

ADVANCED NUMERICAL APPROACHES IN THE DYNAMICS OF RELATIVISTIC FLOWS

**Der Fakultät für Mathematik und Physik
der Gottfried Wilhelm Leibniz Universität Hannover**

zur Erlangung des Grades

Doktor der Naturwissenschaften
Dr. rer. nat.

genehmigte Dissertation
von

David Radice
geboren am 03.06.1984 in Rho (Mi), Italien

2013

Referent: Prof. Bernard F. Schutz
Korreferent: Prof. Jose-Maria Ibañez
Tag der Promotion: 18.11.2013

Abstract

Keywords: relativistic hydrodynamics, numerical methods, general relativity.

Strong gravity and relativistic plasma flows are among the fundamental ingredients powering high-energy astrophysical phenomena such as short and long gamma ray bursts, core-collapse supernovae and relativistic outflows from black-hole accreting systems. General-relativistic hydrodynamics is also essential in modelling the merger of neutron stars binaries and black-hole neutron-star binaries that are among the best sources for future gravitational-wave detectors such as LIGO, Virgo or KAGRA.

Over the past decade, the understanding of these phenomena has benefited significantly from the results obtained through non-linear numerical calculations. Key factors in this progress have been the switch to more advanced numerical schemes that are able to properly treat relativistic shock waves, and the progressive inclusion of more “physics”, such as magnetic fields or realistic equations of state. Following this trend, even better numerical tools and more accurate physical description will be essential to understand these phenomena. This thesis aims at contributing to both of these aspects.

It is in this view that I have developed the first (and so far the only) truly higher-than-second-order fully general-relativistic multi-dimensional hydrodynamics code using state-of-the-art high-resolution shock-capturing finite-difference methods. This code is able to compute gravitational waveforms from binary neutron star merger with superior accuracy and at a reduced cost when compared to traditional codes. The code has also been used to study the statistical properties of relativistic turbulent flows of an ultrarelativistic plasma, and I was able to uncover similarities and differences with classical turbulence that could not be observed before with standard codes.

Similarly, I have developed the first general relativistic hydrodynamics code in spherical symmetry based on discontinuous Galerkin methods. These constitute a new family of methods that have gained a lot of attention recently due to their superior accuracy and scaling properties. Using a one-dimensional code I could start evaluating what could be the paradigm for the next-generation of numerical codes in computational relativistic astrophysics. In particular, I showed that these methods are able to properly treat relativistic shock waves and deliver spectral-like accuracy in smooth flow regions.

In addition, I started the development of a new multi-dimensional, multi-group, multi-angle, velocity dependent, general-relativistic, full-Boltzmann solver for the solution of the radiative transfer equation. This code is also based on discontinuous Galerkin schemes and uses the recently proposed filtered-spherical harmonics method to deal with angular dependence. In the first stage in this effort, which does not include velocity-dependence and general relativistic gravity yet, I have improved the original method by making use of new filters and by demonstrating the great potentialities of this approach.

Zusammenfassung

Schlagworte: relativistische Hydrodynamik, numerische Methoden, Allgemeine Relativitätstheorie.

Starke Gravitation und relativistische Plasmaströme sind mitverantwortlich für hoch energetische astrophysikalische Phänomene wie kurze und lange Gammastrahlenblitze (gamma-ray bursts), Kernkollaps-Supernovae und relativistische Plasmaausströme von Akkretionsvorgängen um Schwarze Löcher. Die allgemein-relativistische Hydrodynamik ist ebenfalls unabdingbar zur Beschreibung der Verschmelzung zweier Neutronensterne oder eines Neutronensterns mit einem Schwarzen Loch in Doppelsternsystemen. Diese Systeme gehören zu den vielversprechendsten Signalquellen für zukünftige Gravitationswellendetektoren wie LIGO, Virgo oder KAGRA.

In der letzten Dekade wurde das Verständnis dieser Phänomene erheblich durch Resultate von nicht-linearen numerischen Simulationen vorangetrieben. Eine Schlüsselrolle nahm hierbei der Übergang zu fortgeschritteneren numerischen Verfahren ein, welche in der Lage sind, relativistische Schockwellen richtig zu behandeln. Ebenso wichtig war die schrittweise Einbeziehung von mehr "Physik", wie zum Beispiel Magnetfelder und realistische Zustandsgleichungen. Gemäß diesem Trend werden noch bessere numerische Verfahren und eine noch genauere physikalische Beschreibung erforderlich sein, um diese Phänomene letztlich zu verstehen. Die vorliegende Arbeit setzt sich zum Ziel, zu beiden genannten Aspekten beizutragen.

Zu diesem Zweck habe ich den ersten (und bislang einzigen) allgemein-relativistischen, multidimensionalen Hydrodynamik-Code entwickelt, der modernste high-resolution shock-capturing Finite-Differenzen-Methoden verwendet und eine echt höhere Genauigkeit als zweite Ordnung erreicht. Dieser Code kann Gravitationswellenformen von Doppelneutronensternsystemen (insbesondere bei der Verschmelzung der Sterne) mit höherer Genauigkeit als bisherige Codes bei gleichzeitig reduziertem Rechenaufwand berechnen. Dieser Code wurde ebenso verwendet um die statistischen Eigenschaften von relativistisch-turbulenten Strömungen eines ultra-relativistischen Plasmas zu untersuchen, und so Gemeinsamkeiten und Unterschiede zur klassischen Turbulenz aufzudecken, was zuvor mit bisherigen Codes nicht möglich war.

Darüber hinaus habe ich den ersten allgemein-relativistischen Hydrodynamik-Code in sphärischer Symmetrie entwickelt, der auf diskontinuierlichen Galerkin-Methoden beruht. Diese bilden eine neue Familie von numerischen Verfahren, die in letzter Zeit aufgrund ihrer höheren Genauigkeit und besseren Skalierungseigenschaften viel Aufmerksamkeit erfahren haben. Durch die Verwendung eines solchen eindimensionalen Codes konnte ich beginnen herauszufinden, was das Paradigma für die nächste Generation an numerischen Codes für die relativistische Astrophysik werden könnte. Speziell habe ich gezeigt, dass diese Verfahren in der Lage sind, relativistische Schockwellen richtig zu behandeln und in glatten Strömungsregionen eine Genauigkeit ähnlich der von spektralen Methoden zu erreichen.

Ferner habe ich mit der Entwicklung eines neuen multidimensionalen, geschwindigkeits- und winkelabhängigen, allgemein-relativistischen Multi-Group Boltzmann Solvers zur numerischen Lösung der Strahlungstransportgleichung begonnen. Dieser Code basiert ebenfalls auf diskontinuierlichen Galerkin-Verfahren und verwendet die kürzlich vorgeschlagene filtered-spherical harmonics Methode, um die Winkelabhängigkeit zu berücksichtigen. Als erstes Ziel in diesem Bestreben konnte ich durch die Verwendung von neuen Filtern die bisherigen Methoden in einer noch geschwindigkeitsunabhängigen und nicht-allgemein-relativistischen Version verbessern und die vielversprechenden Möglichkeiten dieser Herangehensweise aufzeigen.

Acknowledgements

First of all I would like to thank Viktoriya for her love. Without her writing this thesis would have been meaningless. Then I would like to thank my parents, my brother and my sister for their support and patience. Also, and in a special way, my child Margarita for carefully going with me over all the gravitational waveforms and many of the plots shown in this thesis and patiently explaining me which of them are “gut” and which are “blöde”.

I am deeply indebted to my advisor Luciano Rezzolla for having guided me through my PhD studies: the first time we met I told him that I was looking for new challenges. Indeed he found many for me! He gave me enough freedom to explore many of the open problems in numerical relativity and, at the same time, full support every time I needed it.

A sincere gratitude goes also to the people with whom I worked and discussed most closely in the past few years: Ernazar Abdikamalov, Daniela Alic, Riccardo Ciolfi, Filippo Galeazzi, Ian Hawke, Wolfgang Kastaun, Gianmario Manca, Christian Ott, Erik Schnetter, Bernard Schutz, Kentaro Takami, Aaryn Tonita and Olindo Zanotti.

It is a pleasure to acknowledge Daniel Siegel for helping me with the German abstract of this thesis, Francesco Pannarale for providing the Taylor-T4 waveform, Kentaro Takami for providing the perturbative eigenfrequencies of the TOV model we evolve in Chapter 6 and Cecilia Chirenti and Shin’ichirou Yoshida for sharing with me their perturbation theory codes.

I would also like to thank all the members of the astrophysical relativity group at AEI for the helpful discussions we had in the past few years and, in a particular way, Eloisa Bentivegna, Giovanni Corvino, Kyriaki Dionysopoulou, Abraham Harte, Ian Hinder, Thorsten Kellermann, Jose-Luis Jaramillo, Bruno Giacomazzo, Philipp Moesta, Bruno Mundim, Constanze Roedig, Christian Reisswig, Jocelyn Read, Alberto Sesana, Barry Wardell and Burkhard Zink.

It is also a pleasure to thank “my student”, Massimiliano Leoni, for the many pleasant discussions on numerical methods, hydrodynamics and math in general.

Finally, last but not least, I would like to thank you, my reader. If you were reading this page looking for your name and you could not find it, then chances are high that I really meant to thank you, but I forgot. If not and you just happen to be reading my thesis for some reason I hope you will not have a too hard time disentangling my notation. Either way, thank you for your patience.

Contents

Abstract	iii
Zusammenfassung	iv
Acknowledgements	vi
I Background	1
1 Introduction	3
2 General-Relativistic Hydrodynamics	7
2.1 The Cauchy Problem in General Relativity	7
2.1.1 The Einstein Field Equations	7
2.1.2 From Spacetime to Space and Time	9
2.1.3 The ADM Formulation	12
2.1.4 Strongly Hyperbolic Formulations of the Einstein Equations	16
2.1.5 Gauge conditions	19
2.2 The Equations of General-Relativistic Hydrodynamics	20
2.2.1 Kinematics of a Relativistic Fluid	21
2.2.2 Dynamics of a Relativistic Fluid	24
2.2.3 Conservative Formulations	25
2.3 The General-Relativistic Boltzmann Equation	27
2.3.1 The geometry of the tangent bundle	28
2.3.2 The Liouville Theorem	30
2.3.3 The Boltzmann equation	33
2.3.4 From the Boltzmann Equation to the Euler Equation . . .	34
3 Numerical Approximation of Conservation Laws	37
3.1 Theoretical Background	37
3.1.1 Conservation Laws	38
3.1.2 Consistency, Stability and Convergence	39
3.1.3 Non-Linear Equations and Non-Linear Stability	41
3.2 Finite-Volume Methods	44
3.2.1 The Godunov Method	44
3.2.2 TVD Finite-Volume Methods	48
3.2.3 Higher-Order Finite-Volume Methods	49
3.3 Central Methods	52
3.4 Finite-Difference Methods	54

3.5	Discontinuous Galerkin Methods	57
3.5.1	Runge-Kutta Discontinuous-Galerkin Methods	60
II	High-Order Methods for Relativistic Hydrodynamics	65
4	Finite-Differencing Methods: Flat Spacetimes	67
4.1	Introduction	67
4.2	The THCCode	68
4.2.1	Newtonian Hydrodynamics	68
4.2.2	Special-relativistic hydrodynamics	68
4.3	Numerical tests	71
4.3.1	Newtonian hydrodynamics	71
4.3.2	Special-relativistic hydrodynamics	74
4.4	The relativistic Kelvin-Helmholtz instability in 3D	91
4.4.1	The linear evolution of the instability	92
4.4.2	The non-linear evolution of the instability	94
4.5	Driven Relativistic Turbulence	99
4.5.1	Introduction	99
4.5.2	Model and method	99
4.5.3	Basic flow properties	101
4.5.4	Universality	102
4.5.5	Intermittency	104
4.6	Conclusions	105
5	Finite-Differencing Methods: General Spacetimes	109
5.1	Introduction	109
5.2	WhiskyTHC	110
5.2.1	Numerical Methods	110
5.2.2	Atmosphere Treatment	111
5.3	Single Neutron Stars	116
5.3.1	Linear Oscillations: Cowling Approximation	116
5.3.2	Linear Oscillations: Full-GR	120
5.3.3	Non-linear Oscillations: the Migration Test	122
5.3.4	Gravitational Collapse to Black-Hole	125
5.4	Binary Neutron Stars	127
5.4.1	Small Separation	128
5.4.2	Large Separation	135
5.5	Conclusions	142
6	Discontinuous Galerkin Methods	145
6.1	Introduction	145
6.2	Discontinuous Galerkin methods for general relativistic hydrodynamics	146
6.2.1	Weak formulation of the equations of relativistic hydrodynamics	147
6.2.2	Spacetime discontinuous Galerkin formulation	149
6.2.3	Discontinuous Galerkin formulation in the 3 + 1 split	152
6.2.4	Discontinuous Galerkin formulation in spherical symmetry	154
6.3	The EDGES code	156

6.3.1	The DG equations in a fully discrete form	156
6.3.2	Coupling with the spacetime	158
6.3.3	Limiters, spectral viscosity and spectral filtering	159
6.3.4	Treatment of low-density regions	165
6.4	Numerical tests	168
6.4.1	Shock tubes	168
6.4.2	Spherical shock reflection	170
6.4.3	Spherical accretion onto a Schwarzschild black hole	172
6.4.4	Linear oscillations of spherical stars	173
6.4.5	Nonlinear oscillations of spherical stars: the migration test	179
6.4.6	Gravitational collapse of unstable spherical stars	183
6.5	Conclusions	185
III	Relativistic Radiation Transport	187
7	The Filtered Spherical Harmonics Method	189
7.1	Introduction	189
7.2	The relativistic Boltzmann equation	192
7.2.1	The distribution function for radiation	192
7.3	The Charon Code	194
7.3.1	Frequency discretization	194
7.3.2	Real spherical harmonics	195
7.3.3	Angular discretization	197
7.3.4	The multi-group P_N scheme	197
7.3.5	Spatial discretization	199
7.3.6	Time discretization	201
7.3.7	Filtering	202
7.4	Tests	205
7.4.1	1D diffusion of a step function	205
7.4.2	1D diffusion of a sine wave	207
7.4.3	The 2D line-source problem	208
7.4.4	A lattice problem	213
7.4.5	3D Homogeneous sphere	215
7.5	Conclusions	220
IV	Conclusions	221
8	Conclusions	223
	Bibliography	227
	Curriculum Vitæ	255
	Publications	257

Part I

Background

Chapter 1

Introduction

Numerical relativistic hydrodynamics has come a long way since the pioneering works by [213] and [333] and it is now playing a central role in the modelling of systems involving strong gravity and/or flows with high Lorentz factors. Examples of applications are relativistic jets, core-collapse supernovae, the merger of compact binaries and the study of gamma-ray bursts, see [209] and [130] for a complete overview.

In all of these areas progress has been continuous over the past few years to the point that relativistic computational fluid dynamics is starting to provide a realistic description of many relativistic-astrophysics scenarios, see, *e.g.*, [274]. Key factors in this progress have been the switch to more advanced and accurate numerical schemes, and in particular the adoption of high resolution shock capturing (HRSC) schemes [210, 288, 36, 116, 14, 28] and the progressive inclusion of more “physics” for a more accurate description of the different scenarios. Examples of the latter are the inclusion of magnetic fields [182, 110, 138, 184, 117, 234, 20, 145], the use of realistic tabulated equations of state, see, *e.g.*, [292], and the description of radiative processes [127, 241, 291, 337].

We expect that both improved physical models and better numerical techniques will be key elements in the future generation of codes for relativistic astrophysics. On the one hand it is necessary to take into account many physical phenomena that are currently oversimplified and, on the other hand, higher accuracy is necessary to make quantitative predictions even in the case where simplified models are used to describe the objects of study. For example, in the case of inspiralling binary neutron stars, waveforms that are sufficiently accurate for gravitational-waves templates are just now becoming available and only in the simple case of polytropic stars [31, 32, 48]. Clearly, even higher accuracy will be required as more realistic equations of state are considered or better characterisations of the tidal effects are explored [26, 46].

On the one hand, the development of more accurate numerical tools for relativistic hydrodynamics is an active and lively field of research. Most of the effort has been directed towards the development of high-order finite-volume [318, 120] and finite-difference [339, 111] schemes, but many alternative approaches have been also proposed, including finite-element methods [206, 220], spectral methods [155], smoothed-particle-hydrodynamics [302, 279] and discontinuous Galerkin methods [121].

