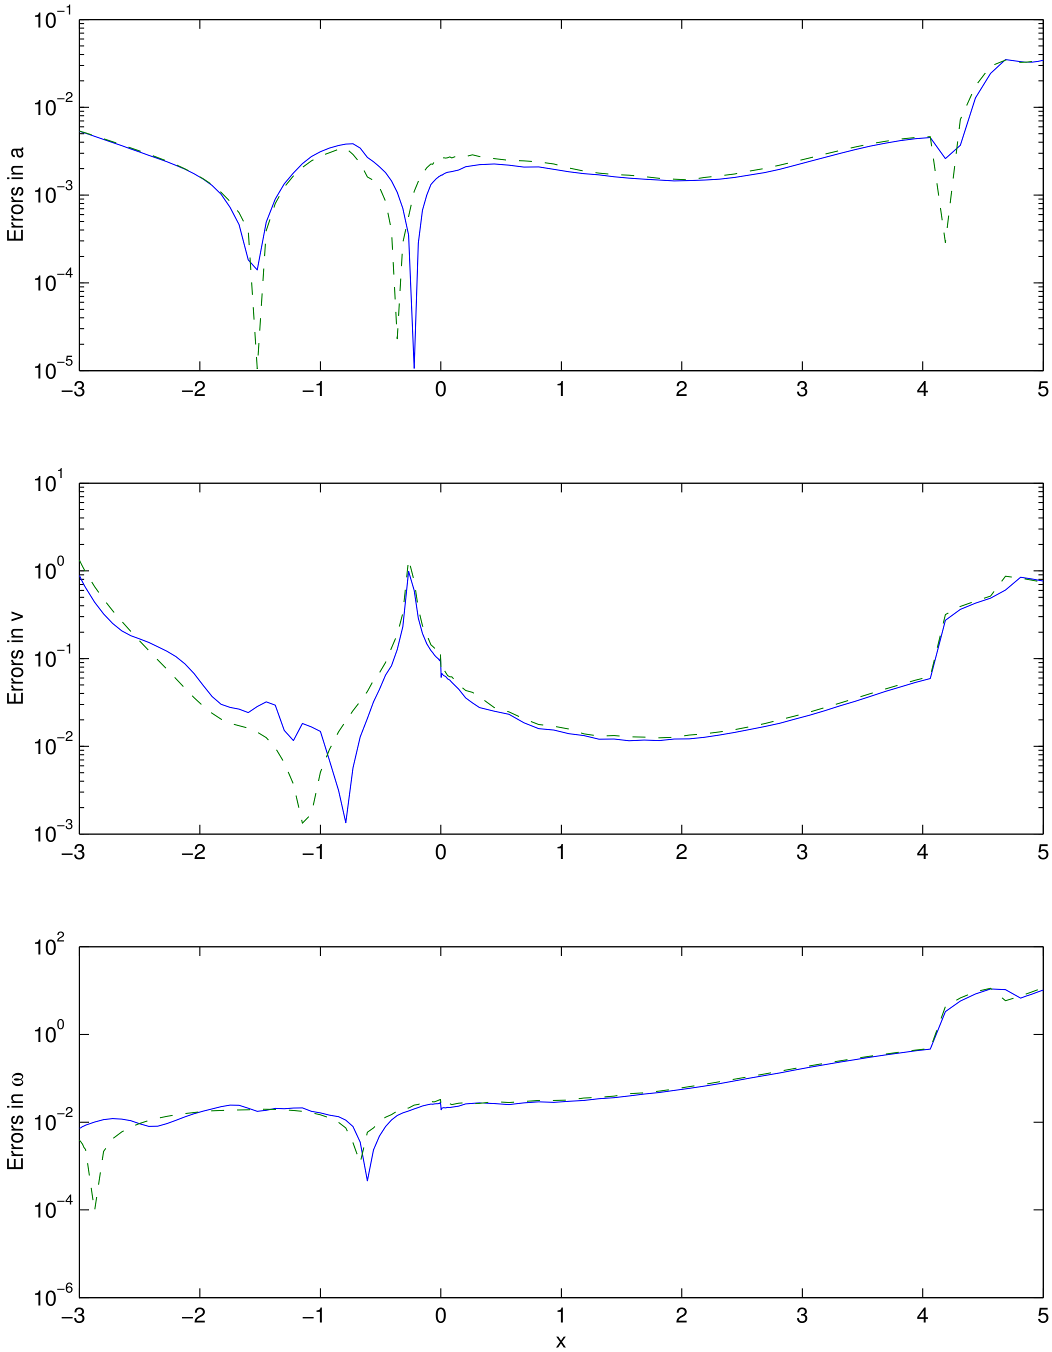

Figure 5.23: Approximate critical solutions compared to the Evans–Coleman self-similar solution. The deviation of the metric function a, fluid velocity v and combined variable $\omega = 4\pi r^2 a^2 \rho$ from the Evans–Coleman solution are shown. The radial coordinate is logarithmic. The solution found from the centred Gaussian is the blue solid line, and the off centred Gaussian is the green dashed line. In the interior the solutions agree well given the interpolations necessary to compare the data.

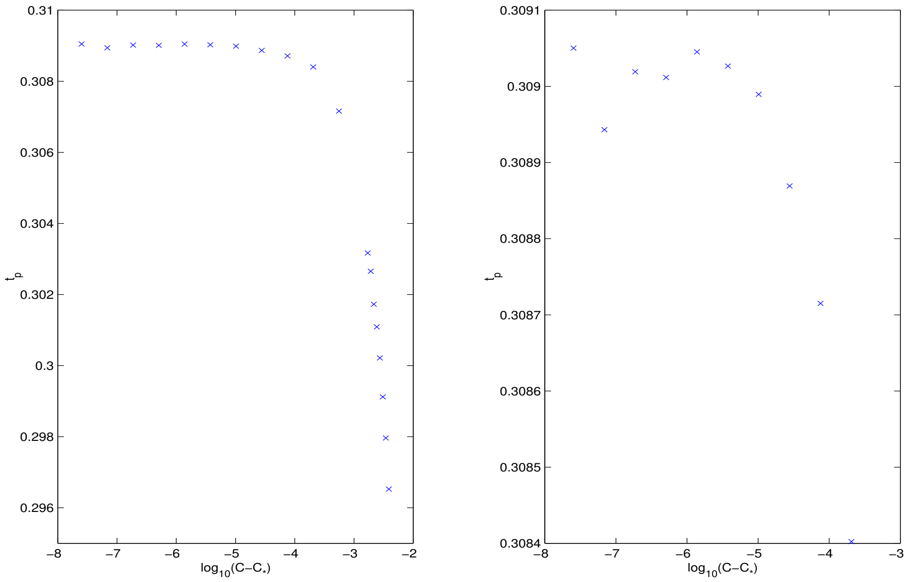

Figure 5.24: Typical results for the time spent near the critical solution t_p. The scaling law (5.2) for t_p clearly does not hold. Although it is possible we are considering the wrong time coordinate, it appears from the right hand plot that near the critical value t_p is no longer monotonic.

the critical solution. The time coordinate used should be that generated by the Killing vector of the critical solution. Without clear knowledge of what that should be, we have estimated t_c by the amount of coordinate time that the peak of the metric function a is within 1% of the Evans–Coleman solution. Typical results are shown in figure 5.24. It is clear that the scaling does not hold in this time coordinate. We also note that t_p appears to stop being monotonic near the black hole threshold. This strongly suggest that the behaviour is not type I.

Comparing our results with the results of Niemeyer and Jedamzik (1999a) is not simple. To calculate the equivalent of the density contrast indicator δ by just integrating the extra mass on the slice we must make a number of assumptions. Firstly, we assume that before the initial time t_0 the spacetime was exactly FRW (in other words the perturbation appears "as if by magic" at time t_0). Secondly, we assume that the Hubble horizon is unchanged by the appearance of the perturbation. With these assumptions we find the coordinate time τ on the initial surface of constant t and evolve it along with the rest of the variables. This allows us to find the data on a surface of constant τ. We then assume that as none of the matter in the perturbation leaves the domain the indicator is unchanged by this evolution, allowing us to calculate the density contrast of the perturbation.

A second method of calculating δ is to note that it is equivalent to $m_{ptb}/m_{FRW} - 1$ evaluated at the horizon, using the covariant definition given in section 5.2. As they both diverge independently at the horizon we must instead evaluate this inside the horizon. We then check that the radial derivative of the density and hence the integrand is suitably small. Nonetheless both of these methods are only rough approximations to the true value. However, as the agreement is reasonable we assume that the results are sufficiently accurate for our purposes.

Using the first method we find that for the off-centred Gaussian the density contrast indicator near critical is $\delta_c \approx 0.82$, whilst for the centred Gaussian the threshold value is $\delta_c \approx 0.36$. The second method gives $\delta_c \approx 0.86$ and $\delta_c \approx 0.37$ respectively. Although the same order of magnitude as the value $\delta_c \approx 0.7$ found by Niemeyer and Jedamzik (1999a), we are expecting $\delta_c \leqslant 1$ anyway. It is clear that δ_c is not an invariant for different families of initial data.

5.10 Conclusions

The numerical evidence presented here shows that we do not have any of the standard critical collapse scenarios. Due to the careful testing of the numerical methods and codes in chapters 2, 3 and section 5.8, and that our method is based on Neilsen and Choptuik (2000) that did find standard critical collapse, we are confident that the numerical method is not the reason for this.