On the other hand, one of the main missing ingredients in computational relativistic astrophysics is an accurate treatment of radiation-transport: many phenomena in the universe involve the transport of radiation and need to be modeled with radiation-transport techniques that are as accurate as possible to maximize the match with observations. Examples are nova and supernova explosions, gamma-ray bursts, star or planet formation, luminous blue variable outbursts, stellar winds, etc. In these examples, radiation plays a major role in exchanging energy and/or momentum between different parts of the system. In most of these cases, the radiation is composed of photons, but the radiation can also be composed of neutrinos in studies of core-collapse supernova explosion mechanisms (see, *e.g.*, [49, 175]) or in modeling the torus orbiting the black hole produced in the merger of neutron stars (see, *e.g.*, [280]).

The aim of this dissertation is to contribute new developments on both of these fronts: improved numerical methods and a better physical description, with the inclusion of radiation transport. On the numerical side we present the first higher-than-second-order multidimensional fully-general relativistic hydrodynamics code, which we demonstrate in the case of the inspiral and merger of binary neutron star in quasi-circular orbit. As an example illustrating the potential of our methods, we present a study of the statistical properties of relativistic turbulence. We also present the first discontinuous Galerkin general relativistic hydrodynamics code in 1D/spherical symmetry. On the physics side we present work done towards the creation of the first, genuinely multi-dimensional, general-relativistic radiation transport code. In all cases we report the results from a series of stringent tests to demonstrate the potential of the approaches we propose.

The rest of this thesis is organized as follows.

- In Chapter 2 we introduce the equations of general relativity and, in particular, the so-called “3+1” formalism used in numerical relativity to rewrite the Einstein field equations as a set of hyperbolic evolution equations and constraints. Secondly we present the equations of general relativistic hydrodynamics emphasizing the aspects relevant for our presentation, such as the so-called “Valencia formulation”. Finally we present a detailed derivation of the Boltzmann equation in general-relativity as this equation is at the basis of both the kinetic theory of gases and of the standard treatment of radiation. We will also show how the equations of relativistic hydrodynamics can be derived from the Boltzmann equation.
- In Chapter 3 we recall some of the theoretical background needed in the rest of the thesis. In particular we describe the mathematical theory of conservation laws and of their numerical approximation. We also give a short review of the most commonly adopted numerical schemes for conservation laws. Finally we discuss the strengths and weaknesses of each scheme for general problems and, in particular, for the case of the relativistic hydrodynamics equations.
- In Chapter 4 we present THC: a new code for Newtonian and special relativistic hydrodynamics employing state-of-the-art finite-differencing high-resolution shock-capturing methods. We present the results obtained in a representative number of test cases both in Newtonian and in special-relativistic hydrodynamics. As examples of possible applications

we study the linear and non-linear development of the Kelvin-Helmholtz instability in 2D and 3D and the statistical properties of driven relativistic turbulence.

- In Chapter 5 we present the extension of THC to the general relativistic case: `WhiskyTHC`. We describe in detail some of the key issues in the development of high-order general relativistic hydrodynamics code, in particular for what concerns the treatment of fluid–vacuum interfaces. We present the results obtained in a series of classical tests involving the linear and non-linear evolution of isolated, non-rotating stars. Finally we demonstrate the superior accuracy of our new code in the calculation of gravitational radiation in the inspiral and merger of binary neutron stars in quasi-circular orbits.
- In Chapter 6 we present `EDGES`: a new code for general relativistic hydrodynamics in 1D/spherical symmetry using discontinuous Galerkin methods. First of all we present a detailed derivation of the weak formulation of the equations of relativistic hydrodynamics. Next we discuss our implementation and present the results obtained in a series of tests in 1D and in spherical-symmetry, showing the great potential of discontinuous Galerkin methods.
- In Chapter 7 we present `Charon`: a new code for multi-dimensional radiation transport based on the filtered spherical-harmonics approach. We discuss some of the improvements we have made in the original method and present the results obtained in a series of very stringent tests showing that filtered spherical-harmonics methods represent a viable strategy for the modeling of multi-dimensional, general relativistic, radiative transfer.
- Finally Chapter 8 is dedicated to discussion and conclusions.

In the following we denote vectors with $\vec{u}$ and one-forms with $\underline{\alpha}$. Tensors are written with bold characters. Unless otherwise specified, greek indices will run over 0, 1, 2, 3, the indices i, j, k, l will run over 1, 2, 3 and capitalized latin indices A, B, C, D will run over 0, 1, 2, 3, 4, 5, 6, 7. Finally we use a system of units in which $M_\odot = c = G = 1$, unless explicitly stated.

Chapter 2

General-Relativistic Hydrodynamics

In this chapter we sketch a number of important ideas that make up the mathematical background of this thesis. We do not attempt to give a comprehensive survey of the topics we present, but we emphasise the aspects that are relevant for our discussion.

This chapter is divided in three parts. In Section 2.1 we recall the $3 + 1$ formalism used in numerical relativity to recast the Einstein equation in a set of evolutionary equations and constraints. In Section 2.2 we present the equations of general relativistic hydrodynamics (GRHD). The emphasis here is posed on their interpretation as balance law, as this is the point of view adopted in their numerical approximation. Finally in Section 2.3 we derive the general relativistic Boltzmann equation and show how the equations of relativistic hydrodynamics can be derived from the kinetic theory.

2.1 The Cauchy Problem in General Relativity

In this section we recall some basic concepts about the initial-value formulation of the Einstein equations of general relativity. First of all we start from the basics of general relativity, in order to introduce our notation. We briefly mention the derivation of the Einstein field equations and then we show how these can be split in a set of evolutionary equations and constraints using the Arnowitt, Deser and Misner, or ADM, formalism. Finally we comment on the stability of the ADM equations, on the need for strongly-hyperbolic formulations of the Einstein equations, and on the choice of gauge conditions commonly used to evolve spacetimes with singularities. Our treatment is sketched on the basis of the one found in [23, 191, 229, 330, 253, 38], we refer to these sources for a more extended treatment.

2.1.1 The Einstein Field Equations

In general relativity the spacetime is represented by a Lorentzian manifold, *i.e.*, a tuple $(\mathcal{M}, g)$, where $\mathcal{M}$ is a smooth manifold (at least of class C^2) and g is

a C^2 Lorentzian metric. The geometry of space and time is prescribed by the Einstein field equations.

Excursus: Lagrangian Field Theory

For technical reasons it is useful to introduce the totally antisymmetric symbol, α , which, in any coordinate patch $\{x^\mu\}$, is set equal to

$$\alpha = dx^0 \wedge dx^1 \wedge dx^2 \wedge dx^3, \quad (2.1)$$

so that the proper volume pseudo-form of the spacetime, which we denote as ε or Vol_x^4 (we use the second notation to avoid confusion when dealing with multiple volume pseudo-forms), is simply $\varepsilon = \sqrt{-g} \alpha$, where g is the determinant of the spacetime metric.

With this notation a Lagrangian field theory on the spacetime $(\mathcal{M}, g)$ is described by the action¹

$$S = \int_{\mathcal{M}} \Lambda(q, \nabla q) \alpha, \quad (2.2)$$

where q are a set of (tensorial) generalized coordinates for the fields described by the theory, ∇ is the Levi-Civita connection and Λ is a scalar density, *i.e.*,

$$\Lambda(q, \nabla q) = \mathcal{L}(q, \nabla q) \sqrt{-g}, \quad (2.3)$$

for some scalar quantity $\mathcal{L}(q, \nabla q)$. The field equations for a Lagrangian theory are obtained by requiring the action S to be stationary with respect to variations δq with compact support:

$$0 = \delta S = \int_{\mathcal{M}} \left(\frac{\partial \Lambda}{\partial q} - \nabla \frac{\partial \Lambda}{\partial(\nabla q)} \right) \cdot \delta q \alpha. \quad (2.4)$$

From the arbitrariness of δq one gets the Euler-Lagrange equations

$$\frac{\partial \Lambda}{\partial q} - \nabla \frac{\partial \Lambda}{\partial(\nabla q)} = 0. \quad (2.5)$$

The Hilbert Action

The Einstein field equations can also be derived within the Lagrangian field-theory framework. In particular the action describing them can be written as

$$S = S_g + S_m, \quad (2.6)$$

where S_g is the action of the gravitational field, the so-called *Hilbert action*, which is, in our system of units,

$$S_g = \frac{1}{16\pi} \int_{\mathcal{M}} R \sqrt{-g} \alpha, \quad (2.7)$$

where R is the Ricci scalar and S_m is the action of all the other “matter” fields, which we leave unspecified. The Euler-Lagrange equations can be obtained

¹Here we are implicitly requiring the derivatives of Λ to be tensor field densities.

by varying the action with respect to the inverse metric, which we write in components as $g^{\mu\nu}$.

The variation of the gravitational sector of the action gives (after a lengthy calculation)

$$\delta S_g = -\frac{1}{16\pi} \int_{\mathcal{M}} \left(R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R \right) \delta g^{\mu\nu} \sqrt{-g} \, \alpha, \quad (2.8)$$

where we have introduced the Ricci tensor $R_{\mu\nu}$. This can be written in more compact form with the introduction of the Einstein tensor,

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R, \quad (2.9)$$

as

$$\delta S_g = -\frac{1}{16\pi} \int_{\mathcal{M}} G_{\mu\nu} \delta g^{\mu\nu} \sqrt{-g} \, \alpha. \quad (2.10)$$

We can write the variation of the matter sector as

$$\delta S_m = \frac{1}{2} \int_{\mathcal{M}} T_{\mu\nu} \delta g^{\mu\nu} \sqrt{-g} \, \alpha, \quad (2.11)$$

where we have defined the densitized stress-energy tensor density of the matter, $\sqrt{-g} T_{\mu\nu}$, as the functional derivative of the matter action with respect to the inverse metric:

$$\sqrt{-g} T_{\mu\nu} = 2 \frac{\delta S_m}{\delta g^{\mu\nu}}. \quad (2.12)$$

Thanks to the arbitrariness of $\delta g^{\mu\nu}$ we obtain the Einstein equations

$$G = 8\pi T. \quad (2.13)$$

2.1.2 From Spacetime to Space and Time

The Einstein field equations (2.13) constitute a set of ten non-linear partial differential equations determining the metric on the whole spacetime or on a domain $\Omega \subset \mathcal{M}$, once appropriate boundary conditions are given on $\partial\Omega$. It is clear, however, that their numerical solution in this form, *i.e.*, as a boundary-value problem with data on $\partial\Omega$, is feasible only in particular cases, for instance if the spacetime is assumed to be stationary. More generally one would like to be able to specify initial data on an appropriated spacelike or null hypersurface, $\Sigma \subset \mathcal{M}$, and then “evolve it in time” to reconstruct the geometry of the whole spacetime. It turns out that this is possible for a whole class of spacetimes, called strongly-hyperbolic, that is spacetimes which admit a foliation of the form $\mathcal{M} = \Sigma \times \mathbb{R}$, where the leaves of the foliation, Σ_t are Cauchy hypersurfaces, *i.e.*, spacelike or null hypersurfaces whose past and future domains of influence cover the whole spacetime.

Spacelike Foliations

Here we consider the case in which the Σ_t are spacelike. In this case we can introduce a global, smooth, “time function”, t , such that

$$\Sigma_\tau = \{x^\alpha \in \mathcal{M}: t(x^\alpha) = \tau\}, \quad (2.14)$$




Figure 2.1: 3 + 1 decomposition of spacetime into space and time

i.e., the leaves of the foliation are the level-sets of the function t . Let $\vec{t}$ be a vector such that $\langle \nabla t, \vec{t} \rangle = 1$, so that we can interpret it as being the “time-flow vector field”. In particular the Lie derivative along $\vec{t}$ of any tensor quantity, q , gives the rate of change of this quantity as we go “from one leaf of the foliation to the next one” and can be interpreted as the “time derivative”.

The vector $\vec{t}$ can be uniquely decomposed into a part which is parallel to the future-oriented unit-normal of Σ_t , $\vec{n}$, and a part which lies in the tangent plane to Σ_t , $\vec{\beta}$, as

$$\vec{t} = \alpha \vec{n} + \vec{\beta}. \quad (2.15)$$

In the previous equation α is the so-called *lapse-function* and gives the rate of change of the proper time measured by an observer, called Eulerian or normal observer, moving with four-velocity $\vec{n}$ with respect to the “coordinate time” t . $\vec{\beta}$ is called the *shift vector* and can be interpreted as being (minus) the coordinate velocity of the normal observer. The geometrical interpretation of lapse and shift is shown in Figure 2.1.

The spacetime metric g induces a Riemannian metric, which we call the spatial metric, on Σ_t given by

$$\gamma = g + \underline{n} \otimes \underline{n}. \quad (2.16)$$

The associated Levi-Civita connection, D , is simply obtained by projecting the connection ∇ on the tangent space to Σ_t . *I.e.*, in component form

$$D_\mu T^{\alpha_1 \dots \alpha_k}_{\beta_1 \dots \beta_l} = \gamma^{\alpha_1}_{\gamma_1} \dots \gamma^{\alpha_k}_{\gamma_k} \gamma^{\delta_1}_{\beta_1} \dots \gamma^{\delta_l}_{\beta_l} \gamma^\nu_\mu \nabla_\nu T^{\gamma_1 \dots \gamma_k}_{\delta_1 \dots \delta_l}. \quad (2.17)$$

Things are simplified if we work in a coordinate frame adapted to the split, *i.e.*, $\{t, x^i\}$, with $\vec{\partial}_i \cdot \vec{n} = 0$. In this coordinate system ∇t is simply dt and $\vec{t} = \vec{\partial}_t$.

The components of the spacetime metric can also be written very easily in terms of the ones of the spatial metric, the lapse function and the shift vector: remembering that

$$g_{\mu\nu} = \vec{\partial}_\mu \cdot \vec{\partial}_\nu, \quad (2.18)$$

we obtain for the spatial components

$$g_{ik} = \vec{\partial}_i \cdot \vec{\partial}_k = \gamma_{ik}, \quad (2.19)$$

while, for the time component we have²

$$g_{tt} = \vec{\partial}_t \cdot \vec{\partial}_t = \vec{t} \cdot \vec{t} = (\alpha \vec{n} + \vec{\beta}) \cdot (\alpha \vec{n} + \vec{\beta}) = -(\alpha^2 - \vec{\beta} \cdot \vec{\beta}) \quad (2.20)$$

and, finally, the mixed components are

$$g_{ti} = \vec{\partial}_t \cdot \vec{\partial}_i = \vec{t} \cdot \vec{\partial}_i = (\alpha \vec{n} + \vec{\beta}) \cdot \vec{\partial}_i = \beta_i. \quad (2.21)$$

In conclusion the line-element can be written as

$$ds^2 = -(\alpha^2 - \beta_i \beta^i) dt^2 + 2\beta_i dx^i dt + \gamma_{ik} dx^i dx^k, \quad (2.22)$$

i.e., in a form which we can interpret as the four-dimensional version of the Pythagorean theorem.

Exterior Curvature and Constraint Equations

We define the second fundamental form, or extrinsic curvature of Σ_t

$$K_{\mu\nu} = -\gamma^\alpha{}_\mu \nabla_\alpha n_\nu = -\frac{1}{2} \mathcal{L}_{\vec{n}} \gamma_{\mu\nu}, \quad (2.23)$$

which we can interpret as the rate of change of the three-metric as measured by the Eulerian observer, or as the change of the normal vector $\vec{n}$ under parallel transport on Σ_t . We notice that $\mathbf{K}$ is a purely spatial tensor, that is

$$K_{\mu\nu} n^\nu = 0. \quad (2.24)$$

The extrinsic curvature is related to the four-dimensional Ricci tensor by the Codazzi equations

$$D_\beta K - D_\alpha K^\alpha{}_\beta = R_{\gamma\delta} n^\delta \gamma^\gamma{}_\beta, \quad (2.25)$$

where K is the trace of $\mathbf{K}$.

Similarly the three-dimensional Riemann tensor, ${}^{(3)}R_{\alpha\beta\gamma}{}^\delta$, is related to the four-dimensional one and to the extrinsic curvature by the Gauss equation

$${}^{(3)}R_{\alpha\beta\gamma}{}^\delta = \gamma^\mu{}_\alpha \gamma^\nu{}_\beta \gamma^\lambda{}_\gamma \gamma^\delta{}_\sigma R_{\mu\nu\lambda}{}^\sigma - K_{\alpha\gamma} K_{\beta}{}^\delta + K_{\beta\gamma} K_{\alpha}{}^\delta. \quad (2.26)$$

If we substitute the Einstein equations into (2.25) we get the so-called *momentum constraint*

$$D_\beta K - D_\alpha K^\alpha{}_\beta = -8\pi \gamma^\alpha{}_\beta n^\gamma T_{\alpha\gamma} =: 8\pi j_\beta, \quad (2.27)$$

²Note that, since $\vec{\beta}$ is a spatial vector, $\vec{\beta} \cdot \vec{\beta} = \gamma_{ik} \beta^i \beta^k$.

where j^α is called ADM momentum density.

In the same way we obtain from (2.26) the so-called *Hamiltonian constraint*

$${}^{(3)}R + K^2 - K_{\alpha\beta}K^{\alpha\beta} = 2G^{\alpha\beta}n_\alpha n_\beta = 16\pi n_\alpha n_\beta T^{\alpha\beta} =: 16\pi E, \quad (2.28)$$

where E is called ADM energy density.

Equations (2.27) and (2.28) are called *constraint equations* because they constitute a set of elliptic equations which have to be satisfied on each leaf of the foliation. It is possible to show that, if the constraints are satisfied on an initial slice Σ_0 , they will be satisfied at any later time, *i.e.*, the Einstein equations preserve the constraints. We further discuss the nature of these constraints in the context of the ADM formalism.

2.1.3 The ADM Formulation

Now that we have introduced the so-called 3 + 1 formalism for the foliation of spacetime into spacelike hypersurfaces, we sketch the derivation of the ADM formulation of general-relativity as a Cauchy problem.

Excursus: Hamiltonian Field Theory

We consider a Lagrangian field theory described by a Lagrangian density Λ which is a function of the generalized coordinates of the field, $\mathbf{q}$ and their first (covariant) derivatives $\nabla\mathbf{q}$.

First of all it is convenient to decompose α as

$$\alpha = dt \wedge {}^{(3)}\alpha, \quad (2.29)$$

where ${}^{(3)}\alpha$ is the antisymmetric symbol on Σ_t , *i.e.*, in any coordinate patch $\{t, x^i\}$ adapted to the foliation,

$${}^{(3)}\alpha = dx^1 \wedge dx^2 \wedge dx^3. \quad (2.30)$$

Secondly we define the “time derivative” of $\mathbf{q}$ as

$$\dot{\mathbf{q}} := \mathcal{L}_t \mathbf{q}. \quad (2.31)$$

The canonical momentum $\mathbf{p}$ associated with the generalized coordinates of the system, $\mathbf{q}$, is defined as

$$\mathbf{p} = \frac{\partial \Lambda}{\partial \dot{\mathbf{q}}}. \quad (2.32)$$

If it is possible to express $\dot{\mathbf{q}}$ and $\nabla\mathbf{q}$ in terms of $\mathbf{q}$ and $\mathbf{p}$, then we can define the Hamiltonian density

$$\mathcal{H}(\mathbf{q}, \mathbf{p}) = \mathbf{p} \cdot \dot{\mathbf{q}} - \Lambda(\mathbf{q}, \nabla\mathbf{q}) \quad (2.33)$$

and the quantity

$$J = \int_0^t H(\mathbf{q}, \mathbf{p}) dt := \int_0^t dt \int_\Sigma \mathcal{H}(\mathbf{q}, \mathbf{p}) {}^{(3)}\alpha. \quad (2.34)$$

If we vary J with respect to variations $\delta \mathbf{p}$ and $\delta \mathbf{q}$ of compact support we get

$$\begin{aligned}\delta J &= \int_0^t \delta H(\mathbf{q}, \mathbf{p}) dt \\ &= \int_0^t dt \int_{\Sigma} (\mathbf{p} \cdot \delta \dot{\mathbf{q}} + \dot{\mathbf{q}} \cdot \delta \mathbf{p})^{(3)} \alpha - \delta S,\end{aligned}\quad (2.35)$$

from the stationarity of the action, $\delta S = 0$, we get, integrating by parts,

$$\int_0^t \delta H(\mathbf{q}, \mathbf{p}) dt = \int_0^t \int_{\Sigma} (-\dot{\mathbf{p}} \cdot \delta \mathbf{q} + \dot{\mathbf{q}} \cdot \delta \mathbf{p})^{(3)} \alpha \quad (2.36)$$

and from the arbitrariness of $\delta \mathbf{p}$ and $\delta \mathbf{q}$ we get the so-called Hamilton equations

$$\dot{\mathbf{q}} = \frac{\delta H}{\delta \mathbf{p}}, \quad \dot{\mathbf{p}} = -\frac{\delta H}{\delta \mathbf{q}}. \quad (2.37)$$

In conclusion, the Hamiltonian formalism provides a natural framework that can be used to obtain the field-equations in a form suitable for the solution of the Cauchy problem: once initial data is given on a Cauchy hypersurface for $\mathbf{q}$ and $\mathbf{p}$ one can evolve them according to (2.37) and obtain their value in the whole spacetime.

The Hamiltonian Formulation of the Einstein Equations

The initial-value problem for the Einstein field equations can also be obtained within the Hamiltonian framework, which is sketched here. For a more complete treatment, including the treatment of boundary terms see [253].

We begin by rewriting the scalar curvature, R , as

$$R = 2(G_{\mu\nu} n^\mu n^\nu - R_{\mu\nu} n^\mu n^\nu). \quad (2.38)$$

From (2.28) we have

$$2G_{\mu\nu} n^\mu n^\nu = {}^{(3)}R + K^2 - K_{\mu\nu} K^{\mu\nu}. \quad (2.39)$$

We can also rewrite $R_{\mu\nu} n^\mu n^\nu$ as

$$\begin{aligned}R_{\mu\nu} n^\mu n^\nu &= R_{\mu\gamma\nu}{}^\gamma n^\mu n^\nu \\ &= -n^\mu (\nabla_\mu \nabla_\gamma - \nabla_\gamma \nabla_\mu) n^\gamma \\ &= K^2 - K_{\mu\nu} K^{\mu\nu} - \nabla_\mu (n^\mu \nabla_\nu n^\nu) + \nabla_\mu (n^\nu \nabla_\nu n^\mu).\end{aligned}\quad (2.40)$$

The last two terms are total divergences, so they are not important in the case of variations with compact support, such as the ones we are interested in, thus we are able to write the Lagrangian density as

$$\Lambda = \frac{1}{16\pi} \alpha ({}^{(3)}R + K_{\mu\nu} K^{\mu\nu} - K^2) \sqrt{\gamma} + \Lambda_m, \quad (2.41)$$

where Λ_m is the Lagrangian density of the matter and $\sqrt{\gamma} = \frac{\sqrt{-g}}{\alpha}$ is the square root of the determinant of the three-metric.

In order to compute the canonical momentum $\mathbf{p}$ we notice that

$$K_{\mu\nu} = -\frac{1}{2}\mathcal{L}_{\vec{n}}\gamma_{\mu\nu} = -\frac{1}{2\alpha}(\partial_t\gamma_{\mu\nu} - D_\mu\beta_\nu - D_\nu\beta_\mu), \quad (2.42)$$

so that, if we assume that the matter Lagrangian does not depend explicitly on $\dot{\gamma}_{\mu\nu}$, we get

$$p^{\mu\nu} = \frac{\partial\Lambda}{\partial\dot{\gamma}_{\mu\nu}} = \frac{\sqrt{\gamma}}{16\pi}(K\gamma^{\mu\nu} - K^{\mu\nu}). \quad (2.43)$$

The canonical momenta associated with α and $\vec{\beta}$ vanish identically due to the fact that the lapse function and the shift vector are not real dynamical degrees of freedom of the gravitational field, but are associated with the gauge freedom in the choice of the foliation of $\mathcal{M}$.

Finally we can write the Hamiltonian density

$$\begin{aligned} \mathcal{H} = \frac{\gamma^{1/2}}{16\pi} \left\{ \alpha \left[-^{(3)}R + \gamma^{-1}p^{\mu\nu}p_{\mu\nu} - \frac{1}{2}\gamma^{-1}p^2 \right] \right. \\ \left. + 2\beta_\nu \left[D_\mu(\gamma^{-1/2}p^{\mu\nu}) \right] - 2D_\mu(\gamma^{-1/2}\beta_\nu p^{\mu\nu}) \right\} - \Lambda_m, \end{aligned} \quad (2.44)$$

where p is the trace of $\mathbf{p}$. The last term in brackets is a boundary term and gives no contribution in the case of compact support variations. We can also clearly see that α and β_ν behave as Lagrange multipliers enforcing two constraints³:

$$0 = 16\pi\gamma^{-1/2}\frac{\delta H}{\delta\alpha} = -^{(3)}R + \gamma^{-1}p^{\mu\nu}p_{\mu\nu} - \frac{1}{2}\gamma^{-1}p^2 + 16\pi T^{\mu\nu}n_\mu n_\nu \quad (2.45)$$

and

$$0 = 16\pi\gamma^{-1/2}\frac{\delta H}{\delta\beta_\mu} = -D_\nu(\gamma^{-1/2}p^{\mu\nu}) + 8\pi\gamma^\mu{}_\nu n_\gamma T^{\nu\gamma}, \quad (2.46)$$

which are easily recognizable as being the constraints (2.27) and (2.28). In the previous we have also used the fact that

$$\frac{\delta S_m}{\delta\beta_\mu} = \frac{\delta S_m}{\delta g_{\mu 0}} = \frac{1}{2}\sqrt{-g}T^{\mu 0} = -\frac{1}{2}\sqrt{\gamma}T^{\mu\nu}n_\nu \quad (2.47)$$

and

$$\frac{\delta S_m}{\delta\alpha} = -2\alpha\frac{\delta S_m}{\delta g_{00}} = -\alpha\sqrt{-g}T^{00} = -\alpha^2\sqrt{\gamma}T^{00} = -\sqrt{\gamma}T^{\mu\nu}n_\mu n_\nu. \quad (2.48)$$

Finally, the reader might be puzzled by the presence of the projector $\gamma^\mu{}_\nu$ in (2.46). The reason why this term is necessary is the following. When we write $\delta H/\delta\beta_\mu = 0$, what we really mean is that the action of the Frechét differential of the functional H , dH , on all compactly-supported variations, $\delta\beta_\mu$, is zero:

$$0 = 16\pi\gamma^{-1/2}\langle dH, \delta\beta_\mu \rangle = \delta\beta_\mu \left[-D_\nu(\gamma^{-1/2}p^{\mu\nu}) + 8\pi n_\gamma T^{\mu\gamma} \right], \quad (2.49)$$

³Note that H must be stationary with respect to variations of a quantity whose canonical momentum vanishes.

now since $\delta\beta_\mu$ is purely spatial only the spatial part of the term in square bracket is actually constrained to be zero by the arbitrariness of $\delta\beta_\mu$. This is expressed by the presence of the projector operator in (2.46).

As we can see the momentum and Hamiltonian constraints are associated respectively with the coordinate freedom on Σ and with the choice of the foliation of $\mathcal{M}$.

Using the Hamiltonian formalism we find the evolution equations for the metric:

$$\dot{\gamma}_{\mu\nu} = \frac{\delta H}{\delta p^{\mu\nu}} = 2\gamma^{-1/2}\alpha\left(p_{\mu\nu} - \frac{1}{2}\gamma_{\mu\nu}p\right) - 2D_{(\mu}\beta_{\nu)}, \quad (2.50)$$

where we used the notation $T_{(\mu\nu)} = \frac{1}{2}(T_{\mu\nu} + T_{\nu\mu})$ and we have assumed that the matter Lagrangian do not depend explicitly on $p^{\mu\nu}$.

In the same way we find the evolution equations for the canonical momentum

$$\begin{aligned} \dot{p}^{\mu\nu} = -\frac{\delta H}{\delta \gamma_{\mu\nu}} = & +\alpha\gamma^{1/2}\left({}^{(3)}R^{\mu\nu} - \frac{1}{2}{}^{(3)}R\gamma^{\mu\nu}\right) \\ & - \frac{1}{2}\alpha\gamma^{-1/2}\gamma^{\mu\nu}\left(p_{\gamma\delta}p^{\gamma\delta} - \frac{1}{2}p^2\right) \\ & + 2\alpha\gamma^{-1/2}\left(p^{\mu\gamma}p^\nu{}_\gamma - \frac{1}{2}pp^{\mu\nu}\right) \\ & - \gamma^{1/2}\left(D^\mu D^\nu\alpha - \gamma^{\mu\nu}D^\gamma D_\gamma\alpha\right) \\ & - \gamma^{1/2}D_\gamma\left(\gamma^{-1/2}\beta^\gamma p^{\mu\nu}\right) \\ & + 2p^{\gamma(\mu}D_\gamma\beta^{\nu)} + 8\pi\alpha\gamma^{1/2}S^{\mu\nu}, \end{aligned} \quad (2.51)$$

where $S^{\mu\nu} = \gamma^\mu{}_\alpha\gamma^\nu{}_\beta T^{\alpha\beta}$ and we have used the fact that

$$\frac{\delta S_m}{\delta \gamma_{ik}} = \frac{\delta S_m}{\delta g_{ik}} = \frac{1}{2}\sqrt{-g}T^{ik}. \quad (2.52)$$

Equations (2.45), (2.46), (2.50) and (2.51) are the so-called ADM system. We can also rewrite these equations in terms of the three-metric and the extrinsic curvature to get

$$(\partial_t - \mathcal{L}_{\vec{\beta}})\gamma_{ik} = -2\alpha K_{ik}; \quad (2.53a)$$

$$\begin{aligned} (\partial_t - \mathcal{L}_{\vec{\beta}})K_{ik} = & -D_i D_k \alpha + \alpha(R_{ik} - 2K_{ij}K^j{}_k + KK_{ik}) \\ & - 8\pi\alpha\left(S_{ik} - \frac{1}{2}\gamma_{ik}(S - E)\right); \end{aligned} \quad (2.53b)$$

$${}^{(3)}R + K^2 - K_{ik}K^{ik} = 16\pi E; \quad (2.53c)$$

$$D_i K - D_k K^k{}_i = 8\pi j_i, \quad (2.53d)$$

where $S = \gamma^{ij}S_{ij}$.

These equations, known as the ADM equations, provide the wanted initial-value problem for the Einstein equations. Initial data can be given on a Cauchy hypersurface by specifying a spatial metric and its extrinsic curvature, subject to the constraints equations (2.53c) and (2.53d). Finally the metric on the whole spacetime can be obtained by evolving the equations (2.53a) and (2.53b).

2.1.4 Strongly Hyperbolic Formulations of the Einstein Equations

The ADM system in its original form (2.53) is never used in recent numerical relativity codes as it turns out to be unstable even for the evolution of perturbed Minkowsky spacetimes⁴. The reasons for that are not yet fully understood [38], but it appears clear that most of the issues are due to the fact that the ADM system is only weakly hyperbolic and this leads to numerical problems as the errors tend to couple with zero-velocity modes. This was shown in Alcubierre et al. [10] where the authors studied the linearized ADM system on the Minkowsky background. An intuitive picture of why zero-speed modes can lead to instabilities can be gained by studying the toy model equation

$$\partial_t u - \lambda \partial_x u = \delta u, \quad (2.54)$$

where δu , with $0 < \delta \ll 1$, can be thought of as a source of numerical errors. When λ goes to 0, u grows exponentially, signaling the instability of the system.

The solution to this problem was found by looking at alternative formulations of the initial-value problem for the Einstein equations. Among the different formulations some of the most successful were the Baumgarte, Shapiro, Shibata, Nakamura, Oohara and Kojima formulation, BSSNOK (often also referred to as BSSN or conformal traceless formalism) [233, 294, 37], the generalized-harmonic formulation [133, 200, 199] and the Z4 formulation [55, 47, 281, 332, 12]. We do not discuss all these formulations in detail as this would take us too far from the goals of this thesis. We only point out that each of these formulations has its strengthes and weaknesses and that the search for new and improved formulations of the Einstein equations is still a lively research topic. Here we sketch the BSSNOK and the conformal-covariant variant of the Z4 formulation (also called CCZ4) as these are relevant for the work presented here.

The BSSNOK Formulation

The BSSNOK formulation is a modification of the original ADM formulation, where auxiliary variables are introduced to obtain a strongly-hyperbolic system.

First of all the metric is split into a conformal metric

$$\tilde{\gamma}_{ij} = e^{-4\phi} \gamma_{ij} \quad (2.55)$$

and a conformal factor

$$\phi = \frac{1}{12} \log \gamma. \quad (2.56)$$

Then the extrinsic curvature is split into its trace K and a traceless part:

$$\tilde{A}_{ij} = e^{-4\phi} \left(K_{ij} - \frac{1}{3} \gamma_{ij} K \right). \quad (2.57)$$

Finally one defines

$$\Gamma^i = \gamma^{jk} \Gamma^i_{jk}, \quad \tilde{\Gamma}^i = \tilde{\gamma}^{jk} \tilde{\Gamma}^i_{jk}, \quad (2.58)$$

⁴An important exception is the case of spherically symmetric spacetimes.

where Γ^i_{jk} is the Levi-Civita connection associated with γ_{ij} .