Firstly we note that there is evidence of type II collapse but that it breaks down near the threshold. This evidence is clearest in figure 5.21 where the universality of the majority of the critical solution is clear, as is the close agreement with the Evans–Coleman solution. We also see scaling behaviour for around 3 orders of magnitude of $p - p_*$ both in the black hole mass and dimensionful quantities such as the maximum of the density. The critical exponent $\gamma \simeq 0.36$ is about the expected value. However, near the threshold of black hole formation there is a finite minimum mass for the black holes. This is also seen in the maximum of the density which also fails to scale correctly near the threshold.

Aside from the mass gap there is no evidence for type I behaviour whatsoever. The intermediate solution is definitely not static or periodic. Nor does the time spent near the critical solution appear to scale in the expected way. Instead we conclude that the behaviour is approximately type II but that scaling fails near the threshold. The best way to check that the behaviour really is type two would be to compare the difference between the subcritical and supercritical intermediate solutions and the Evans–Coleman solution for the same values of $|p - p_*|$. The growing mode should have the same profile but different signs. This was attempted, but meaningful results were not found. The most likely reason for this is insufficient resolution of the intermediate solutions and insufficiently accurate solution of the ordinary differential equations giving the Evans–Coleman solution, combined with the interpolations required for the comparison. However all other evidence clearly points to type II behaviour.

I believe that the appearance of the shock leads to the failure of the scaling behaviour. Here I do not mean the accretion of material in a shock wave as discussed in section 5.9.2. Although that will lead to a minimum mass regardless of the black hole mass, the minimum mass discussed here appears to be independent of the accretion. First we shall look at why a shock could lead to the failure of critical behaviour, and then we shall see that it is possible.

The standard analytical picture of type II critical collapse relies on the existence of a self-similar intermediate solution. Linear perturbations of this are tuned away, with the only growing mode giving the critical exponent γ. However when a shock forms these perturbation equations are no longer valid as the differentiability is lost. A shock wave is not a linear perturbation of the critical solution and so it may not be possible to tune it away.

The appearance of a shock also causes the standard dynamical systems picture to fail. The flow on the phase space is given by the flow map f which is derived from the equations of motion of the system. Fixed points in the phase space are fixed points of f. The stable manifold of a point p is given by

$$W^s(p) = \left\{ q : \lim_{n \to \infty} f^n(q) = p \right\}. \tag{5.37}$$

The unstable manifold is given by considering f^{-n}. Implicit in this definition of the unstable manifold is that the flow map f is invertible. However at a shock entropy increases and information is lost; the situation is irreversible. This means that the flow map cannot be uniquely inverted and it is no longer possible to define the unstable manifold and talk of the growing modes.

Finally we note that in the construction of the self similar solution it is not necessary to consider boundary conditions except for regularity at the origin and at the sonic point. However as soon as a shock appears the solution across the shock is forced to obey the Rankine-Hugoniot conditions and the "interior" is coupled to the "exterior". This again indicates the likelihood that the appearance of a shock would perturb the self similar solution even at the threshold.

The appearance of a shock may perturb the self similar solution in such a way that infinitesimally small black holes cannot be formed. But is this causally possible for the initial data used here? It is clear that the shock wave cannot directly affect the solution as it does not fall in until after the black hole has formed. But this is not necessary for the threshold solution to be perturbed. All that is required is that the shock wave forms inside the past domain of dependence of the accumulation point of the self similar solution. Then the perturbations could not necessarily be tuned away, as is outlined above.

For the initial data used here it would appear to be the case that the shock wave does form within the past domain of dependence of the accumulation point. However, it is very close to the boundary of the domain. This may explain why the minimum mass is so small. The accumulation point for the centred Gaussian profile is at $t_a \approx 2.5$ and for the off-centre profile at $t_a \approx 4.1$. In each case as the threshold solution does not collapse to a small black hole some approximation is necessary. This is done from comparing the real radius r and the coordinate $x = log(r/(t_a - t))$ in the comparison of the self-similar Evans–Coleman solution and the intermediate solutions obtained in the collapse simulations. The shock wave appears at $r \approx 0.32$ and $t \approx 1.92$ for the centred profile and at $r \approx 0.41$ and $t \approx 3.35$ for the off-centre profile. As the critical solution should be self similar the past characteristic of the accumulation point has a constant velocity $\xi_a = 1/\sqrt{3}$. This means that the past domain of dependency of the accumulation point extends out to $r \approx 0.33$ for the centred profile and $r \approx 0.43$ for the off-centre profile.

A remaining question is the universality of this result. Figure 5.21 also shows that the position and strength of the shock wave is not universal. However the minimum mass found is remarkably similar, as is the time and position that the shock forms relative to the past horizon of the accumulation point of the self similar threshold solution. This might be just due to our initial data. Because of the horizon in our coordinates the profiles chosen were all quite narrow. This makes it more likely that a shock would form. However, it is possible that the outer boundary not being asymptotically flat also has an effect. A conclusive answer to this is unlikely to be found without using a coordinate system that is regular at the horizon but allows for the numerical evolution of shocks. A "matching" method has been considered, where comoving coordinates are used near the horizon and outside, and the algorithm above inside the horizon, but this has not been implemented due to the inherent computational complexity.

Aside from critical collapse there is also the astrophysical questions. The results on accretion indicate that even if the initial data were such that an infinitesimally small black hole formed, it would soon accrete a finite amount of matter. Assuming that $10^{-4}M_h$ is a reasonable order of magnitude estimate for realistic initial data, this corresponds to $10^{-2}M_{\odot}t$, where t is the formation time and $M_{\odot}$ a solar mass. A black hole formed from stellar collapse has a minimum mass of the order of a solar mass. So in order to be able to distinguish a black hole formed in the early universe it would need to form well before $t = 10^2$ seconds, which is quite early in the radiation dominated era.

Our results also indicate that the critical value of the extra mass δ_c is dependent on the particular family of initial data, contrary to the results of Niemeyer and Jedamzik (1999a).

Chapter 6

Conclusion

The aims of this thesis have been to look at the best way to solve problems in relativity involving fluids numerically and to apply these methods to the study of critical collapse and primordial black holes. The first three chapters looked at the numerical methods in a wide variety of test cases. The final chapter contained the application to critical collapse.

One very specific aim was to find the advantages of using the ultrarelativistic equation of state over the ideal gas equation of state. The results of chapters 2 and 3 show that there are clear advantages to the simpler ultrarelativistic approximation. Firstly, it is more efficient to solve the corresponding Riemann problem, especially in the general case where tangential velocities are included, as shown in sections 2.6 and 2.9.3. Secondly, the problems of accurately capturing discontinuities occur at contact and tangential surfaces, which do not appear in one dimensional ultrarelativistic hydrodynamics. In higher dimensions the improvements in accuracy will occur only with numerical methods that are accurate at the tangential discontinuity.

Another specific aim was to try and find the "best" numerical method to use in relativistic hydrodynamics. This will of course depend on the problem under consideration and the computational resources available. However, from the results given in the chapters 2 and 3 the advantages of the ENO methods with accurate Riemann solvers are clear. The advantages of the exact Riemann solver are probably insufficient to justify using it for anything other than validation of results found by other methods. The advantages of the very high (fourth and fifth) order accurate ENO schemes are more obvious, but probably insufficient to justify their use except in very specific problems. For these reasons we concentrated on the use of a third order ENO method with a Marquina solver in the final chapter.

The most surprising result in this work occurs when these methods are applied to the problem of critical collapse on an FRW background. We do not see the standard type I or type II collapse over the entire range of parameters. The only possible reason for believing that we are seeing type I behaviour is that there is a minimum black hole mass. However, the good agreement between the intermediate solution which is (almost) universal, and the Evans–Coleman solution that is central to type II perfect fluid collapse, combined with the scaling behaviour for the black hole mass and other dimensionful quantities over large ranges of the parameter space, indicate that the underlying behaviour is type II. However, this behaviour does not hold near the threshold of black hole formation.

This result differs from previous numerical and analytic results such as Evans and Coleman (1994), Koike et al. (1995), Neilsen and Choptuik (2000). The clearest feature that might cause our results to differ from previous results is the appearance of the shock wave. This breaks self similarity, causing the analytical approximations to fail. We have seen that the shock is causally connected to the accumulation point of the critical solution, under the assumption that it exists. It also breaks the universality of the intermediate solution, even if only very weakly. All these points indicate how the shock could cause the phenomena that we see.

The question that immediately arises is why do we see this phenomena when other simulations, notably Neilsen and Choptuik (2000) do not? As noted in chapter 2, shock waves are in some sense generic in hydrodynamics. However in the situation considered here there are a number of features that could cause the difference. First, our boundary conditions are not asymptotically flat. This may lead to stronger shocks forming. Second, there is a considerable amount of matter on the background already. Interaction with this matter is more likely to lead to strong shocks. Thirdly, due to our coordinate system which fails at the horizon we have taken very narrow and steep initial data. This may increase the likelihood of shock formation. It is possible that different initial data on the same spacetime may not form a shock inside the past domain of dependence of the accumulation point, leading to true type II behaviour. We were unable to broaden our initial data to find this without introducing errors at the outer boundary due to discrepancies between the interior and the assumption of exact FRW in the expanding grid.

For the astrophysical questions the precise nature of the critical behaviour is perhaps not so important compared to other results. Firstly, the result that the amount of matter accreted will be approximately independent of the mass of the black hole if the black hole is near the threshold. Secondly, the result that the indicator δ_c, characterizing the additional mass in the spacetime at the threshold of black hole formation, is not a universal value. These results unfortunately make it more difficult to calculate the primordial black hole spectrum as attempted by e.g., Green and Liddle (1999). However, calculations with more "realistic" initial data may still give useful information.

Appendix A

Primitive variable methods, applications, and different formulations

A.1 Numerical methods for primitive variables

A.1.1 Primitive SLIC

The numerical methods in chapter 3 are all based on a conservative form of the system of hyperbolic partial differential equations. In one dimension this is written in the conservation form

$$\partial_t \mathbf{q} + \partial_x \mathbf{f}(\mathbf{q}) = 0. \tag{A.1}$$

However it may be that the system of equations is not in conservation form. In this situation we would expect any numerical method to give incorrect shock speeds for the reasons outlined in section 2.2.2. However we may still be able to gain some useful information from evolving the equations in primitive form. It may be that shocks do not form. But we may still need a method that can deal with very steep gradients.

The SLIC method given in section 3.7.1 was extended to the more general primitive variable case. This extension was originally performed in order to apply it to the equations found in the Friedrich approach, outlined in section A.2. Here we give them for a general system of hyperbolic partial differential equations and use the method to show explicitly that the conservative form is necessary to evolve a shock correctly.

We shall write the system of equations in the form

$$\mathbf{W}_t + \mathbf{A}(\mathbf{W})\mathbf{S}(\mathbf{W})_x = 0, \tag{A.2}$$

where $\mathbf{W}$ is the vector of basic variables, $\mathbf{S}$ a vector of conserved type and $\mathbf{A}$ a matrix with coefficients possibly dependent on the basic variables $\mathbf{W}$. This reduces to the conservative form when $\mathbf{W} = \mathbf{q}$, $\mathbf{S} = \mathbf{f}$ and $\mathbf{A}$ is the identity matrix. It also reduce to the standard primitive-variable form

$$\mathbf{W}_t + \mathbf{A}(\mathbf{W})\mathbf{W}_x = 0 \tag{A.3}$$

when $\mathbf{S}$ is the identity function.

The primitive variable method is found by expanding the Lax-Friedrichs and Richtmyer fluxes, replacing any terms of the form $\mathbf{f}^L - \mathbf{f}^R$ with terms of the form

$$\mathbf{A}(\mathbf{W}^L, \mathbf{W}^R)(\mathbf{W}^L - \mathbf{W}^R), \tag{A.4}$$

where $\mathbf{A}$ is some average matrix which we will always choose to be

$$\mathbf{A}(\mathbf{W}^L, \mathbf{W}^R) = \mathbf{A}\left(\tfrac{1}{2}(\mathbf{W}^L + \mathbf{W}^R)\right). \tag{A.5}$$

To simplify the presentation we introduce the notation

$$\mathcal{A}(\mathbf{W}_L, \mathbf{W}_R) = \mathbf{A}\left(\tfrac{1}{2}(\mathbf{W}_L + \mathbf{W}_R)\right)\left(\mathbf{S}(\mathbf{W}_L) - \mathbf{S}(\mathbf{W}_R)\right). \tag{A.6}$$

We then proceed as follows:

1. **Extrapolation** The data are reconstructed by extending the grid point values to an interval $[x_{i-1/2}, x_{i+1/2}]$ around the grid point x_i as

$$\mathbf{W}_i(x) = \mathbf{W}_i^n + \frac{(x - x_i)}{\Delta x}\Delta_i, \tag{A.7}$$

where Δ_i is an estimated slope and which in general can be written

$$\begin{aligned} \Delta_i &= \alpha\Delta_{i-1/2} + (1-\alpha)\Delta_{i+1/2} \\ \Delta_{i-1/2} &= \mathbf{W}_i^n - \mathbf{W}_{i-1}^n && \text{(A.8a)} \\ \Delta_{i+1/2} &= \mathbf{W}_{i+1}^n - \mathbf{W}_i^n && \text{(A.8b)} \end{aligned}$$

where α is a parameter in the range $[0, 1]$ and is arbitrary. We will always choose $\alpha = 1/2$ to give a centred slope. For the TVD version of the algorithm this extension will depend on neighbouring cells through slope limiters (such as Van Leer or Superbee). The values at the cell boundaries are

$$\mathbf{W}_i^R = \mathbf{W}_i^n + \tfrac{1}{2}\Delta_i, \tag{A.9a}$$

$$\mathbf{W}_{i+1}^L = \mathbf{W}_{i+1}^n - \tfrac{1}{2}\Delta_{i+1}. \tag{A.9b}$$

where again the superscript L/R indicates the left/right extension within that computational cell.

2. **Evolution**

The next step is to evolve the cell boundary values forward half a time step:

$$\bar{\mathbf{W}}_i^{L/R} = \mathbf{W}_i^{L/R} + \frac{1}{2}\frac{\Delta t}{\Delta x}\mathcal{A}_i^n(\mathbf{W}_i^L, \mathbf{W}_i^R). \tag{A.10}$$

3. **The FORCE flux**

The FORCE flux is replaced by terms compatible with the differencing. The term which is analogous to the Richtmyer velocity, $\mathbf{W}^{Ri}$, is given by

$$\mathbf{W}_i^{Ri} = \frac{1}{2}\left[(\mathbf{W}_i^R + \mathbf{W}_{i+1}^L) + \frac{\Delta x}{\Delta t}\mathcal{A}(\mathbf{W}_i^R, \mathbf{W}_{i+1}^L)\right]. \tag{A.11}$$

We also write intermediate left and right Lax-Friedrichs terms, $\mathcal{A}^{LF_L}$ and $\mathcal{A}^{LF_R}$ respectively, as

$$\mathcal{A}_{i+1/2}^{LF_L} = \mathcal{A}(\mathbf{W}_i^L, \mathbf{W}_{i+1}^L), \tag{A.12a}$$

$$\mathcal{A}_{i+1/2}^{LF_R} = \mathcal{A}(\mathbf{W}_{i-1}^R, \mathbf{W}_i^R). \tag{A.12b}$$

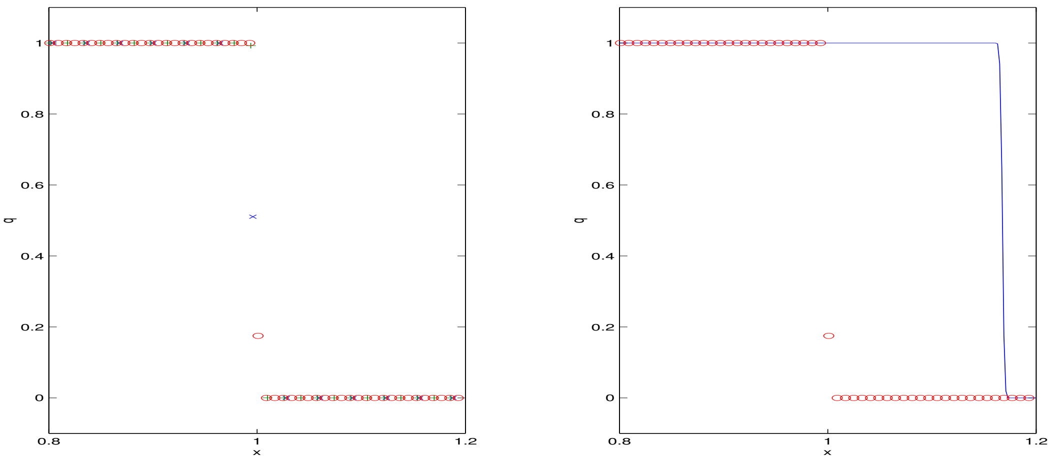

Figure A.1: PSLIC is used to evolve Burgers equation in primitive form. Results are shown for 250, 500 and 1000 points (only 1 in 4 points plotted) in the left plot. The result is clearly stable and convergent on a discontinuity at $x = 1$. However we are modelling the conservation law (A.15) which has a shock at $x = 7/6$, given by the line in the right hand plot. This convergent, stable numerical method produces an unphysical answer as it cannot enforce conservation.

We can then produce new data $\mathbf{W}_i^{n+1}$ at time $t + \Delta t$ using the update formula

$$\begin{aligned}\mathbf{W}_i^{n+1} \quad &= \quad \mathbf{W}_i^n + \frac{1}{2}\frac{\Delta t}{\Delta x}\left[\left(\mathcal{A}_{i+1/2}^{Ri} + \frac{1}{2}\left(\mathcal{A}_{i+1/2}^{LF_L} + \mathcal{A}_{i+1/2}^{LF_R}\right)\right)\right. \\ &\qquad \left. + \frac{1}{2}\frac{\Delta x}{\Delta t}\left(\left(\mathbf{W}_{i-1}^R - \mathbf{W}_i^R\right) - \left(\mathbf{W}_i^L - \mathbf{W}_{i+1}^L\right)\right)\right]. \qquad \text{(A.13)}\end{aligned}$$

This is found by explicitly writing out the FORCE flux terms in the standard update formula and replacing all terms of the conservative form $\mathbf{f}^L - \mathbf{f}^R$ with the primitive form (A.4).