Alternative formulations exist where slightly different variables are used. For instance Campanelli et al. [73] proposed to evolve $\chi = e^{-4\phi}$ instead of ϕ , to take advantage of the fact that χ is always non-singular.

The evolution equations are then obtained after a rather lengthy calculation starting from the ADM equations and of which we only report the final result:

$$(\partial_t - \mathcal{L}_{\vec{\beta}})\tilde{\gamma}_{ij} = -2\alpha\tilde{A}_{ij}, \quad (2.59a)$$

$$(\partial_t - \mathcal{L}_{\vec{\beta}})\phi = -\frac{1}{6}\alpha K, \quad (2.59b)$$

$$(\partial_t - \mathcal{L}_{\vec{\beta}})\tilde{A}_{ij} = e^{-4\phi} \left[-D_i D_j \alpha + \alpha({}^{(3)}R_{ij} - 8\pi S_{ij}) \right]^{\text{TF}} + \alpha(K\tilde{A}_{ij} - 2\tilde{A}_{il}\tilde{A}^l_j), \quad (2.59c)$$

$$(\partial_t - \mathcal{L}_{\vec{\beta}})K = -\gamma^{ij}D_i D_j \alpha + \alpha[\tilde{A}_{ij}\tilde{A}^{ij} + \frac{1}{3}K^2 + 4\pi(E + S)], \quad (2.59d)$$

$$(\partial_t - \mathcal{L}_{\vec{\beta}})\tilde{\Gamma}^i = -2\tilde{A}^{ij}\partial_j \alpha - \frac{2}{3}\tilde{\Gamma}^i \partial_j \beta^j + \frac{1}{3}\tilde{\gamma}^{li}\partial_j \partial_l \beta^j + \tilde{\gamma}^{lj}\partial_l \partial_j \beta^i \quad (2.59e)$$

$$+ 2\alpha(\tilde{\Gamma}^i_{jk}\tilde{A}^{kj} - \frac{2}{3}\tilde{\gamma}^{ij}\partial_j K - 8\pi\tilde{\gamma}^{ij}S_j + 6\tilde{A}^{ij}\partial_j \phi); \quad (2.59f)$$

where $\cdot^{\text{TF}}$ denotes the trace-free part of $\cdot$.

The CCZ4 Formulation

The Z4 formulation can be obtained from the covariant Lagrangian

$$\Lambda = g^{\mu\nu}[R_{\mu\nu} + 2\nabla_\mu Z_\nu] \sqrt{-g} + \Lambda_m, \quad (2.60)$$

by means of a Palatini-type variational principle [56]. The variational principle yields the field equations

$$R_{\mu\nu} + \nabla_\mu Z_\nu + \nabla_\nu Z_\mu = 8\pi \left(T_{\mu\nu} - \frac{1}{2}Tg_{\mu\nu} \right), \quad (2.61)$$

as well as a set of constraints fixing the connection

$$\nabla_\rho g^{\mu\nu} = 0, \quad (2.62)$$

and the algebraic constraint

$$Z_\mu = 0. \quad (2.63)$$

If (2.63) is satisfied then (2.61) and (2.62) reduce to the standard Einstein field equations. Otherwise Z_μ gives a measure of the deviation of the solution from the one of the original Einstein equations. In addition we point out that it is possible to show that the condition that the first derivatives of Z_μ vanish amounts to imposing the ADM momentum and Hamiltonian constraints [57].

The key idea of the Z4 formalism is to develop a set of evolution equation starting from the Lagrangian (2.60), *without explicitly enforcing* (2.63), *i.e.*, treating Z_μ as a new independent variable. The resulting set of equations is then strongly-hyperbolic, *i.e.*, free from the zero-speed modes of the original ADM system, and the solution of the Einstein equations is obtained exploiting the fact that the Z4 evolution system preserves the constraint (2.63), *i.e.*, $\partial_t(Z_\mu) = 0$. In

particular, if the initial-data is constraint satisfying, the Z4 evolution recovers the solution of the Einstein equations, even though Z_μ is an evolved variable. In practice small numerical errors introduce constraint violations during the evolution, for this reason the Z4 system is usually modified, with the addition of terms that cancel out in the case in which the constraints are satisfied, to ensure that eventually constraint violations are propagated away and exponentially damped [160].

The version of Z4 that we employ was recently introduced by Alic et al. [12] and is based on a conformal decomposition of the original Z4 system aiming to incorporate both the advantages of BSSNOK, and in particular the possibility of treating black-holes using the moving-puncture approach (more on this in the next session), with the constraint-damping properties of the original Z4 formulation. The CCZ4 system reads

$$\partial_t \tilde{\gamma}_{ij} = -2\alpha \tilde{A}_{ij}^{\text{TF}} + 2\tilde{\gamma}_{k(i}\partial_{j)}\beta^k - \frac{2}{3}\tilde{\gamma}_{ij}\partial_k\beta^k + \beta^k\partial_k\tilde{\gamma}_{ij}, \quad (2.64a)$$

$$\begin{aligned} \partial_t \tilde{A}_{ij} &= \phi^2 \left[-\nabla_i \nabla_j \alpha + \alpha \left({}^{(3)}R_{ij} + \nabla_i Z_j + \nabla_j Z_i - 8\pi S_{ij} \right) \right]^{\text{TF}} \\ &\quad + \alpha \tilde{A}_{ij} (K - 2\Theta) - 2\alpha \tilde{A}_{il} \tilde{A}_j^l + 2\tilde{A}_{k(i}\partial_{j)}\beta^k \\ &\quad - \frac{2}{3}\tilde{A}_{ij}\partial_k\beta^k + \beta^k\partial_k\tilde{A}_{ij}, \end{aligned} \quad (2.64b)$$

$$\partial_t \phi = \frac{1}{3}\alpha\phi K - \frac{1}{3}\phi\partial_k\beta^k + \beta^k\partial_k\phi, \quad (2.64c)$$

$$\begin{aligned} \partial_t K &= -\nabla^i \nabla_i \alpha + \alpha \left({}^{(3)}R + 2\nabla_i Z^i + K^2 - 2\Theta K \right) + \beta^j \partial_j K - \\ &\quad 3\alpha\kappa_1 (1 + \kappa_2) \Theta + 4\pi\alpha (S - 3E), \end{aligned} \quad (2.64d)$$

$$\begin{aligned} \partial_t \Theta &= \frac{1}{2}\alpha \left(R + 2\nabla_i Z^i - \tilde{A}_{ij}\tilde{A}^{ij} + \frac{2}{3}K^2 - 2\Theta K \right) - Z^i \partial_i \alpha + \\ &\quad \beta^k \partial_k \Theta - \alpha\kappa_1 (2 + \kappa_2) \Theta - 8\pi\alpha E, \end{aligned} \quad (2.64e)$$

$$\begin{aligned} \partial_t \hat{\Gamma}^i &= 2\alpha \left(\tilde{\Gamma}_{jk}^i \tilde{A}^{jk} - 3\tilde{A}^{ij} \frac{\partial_j \phi}{\phi} - \frac{2}{3}\tilde{\gamma}^{ij} \partial_j K \right) + 2\tilde{\gamma}^{ki} \left(\alpha \partial_k \Theta - \Theta \partial_k \alpha \right. \\ &\quad \left. - \frac{2}{3}\alpha K Z_k \right) - 2\tilde{A}^{ij} \partial_j \alpha + \tilde{\gamma}^{kl} \partial_k \partial_l \beta^i + \frac{1}{3}\tilde{\gamma}^{ik} \partial_k \partial_l \beta^l + \frac{2}{3}\tilde{\Gamma}^i \partial_k \beta^k \\ &\quad - \tilde{\Gamma}^k \partial_k \beta^i + 2\kappa_3 \left(\frac{2}{3}\tilde{\gamma}^{ij} Z_j \partial_k \beta^k - \tilde{\gamma}^{jk} Z_j \partial_k \beta^i \right) + \beta^k \partial_k \hat{\Gamma}^i \\ &\quad - 2\alpha\kappa_1 \tilde{\gamma}^{ij} Z_j - 16\pi\alpha \tilde{\gamma}^{ij} S_j, \end{aligned} \quad (2.64f)$$

where Θ is the projection of the Z4 four-vector along the normal direction, $\Theta := n_\mu Z^\mu = \alpha Z^0$.

The three-dimensional Ricci tensor ${}^{(3)}R_{ij}$ is split into a part containing conformal terms $\tilde{R}_{ij}^\phi$ and another one containing space derivatives of the conformal metric $\tilde{R}_{ij}$, defined as

$$\tilde{R}_{ij} = -\frac{1}{2}\tilde{\gamma}^{lm}\partial_l\partial_m\tilde{\gamma}_{ij} + \tilde{\gamma}_{k(i}\partial_{j)}\tilde{\Gamma}^k + \tilde{\Gamma}^k\tilde{\Gamma}_{(ij)k} + \tilde{\gamma}^{lm}\left[2\tilde{\Gamma}_{l(i}^k\tilde{\Gamma}_{j)km} + \tilde{\Gamma}_{im}^k\tilde{\Gamma}_{kjl}\right], \quad (2.65a)$$

$$\tilde{R}_{ij}^\phi = \frac{1}{\phi^2}\left[\phi\left(\tilde{\nabla}_i\tilde{\nabla}_j\phi + \tilde{\gamma}_{ij}\tilde{\nabla}^l\tilde{\nabla}_l\phi\right) - 2\tilde{\gamma}_{ij}\tilde{\nabla}^l\phi\tilde{\nabla}_l\phi\right]. \quad (2.65b)$$

The evolution variable Z_i of the original Z4 formulation is now included in

the $\hat{\Gamma}^i$ variable of the CCZ4 formulation

$$\hat{\Gamma}^i := \tilde{\Gamma}^i + 2\tilde{\gamma}^{ij}Z_j, \quad (2.66)$$

where

$$\tilde{\Gamma}^i := \tilde{\gamma}^{jk}\tilde{\Gamma}_{jk}^i = \tilde{\gamma}^{ij}\tilde{\gamma}^{kl}\partial_l\tilde{\gamma}_{jk}. \quad (2.67)$$

Finally κ_1 and κ_2 are constants associated with the constraint damping terms and κ_3 is an extra constant used to select among different variants of the formulation.

A technical point to keep in mind is the following. In the BSSNOK formulation the Hamiltonian constraint is used to eliminate the Ricci scalar from the right-hand-side (RHS) of K . For this reason, in the case in which the constraint violation is non-zero, the trace of the extrinsic curvature used in the CCZ4 formulation might be different from the one used in BSSNOK and in particular one has

$$K^{\text{BSSNOK}} = K - 2\Theta. \quad (2.68)$$

2.1.5 Gauge conditions

The ADM equations (as well as the BSSNOK or CCZ4 ones) do not prescribe the evolution of α and β^i so, in order to successfully evolve these systems, it is necessary to choose a prescription for the lapse and the shift. This is not surprising as the prescription of the spacetime foliation is clearly a gauge.

The right choice of the prescription for the lapse, the *slicing condition*, and for the shift, the *spatial gauge condition*, is critical for the stable evolution of non-trivial spacetimes and there is indeed a whole branch of numerical relativity that is dedicated to the study of suitable gauge conditions. Here we are going to present only a few concepts needed in our case.

Slicing conditions

The most simple choice is *geodesic slicing*: $\alpha = 1$ and $\beta^i = 0$. But this has obvious limitations especially in the case of singularities and is never used.

Another classical slicing condition is the maximal slicing condition, $K = 0$, that implies:

$$D^i D_i \alpha = \alpha [K_{ij} K^{ij} + 4\pi(e + S)]. \quad (2.69)$$

This slicing condition has the interesting property of being *singularity-avoiding*, that is the lapse collapses to zero near singularities so that spacetime singularities are not met by the slicing and thus the numerical grid. For example maximal slicing in the case of a Schwarzschild black hole has been studied by Geyer and Herold [142]. They showed in particular that the spacetime slices converge to the maximal hypersurface $r = 3/2M$, remaining at finite distance from the singularity.

The maximal slicing has interesting mathematical characteristics, but involves the solution of elliptic equations at each time step and thus it is too expensive for practical applications. For this reason Bona and Masso [54] developed a new class of slicing conditions in the form of hyperbolic equations for

the lapse, that are able to mimic the behaviour of the maximal slicing condition⁵:

$$(\partial_t - \beta^i \partial_i) \alpha = -\alpha^2 f(\alpha) K, \quad (2.70)$$

where $f(\alpha)$ is a positive function. When $f = 1$, one gets geodesic slicing, while with $f \rightarrow \infty$ one formally gets the maximal slicing [38].

A popular choice is $f(\alpha) = 2/\alpha$, the so-called “1 + log” slicing. The name is a consequence of the fact that, when $\beta^i = 0$, (2.70) can be integrated to obtain

$$\alpha = 1 + \log \gamma. \quad (2.71)$$

This condition has excellent singularity avoiding properties because $f \rightarrow \infty$ near singularities so that it behaves like a maximal slicing condition [38].

Spatial gauge conditions

Gauge conditions for the shift function are also very important, especially in the case of compact binaries where proper spatial gauge conditions have been a key component for successful stable simulations.

A classical shift condition is the so called *Gamma-driver* condition [11]:

$$\partial_t \beta^i = \frac{3}{4} \alpha B^i, \quad (2.72)$$

$$\partial_t B^i = \partial_i \tilde{\Gamma}^i - \eta B^i; \quad (2.73)$$

where η is a dumping coefficient that must be tuned for optimal performance.

This shift condition is the hyperbolic version of an elliptic gauge condition called “minimal distortion” because it tried to minimize the stretching of the coordinates, in a fashion similar to the Bona-Masso slicing conditions for the lapse [38], and it is intended to avoid the development of large shears in the metric near singularities. The Gamma-driver condition has been successfully used in the evolution of single moving black holes, but has a zero-speed mode which, as we have seen, can couple with the numerical errors and destabilize the system [326].

Van Meter [326] carried out an extensive study of various prescriptions looking for a gauge condition not having zero or small speed modes. They proposed a modified Gamma-driver (or moving-puncture) gauge:

$$(\partial_t - \beta^j \partial_j) \beta^i = \frac{3}{4} B^i, \quad (2.74a)$$

$$(\partial_t - \beta^j \partial_j) B^i = (\partial_t - \beta^j \partial_j) \tilde{\Gamma}^i - \eta B^i; \quad (2.74b)$$

which was successfully used to evolve binary black holes [73].

2.2 The Equations of General-Relativistic Hydrodynamics

In this section we introduce the equations of general relativistic hydrodynamics. These are the equations describing the flow of a fluid on a Lorentzian manifold

⁵Note that this becomes $(\partial_t - \beta^i \partial_i) \alpha = \alpha^2 f(\alpha) (K - 2\Theta)$ in the CCZ4 formulation.

coupled to the Einstein equations describing how the geometry of the spacetime changes in response to the changes in the fluid. We start from the kinematics of a fluid, *i.e.*, from the mathematical tools necessary to describe the state and the motion of the fluid. Then we present the equations of motion for perfect fluids, *i.e.*, in the absence of thermal conduction or viscosity. Finally we recall the so-called “Valencia formulation” of the hydrodynamic equations, which is the one most commonly employed in numerical relativity. Our presentation loosely follows the ones by [229, 290, 158, 17, 275] and we refer the reader to those sources for a more extended treatment.

2.2.1 Kinematics of a Relativistic Fluid

Broadly speaking a fluid in Newtonian physics is something that *flows*, *i.e.*, a system whose thermodynamical state is entirely determined as soon as quantities like its mass, momentum and energy densities are known and whose dynamics is described in terms of the respective flows. In special and general relativity quantities such as the mass or the energy density of a fluid at a particular point are ill-defined, as their value would depend on the observer, *i.e.*, there is no such a thing as a thermodynamical state of a fluid upon which all the observers can agree. For this reason, while Newtonian hydrodynamics describes the evolution of these quantities in terms of their flows, in relativistic hydrodynamics this is not possible. In special and general relativity the fundamental quantities describing a fluid are not scalar or vector fields determining the state of the fluid, but directly their flow, in space and time. In this sense flows and flux-conservative formulations of the equations of hydrodynamics play an even more fundamental role in relativity than in classical physics.