A.1.2 Burgers equation in primitive form

As mentioned in section 2.2.2, a hyperbolic system in primitive form need not have a unique conserved form. This may be important for numerical calculations when discontinuities appear. Figure A.1 shows a stable, convergent numerical calculation that clearly converges to the wrong solution. Here, PSLIC was used to evolve Burgers equation in primitive variable form, equation (2.11), with the initial data

$$q = \begin{cases} 1 & x < 0.5 \\ 0 & x > 0.5 \end{cases} \tag{A.14}$$

on the domain $x \in [0, 2]$. The "correct" conserved form from with the exact solution was found was chosen to be the alternative form given in section 2.2.2,

$$\partial_t \left(\tfrac{1}{2} q^2\right) + \partial_x \left(\tfrac{1}{3} q^3\right) = 0. \tag{A.15}$$

Here an incorrect numerical solution was found because conservation was ignored. If this is considered to be a constraint on the system then convergence testing against this constraint would show the failure of the numerical method. However, sometimes the failings of a numerical method may be due to different underlying assumptions. In an example given by LeVeque (1998) of the evolution of a very strong Chapman-Jouguet detonation, the failure of the method to reproduce (even approximately) the correct solution is due to the way that the source terms are treated in the finite volume formulation. The material ignites once the temperature reaches a certain point. However this causes the material across an entire cell to ignite. Near the shock, the physical solution is to enforce that only the material heated above ignition temperature by the shock wave should ignite.

This emphasises that knowledge of the underlying physics is essential to compute physical solutions. Convergence testing is an essential tool, but to ensure physical answers we must test against all the possible constraints. If complete knowledge of the physics is not available, then it is best to implement as many different numerical methods as possible. In this case, results that can be reproduced by different numerical methods with different underlying assumptions are more believable. The areas where the methods differ will also give information about the physics, by looking at where the underlying physical assumptions of the methods differ.

A.2 Friedrich's equations

A.2.1 Equations

Einstein's equations coupled to a perfect fluid have been written in a symmetric hyperbolic form by Friedrich (1998). The characteristic speeds of this formulation are all physical, that is either the speed of light or the speed of sound of the fluid. An attempt was made to solve the problem of primordial black hole formation using this system of equations. Here the equations and methods used are outlined along with possible reasons why it failed.

The spherically symmetric line element is written in the form

$$\mathrm{d}s^2 = -n(t,r)^2\mathrm{d}t^2 + b(t,r)^2\mathrm{d}r^2 + c(t,r)^2\mathrm{d}\Omega^2. \tag{A.16}$$

Note that the coordinates (t, r) do not correspond to those used in chapter 5. In the Friedrich formulation it is assumed that observers are comoving with the fluid. A number of auxilliary quantities, some of them combinations of derivatives of the metric components, are taken to be the dependent variables. For this line element the system of evolution equations reduces to

$$\partial_t b = nbc_1 \tag{A.17a}$$

$$\partial_t c = ncc_2 \tag{A.17b}$$

$$\partial_t a - c_s^2 \frac{n}{b} \partial_r c_1 = n \left[a \left(c_s^2 (c_1 + 2c_2) - c_1 + 2c_s^2 g (c_1 - c_2) \right) \right] \tag{A.17c}$$

$$\partial_t c_1 - \frac{n}{b} \partial_r a = n \left[a^2 - w - c_1^2 + c_2^2 - \tfrac{2}{3}\rho \right] \tag{A.17d}$$

$$\partial_t c_2 = n \left[ag + \tfrac{1}{2}(w - 3c_2^2) - \tfrac{1}{6}\rho \right] \tag{A.17e}$$

$$\partial_t y = n(a - y)c_2 \tag{A.17f}$$

$$\partial_t \rho = -\tfrac{4}{3} n (c_1 + 2c_2)\rho \tag{A.17g}$$

where

$$w = y(y + 2/c) \tag{A.18}$$

and

$$g = y + 1/c. \tag{A.19}$$

The constraints reduce to

$$\partial_r n = abn \tag{A.20a}$$

$$\partial_r c = bgc \tag{A.20b}$$

$$\partial_r c_2 = (c_1 - c_2)bg \tag{A.20c}$$

$$\partial_r g = b\left\{c_1 c_2 + \tfrac{1}{2}(c_2^2 + c^{-2} - \rho - 3g^2)\right\} \tag{A.20d}$$

$$c_s^2 \partial_r \rho = -ab(1 + c_s^2)\rho. \tag{A.20e}$$

We are free to choose the density ρ and the metric functions b and c. The other auxilliary variables are then analytically soluble by quadratures.

During the evolution the lapse n is either found from equations (A.20a) and (A.20e) and a suitable boundary condition, or by the evolution equation

$$\partial_t n = n\left(\tfrac{1}{3}n(c_1 + 2c_2) - \frac{1}{2t}\right). \tag{A.21}$$

The evolution equations containing spatial derivatives (A.17d)-(A.17e) are not in conservation form. A conservation form can be found, which is

$$\partial_t(ba) - c_s^2\partial_r(nc_1) = 2c_s^2 nb\left[ac_2 + (c_1 - c_2)g\right] \tag{A.22a}$$

$$\partial_t(bc_1) - \partial_r(na) = nb\left[c_2^2 - w - \frac{2}{3}\rho\right]. \tag{A.22b}$$

It is not clear that this conservation form has any physical meaning, and so shocks may not be evolved correctly if they appear.

A.2.2 Evolution of density fluctuations

In this coordinate system the spatially flat radiation dominated FRW spacetime has the form

$$\begin{aligned}
\rho_{FRW} &= \frac{3}{4t^2}, n_{FRW} = 1, b_{FRW} = t^{1/2}, c_{FRW} = rt^{1/2} \qquad &\text{(A.23a)}\\
(c_1)_{FRW} &= (c_2)_{FRW} = \frac{1}{2t} &\text{(A.23b)}\\
y_{FRW} &= a_{FRW} = 0. &\text{(A.23c)}
\end{aligned}$$

Density fluctuations identical to those of Niemeyer and Jedamzik (1999a) were used with the assumption that the metric variables b and c would initially be unchanged.

A number of different numerical methods were implemented to evolve this system of equations. These included SLIC and the slope limited method of lines approach outlined in chapter 3. As the characteristic speeds $c_s n/b$ are constant up to a metric factor and this metric factor is continuous an "exact" Riemann solver was implemented using the form for a linear system as given in chapter 2. With the exception of the method of lines schemes, the evolution of the source terms used Strang splitting with either second or fourth order Runge–Kutta integrators, or a stiff semi-implicit extrapolation method of Bader and Deuflhard. These methods are given by Press and Flannery (1989).

Evolutions of the FRW spacetime were correct to machine accuracy until the density ρ dropped beneath the truncation error. This is because in these coordinates we are just solving a simple system of ordinary differential equations. For weak fluctuations evolutions of thousands of units of coordinate time were performed with no stability problems and with the constraint violations convergent and small (on the order of 10^{-5}). However for larger fluctuations discontinuities were starting to form in the density. The evolutions continued until the discontinuity reached the origin; on reflecting from the origin oscillations appeared which quickly caused the evolution to fail.

A number of different boundary conditions and numerical methods were implemented to try and alleviate this problem. None succeeded. In many ways this is not surprising. Firstly we are evolving a which is proportional to the spatial derivative of the density; in other words we are taking the second derivative of a discontinuous function. Secondly we note that the lapse will also become discontinuous when the density does, leading to an unphysical coordinate discontinuity. Finally we note that although we have a conservation form for our equations, (A.22), we do not know if it is physical or not so should not expect to be able to evolve a shock correctly.

A.3 A null formulation

The equations used in chapter 5 use a timelike coordinate t. The slice stretching needed to stop the singularity of the black hole forming on a computational slice seemed to cause numerical problems near the apparent horizon when AMR was used, as explained in section 4.4. As an alternative an outgoing null coordinate was considered. As such a coordinate cannot penetrate the apparent horizon it was thought that these numerical problems might be avoided.

However early attempts to implement this approach showed that while it worked on the FRW spacetime, once a perturbation was introduced it was not as useful. This is because at a fast moving shock such as is shown in section 5.9.1 the characteristic velocities need not be restricted to $(-1,1)$ but could diverge, leading to arbitrarily small timesteps. A few attempts were made to implement *implicit* methods in order to circumvent this restriction but these failed even on simple test problems.

Instead an adaptive coordinate was considered. This coordinate would be outgoing null at the origin to avoid the problems with singularity formation, but timelike further away. Explicitly the metric is

$$\mathrm{d}s^2 = -\alpha^2(T,r)(1-h(T,r))\mathrm{d}T^2+2h(T,r)\mathrm{d}T\mathrm{d}r+a^2(T,r)(1+h(T,r))\mathrm{d}r^2+r^2\mathrm{d}\Omega^2. \tag{A.24}$$

In the limit $h=0$ we recover the metric used in chapter 5. In the limit $h=1$ we recover the outgoing null form. Although the equations were written for a general function h depending on both time and space, here we only consider functions of r.

The resulting equations could be expressed in conservation law form with conserved quantities, fluxes and source terms

$$\mathbf{q} = (\sigma + Q(1+h), \sigma - Q(1-h))^T, \tag{A.25a}$$

$$\mathbf{f} = (Q + Y(1+h), Q - Y(1-h))^T, \tag{A.25b}$$

$$\mathbf{s} = (-\alpha a F, \alpha a F)^T, \tag{A.25c}$$

where

$$\begin{aligned}
\sigma &= 4pW^2(1-hv)^2 - (1-h^2)p, && \text{(A.26a)}\\
Q &= 4pW^2 v(1-hv) - hp, && \text{(A.26b)}\\
Y &= \frac{(Q+hp)^2}{\sigma + (1-h^2)p} + p, && \text{(A.26c)}\\
F &= 8\pi r(\sigma Y - Q^2) + (\sigma + 2hQ + (h^2-1)Y)\frac{m}{r^2}. && \text{(A.26d)}
\end{aligned}$$

The metric variable equations are

$$\begin{aligned}
\partial_t a &= -4\pi r \alpha a^2 (Q + hY), && \text{(A.27a)}\\
\partial_r a &= a^3 \left(4\pi r(hQ+\sigma) - \frac{m}{r^2}\right), && \text{(A.27b)}\\
\partial_r \log(\alpha) &= a^2(4\pi r(Y(1-h^2) - hQ) + \frac{m}{r^2}). && \text{(A.27c)}
\end{aligned}$$

A number of different functions h were tested, mainly combinations of exponentials such as

$$h(r) = \tfrac{1}{2}\left(1 - \tanh(A(r - r_0))\right), \qquad \text{(A.28)}$$

or high (sixth or seventh) order polynomials. These functions were enforced to be 1 at the origin and 0 at $r \geqslant 1$ with continuous derivatives and decaying rapidly. These conditions require using the polynomials only where $r \in [0,1]$. In practise the differences between functions were small.

This method was not used for the problem in chapter 5 for a number of reasons. Firstly, it was difficult to find a function h that worked well both on the background FRW spacetime and for perturbations. If the function decayed sufficiently quickly so that the shock wave formed in the region covered by the timelike coordinate where the characteristic speeds were bounded, then accuracy was lost in the region where h varied. Sufficient accuracy could only be recovered by increasing the resolution substantially. Secondly, there were problems at the origin when a perturbation was introduced. This problem appeared to be some sort of instability caused by matter reflecting from the origin. We believe that this is due to the boundaries of computational cells moving at the speed of light, faster than the characteristic speed of the matter. For collapse problems where the velocities are directed inwards a retarded-time null coordinate should always be satisfactory, e.g., Baumgarte et al. (1995), Font and Papadopoulos (2000). For expanding problems where the fluid velocities are directed outwards, an advanced-time null coordinate should be equally satisfactory. Here though, we have both incoming and outgoing flows and so no global null coordinate will work well.

A.4 Convergence tests

This section contains the remaining convergence tables for the tests not included in section 3.10. As all these problems were hydrodynamical the error was measured on the density. This is because the velocity, and hence some of the conserved variables, may cause problems due to being very near zero over large parts of the domain. Also, for the ideal gas cases the density is more sensitive than the pressure as the density will change across the linear (that is, contact and tangential) discontinuities.

For all tests and methods the convergence is poor. This is to be expected with a shock wave in the solution.

Method	Cells	L_1 error	L_1 order	L_2 error	L_2 order	L_∞ error	L_∞ order
SLIC	100	20	–	3.5	–	1.5	–
	200	11	0.91	1.4	1.3	0.65	1.2
	400	5.6	0.94	0.59	1.3	0.25	1.4
WAF	100	9.7	–	2.0	–	1.1	–
	200	5.1	0.92	0.83	1.3	0.47	1.3
	400	2.6	0.95	0.34	1.3	0.18	1.4
Slope-LF	100	25	–	4.1	–	1.8	–
	200	14	0.87	1.7	1.2	0.73	1.3
	400	7.3	0.90	0.72	1.3	0.27	1.4
SLOPE-Roe	100	21	–	3.8	–	1.7	–
	200	12	0.85	1.6	1.2	0.69	1.3
	400	6.3	0.89	0.67	1.3	0.26	1.4
SLOPE-Marquina	100	23	–	4.0	–	1.8	–
	200	13	0.86	1.7	1.2	0.72	1.3
	400	6.9	0.90	0.71	1.3	0.27	1.4
SLOPE-Exact	100	20	–	3.7	–	1.7	–
	200	11	0.84	1.6	1.2	0.68	1.3
	400	6.1	0.89	0.67	1.3	0.26	1.4
ENO3-LF	100	14	–	2.6	–	1.4	–
	200	7.6	0.93	1.0	1.3	0.54	1.3
	400	3.9	0.95	0.43	1.3	0.20	1.4
ENO3-Exact	100	12	–	2.3	–	1.2	–
	200	6.1	0.92	0.97	1.3	0.51	1.3
	400	3.2	0.95	0.40	1.3	0.19	1.4
ENO4-LF	100	12	–	2.1	–	1.2	–
	200	6.0	0.94	0.87	1.3	0.47	1.3
	400	3.3	0.89	0.37	1.2	0.18	1.4
ENO4-Exact	100	9.3	–	1.9	–	1.1	–
	200	4.8	0.95	0.79	1.3	0.44	1.3
	400	2.5	0.93	0.33	1.3	0.17	1.4
ENO5-LF	100	10	–	1.9	–	1.1	–
	200	5.4	0.90	0.81	1.2	0.46	1.3
	400	4.3	0.36	0.47	0.78	0.18	1.3
ENO5-Exact	50	15	–	4.3	–	2.7	–
	100	8.3	0.89	1.8	1.3	1.0	1.4
	200	4.4	0.92	0.72	1.3	0.41	1.3

Table A.1: Convergence tests for the one dimensional ultrarelativistic hydrodynamics test. Convergence is not second order for any method because of the appearance of the shock and the nonlinearity of the problem. ENO5-Exact uses a different number of cells as the 400 cell test failed to complete. Errors are measure in the density only, to avoid spurious results as the velocity is near zero.

Method	Cells	L_1 error	L_1 order	L_2 error	L_2 order	L_∞ error	L_∞ order
SLIC	100	1.1e-2	–	2.1e-3	–	1.0e-3	–
	200	6.6e-3	0.78	1.1e-3	0.97	4.6e-4	1.2
	400	4.4e-3	0.60	6.8e-4	0.68	3.0e-4	0.62
WAF	100	5.8e-3	–	1.3e-3	–	7.2e-4	–
	200	3.5e-3	0.74	7.6e-4	0.80	4.3e-4	0.75
	400	2.6e-3	0.43	5.7e-4	0.41	3.5e-4	0.29
Slope-LF	100	1.3e-2	–	2.3e-3	–	9.9e-4	–
	200	7.9e-3	0.75	1.2e-3	0.92	4.8e-4	1.0
	400	5.1e-3	0.63	7.5e-4	0.71	3.2e-4	0.57
SLOPE-Roe	100	1.2e-2	–	2.1e-3	–	9.1e-4	–
	200	7.1e-3	0.74	1.1e-3	0.90	4.8e-4	0.93
	400	4.6e-3	0.62	7.1e-4	0.67	3.4e-4	0.50
SLOPE-Marquina	100	1.3e-2	–	2.2e-3	–	1.0e-3	–
	200	7.5e-3	0.75	1.2e-3	0.93	4.5e-4	1.2
	400	4.9e-3	0.62	7.2e-4	0.70	3.2e-4	0.49
SLOPE-Exact	100	1.2e-2	–	2.1e-3	–	9.1e-4	–
	200	7.0e-3	0.74	1.1e-3	0.89	4.8e-4	0.92
	400	4.6e-3	0.61	7.1e-4	0.66	3.4e-4	0.51
ENO3-LF	100	9.1e-3	–	1.8e-3	–	9.1e-4	–
	200	5.2e-3	0.82	9.4e-4	0.92	4.4e-4	1.0
	400	3.4e-3	0.58	6.2e-4	0.60	3.3e-4	0.44
ENO3-Exact	100	7.8e-3	–	1.6e-3	–	8.2e-4	–
	200	4.5e-3	0.79	8.5e-4	0.88	4.3e-4	0.94
	400	3.1e-3	0.54	5.9e-4	0.53	3.4e-4	0.33
ENO4-LF	100	7.5e-3	–	1.5e-3	–	8.1e-4	–
	200	4.2e-3	0.83	8.3e-4	0.90	4.2e-4	0.95
	400	2.9e-3	0.55	5.7e-4	0.53	3.3e-4	0.33
ENO4-Exact	100	6.5e-3	–	1.4e-3	–	7.3e-4	–
	200	3.8e-3	0.79	7.7e-4	0.83	4.0e-4	0.86
	400	2.7e-3	0.50	5.6e-4	0.46	3.4e-4	0.23
ENO5-LF	100	6.6e-3	–	1.4e-3	–	7.5e-4	–
	200	3.8e-3	0.79	7.8e-4	0.85	4.1e-4	0.89
	400	2.7e-3	0.50	5.6e-4	0.49	3.3e-4	0.29
ENO5-Exact	100	6.1e-3	–	1.3e-3	–	6.9e-4	–
	200	3.6e-3	0.77	7.3e-4	0.81	4.0e-4	0.78
	400	2.6e-3	0.44	5.5e-4	0.42	3.5e-4	0.20

Table A.2: Convergence tests for the one dimensional ideal gas hydrodynamics test. Convergence is not second order for any method because of the appearance of the shock and the nonlinearity of the problem. Errors measured on the density.