First of all, in classical hydrodynamics the mass density (or simply density) is described by a scalar quantity, usually denoted by ρ . In the relativistic case one could be tempted to define, as in the Newtonian case, $\rho := mn$, where m is the rest-mass of a single particle⁶ and n is the number density of particles, *i.e.*, the number of particles in a volume d^3x so that

$$\int_V \rho d^3x = m \int_V n d^3x =: mN \quad (2.75)$$

would give the total mass of fluid enclosed in a volume V . The problem with this definition is that d^3x is an observer dependent quantity. This means that two observers in relative motion with respect to each other will not agree on the measure of n (but will agree on the measure of N). One could bypass the problem by defining the density as being the one measured in a particular frame, for instance the one instantaneously comoving (or simply comoving) with the fluid. While the density in the comoving frame is an interesting quantity, it is clear that, to obtain a covariant description of the dynamics of the fluid, it is necessary to use a formulation in terms of invariant quantities, from which frame-dependent quantities can be constructed once an observer is selected.

A covariant description of the mass density of a fluid can be done once one realizes that the truly fundamental quantity is not the mass density, but its flow in spacetime, *i.e.*, a three pseudo-form $\boldsymbol{\rho}$, whose value on any three dimensional submanifold gives the flow of mass transverse to the submanifold. In particular

⁶For simplicity we assume the fluid to be composed of identical particles.

if Σ is a spacelike hypersurface with normal $\vec{n}$, which we take future-oriented, then the density as measured by an observer with four-velocity $\vec{n}$ is given by

$$\int_{\Sigma} \rho. \quad (2.76)$$

The mass-flow measured by the Eulerian observer across any spacelike surface $\Omega \subset \Sigma$ is given by

$$\int_{\Omega} i_{\vec{n}} \rho, \quad (2.77)$$

where $i_{\vec{n}} \rho$ denotes the interior product between the vector $\vec{n}$ and the form ρ , *i.e.*, $i_{\vec{n}} \rho$ is the two-form $\rho(\vec{n}, \cdot, \cdot)$.

The statement that the number of particles of the fluid is conserved is expressed by the fact that the exterior differential of ρ vanishes

$$d\rho = 0, \quad (2.78)$$

or equivalently by the statement that, for any sufficiently regular, open set $\Omega \subset \mathcal{M}$, we have

$$\int_{\partial\Omega} \rho = \int_{\Omega} d\rho = 0, \quad (2.79)$$

which means that the total number of particles is conserved, since the net flow across any regular surface enclosing a four-dimensional open set is zero.

On any pseudo-Riemannian manifold, associated with any three form there exists a vector field, such that the three form can be interpreted as describing the flux of this field in a sense that we make more precise in a moment. The vector field associated with the density three-form is usually denoted with $\vec{J}$ and called the rest-mass density four-vector. It is constructed by requiring that

$$\rho = i_{\vec{J}} \text{Vol}_x^4, \quad (2.80)$$

where Vol_x^4 is the volume pseudo-form of the spacetime, *i.e.*, in any local coordinate patch we have

$$\text{Vol}_x^4 = \sqrt{-g} dx^0 \wedge dx^1 \wedge dx^2 \wedge dx^3, \quad (2.81)$$

where g is the determinant of the spacetime metric. More precisely $\vec{J}$ is constructed from the one-form $\underline{J} = \star \rho$, where $\star$ is the Hodge dual operator, see *e.g.*, [132], by raising its indices, *i.e.*, $\vec{J} = \# \underline{J}$.

The statement that ρ represents the flux of $\vec{J}$ reads

$$\int_{\Sigma} \rho = - \int_{\Sigma} \vec{J} \cdot \vec{n} \text{Vol}_x^3, \quad (2.82)$$

where $\vec{n}$ is the future-oriented unit-timelike normal to Σ and Vol_x^3 is the intrinsic volume form on the submanifold Σ , *i.e.*, $\text{Vol}_x^3 = i_{\vec{n}} \text{Vol}_x^4$.

More generally the flux associated with a flow defined by a vector field, $\vec{X}$, across a hypersurface, Σ , transverse to it and with normal $\vec{v}$ (with appropriate

sign depending on the signature of the metric and on Σ), is given by⁷

$$\int_{\Sigma} \star \underline{X} = \int_{\Sigma} i_{\vec{X}} \text{Vol}^n = \int_{\Sigma} i_{\vec{X}} [\underline{v} \wedge \text{Vol}^{n-1}] = \int_{\Sigma} \vec{X} \cdot \vec{v} \text{Vol}^{n-1}. \quad (2.83)$$

It is easy to see that, since the flux of particles across any future-oriented spacelike hypersurface must be positive or null, $\vec{J}$ must be a timelike or null vector. In the former case it can be decomposed in a unique way as

$$\vec{J} =: \rho \vec{u} \quad (2.84)$$

where ρ is a scalar that we can interpret as being the comoving density, *i.e.*, the density in the comoving frame, and $\vec{u}$ is a unit-timelike vector that we can interpret as being the fluid four-velocity.

The statement of mass conservation, as expressed in terms of $\vec{J}$, becomes, in any coordinate frame, the well known continuity equation:

$$0 = \nabla_{\mu} J^{\mu} = \frac{1}{\sqrt{-g}} \partial_{\mu} [\sqrt{-g} \rho u^{\mu}]. \quad (2.85)$$

Similarly energy and momentum of a fluid can be defined, using the Cartan formalism, as a three pseudo-form valued form, *i.e.*, a one-form that, when acting on the vector $\vec{\partial}_{\mu}$ returns a three pseudoform representing the flow of the μ -momentum of the fluid, $\tau \otimes \underline{\alpha}$

$$\tau_{\mu} := [\tau \otimes \underline{\alpha}](\vec{\partial}_{\mu}). \quad (2.86)$$

More simply, exploiting the equivalence between three-forms and vectors, we can describe energy and momentum flows using a mixed tensor T such that

$$T^{\mu}_{\nu} = T(\text{dx}^{\mu}, \partial_{\nu}) = \text{"Flow of the } \nu \text{ momentum across the volume element perpendicular to } \text{dx}^{\mu}\text{"}. \quad (2.87)$$

It is not difficult to show that with this definition of the stress energy tensor coincides with the definition of stress-energy tensor given by (2.12).

As a consequence of the Bianchi identities the divergence of the stress-energy tensor of the fluid must vanish if the Einstein equation are satisfied:

$$0 = \nabla_{\mu} T^{\mu}_{\nu} = \frac{1}{\sqrt{-g}} \partial_{\mu} (\sqrt{-g} T^{\mu}_{\nu}) - \Gamma^{\alpha}_{\mu\nu} T^{\mu}_{\alpha}. \quad (2.88)$$

Although these equations are often described as a statement of conservation of the energy and momentum of the fluid, we must remark that (2.88) is not properly describing the conservation of energy and momentum of the fluid. Actually these quantities are not, in general, conserved. Conservation of the ν -momentum would, instead, read

$$d\tau_{\nu} = 0 \quad (2.89)$$

and it is easy to show that (2.88) and (2.89) agree, for a fixed ν , only if $\vec{\partial}_{\nu}$ is a Killing vector.

⁷We remember that the volume form of the spacetime can be decomposed as $\text{Vol}_x^4 = \underline{v} \wedge \text{Vol}_x^3$.

2.2.2 Dynamics of a Relativistic Fluid

So far we have been concerned only with the description of the kinematics of a relativistic fluid. We derived the equations (2.85) and (2.88) describing mass, energy and momentum “conservation”, but clearly these are not enough to fully determine the motion of the fluid. This is not surprising since we have not yet specified any of the physical properties of the fluid that we want to describe. In particular, in this thesis we are concerned with *perfect fluids*, *i.e.*, fluids where, in the comoving frame,

1. there is no heat conduction;
2. there is no viscosity.

The absence of heat conduction means that the medium is in local thermodynamical equilibrium (LTE). Viscosity is neglected simply because it is expected to be very small for the systems of interest, even when compared with the “numerical viscosity”, and because the correct mathematical formulation of viscous and/or thermally conducting fluids in general-relativity is still the object of research, see *e.g.*, [17] and references therein.

The stress-energy tensor of a perfect fluid is easily calculated in the comoving frame with the fluid. To do that we construct an orthonormal tetrad $\{\vec{u}, \vec{e}_i\}$ comoving with the fluid, where $\vec{u}$ is the four-velocity of the fluid, as defined in (2.84) $\vec{u} \cdot \vec{e}_i = 0$ and $\vec{e}_i \cdot \vec{e}_k = \delta_{ik}$. We also introduce the dual basis $\{\underline{u}, \underline{e}^i\}$. At this point the $T(\underline{u}, \vec{u})$ component of the stress-energy tensor is the energy-density in the rest-frame of the fluid, which we denote as e . The mixed components of the stress energy tensor, $T(\underline{u}, \vec{e}_i)$, represent the energy flowing transverse to the four-velocity of the fluid and are zero since we assumed that there is no heat-conduction. Finally the spatial components of the stress-energy tensor $T(\underline{e}^i, \vec{e}_k)$ represent the k component of the force exchanged across the surface element orthogonal to $\underline{e}^i$. Since we neglect viscosity and since the stress-energy tensor must be invariant with respect to rotations of the triad $\{\vec{e}_i\}$, we have

$$T(\underline{e}^i, \vec{e}_k) = p \delta^i_k, \quad (2.90)$$

for some scalar p , which we call pressure. Collecting all these considerations together we get the following form for the stress-energy tensor

$$T = (e + p) \vec{u} \otimes \underline{u} + p \delta. \quad (2.91)$$

It is customary to define the specific enthalpy of the fluid $h = 1 + e + p/\rho$, where e is the specific internal energy, *i.e.*, $e = \rho(1 + \epsilon)$, so that the stress-energy tensor is written as

$$T = \rho h \vec{u} \otimes \underline{u} + p \delta. \quad (2.92)$$

The equations (2.85) and (2.88) together with the form of the stress-energy tensor (2.92) are still not sufficient to determine the motion of the fluid. In order to do that it is necessary to also prescribe an equation of state for the fluid, *i.e.*, a relation between pressure, internal energy and density.

Typical equations of state that we use throughout this work are:

1. the ideal-gas, or gamma-law equation of state (EoS)

$$p = (\Gamma - 1) \rho \epsilon, \quad (2.93)$$

where Γ is the polytropic index of the gas;

2. the polytropic EoS

$$p = K \rho^\Gamma, \quad (2.94)$$

which is just the ideal-gas EoS restricted to the case of isoentropic flows;

3. the ultrarelativistic EoS

$$p = (\Gamma - 1) e, \quad (2.95)$$

which can be interpreted as the ultra-high temperature limit of the ideal-gas EoS.

Note that in the case of the polytropic EoS the energy equation becomes redundant, while in the case of the ultrarelativistic EoS it is the continuity equation which becomes redundant. For these reasons the former EoS is often employed in situations where the flow is isoentropic to simplify the evolution, while the former can be used to described a hot, optically thick pair-plasma, since pair-creation is automatically taken into account. More on the properties of these EoS can be found, for instance, in [18].

With the addition of an equation of state, equations (2.59), (2.85) and (2.88) together with the form of the stress-energy tensor (2.92) form a hyperbolic system of equations that can be evolved, once initial data is prescribed on a Cauchy surface, to obtain both the geometry of the spacetime and the dynamics of the matter in it.

2.2.3 Conservative Formulations

The first attempts to solve the equations of general relativistic hydrodynamics date back to the seminal works by May and White [213] and by Wilson [333]. In these investigations, the approach was to cast the relativistic hydrodynamics equations as non-linear advection-like equations in a form that resembles the Newtonian Euler equations. These were then solved using finite-difference (FD) schemes, stabilized using a combination of upwinding and artificial-viscosity methods to avoid excessive oscillations at shocks (see [130] for a comprehensive list of references). Although these methods allowed to perform the first numerical studies in general relativistic hydrodynamics, they also had several limitations, such as the difficulty of tuning the artificial viscosity to avoid excessive smearing of the shock fronts or, most importantly, the limitation to mildly relativistic flows, *i.e.*, with Lorentz factor $W \lesssim 2$ [130].

A major leap forward in numerical relativistic hydrodynamics took place when it was realized that the major problem behind Wilson's approach was the use of a formulation which breaks the conservative nature of the equations [210]. This realization lead to the formulation of the equations of relativistic hydrodynamics in conservation form, *i.e.*, as a system of equations of the type

$$\frac{\partial \mathbf{F}^0(\mathbf{u})}{\partial t} + \frac{\partial \mathbf{F}^i(\mathbf{u})}{\partial x^i} = \mathbf{S}(\mathbf{u}), \quad (2.96)$$

where $\mathbf{u}$ is a “vector” of *primitive quantities*, such as the rest-mass density or the specific internal energy, $\mathbf{F}^0$ is a “vector” of *conserved quantities*⁸ and $\mathbf{F}^i$ and $\mathbf{S}$ are their *fluxes* and *sources* respectively.

⁸These are not strictly conserved if $\mathbf{S}$ is not zero.

In particular the “Valencia formulation” [36], which is adapted for the use in conjunction with the 3 + 1 formalism outlined above, was very successfully used in conjunction with special numerical methods developed in the context of conservation laws, such as finite-volume (FV) and FD high-resolution shock capturing (HRSC) methods, which we discuss in more detail in Chapter 3. These methods were shown to be able to handle ultra-relativistic flows and to sharply resolve shocks without spurious oscillations or the need for artificial viscosity. For these reasons they have been the key ingredient in a number of recent achievements of numerical relativistic hydrodynamics and magnetohydrodynamics (MHD; see, *e.g.*, [146, 274] and references therein). Other conservative formulations of the equations of general relativistic hydrodynamics are available, for instance the one proposed by Papadopoulos and Font [248], but, here, we restrict ourselves to the Valencia formulation as it is the one relevant for this thesis.

First of all we decompose the four-velocity $\vec{u}$ in a part which is parallel to $\vec{n}$ and a part perpendicular to it, *i.e.*, purely spatial

$$\vec{u} = (-\vec{u} \cdot \vec{n})(\vec{n} + \vec{v}), \quad (2.97)$$

where $-\vec{u} \cdot \vec{n}$ is obviously the Lorentz factor of the fluid as measured by the Eulerian observer, which we denote as W , and $\vec{v}$ is the three-velocity of the fluid as measured by the Eulerian observer,

$$\vec{v} = \frac{\vec{u}}{W} - \vec{n}, \quad (2.98)$$

or, in component form

$$v^i = \frac{u^i}{W} + \frac{\beta^i}{\alpha}, \quad v_i = \frac{u_i}{W}. \quad (2.99)$$

The continuity equation is then easily written in conservative form

$$0 = \nabla_\mu J^\mu = \frac{1}{\sqrt{-g}} \partial_t [\sqrt{\gamma} \rho W] + \frac{1}{\sqrt{-g}} \partial_i [\sqrt{\gamma} \rho (\alpha v^i - \beta^i)] \quad (2.100)$$

with conserved density $D = \rho W = -\vec{j} \cdot \vec{n}$.

The energy and momentum equations can be written in conservation form noting that [275]

$$\nabla_\mu [T^\mu_\nu p^\nu] = T^\mu_\nu \nabla_\mu p^\nu, \quad (2.101)$$

for any vector field $\vec{p}$. The Valencia formulation is then obtained by letting $\vec{p}$ vary among the vectors $\{\vec{n}, \vec{\partial}_i\}$. In this way one obtains four conservation laws with conserved quantities, $T^0_\nu p^\nu$,

$$S_i = \alpha T^0_\nu (\partial_i)^\nu = -\mathbf{T}(\vec{n}, \vec{\partial}_i), \quad E = \alpha T^0_\nu n^\nu = \mathbf{T}(\vec{n}, \vec{n}), \quad (2.102)$$

and associated fluxes $T^i_\nu p^\nu$ and sources $T^\mu_\nu \nabla_\mu p^\nu$.

In addition we introduce the quantity $\tau = E - D$, that is the difference between the total internal energy density and the rest-mass density, which we evolve in place of E to avoid large errors in the internal energy density in the case in which this is small compared to the rest-mass density.