Method	Cells	L_1 error	L_1 order	L_2 error	L_2 order	L_∞ error	L_∞ order
SLIC	100	0.20	–	8.1e-2	–	6.4e-2	–
	200	0.17	0.20	5.9e-2	0.46	3.9e-2	0.73
	400	0.15	0.16	4.5e-2	0.39	2.0e-2	0.96
Slope-LF	100	0.22	–	8.7e-2	–	6.9e-2	–
	200	0.20	0.14	6.6e-2	0.40	4.5e-2	0.64
	400	0.19	0.10	5.1e-2	0.36	2.2e-2	1.0
SLOPE-Roe	100	0.22	–	8.7e-2	–	7.0e-2	–
	200	0.21	0.12	6.7e-2	0.39	4.5e-2	0.64
	400	0.19	0.11	5.2e-2	0.35	2.2e-2	1.0
SLOPE-Marquina	100	0.23	–	8.8e-2	–	7.0e-2	–
	200	0.21	0.15	6.6e-2	0.41	4.5e-2	0.66
	400	0.19	0.097	5.2e-2	0.34	2.2e-2	1.0
SLOPE-Exact	100	0.22	–	8.7e-2	–	7.0e-2	–
	200	0.20	0.13	6.6e-2	0.40	4.5e-2	0.64
	400	0.19	0.095	5.2e-2	0.34	2.2e-2	1.0
ENO3-LF	100	0.18	–	7.7e-2	–	6.3e-2	–
	200	0.15	0.22	5.4e-2	0.49	3.7e-2	0.76
	400	0.15	0.051	4.3e-2	0.34	1.9e-2	0.93
ENO3-Exact	100	0.18	–	7.9e-2	–	6.4e-2	–
	200	0.15	0.24	5.5e-2	0.50	3.7e-2	0.77
	400	0.15	0.067	4.3e-2	0.35	2.0e-2	0.93
ENO4-LF	100	0.15	–	7.3e-2	–	6.0e-2	–
	200	0.14	0.13	5.4e-2	0.45	3.6e-2	0.73
	400	0.14	-0.021	4.3e-2	0.33	1.9e-2	0.90
ENO4-Exact	100	0.16	–	7.5e-2	–	6.1e-2	–
	200	0.14	0.18	5.4e-2	0.48	3.6e-2	0.76
	400	0.14	-0.027	4.3e-2	0.34	2.0e-2	0.88
ENO5-LF	50	0.20	–	9.6e-2	–	7.9e-2	–
	100	0.15	0.45	7.0e-2	0.45	5.7e-2	0.47
	200	0.13	0.11	5.2e-2	0.44	3.5e-2	0.70

Table A.3: Convergence tests for the one dimensional ideal gas blast wave test. Convergence is not second order for any method because of the appearance of the shock and the nonlinearity of the problem. The error is measured on the density. ENO5 uses different resolutions due to problems with the higher resolution evolution.

Method	Cells	L_1 error	L_1 order	L_2 error	L_2 order	L_∞ error	L_∞ order
WAF	50	9.4	–	2.4	–	1.5	–
	100	5.1	0.87	1.0	1.3	0.55	1.4
	200	3.0	0.76	0.43	1.2	0.22	1.3
Slope-LF	50	26	–	5.6	–	2.8	–
	100	15	0.84	2.4	1.2	1.0	1.4
	200	8.6	0.77	1.0	1.2	0.40	1.4
SLOPE-Roe	50	23	–	5.3	–	2.7	–
	100	13	0.84	2.3	1.2	1.0	1.4
	200	7.4	0.78	0.97	1.2	0.39	1.4
SLOPE-Marquina	50	24	–	5.5	–	2.8	–
	100	13	0.84	2.4	1.2	1.0	1.4
	200	7.7	0.81	1.0	1.2	0.40	1.4
ENO2-LF	50	26	–	5.4	–	2.8	–
	100	14	0.86	2.3	1.2	1.0	1.4
	200	7.7	0.88	0.99	1.2	0.40	1.4
ENO2-Exact	50	21	–	4.9	–	2.4	–
	100	12	0.81	2.1	1.2	0.94	1.4
	200	6.7	0.81	0.91	1.2	0.37	1.4
ENO3-LF	50	17	–	3.7	–	2.2	–
	100	8.7	0.94	1.5	1.3	0.79	1.5
	200	4.6	0.90	0.63	1.3	0.30	1.4
ENO3-Exact	50	13	–	3.4	–	1.9	–
	100	7.3	0.87	1.4	1.3	0.72	1.4
	200	4.0	0.86	0.58	1.3	0.27	1.4
ENO4-LF	50	13	–	3.1	–	1.9	–
	100	6.9	0.95	1.3	1.3	0.67	1.5
	200	3.7	0.89	0.52	1.3	0.26	1.4
ENO4-Exact	50	11	–	2.9	–	1.7	–
	100	5.9	0.88	1.2	1.3	0.61	1.5
	200	3.3	0.86	0.47	1.3	0.23	1.4
ENO5-LF	50	12	–	2.9	–	1.9	–
	100	6.3	0.95	1.2	1.3	0.62	1.6
	200	3.4	0.88	0.47	1.3	0.24	1.4
ENO5-Exact	50	9.7	–	2.6	–	1.6	–
	100	5.3	0.88	1.1	1.3	0.56	1.5
	200	2.9	0.85	0.43	1.3	0.22	1.4

Table A.4: Convergence tests for the ultrarelativistic hydrodynamics test in more than one dimension. Convergence is not second order for any method because of the appearance of the shock and the nonlinearity of the problem. The error is measured on the density.

Method	Cells	L_1 error	L_1 order	L_2 error	L_2 order	L_∞ error	L_∞ order
SLIC	50	2.0e-2	–	4.4e-3	–	1.5e-3	–
	100	1.7e-2	0.22	3.1e-3	0.50	1.1e-3	0.48
	200	1.5e-2	0.21	2.2e-3	0.52	5.0e-4	1.1
Slope-LF	50	2.4e-2	–	5.1e-3	–	1.7e-3	–
	100	1.8e-2	0.42	3.3e-3	0.65	1.1e-3	0.7
	200	1.6e-2	0.21	2.2e-3	0.54	5.2e-4	1.1
SLOPE-Roe	50	2.0e-2	–	4.5e-3	–	1.5e-3	–
	100	1.7e-2	0.24	3.1e-3	0.54	9.9e-4	0.59
	200	1.5e-2	0.16	2.2e-3	0.50	4.7e-4	1.1
SLOPE-Marquina	50	2.2e-2	–	4.9e-3	–	1.8e-3	–
	100	1.7e-2	0.35	3.2e-3	0.62	1.1e-3	0.71
	200	1.5e-2	0.16	2.2e-3	0.52	5.2e-4	1.1
ENO3-LF	50	2.0e-2	–	4.5e-3	–	1.5e-3	–
	100	1.7e-2	0.27	3.2e-3	0.52	9.8e-4	0.65
	200	1.5e-2	0.16	2.2e-3	0.50	4.8e-4	1.0
ENO3-Marquina	50	1.9e-2	–	4.4e-3	–	1.5e-3	–
	100	1.6e-2	0.24	3.1e-3	0.49	9.8e-4	0.62
	200	1.4e-2	0.14	2.2e-3	0.49	4.8e-4	1.0
ENO4-LF	50	1.8e-2	–	4.3e-3	–	1.6e-3	–
	100	1.6e-2	0.22	3.1e-3	0.48	9.0e-4	0.83
	200	1.4e-2	0.13	2.2e-3	0.49	4.8e-4	0.90
ENO4-Marquina	50	1.7e-2	–	4.2e-3	–	1.6e-3	–
	100	1.5e-2	0.19	3.1e-3	0.45	9.0e-4	0.81
	200	1.4e-2	0.10	2.2e-3	0.48	4.8e-4	0.91
ENO5-LF	50	1.7e-2	–	4.2e-3	–	1.6e-3	–
	100	1.5e-2	0.20	3.1e-3	0.46	9.1e-4	0.82
	200	1.4e-2	0.12	2.2e-3	0.49	4.9e-4	0.91
ENO5-Marquina	50	1.6e-2	–	4.2e-3	–	1.6e-3	–
	100	1.5e-2	0.18	3.1e-3	0.45	9.2e-4	0.81
	200	1.4e-2	0.08	2.2e-3	0.48	4.9e-4	0.92

Table A.5: Convergence tests for the ideal gas hydrodynamics test in more than one dimension. Convergence is not second order for any method because of the appearance of the shock and the nonlinearity of the problem. The error is measured on the density.

Bibliography

A. M. Abrahams and C. R. Evans. Critical behavior and scaling in vacuum axisymmetric gravitational collapse. *Phys. Rev. Lett.*, **70**:2980–2983, 1993.

C. Alcock, R. A. Allsman, D. Alves, T. S. Axelrod, A. C. Becker, D. P. Bennett, K. H. Cook, K. C. Freeman, K. Griest, J. Guern, M. J. Lehner, S. L. Marshall, B. A. Peterson, M. R. Pratt, P. J. Quinn, A. W. Rodgers, C. W. Stubbs, and W. Sutherland. The MACHO Project: Limits on planetary mass dark matter in the galactic halo from gravitational microlensing. *Astrophys. J.*, **471**:774–782, 1996. (MACHO Collaboration).

A. S. Almgren, J. B. Bell, P. Colella, L. H. Howell, and M. L. Welcome. A conservative adaptive projection method for the variable density incompressible Navier-Stokes equations. *J. Comput. Phys.*, **142**:1–46, 1998.

M. A. Aloy, J. A. Pons, and J. M. Ibáñez. An efficient implementation of flux formulae in multidimensional relativistic hydrodynamical codes. *Comput. Phys. Commun.*, **120**:115–121, 1999.

T. W. Baumgarte, S. L. Shapiro, and S. A. Teukolsky. Computing supernova collapse to neutron stars and black holes. *Astrophys. J.*, **443**:717–734, 1995.

J. Bell, M. Berger, J. Saltzman, and M. Welcome. Three-dimensional Adaptive Mesh Refinement for hyperbolic conservation laws. *SIAM J. Sci. Comput.*, **15**:127–238, 1994.

M. J. Berger and J. Oliger. Adaptive Mesh Refinement for hyperbolic partial differential equations. *J. Comput. Phys.*, **53**:484–512, 1984.

G. V. Bicknell and R. N. Henriksen. Formation of primordial black holes. *Astrophys. J.*, **232**:670–682, 1979.

P. R. Brady, C. M. Chambers, and S. M. C. V. Gonçalves. Phases of massive scalar field collapse. *Phys. Rev.*, **D56**:6057–6061, 1997.

D. Brown, G. Cheshire, W. Henshaw, and D. Quinlan. Overture: An object-oriented software system for solving partial differential equations in serial and parallel environments. In *Proceedings of the SIAM Parallel Conference*. Minneapolis, 1997.

B. Brügmann. Adaptive mesh and geodesically sliced Schwarzschild spacetime in 3+1 dimensions. *Phys. Rev.*, **D54**:7361–7372, 1996.

Cactus. Cactus code. http://www.cactuscode.org, 2001.

A. H. Cahill and M. E. Taub. Spherically symmetric similarity solutions of the Einstein field equations for a perfect fluid,. *Commun. Math. Phys.*, **21**: 1, 1971.

B. J. Carr. The primordial black hole mass spectrum. *Astrophys. J.*, **201**: 1–19, 1975.

B. J. Carr and A. A. Coley. A complete classification of spherically symmetric perfect fluid similarity solutions. *Phys. Rev.*, **D62**:044023, 2000.

B. J. Carr, A. A. Coley, M. Goliath, U. S. Nilsson, and C. Uggla. Critical phenomena and a new class of self-similar spherically symmetric perfect-fluid solutions. *Phys. Rev.*, **D61**:081502, 2000.

B. J. Carr and S. W. Hawking. Black holes in the early universe. *Mon. Not. R. astr. Soc.*, **168**:399–416, 1974.

M. W. Choptuik. *A Study of Numerical Techniques for Radiative Problems in General Relativity.* PhD thesis, University of British Columbia, 1986.

M. W. Choptuik. Consistency of Finite-Difference Solutions of Einstein's Equations. *Phys. Rev.*, **D44**:3124–3135, 1991.

M. W. Choptuik. "Critical" Behaviour in Massless Scalar Field Collapse. In R. D'Inverno, editor, *Approaches to Numerical Relativity.* Cambridge University Press, 1992.

M. W. Choptuik, T. Chmaj, and P. Bizon. Critical behavior in gravitational collapse of a Yang-Mills field. *Phys. Rev. Lett.*, **77**:424, 1996.

D. Christodoulou. A mathematical theory of gravitational collapse. *Commun. Math. Phys.*, **109**:613–647, 1987.

D. Christodoulou. The formation of black holes and singularities in spherically symmetric gravitational collapse. *Commun. Pure App. Math.*, **44**: 339–373, 1991.

D. Christodoulou. Bounded variation solutions of the spherically symmetric Einstein-scalar field equations. *Commun. Pure App. Math.*, **46**:1131–1220, 1993.

W. Dai and P.R . Woodward. An iterative Riemann solver for relativistic hydrodynamics. *SIAM J. Sci. Comput.*, **18**:982–995, 1997.

A. Dolezal and S. S. Wong. Relativistic hydrodynamics and essentially non-osciallatory shock capturing schemes. *J. Comput. Phys.*, **120**:266–277, 1995.

R. Donat, J. A. Font, J. M. Ibáñez, and A. Marquina. A flux-split algorithm applied to relativistic flows. *J. Comput. Phys.*, **146**:58–81, 1998.

R. Donat and A. Marquina. Capturing shock reflections: An improved flux formula. *J. Comput. Phys.*, **125**:42–58, 1996.

B. Einfeldt. On Godunov-type methods for gas dynamics. *SIAM J. Numer. Anal.*, **25**:294–318, 1988.

G. F. R. Ellis and M. Bruni. Covariant and gauge-invariant approach to cosmological density fluctuations. *Phys. Rev.*, **40**:1804–1818, 1989.

F. Eulderink and G. Mellema. General relativistic hydrodynamics with a Roe solver. *Astron. Astrophys. Suppl.*, **110**:587–623, 1995.

C. R. Evans and C. S. Coleman. Critical phenomena and self-similarity in the gravitational collapse of radiation fluid. *Phys. Rev. Lett.*, **72**:1782, 1994.

R. Fedkiw, B. Merriman, R. Donat, and S. Osher. The penultimate scheme for systems of conservation laws: Finite difference ENO with Marquina's flux splitting. In M. Hafez, editor, *Progress in Numerical Solutions of Partial Differential Equations.* Arcachon, France, 1998.

J. A. Font and P. Papadopoulos. Relativistic hydrodynamics on spacelike and null surfaces: Formalism and computations of spherically symmetric spacetimes. *Phys. Rev.*, **D61**:024015, 2000.

J.A. Font, N. Stergioulas, and K. D. Kokkotas. Non-linear hydrodynamical evolution of rotating relativistic stars: Numerical methods and code tests. *Mon. Not. R. astr. Soc.*, **313**:678–688., 2000.

H. Friedrich. Evolution equations for gravitating ideal fluid bodies in general relativity. *Phys. Rev.*, **D57**:2317, 1998.

B. Fryxell, K. Olson, P. Ricker, F. Timmes, M. Zingale, D. Lamb, P. MacNeice, R. Rosner, and H. Tufo. FLASH: An Adaptive Mesh hydrodynamics code for modeling astrophysical thermonuclear flashes. *Astrophys. J. Suppl.*, **131**:273–334, 2000.

D. Garfinkle. Choptuik scaling in null coordinates. *Phys. Rev.*, **D51**:5558–5561, 1995.

D. Garfinkle and G. C. Duncan. Scaling of curvature in sub-critical gravitational collapse. *Phys. Rev.*, **D58**:064024, 1998.

S. Godunov. Finite difference methods for numerical computation of discontinuous solutions of the equations of fluid dynamics. *Mat. Sbornik*, **47**: 3–271, 1959.

S. Gottlieb and C. Shu. Total Variation Diminishing Runge-Kutta schemes. *Math. Comp.*, **67**:73–85, 1998.

A. M. Green and A. R. Liddle. Critical collapse and the primordial black hole initial mass function. *Phys. Rev.*, **D60**:063509, 1999.

C. Gundlach. Nonspherical perturbations of critical collapse and cosmic censorship. *Phys. Rev.*, **D57**:7075–7079, 1998.

C. Gundlach. Critical phenomena in gravitational collapse. *Living Rev. Relativity*, **4**, 1999. [Article in online journal], cited on 31/07/01, http://www.livingreviews.org/Articles/Volume2/1999-4gundlach/index.html.

R. S. Hamadé, J. H. Horne, and J. M. Stewart. Continuous self-similarity and S-duality. *Class. Quantum Grav.*, **13**:2241–2254, 1996.

R. S. Hamadé and J. M. Stewart. The spherically symmetric collapse of a massive scalar field. *Class. Quantum Grav.*, **13**:497–512, 1996.

E. R. Harrison. Fluctuations at the threshold of classical cosmology. *Phys. Rev.*, **D1**:2726–2730, 1970.

A. Harten, B. Engquist, S. Osher, and S. R. Chakravarthy. Uniformly high order accurate essentially non-oscillatory schemes, III. *J. Comput. Phys.*, **71**:231–303, 1987.

S. D. Hern. *Numerical Relativity and Inhomogeneous Cosmologies*. PhD thesis, University of Cambridge, 1999.

S. D. Hern and J. M. Stewart. The Gowdy T^3 cosmologies revisited. *Class. Quantum Grav.*, **15**:1581, 1998.

V. Husain and M. Olivier. Scalar field collapse in three-dimensional AdS spacetime. *Class. Quantum Grav.*, **18**:L1–10, 2001.

G. Jiang and C. Shu. Efficient implementation of Weighted ENO schemes. *J. Comput. Phys.*, **126**:202–228, 1996.

T. Koike, T. Hara, and S. Adachi. Critical behaviour in gravitational collapse of radiation fluid - A renormalization group (linear perturbation) analysis. *Phys. Rev. Lett.*, **74**:5170, 1995.