Collecting everything together we can write the general relativistic hydrodynamics equations in the form

$$\frac{1}{\sqrt{-g}} \left[\frac{\partial \sqrt{\gamma} F^0(\mathbf{u})}{\partial t} + \frac{\partial \sqrt{-g} F^i(\mathbf{u})}{\partial x^i} \right] = \mathbf{S}(\mathbf{u}), \quad (2.103)$$

with primitive quantities

$$\mathbf{u} = [\rho, v_i, \epsilon] \quad (2.104)$$

and conservative quantities

$$\mathbf{F}^0(\mathbf{u}) = [D, S_j, \tau] = [\rho W, \rho h W^2 v_j, \rho h W^2 - p - \rho W]. \quad (2.105)$$

The fluxes are

$$\mathbf{F}^i(\mathbf{u}) = \left[D \left(v^i - \frac{\beta^i}{\alpha} \right), S_j \left(v^i - \frac{\beta^i}{\alpha} \right) + p \delta_j^i, \tau \left(v^i - \frac{\beta^i}{\alpha} \right) + p v^i \right] \quad (2.106)$$

and the sources are

$$\mathbf{S}(\mathbf{u}) = \left[0, T^{\mu\nu} \left(\frac{\partial g_{\nu j}}{\partial x^\mu} - \Gamma_{\nu\mu}^\delta g_{\delta j} \right), \alpha \left(T^{\mu 0} \frac{\partial \log \alpha}{\partial x^\mu} - T^{\mu\nu} \Gamma_{\nu\mu}^0 \right) \right]^T. \quad (2.107)$$

The Valencia formulation casts the equations of general relativistic hydrodynamics in a form akin to the one usually employed in the numerical solution of the equations of gas-dynamics in the Newtonian case. There are, however, at least two important differences with respect to the classical case that appear at a first inspection of the equations. The first one is that it is not possible to find an explicit inverse relation between $\mathbf{u}$ and $\mathbf{q}$. As a consequence the primitive variables need to be recovered from the conservative ones by means of a numerical root-finding algorithm (more on this in Chapters 6, 4 and 5). The second one is that the Lorentz factor couples the equation for the momenta in the different directions in a way which has no classical counterpart [257, 272, 273, 15]. For instance the presence of a tangential velocity can change completely the dynamics of a shock wave. This makes the equations of general relativistic hydrodynamics even more challenging to solve [224, 339].

2.3 The General-Relativistic Boltzmann Equation

The foundations of the Boltzmann equation in special relativity were laid by Synge [311], while the extension to general relativity was suggested by Tauber and Weinberg [315] and Chernikov [84]. The relativistic Boltzmann equation was first applied to the study of a simple relativistic gas [174], and later to the study of transient relativistic thermodynamics [173], radiative transfer [201] [321] in core-collapse supernovae [67], and in many other scenarios (see *e.g.*, [79] and references therein).

Unfortunately the current literature on the general relativistic Boltzmann equation is rather sparse and a number of subtly different conventions have been used in different works, to the point that it is often not simple to compare the different formalisms. For this reason, very recently, Debbasch and van

Leuwwen [108, 109] tried to put order in the current understanding of the equation by publishing its detailed derivation. Unfortunately their work is overly concerned with the algebraic aspects of the derivation, hiding the simple geometrical interpretation proposed by Lindquist [201] or Ehlers [124] and based on earlier work on the Riemannian structure of tangent bundles by Sasaki [284] [285]. For this reason we think that it is useful to present a detailed derivation of the general relativistic Boltzmann equation, using modern differential geometry notation.

2.3.1 The geometry of the tangent bundle

Let $(\mathcal{M}, g_{\alpha\beta})$ be a C^2 spacetime, $T\mathcal{M}$ be its tangent bundle and $\pi: T\mathcal{M} \rightarrow \mathcal{M}$ be the projection map.

Extended coordinates on $T\mathcal{M}$

In every coordinate patch U , $\{x^\alpha\}$, of $\mathcal{M}$, for every point $q \in U$ we denote a vector in $T_q\mathcal{M}$ as

$$\vec{p} = p^\alpha \frac{\partial}{\partial x^\alpha} \quad (2.108)$$

and its dual as

$$\underline{p} = p_\alpha dx^\alpha := g_{\alpha\beta} p^\beta dx^\alpha. \quad (2.109)$$

We also introduce a coordinate patch TU , $\{z^A\}$, $A = 0, 1, \dots, 7$, of $T\mathcal{M}$ as:

$$z^\alpha = x^\alpha, \quad z^{\alpha+4} = p^\alpha. \quad (2.110)$$

A coordinate change $\hat{x}^\alpha = \hat{x}^\alpha(x)$ on $\mathcal{M}$ induces an *extended coordinate transformation*

$$\hat{x}^\mu = \hat{x}^\mu(x), \quad \hat{p}^\mu = \frac{\partial \hat{x}^\mu}{\partial x^\nu} p^\nu, \quad (2.111)$$

this corresponds to the Jacobian matrix:

$$\frac{\partial \hat{z}^A}{\partial z^B} = \begin{pmatrix} \frac{\partial \hat{x}^\alpha}{\partial x^\beta} & 0 \\ \frac{\partial^2 \hat{x}^\alpha}{\partial x^\beta \partial x^\gamma} p^\gamma & \frac{\partial \hat{x}^\alpha}{\partial x^\beta} \end{pmatrix}. \quad (2.112)$$

Vectors on $T\mathcal{M}$

Let $b \in TU$ and $T_b T\mathcal{M}$ be the tangent space to $T\mathcal{M}$ at b . Then the push-forward [132] of π , π_* acts on the natural basis as:

$$\pi_* \left[\frac{\partial}{\partial x^\alpha} \right] = \frac{\partial}{\partial x^\alpha}, \quad \pi_* \left[\frac{\partial}{\partial p^\alpha} \right] = 0, \quad (2.113)$$

the pull-back acts as

$$\pi^* dx^\alpha = dx^\alpha. \quad (2.114)$$

Let $\vec{X} \in TTM$ be a vector field in some neighbour of the point b . b is associated with a point q of $\mathcal{M}$ and a vector $\vec{x} \in T_q\mathcal{M}$. The flow of b generated

by $\vec{X}$: $b(\lambda)$ is also associated with a one parameter family of points of $\mathcal{M}$, $q(\lambda)$, and a one-parameter family of vectors in $T\mathcal{M}$, $\vec{x}(\lambda)$. We say that $\vec{X}$ is *vertical* if the point $q(\lambda)$ is constant along the flow. Similarly, we say that $\vec{X}$ is *horizontal* if $\vec{x}(\lambda)$ “constant” along the flow, in the sense that $\vec{x}(\lambda)$ is just $\vec{x}$ parallel transported to $q(\lambda)$.

Obviously there is no unique way of giving a precise definition of “parallel transport” and this non-uniqueness is solved by choosing a linear connection, ∇ , of $\mathcal{M}$. From an abstract point of view, the choice of a linear connection is equivalent, in each point b , to the choice of two vector spaces $\mathcal{O}_b$ and $\mathcal{V}_b$ of horizontal and vertical vectors such that:

$$T_b T\mathcal{M} = \mathcal{O}_b \oplus \mathcal{V}_b. \quad (2.115)$$

Given a vector $\vec{x} \in T_q \mathcal{M}$, we can now define its *lift* as the unique horizontal vector, $\vec{X} \in T_b T\mathcal{M}$, whose projection is $\vec{x}$.

We can define a new vector basis, called the *connection basis*, adapted to the split (2.115): $\{D/\partial z^A\} := \{D/\partial x^\alpha, \partial/\partial p^\alpha\}$, where

$$\frac{D}{\partial x^\alpha} := \frac{\partial}{\partial x^\alpha} - \Gamma_{\alpha\gamma}^\beta p^\gamma \frac{\partial}{\partial p^\beta}. \quad (2.116)$$

A similar construction can be carried over for differential forms defined on the tangent bundle to compute the dual basis $\{Dz^A\} := \{dx^\alpha, Dp^\alpha\}$ using the pull-back π^* and it is easy to show that

$$Dp^\alpha = dp^\alpha + \Gamma_{\beta\gamma}^\alpha p^\gamma dx^\beta. \quad (2.117)$$

Metric on $T\mathcal{M}$

First of all, we notice that:

$$\frac{\partial^2 \hat{x}^\mu}{\partial x^\nu \partial x^\lambda} p^\lambda = \hat{\Gamma}_{\delta\lambda}^\mu p^\lambda \frac{\partial \hat{x}^\delta}{\partial x^\nu}. \quad (2.118)$$

We assume that, for any point, $b \in T\mathcal{M}$, there exist an open set, $TU \ni b$, and a coordinate system defined over TU such that

$$G_{AB} = (\boldsymbol{\eta} \otimes \boldsymbol{\eta})_{AB}, \quad (2.119)$$

where $\boldsymbol{\eta} = \text{diag}(-1, 1, 1, 1)$. Let $\hat{x}^A$ denotes a generic coordinate system in TU , then the metric in this coordinate system can be written as:

$$\hat{G}_{\mu\nu} = \frac{\partial \hat{x}^\alpha}{\partial x^\mu} \frac{\partial \hat{x}^\beta}{\partial x^\nu} \eta_{\alpha\beta} + \frac{\partial \hat{x}^\alpha}{\partial x^\mu} \hat{\Gamma}_{\alpha\lambda}^\gamma p^\lambda \frac{\partial \hat{x}^\beta}{\partial x^\nu} \hat{\Gamma}_{\beta\xi}^\delta p^\xi \eta_{\gamma\delta}; \quad (2.120)$$

$$\hat{G}_{\mu \nu+4} = \frac{\partial \hat{x}^\alpha}{\partial x^\mu} \frac{\partial \hat{x}^\gamma}{\partial x^\nu} \hat{\Gamma}_{\gamma\lambda}^\beta p^\lambda \eta_{\alpha\beta}; \quad (2.121)$$

$$\hat{G}_{\mu+4 \nu+4} = \frac{\partial \hat{x}^\alpha}{\partial x^\mu} \frac{\partial \hat{x}^\beta}{\partial x^\nu} \eta_{\alpha\beta}. \quad (2.122)$$

The line-element is then

$$\begin{aligned} dS^2 &= \hat{G}_{AB} d\hat{z}^A d\hat{z}^B = \hat{g}_{\mu\nu} d\hat{x}^\mu d\hat{x}^\nu + \hat{g}_{\mu\nu} [dp^\mu + \hat{\Gamma}_{\alpha\beta}^\mu p^\beta dx^\alpha] [dp^\nu + \hat{\Gamma}_{\alpha\beta}^\nu p^\beta dx^\alpha] \\ &= \hat{g}_{\mu\nu} d\hat{x}^\mu d\hat{x}^\nu + \hat{g}_{\mu\nu} D\hat{x}^\mu D\hat{x}^\nu. \end{aligned} \quad (2.123)$$

It can be shown that $|\det \mathbf{G}| = g^2$ as the transformation from the natural frame to the connection frame is unimodular [201]. Thus the volume pseudo-form on $T\mathcal{M}$ is, in the coordinate patch TU ,

$$\begin{aligned} \text{Vol}^8 &:= -g \, dx^0 \wedge dx^1 \wedge \dots \wedge dp^3 =: -g \, d^4x d^4p, \\ &:= -g \, dx^0 \wedge dx^2 \wedge \dots \wedge Dp^3 =: -g \, d^4x D^4p. \end{aligned} \quad (2.124)$$

2.3.2 The Liouville Theorem

We define an intrinsic 1-form on $T\mathcal{M}$, $\underline{\lambda} \in T^*T\mathcal{M}$, as follows. Consider a point $A \in T\mathcal{M}$. A is associated with a point $q \in \mathcal{M}$ and a 1-form, $\underline{\alpha} \in T_q^*\mathcal{M}$. We define $\underline{\lambda}$ in A as the unique 1-form such that $\underline{\lambda} = \pi^*\underline{\alpha}$. This is the so-called *Poincaré* 1-form and, in a local coordinate neighbor, TU , is given by:

$$\underline{\lambda} = p_\alpha dx^\alpha. \quad (2.125)$$

The associated vector,

$$\vec{\lambda} = p^\alpha \frac{D}{dx^\alpha} = p^\alpha \frac{\partial}{\partial x^\alpha} - p^\alpha \Gamma_{\alpha\gamma}^\beta p^\gamma \frac{\partial}{\partial p^\beta}, \quad (2.126)$$

is the so-called *geodesic flow* field. The name comes from the fact that its flow is the phase-space flow of a particle moving along geodesics.

Given a point $q \in U$, we define the mass-shell as the set

$$\mathcal{S}_m = \{p^\alpha \in T_q\mathcal{M}: p_\mu p^\mu + m^2 =: f(p) = 0\}. \quad (2.127)$$

In the case in which $m \neq 0$, the normal to the mass-shell is given by

$$\underline{\pi} := \frac{1}{2m} df, \quad df = \frac{\partial f}{\partial x^\mu} dx^\mu + \frac{\partial f}{\partial p^\mu} dp^\mu = 2p_\mu dp^\mu, \quad (2.128)$$

while for massless particles the normal is simply defined as $\underline{\pi} := \frac{1}{2} df$. Finally we introduce the form $\underline{\nu}$ as being the only vertical form whose restriction on $T_q\mathcal{M}$ is equal to $\underline{\pi}$. This is given by

$$\underline{\nu} = \frac{1}{m} p_\alpha Dp^\alpha. \quad (2.129)$$

in the case where $m \neq 0$ and $\underline{\nu} = p_\alpha Dp^\alpha$ in the massless case.

It easy to see from its definition that $\underline{\nu}$ is irrotational, *i.e.*, $d\underline{\nu} = 0$, since, for instance, in the $m \neq 0$ case (the massless case is analogous), it can be written as

$$\underline{\nu} = \frac{1}{2m} \frac{Df}{Dp^\alpha} Dp^\alpha. \quad (2.130)$$

It can be proven [284] that $\underline{\lambda}$ is incompressible, *i.e.*, $0 = d^*\underline{\lambda} := \star d \star \underline{\lambda}$. Moreover, both are harmonic in the sense of Kodaira:

$$0 = \Delta \underline{\lambda} := [dd^* + d^*d]\underline{\lambda}, \quad 0 = \Delta \underline{\nu}. \quad (2.131)$$

In our case we are interested in measuring the density of states, in phase space, of particles moving along geodesics and with velocities on the mass shell. Remembering (2.83), we are led to define the six-form

$$\omega = \star(\underline{v} \wedge \underline{\lambda}) = i_{\vec{\lambda}} i_{\vec{v}} \text{Vol}^8. \quad (2.132)$$

This can be written more explicitly, using the definition of Vol^8 , as

$$\omega = i_{\vec{\lambda}} \left[i_{\vec{v}} (\text{Vol}_x^4 \wedge \text{Vol}_p^4) \right] = i_{\vec{\lambda}} [\text{Vol}_x^4 \wedge \text{Vol}_p^3]; \quad (2.133)$$

where we defined, in TU ,

$$\text{Vol}_x^4 := \sqrt{-g} \, dx^0 \wedge dx^1 \wedge dx^2 \wedge dx^3 \quad (2.134)$$

and

$$\text{Vol}_p^4 := \sqrt{-g} \, Dp^0 \wedge Dp^1 \wedge Dp^2 \wedge Dp^3, \quad (2.135)$$

where $\text{Vol}_p^3 := i_{\vec{v}} \text{Vol}_p^4$ is the volume form on the mass shell, $\mathcal{S}_m$. In coordinates adapted to the mass shell, where $\underline{v} = \frac{p_0}{m} Dp^0$ in the $m \neq 0$ case and $\underline{v} = p_0 Dp^0$ in the massless case, this reads

$$\text{Vol}_p^3 = \frac{\sqrt{-g}}{-p_0} Dp^1 \wedge Dp^2 \wedge Dp^3. \quad (2.136)$$

Notice that the last expression is valid also in the case $m = 0$.

In the context of the usual ADM foliation of the spacetime: $\mathcal{M} = \mathbb{R} \times \Sigma$ in $x^0 = \text{const}$ hypersurfaces with normal $\underline{n} = -\alpha dx^0$, α being the lapse function, the expression for ω can be simplified. We can split Vol_x^4 as

$$\text{Vol}_x^4 = -\underline{n} \wedge \text{Vol}_x^3, \quad (2.137)$$

where

$$\text{Vol}_x^3 = i_{\vec{n}} \text{Vol}_x^4 = \sqrt{\gamma} \, dx^1 \wedge dx^2 \wedge dx^3 \quad (2.138)$$

and γ is the three-metric induced on the slices. Thus we obtain, in coordinates adapted to the spacetime slicing and the mass shell,

$$\begin{aligned} \omega &= -(\vec{p} \cdot \vec{n}) \frac{1}{-p_0} \sqrt{\gamma} \sqrt{-g} \, dx^1 \wedge dx^2 \wedge dx^3 \wedge Dp^1 \wedge Dp^2 \wedge Dp^3 \\ &= \frac{p^0}{-p_0} |g| \, dx^1 \wedge dx^2 \wedge dx^3 \wedge Dp^1 \wedge Dp^2 \wedge Dp^3. \end{aligned} \quad (2.139)$$

If we consider the action of ω on the six-vectors, $\delta_1 x, \delta_2 x, \delta_3 x, \delta_1 p, \delta_2 p, \delta_3 p$, chosen so that $\delta_i x$ are tangent vectors to Σ and $\delta_i p$ are tangent vectors to $\mathcal{S}_m$, we

have

$$\begin{aligned}
\omega(\delta_1 x, \dots, \delta_3 p) &= \frac{p^0}{-p_0} |g| \left[dx^1 \wedge dx^2 \wedge dx^3 \right] (\delta_1 x, \delta_2 x, \delta_3 x) \times \\
&\quad \left[Dp^1 \wedge Dp^2 \wedge Dp^3 \right] (\delta_1 p, \delta_2 p, \delta_3 p) \\
&\quad - \frac{p^0}{-p_0} |g| \left[dx^1 \wedge dx^2 \wedge dx^3 \right] (\delta_1 p, \delta_2 p, \delta_3 p) \times \\
&\quad \left[Dp^1 \wedge Dp^2 \wedge Dp^3 \right] (\delta_1 x, \delta_2 x, \delta_3 x) = \\
&= \frac{p^0}{-p_0} |g| \left[dx^1 \wedge dx^2 \wedge dx^3 \right] (\delta_1 x, \delta_2 x, \delta_3 x) \times \\
&\quad \left[Dp^1 \wedge Dp^2 \wedge Dp^3 \right] (\delta_1 p, \delta_2 p, \delta_3 p) = \\
&= \frac{p^0}{-p_0} |g| \left[dx^1 \wedge dx^2 \wedge dx^3 \right] (\delta_1 x, \delta_2 x, \delta_3 x) \times \\
&\quad \left[dp^1 \wedge dp^2 \wedge dp^3 \right] (\delta_1 p, \delta_2 p, \delta_3 p).
\end{aligned} \tag{2.140}$$

This follows from $dx^i(\delta_j p) = 0$ and the relation

$$Dp^i(\delta_j p) = dp^i(\delta_j p) - \Gamma_{\alpha\beta}^i p^\alpha dx^\beta(\delta_j p) = dp^i(\delta_j p). \tag{2.141}$$

For this reason we have, on the space-like hypersurface Σ and on the mass shell:

$$\omega = \frac{p^0}{-p_0} |g| dx^1 \wedge dx^2 \wedge dx^3 \wedge dp^1 \wedge dp^2 \wedge dp^3 =: \Omega. \tag{2.142}$$

We can split Ω into two parts by defining the three-forms:

$$\Lambda = p^0 \sqrt{-g} dx^1 \wedge dx^2 \wedge dx^3 \tag{2.143}$$

and

$$\Pi = \frac{1}{-p_0} \sqrt{-g} dp^1 \wedge dp^2 \wedge dp^3 \tag{2.144}$$

as

$$\Omega = \Lambda \wedge \Pi. \tag{2.145}$$

Note that Λ and Π are intrinsic forms in $T\mathcal{M}$ as they can be written, at any point $q \in \mathcal{M}$, in a coordinate independent way, as⁹:

$$\Lambda = \star_{\mathcal{M}} \underline{\lambda}, \quad \Pi = \star_{T_q \mathcal{M}} \underline{\pi}; \tag{2.146}$$

From the previous definitions it is clear that Λ and Π are the proper geodesics-flux volume form on respectively, the spacelike hypersurface, Σ , in $\mathcal{M}$ and the mass shell, $\mathcal{S}_m$, in the tangent space, $T_q \mathcal{M}$, at a point, $q \in U$.

Let $\mathcal{S}$ be a “phase tube” in the tangent bundle generated by the geodesic flow, $\vec{\lambda}$. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two sections of this tube. It is possible to show that the flux of points in phase space associated with the geodesic flow satisfies

$$\int_{\mathcal{S}_1} \omega = \int_{\mathcal{S}_2} \omega, \tag{2.147}$$

⁹Note that the tangent space, $T_q \mathcal{M}$, at a point q , is also a manifold with metric $\mathbf{g}$.

which is the relativistic version of the Liouville's Theorem. To prove that we compute the exterior differential of ω :

$$\star d\omega = d^\star(\underline{v} \wedge \underline{\lambda}) = d^\star \underline{v} \wedge \underline{\lambda} + \underline{v} \wedge d^\star \underline{\lambda}, \quad (2.148)$$

using the fact that $dd^\star \underline{v} = 0$ we obtain that $d^\star \underline{v}$ is a constant which we call k and, knowing that $d^\star \underline{\lambda} = 0$, we get

$$d\omega = -k(\star \underline{\lambda}). \quad (2.149)$$

Since $\star \underline{\lambda}$ is zero along the phase tube, as $\star \underline{\lambda}[\vec{\lambda}, \dots] = 0^{10}$, we get

$$\int_S d\omega = 0. \quad (2.150)$$

Finally using Stoke's Theorem and knowing that ω vanishes along the part of the boundary tangent to $\vec{\lambda}$, we get (2.147).

2.3.3 The Boltzmann equation

We introduce a new six-form, μ , whose value on any section of the phase tube, S_1 , is the number of particle phase trajectories crossing it. In the absence of collisions

$$\int_{S_1} \mu = \int_{S_2} \mu. \quad (2.151)$$

Also, since the particles move along geodesics, μ must be proportional to ω ,

$$\mu = F \omega, \quad (2.152)$$

where F is the so-called *invariant distribution function*, i.e., F is the Radon-Nikodym derivative of μ with respect to ω .

Suppose now that the number of phase trajectories changes due to collisions. The change in this number is

$$\delta N = \int_{S_2} \mu - \int_{S_1} \mu = \int_S d\mu = \int_S dF \wedge \omega. \quad (2.153)$$

We notice that

$$dF \wedge \omega = dF \wedge \star(\underline{v} \wedge \underline{\lambda}) = \langle dF, \underline{\lambda} \rangle \star \underline{v} - \langle dF, \underline{v} \rangle \star \underline{\lambda}, \quad (2.154)$$

where $\langle \cdot, \cdot \rangle$ is a scalar product between forms:

$$\langle \alpha, \beta \rangle \text{Vol}^8 := \alpha \wedge \star \beta. \quad (2.155)$$

Finally, as $\star \underline{\lambda}$ vanishes on S we get

$$\delta N = \int_S \langle dF, \underline{\lambda} \rangle \star \underline{v}. \quad (2.156)$$

¹⁰Remember that $\star \underline{\lambda}$ is the volume form of the hypersurfaces orthogonal to $\vec{\lambda}$.

If we denote with

$$\delta N =: \int_S \mathbb{C}[F] \star \underline{\nu} \quad (2.157)$$

the effects of the collisions, we get the Boltzmann equation:

$$\langle dF, \underline{\lambda} \rangle = \mathbb{C}[F], \quad (2.158)$$

or, more explicitly,

$$p^\alpha \frac{\partial F}{\partial x^\alpha} - \Gamma^\gamma_{\alpha\beta} p^\alpha p^\beta \frac{\partial F}{\partial p^\gamma} = \mathbb{C}[F]. \quad (2.159)$$

If we adopt a system of coordinates adapted to it, *i.e.*, $p^0 = p^0(p^i)$, the Boltzmann equation takes the following form [79]

$$p^\alpha \frac{\partial F}{\partial x^\alpha} - \Gamma^i_{\alpha\beta} p^\alpha p^\beta \frac{\partial F}{\partial p^i} = \mathbb{C}[F]. \quad (2.160)$$

We notice that, thanks to the incompressibility of $\underline{\lambda}$, we have

$$\langle dF, \underline{\lambda} \rangle = d^\star[F \underline{\lambda}], \quad (2.161)$$

so that we can rewrite the Boltzmann equation in a form that makes the conservation of the number of particles manifest:

$$d^\star[F \underline{\lambda}] = \mathbb{C}[F]. \quad (2.162)$$

This is the so-called conservative formulation of the Boltzmann equation [77].

We can make this construction more algebraic by introducing the Levi Civita connection, ∇ , in phase space. The incompressibility of $\underline{\lambda}$ can then be written in component form as

$$\nabla_A \lambda^A = 0, \quad (2.163)$$

while the Boltzmann equation reads

$$\lambda^A \partial_A F = \mathbb{C}[F]. \quad (2.164)$$

Using these two relations one easily obtains the conservative formulation of the Boltzmann equation in component form as

$$\nabla_A [F p^A] = \mathbb{C}[F], \quad (2.165)$$

or, more explicitly as

$$\frac{1}{|g|} \frac{\partial}{\partial x^\mu} \left[|g| F p^\mu \right] + \frac{p_0}{|g|} \frac{\partial}{\partial p^k} \left[\frac{|g|}{-p_0} \Gamma^k_{\alpha\beta} p^\alpha p^\beta F \right] = \mathbb{C}[F], \quad (2.166)$$

2.3.4 From the Boltzmann Equation to the Euler Equation

Now that we have derived the Boltzmann equation in general-relativity we are ready to derive the equations of hydrodynamics from the Boltzmann equation. To do that, first, we need to have a way to extract quantities, such as the mass-flux or the energy flux, that are of central importance in relativistic

hydrodynamics, from the kinetic description. The density-flux is easily computed remembering that μ is just the phase-space version of the mass-flux. In particular from its definition it is easy to see the mass flow is simply given by

$$\rho = \int_{S_m} \mu, \quad (2.167)$$

while the rest-mass density four-vector is given by

$$J^\mu = \int_{S_m} F p^\mu \Pi, \quad (2.168)$$

since

$$\int_{\Sigma} (-J^\mu n_\mu) \text{Vol}_x^3 = \int_{\Sigma \times S_m} F p^0 \Pi \text{Vol}_x^3 = \int_{\Sigma \times S_m} F \Omega = \int_{\Sigma} \rho. \quad (2.169)$$

In a similar way the second moments of the distribution function give rise to the stress-energy tensor:

$$T^{\mu\nu} = \int_{S_m} F p^\mu p^\nu \Pi \quad (2.170)$$

and, in particular, the components of the stress energy tensor, in the comoving frame of the fluid, can be interpreted as

$$T^{\mu\nu} = \begin{pmatrix} E & \vec{F} \\ \vec{F} & \mathbf{P} \end{pmatrix}, \quad (2.171)$$

where E is the energy density, $\vec{F}$ is the energy flux and $\mathbf{P}$ is the stress tensor. The particular form of the stress-energy tensor depends on the equilibrium distribution function [79], which in turns depends on the nature of the collisional operator.

Finally the hydrodynamic equations are a consequence of the Liouville Theorem. Let Ψ be a tensorial function of p^i . Then, multiplying both sides of equation (2.165) by Ψ and integrating with respect to Π we obtain

$$\int_{\mathbb{R}^3} \nabla_A [F \lambda^A \Psi] \Pi = \int_{\mathbb{R}^3} \mathbb{C}[F] \Psi \Pi, \quad (2.172)$$

where we have used the fact that $\lambda^A \nabla_A \Psi = 0$, since $\vec{\lambda}$ is the geodesic flow. Finally, if we assume that F decays for large moments, we obtain the transfer equation [174, 79]:

$$\nabla_\mu \int_{\mathbb{R}^3} F \Psi p^\mu \Pi = \int_{\mathbb{R}^3} \mathbb{C}[F] \Psi \Pi, \quad (2.173)$$

In the case of a simple gas it is possible to show that, if Ψ is one of $\{1, p^0, p^1, p^2, p^3\}$, we have [79]

$$\int_{\mathbb{R}^3} \mathbb{C}[F] \Psi \Pi = 0, \quad (2.174)$$

these particular choices of Ψ are called *collisional invariants* and are associated with quantities that are conserved by the collisional operator. In particular

1 is obviously associated with the mass conservation, while p^μ yields energy and momentum conservation. In particular and, as a consequence of equation (2.173), we find that

$$\nabla_\mu J^\mu = 0, \quad \nabla_\nu T^{\mu\nu} = 0. \quad (2.175)$$

Chapter 3

Numerical Approximation of Conservation Laws

In this chapter we present some key aspects concerning the theory of the numerical approximation of conservation laws. Conservation laws are ubiquitous in physics and, not surprisingly, the literature covering their mathematical and numerical treatment is extremely large and several monographs are available on the subject. Here we focus on those aspects that are relevant for our subsequent discussion and the emphasis is put on the presentation of high-order, state-of-the-art numerical methods for the solution of conservation laws. We do not provide a complete description of the properties of all these schemes, as their complete presentation can be easily found elsewhere. Instead, we emphasise the key ideas behind their development and discuss their strengths and weakness in the context of general-relativistic hydrodynamics.

This chapter is divided in four parts. In Section 3.1 we recall some of the most basic results in the theory of conservation laws and their numerical approximation, while in Section 3.2 we present Godunov-like finite-volume schemes. In Section 3.3 we present a description of central methods and, in Section 3.4, we describe high-resolution shock-capturing (HRSC) finite-difference schemes and in particular their high-order variants. Finally, in Section 3.5 we present discontinuous Galerkin methods.

3.1 Theoretical Background

In this section we recall some basic results concerning the mathematical theory of conservation laws and their numerical approximation. First of all we start from the definition of weak and entropic solutions of conservation laws and we state some results concerning existence and uniqueness of entropic solutions for conservation laws. We introduce then the basic concepts of consistency, stability and convergence of the numerical approximation to conservation laws and the Lax-Richtmeyer theorem. Finally we present the extension to the case of non-linear equations. Our treatment is loosely based on the one found in [196, 312] and we refer the reader to those references for a more complete discussion.

3.1.1 Conservation Laws

In this section we consider equations in the form

$$\partial_t \mathbf{u} + \nabla \cdot \mathbf{f}(\mathbf{u}) = 0, \quad (t, x) \in \mathbb{R}_+ \times \mathbb{R}^d, \quad (3.1a)$$

$$\mathbf{u}(0, x) = \mathbf{u}_0(x), \quad x \in \mathbb{R}^d. \quad (3.1b)$$

where $\mathbf{u}$ is a vector of m unknowns, $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^m)$ is a d -dimensional flux and $\mathbf{u}_0 \in [L^\infty(\mathbb{R}^d)]^m$ is the initial data.

It is well known that the solution of (3.1) can develop shocks in finite time, even if the initial data is analytic. For this reason (3.1) has to be interpreted in the sense of the distributions. In particular a function $\mathbf{u} \in [L^\infty(\mathbb{R}_+ \times \mathbb{R}^d)]^m$ is a *weak solution* of (3.1) if, for all test functions $v \in C_0^1(\mathbb{R}^{d+1})$ and $i = 1, 2, \dots, m$, we have

$$\int_0^\infty dt \int_{\mathbb{R}^d} [u^i \partial_t v + \mathbf{f}^i(\mathbf{u}) \cdot \nabla v] dx = \int_{\mathbb{R}^d} u_0^i v dx. \quad (3.2)$$

Unfortunately it is easy to show that even simple scalar conservation laws admit multiple weak solutions. In order to single out the “physically relevant” solution it is necessary to introduce the concept of entropic solutions. A convex function, $\eta(\mathbf{u})$, is said to be an *entropy function* if its Hessian, $\nabla_{\mathbf{u}}^2 \eta$, symmetrizes the spatial Jacobian, $\nabla_{\mathbf{u}} \mathbf{f}^i$:

$$\nabla_{\mathbf{u}}^2 \eta \cdot \nabla_{\mathbf{u}} \mathbf{f}^i = [\nabla_{\mathbf{u}} \mathbf{f}^i]^T \cdot \nabla_{\mathbf{u}} \eta, \quad i = 1, \dots, m. \quad (3.3)$$

In this case, there exists an entropy flux, $\boldsymbol{\psi} = (\boldsymbol{\psi}^1, \dots, \boldsymbol{\psi}^m)$, which is determined by the compatibility relations

$$[\nabla_{\mathbf{u}} \eta]^T \cdot \nabla_{\mathbf{u}} \mathbf{f}^i = [\nabla_{\mathbf{u}} \boldsymbol{\psi}^i]^T, \quad i = 1, \dots, m \quad (3.4)$$

and the tuple $(\eta, \boldsymbol{\psi})$ is called *entropy pair*.

We define as an *entropic solution* a weak solution such that, for any entropy pair, the following holds

$$\partial_t \eta(\mathbf{u}) + \nabla \cdot \boldsymbol{\psi}(\mathbf{u}) \leq 0, \quad (3.5)$$

in the sense of the distributions, *i.e.*, for any non-negative test function $v \in C_0^1(\mathbb{R}_+ \times \mathbb{R}^d)$, we have

$$\int_{\mathbb{R}_+ \times \mathbb{R}^d} [\eta(\mathbf{u}) \partial_t v + \boldsymbol{\psi}(\mathbf{u}) \cdot \nabla v] dt dx \geq 0. \quad (3.6)$$

This condition has a simple interpretation in the scalar case, where it is possible to show that (3.6) is equivalent to the requirement that characteristic lines should impinge into shock waves [196]. This is basically a statement of the irreversibility of the processes leading to the formation of a shock wave and enforces the breaking of the time-symmetry of the problem.