T. Koike, T. Hara, and S. Adachi. Critical behavior in gravitational collapse of a perfect fluid. *Phys. Rev.*, **D59**:104008, 1999.

C. B. Laney. *Computational Gasdynamics.* Cambridge University Press, 1998.

L. Lehner. Numerical relativity: a review. *Class. Quantum Grav.*, **18**:R25–86, 2001.

R. J. LeVeque. *Numerical Methods for Conservation Laws.* Birkhäuser, 1992.

R. J. LeVeque. Nonlinear conservation laws and finite volume methods for astrophysical fluid flow. In O. Steiner and A. Gautschy, editors, *Computational Methods for Astrophysical Fluid Flow.* Springer-Verlag, 1998.

R. Löhner. An adaptive finite element scheme for transient problems in CFD. *Comp. Meth. App. Mech. Eng.*, **61**:323–338, 1987.

P. MacNeice, K. M. Olson, C. Mobarry, R. deFainchtein, and C. Packer. PARAMESH: A parallel adaptive mesh refinement community toolkit. *Computer Physics Comm.*, **126**:330–354, 2000.

D. Maison. Non-universality of critical behaviour in spherically symmetric gravitational collapse. *Phys. Rev.*, **B366**:82, 1996.

J. M. Martí, J. M. Ibáñez, and J. A. Miralles. Numerical relativistic hydrodynamics: Local characteristic approach. *Phys. Rev.*, **D43**:3794–3801, 1991.

J. M. Martí and E. Müller. The analytical solution of the Riemann problem in relativistic hydrodynamics. *J. Fluid Mech.*, **258**:317–333, 1994.

J. M. Martí and E. Müller. Numerical hydrodynamics in Special Relativity. *Living Rev. Relativity*, **2**, 1999. [Article in online journal], cited on 31/7/01, http://www.livingreviews.org/Articles/Volume2/1999-3marti/index.html.

J. M. Martín-García and C. Gundlach. All nonspherical perturbations of the Choptuik spacetime decay. *Phys. Rev.*, **D59**:064031, 1999.

D. K. Nadezhin, I. D. Novikov, and A. G. Polnarev. The hydrodynamics of primordial black hole formation. *Soviet Astr.*, **22**:129–138, 1978.

D. W. Neilsen. *Extremely Relativistic Fluids in Strong-Field General Relativity.* PhD thesis, University of Texas (Austin), 1999.

D. W. Neilsen and M. W. Choptuik. Ultrarelativistic fluid dynamics. *Class. Quantum Grav.*, **17**:733–759, 2000.

J. C. Niemeyer and K. Jedamzik. Near-critical gravitational collapse and the initial mass function of primordial black holes. *Phys. Rev. Lett.*, **80**: 5481–4, 1998.

J. C. Niemeyer and K. Jedamzik. Dynamics of primordial black hole formation. *Phys. Rev.*, **D59**:124013, 1999a.

J. C. Niemeyer and K. Jedamzik. Primordial black hole formation during first-order phase transitions. *Phys. Rev.*, **D59**:124014, 1999b.

M. L. Norman and K. H. A. Winkler. Why ultrarelativistic numerical hydrodynamics is difficult. In M. L. Norman and K. H. A. Winkler, editors, *Astrophysical Radiation Hydrodynamics*. Dordrecht:Reidel, 1986.

I. Olabarrieta and M. W. Choptuik. Critical phenomena at the threshold of black hole formation for collisionless matter in spherical symmetry. (preprint `gr-qc 0107076`), 2001.

P. Papadopoulos, E. Seidel, and L. Wild. Adaptive computation of gravitational waves from black hole interactions. *Phys. Rev.*, **D58**:084002, 1998.

T. Plewa. Adaptive Mesh Refinement for structured grids. Online resource at http://www.camk.edu.pl/ tomek/AMRA/amr.html, 2001.

J. A. Pons, J. A. Font, J. M. Ibáñez, J. M. Martí, and J.A . Miralles. General relativistic hydrodynamics with special relativistic Riemann solvers. *Astron. Astrophys.*, **339**:638–642, 1998.

J. A. Pons, J. M. Martí, and E. Müller. The exact solution of the Riemann problem with non-zero tangential velocities in relativistic hydrodynamics. *J. Fluid Mech.*, **422**:125–139, 2000.

W. H. Press and B. P. Flannery. *Numerical Recipes in Pascal: The Art of Scientific Computing.* Cambridge University Press, 1989.

F. Pretorius and M. W. Choptuik. Gravitational collapse in 2+1 AdS spacetime. *Phys. Rev.*, **D62**:124012, 2000.

J. J. Quirk. A contribution to the great Riemann solver debate. *Int. J. Numer. Methods Fluids*, **18**:555–574, 1994.

C. Renault, C. Afonso, E. Aubourg, P. Bareyre, F. Bauer, S. Brehin, C. Coutures, C. Gaucherel, J. F. Glicenstein, B. Goldman, M. Gros, D. Hardin, J. de Kat, M. Lachieze-Rey, B. Laurent, E. Lesquoy, C. Magneville, A. Milsztajn, L. Moscoso, N. Palanque-Delabrouille, F. Queinnec, J. Rich, M. Spiro, L. Vigroux, S. Zylberajch, R. Ansari, F. Cavalier, F. Couchot, B. Mansoux, M. Moniez, O. Perdereau, J.-P. Beaulieu, R. Ferlet, P. Grison, A. Vidal-Madjar, J. Guibert, O. Moreau, E. Maurice, L. Prevot, C. Gry, S. Char, and J. Fernandez. Observational limits on MACHOs in the Galactic Halo. *Astron. Astrophys.*, **324**:L69, 1997. (EROS Collaboration).

L. Rezzolla, A. M. Abrahams, R. A. Matzner, M. E. Rupright, and S. L. Shapiro. Cauchy-perturbative matching and outer boundary conditions: Computational studies. *Phys. Rev.*, **59**:064001, 1999.

L. Rezzolla and O. Zanotti. An improved exact Riemann solver for relativistic hydrodynamics. Submitted to *J. Fluid Mech.*, (preprint `gr-qc/0103005`), 2001.

P. L. Roe. Approximate Riemann solvers, parameter vectors, and difference schemes. *J. Comput. Phys.*, **43**:357–372, 1981.

J. V. Romero, J. M. Ibáñez, J. M. Martí, and J. A. Miralles. A new spherically symmetric general relativistic hydrodynnamical code. *Astrophys. J.*, **462**: 839–854, 1996.

D. Ryu and T. W. Jones. Numerical magnetohydrodynamics in astrophysics: Algorithm and tests for one-dimensional flow. *Astrophys. J.*, **442**:228–258, 1995.

L. F. Shampine and M. K. Gordon. *Computer Solution of Ordinary Differential Equations: the Initial Value Problem.* Freeman, 1975.

M. Shibata and M. Sasaki. Black hole formation in the Friedmann universe: Formulation and computation in numerical relativity. *Phys. Rev.*, **D60**: 084002, 1999.

C. Shu. Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws. Technical report, NASA ICASE, 1997.

C. Shu and S. Osher. Efficient implementation of essentially non-oscillatory shock-capturing schemes. *J. Comput. Phys.*, **77**:439–471, 1988.

J. Smoller and B. Temple. Global solutions of the relativistic Euler equations. *Commun. Math. Phys.*, **156**:67–99, 1993.

G. A. Sod. A survey of several finite-difference methods for systems of non-linear hyperbolic conservation laws. *J. Comput. Phys.*, **27**:1–31, 1978.

J. Stoer and R. Burlisch. *Introduction to Numerical Analysis.* Springer, second edition, 1993.

E. F. Toro. The weighted average flux method applied to the Euler equations. *Philos. Trans. Roy. Soc. London Ser. A.*, **341**:499–530, 1992.

E. F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics.* Springer-Verlag, 2nd edition, 1999.

E. F. Toro. ADER: High-order non-oscillatory advection schemes. In *Proceedings of the 8th International Conference on Non-linear Hyperbolic Problems.* Magdeburg, Germany, 2000.

E. F. Toro and S. J. Billet. Centred TVD schemes for hyperbolic conservation laws. *IMA J. Numer. Anal.*, **20**:47–79, 2000.

E. F. Toro, M. Spruce, and W. Speares. Restoration of the contact surface in the Harten-Lax-van Leer Riemann solver. *Shock Waves*, **4**:25–34, 1994.

B. van Leer. Towards the ultimate conservative difference scheme. *J. Comput. Phys.*, **32**:101–136, 1979.

R. M. Wald. *General Relativity.* University of Chicago, 1984.

C. M. Will. The confrontation between general relativity and experiment. *Living Rev. Relativity*, **4**, 2001. [Article in online journal], cited on 31/07/01, http://www.livingreviews.org/ Articles/Volume4/2001-4will/index.html.

P. Woodward and P. Colella. The numerical simulation of two-dimensional flow with strong shock waves. *J. Comput. Phys.*, **54**:115–173, 1984.

J.-Y. Yang, M.-H. Chen, I.-N. Tsai, and J.-W. Chang. A kinetic beam scheme for relativistic gas dynamics. *J. Comput. Phys.*, **136**:19–40, 1997.