In the case of scalar conservation laws, *i.e.*, $m = 1$, Kruzkov [188] proved the existence and uniqueness of the entropic solution under very general conditions. Kruzkov theory has been extended to measure-valued solutions by DiPerna [114] and to the case of conservation laws on manifolds by Ben-Artzi and LeFloch [40].

In the case of systems of conservation laws very little is known concerning existence, uniqueness and stability of entropic solutions, especially in the multi-dimensional case. For instance, not even the existence of entropy pairs is guaranteed for general systems of equations. A promising approach is the one based on divergence-measure vector fields by Chen et al. [83]. Within this framework Chen and Frid [82] proved existence, uniqueness and stability of the entropic solution of the Euler equations for a classical ideal-gas of one-dimensional Riemann problems, *i.e.*, problems with initial data in the form

$$\mathbf{u}_0(x) = \begin{cases} \mathbf{u}^L, & \text{if } x < 0; \\ \mathbf{u}^R, & \text{if } x > 0. \end{cases} \quad (3.7)$$

On the other hand, for general equation of state, the existence of a weak solution to the Riemann problem is not even guaranteed [222] (see also [81] for a modern review concerning the mathematical study of the classical Euler equations).

In the relativistic case, the existence of solutions to the Riemann problem was shown, in the case of ultrarelativistic equation of state, by Smoller and Temple [304] using Glimm's method [147].

For general systems of conservation laws the existence of weak solutions to the Riemann problem was proven by Lax [192] for initial data with small enough initial jump, in the case of *strictly hyperbolic systems*, *i.e.*, when $\nabla_{\mathbf{u}} \mathbf{f}$ has a complete set of real eigenvalues and eigenvectors.

3.1.2 Consistency, Stability and Convergence

We now turn to the theory related to the numerical approximation of conservation laws. Since the non-linear theory is basically known only in the scalar case, here we focus only on the case in which $m = 1$, that is, we consider the problem

$$\partial_t u + \nabla \cdot \mathbf{f}(u) = 0, \quad (t, x) \in \mathbb{R}_+ \times \mathbb{R}^d \quad (3.8a)$$

$$u(0, x) = u_0(x), \quad x \in \mathbb{R}^d, \quad (3.8b)$$

where u is now just a scalar function.

For the purpose of our discussion it is useful to introduce the following notation. We can view the solution of (3.8) as a sequence of bounded functions $u(t, \cdot) \in L^\infty(\mathbb{R}^d)$, *i.e.*, as a “curve” in the (infinite dimensional) vector space $L^\infty(\mathbb{R}^d)$, and consider it only as a function of time, $u(t)$, with values in $L^\infty(\mathbb{R}^d)$. Since the “curve” is itself bounded in $L^\infty(\mathbb{R}^d)$, we can write $u(t) \in L^\infty[\mathbb{R}_+; L^\infty(\mathbb{R}^d)]$. In particular we can interpret (3.8) as being an (infinite) system of ordinary differential equations (ODEs), which we write symbolically as

$$\frac{du(t)}{dt} = \mathcal{L}[u(t)], \quad u(0) = u_0, \quad (3.9)$$

where $\mathcal{L}(\cdot)$ is an operator associated with $-\nabla \cdot \mathbf{f}(\cdot)$. We point out that, again, (3.9) has to be intended in the sense of the distributions since $u(t)$ is not a smooth function of time. Furthermore we warn the reader that a careful mathematical treatment would require us to address the rather non-trivial question of what

we actually mean when we require $u(0) = u_0$, since the restriction of general functions in $L^\infty[\mathbb{R}_+; L^\infty(\mathbb{R}^d)]$ to a set of zero Lebesgue measure is not defined. We just mention that this question is addressed in [188]. For our purposes we can assume $u(t)$ to be a smooth function of time, because this is the case which can actually be treated numerically and ignore these technical aspects.

With this notation a numerical approximation of (3.8) is a finite-dimensional approximation of (3.9) depending on some discretization parameter Δ , which we write as

$$\frac{du^\Delta(t)}{dt} = L^\Delta[u^\Delta(t)], \quad u^\Delta(0) = P^\Delta[u_0], \quad (3.10)$$

where u^Δ and L^Δ are approximations of u and $\mathcal{L}$, i.e., $u^\Delta \simeq u$, $L^\Delta \simeq \mathcal{L}$ and P^Δ is some (possibly non-linear) projection operator. The error associated with P^Δ is usually negligible with respect to the other errors due to the discretization of the conservation law and for this reason we neglect it, i.e., we set $u^\Delta(0) = u_0$.

We define the *local truncation error* of the scheme as the quantity

$$r^\Delta(t) = L^\Delta[u(t)] - \mathcal{L}[u(t)], \quad (3.11)$$

where $u(t)$ is the exact solution to (3.9) and we say that a scheme is *consistent* if the truncation error converges to zero as $\Delta \rightarrow 0$ in some (problem- and method-dependent) norm for all possible initial data u_0 ¹. In particular we say that a scheme is of order r if

$$\|r^\Delta(t)\| = O(\Delta^r). \quad (3.12)$$

A scheme is said to be *stable* if the norm of L^Δ is limited:

$$\|L^\Delta\| := \sup_{v \neq 0} \frac{\|L^\Delta(v)\|}{\|v\|} \leq C, \quad (3.13)$$

where $C \geq 0$ is a constant independent from v .

Finally we say that a numerical scheme is *convergent* if

$$\lim_{\Delta \rightarrow 0} \|u^\Delta(t) - u(t)\| = 0, \quad \text{a.e. } t \in \mathbb{R}_+. \quad (3.14)$$

In the case of linear equations, consistency, stability and convergence are related by the famous Lax-Richtmeyer equivalence theorem [193] stating that the numerical approximation of well-posed problems is convergent if and only if the scheme is stable and consistent. Furthermore if a scheme is of order r then

$$\|u^\Delta(t) - u(t)\| = O(\Delta^r). \quad (3.15)$$

In the non-linear case stability and consistency are not enough to guarantee convergence and a stronger notion of stability, the so-called *non-linear stability*, is required.

¹Note that the choice of the norm in which the error is measured could restrict the range of the allowable initial data.

3.1.3 Non-Linear Equations and Non-Linear Stability

In order to develop the non-linear theory it is necessary to introduce additional notation. So far we introduced an operator $\mathcal{L}$ representing the spatial part of the conservation law and its approximation, L^Δ , representing the spatial-discretization. In so doing we ignored the fact that the system of ODEs (3.10) needs to be also solved numerically and this introduces an extra source of error, *i.e.*, we worked with schemes in their *semi-discrete* form. This was done because the truncation error associated with the time-discretization is often very small with respect to the one associated with L^Δ in the case of non-linear conservation laws and because, when discussing the different approaches to actually construct L^Δ , we often refer to the discretization error associated with this operator. On the other hand, while discussing the non-linear stability and convergence of numerical schemes in the non-linear case, the properties of the time-discretization play a very important role. For this reason it is useful to work with *fully discrete* schemes.

To this end we introduce a one-parameter family of evolution operators $\{\mathcal{T}_s\}_{s \in \mathbb{R}_+}$, such that

$$u(t) = \mathcal{T}_t(u_0), \quad \mathcal{T}_s \circ \mathcal{T}_t = \mathcal{T}_{t+s}, \quad (3.16)$$

i.e., $\mathcal{T}_t(u_0)$ gives the solution to (3.8) with initial data u_0 at the time t and $\{\mathcal{T}_s\}_{s \in \mathbb{R}_+}$ equipped with the operation of composition forms a semi-group. The existence of these operators is shown in [188]. We also introduce their discrete version $\{T_{k\Delta t}^\Delta\}_{k \in \mathbb{N}}$, such that

$$u^\Delta(k\Delta t) = T_{k\Delta t}^\Delta(u_0), \quad T_s^\Delta \circ T_t^\Delta = T_{t+s}^\Delta, \quad (3.17)$$

i.e.,

$$u^\Delta(t + \Delta t) = T_{\Delta t}^\Delta[u^\Delta(t)] \quad (3.18)$$

represents the numerical scheme in its fully-discrete form.

Notice that we are again assuming that T depends on a single spatial discretization parameter Δ . This is done because the time-discretization is usually linked to the spatial one by the Courant–Friedrichs–Lewy (CFL) condition, which is basically the linear-stability condition for the chosen time-integrator.

The definitions of the previous section are easily translated to the case of fully-discrete schemes. The truncation error is defined as

$$r^\Delta(t) = T_{\Delta t}^\Delta[u(t)] - u(t + \Delta t) = T_{\Delta t}^\Delta[u(t)] - \mathcal{T}_{\Delta t}[u(t)], \quad (3.19)$$

where $u(t)$ is the exact solution to (3.8). A fully-discrete scheme is said to be consistent if, in some norm $\|\cdot\|$, $\|r^\Delta(t)\| \rightarrow 0$ as $\Delta \rightarrow 0$. In particular we say that the scheme is of order r if $\|r^\Delta(t)\| = \mathcal{O}(\Delta^r)$. Finally a fully-discrete scheme is said to be linearly stable, or simply stable, if

$$\|T_{\Delta t}^\Delta\| \leq C, \quad (3.20)$$

for some constant $C \geq 0$.

The starting point for the construction of numerical schemes in the non-linear case is the celebrated Lax-Wendroff theorem [194] stating that *if* the

numerical approximation, u^Δ , to (3.9), obtained with a *conservative*, *i.e.*, such that

$$\int_{\mathbb{R}^d} T_s^\Delta(v) \, dx = \int_{\mathbb{R}^d} v \, dx, \quad \forall v \in L^1(\mathbb{R}^d), \quad (3.21)$$

and consistent scheme convergences strongly, *i.e.*, in L^1 -norm, to some function u , then u is a weak solution of (3.8). The key idea is then to find some sufficient condition to ensure the convergence of the scheme, because then the Lax-Wendroff theorem will guarantee that the obtained solution will be a weak solution of (3.8). Note that the Lax-Wendroff theorem does not guarantee that the found solution is the entropic one. For that it is necessary to show that the numerical scheme is also satisfying the entropy inequality.

In order to derive a sufficient convergence condition, we introduce the concept of total-variation of a function in a domain Ω ,

$$\text{TV}(v; \Omega) := \sup \left\{ \int_{\Omega} v \nabla \cdot \boldsymbol{\varphi} \, dx : \boldsymbol{\varphi} \in [C_0^1(\Omega)]^d, \|\boldsymbol{\varphi}\|_{L^\infty(\Omega)} \leq 1 \right\}. \quad (3.22)$$

It is possible to show that (*e.g.*, [16]) if $\text{TV}(v; \Omega) < \infty$, then²

$$\text{TV}(v; \Omega) = \int_{\Omega} |Dv| \, dx, \quad (3.23)$$

where Du is the gradient of u in the sense of the distributions. The space of all functions in $L^1(\Omega)$ with finite total-variation is an (infinite dimensional) vector space denoted with $\text{BV}(\Omega)$, which is a Banach space with respect to the norm

$$\|v\|_{\text{BV}(\Omega)} = \int_{\Omega} (|v| + |Dv|) \, dx. \quad (3.24)$$

The reason why BV spaces are interesting for us is that, first of all, Conway and Smoller [97] proved that, if the initial data $u_0 \in \text{BV}(\mathbb{R}^d)$, then $u \in \text{BV}(\mathbb{R}_+ \times \mathbb{R}^d)$ and, for any $t \geq 0$, $u(t) \in \text{BV}(\mathbb{R}^d)$, *i.e.*, BV-spaces have an invariance property³. Secondly; bounded sets in $\text{BV}(\Omega)$ are sequentially compact in $L_{\text{loc}}^1(\Omega)$, *e.g.*, [16], *i.e.*, any bounded sequence of functions, $\{v_n\} \in \text{BV}(\Omega)$, has a subsequence converging in the L^1 -norm to a function in $L_{\text{loc}}^1(\Omega)$.

The idea is then to build numerical schemes able to mimic the BV-invariance of exact conservation laws. If these schemes are L^1 -stable and consistent, then the sequence of solutions produced by such a schemes, as we decrease Δ , will converge in the L^1 -norm to a weak solution of (3.8) (thanks to the Lax-Wendroff Theorem)⁴. Such schemes are said to be *TV-stable*. In particular a scheme is

²More precisely $\text{TV}(v; \Omega) = |Dv|(\Omega)$ where $|\cdot|$ denotes the total variation of the vector-valued measure Dv .

³Note that this is valid only for scalar conservation laws. System of conservation laws are known to have solutions which are not BV even for BV initial data. An example is given by the Euler equation in Lagrangian coordinates: the solution can spontaneously develop vacuum regions for initial data with sufficiently large total variation. In this case the solution is not BV, but rather an L^1 function or a Radon measure [81].

⁴More precisely $\{u^\Delta\}$ will be an union of subsequences each converging to a (possibly different) weak solution. Suppose that this is not the case. Then there must exist a sequence $\{u^{\Delta_i}\}$ and $\epsilon > 0$ such that, $\Delta_i \rightarrow 0$ when $i \rightarrow \infty$ and

$$\text{dist}(W, u^{\Delta_i}) > \epsilon, \quad \forall i, \quad (3.25)$$

TV-stable if, for all initial data $u_0 \in L^\infty(\mathbb{R}^d) \cap BV(\mathbb{R}^d)$, there exist $\Delta_0 > 0$ and $C \geq 0$ such that

$$\text{TV}(T_s^\Delta(u_0); \mathbb{R}^d) \leq C, \quad \forall \Delta < \Delta_0. \quad (3.27)$$

The main class of schemes having this property is the one of *total-variation diminishing* (TVD) schemes, *i.e.*, schemes such that

$$\text{TV}(T_{\Delta t}^\Delta(v); \mathbb{R}^d) \leq \text{TV}(v; \mathbb{R}^d), \quad \forall v \in L^\infty(\mathbb{R}^d) \cap BV(\mathbb{R}^d). \quad (3.28)$$

The first schemes that were proven to be convergent in the case of non-linear scalar conservation laws are the so called *monotone* schemes, *i.e.*, schemes such that

$$u \geq v \quad \text{a.e.} \quad \implies \quad T_s^\Delta(u) \geq T_s^\Delta(v) \quad \text{a.e.} \quad (3.29)$$

Crandall and Tartar [101] proved that if T_s^Δ is conservative and monotone, then it is an L^1 -contraction, *i.e.*, satisfying the strong stability condition

$$\|T_{\Delta t}^\Delta(u) - T_{\Delta t}^\Delta(v)\|_{L^1(\mathbb{R}^d)} \leq \|u - v\|_{L^1(\mathbb{R}^d)}, \quad \forall u, v \in L^1(\mathbb{R}^d). \quad (3.30)$$

In particular this implies that the method is TVD, *e.g.*, [196]. Using this fact Crandall and Majda [100] proved that the numerical solution obtained with a conservative, consistent and stable monotone scheme converges to a weak solution of (3.8). Furthermore they also showed that these schemes satisfy an entropy inequality and, for this reason, the solution obtained with a conservative, consistent and stable monotone scheme converge to the (unique) entropic solution of the problem (3.8). Unfortunately Harten et al. [161] proved that monotone schemes are only at most first-order accurate. This means that, while monotonicity is a condition that ensures the non-linear stability of the scheme, it is also too strict a requirement if higher order methods are to be designed.

To achieve higher-order accuracy schemes, high-resolution shock-capturing (HRSC) schemes have been developed that are TVD, in the one-dimensional case, but not monotone. These schemes achieve high-order of accuracy and non-linear stability with the use of a non-linear dissipation mechanism that reduces the local convergence order to first order in the vicinity of discontinuities to avoid the appearance of spurious oscillations. Unfortunately, while in the one-dimensional case TVD schemes can be designed for any formal order of accuracy, Goodman and LeVeque [151] proved that, in the multi-dimensional case, TVD schemes are necessarily at most first-order accurate⁵. For this reason, in the quest for higher-order schemes, even weaker non-linear stability conditions have been proposed, even though the convergence of schemes satisfying these conditions has not yet been proven. One important example is the case

where W is the set of all weak solutions of (3.8). On the other hand since $\{u^{\Delta_i}\}$ is in a sequentially compact set, it must contain a convergent subsequence $\{u^{\Delta_{i_k}}\}$ converging to an element v . Because of the Lax-Wendroff theorem we necessarily have $v \in W$, so that one has

$$\text{dist}(W, u^{\Delta_{i_k}}) \rightarrow 0, \quad (3.26)$$

contradicting (3.25).

⁵Note, however, that many classical multi-dimensional schemes built in a dimensionally split fashion from one-dimensional TVD schemes are actually not TVD and for this reason they are not restricted to first-order accuracy.

of schemes satisfying a maximum-principle, *i.e.*, such that if $m < u_0 < M$ then $u^\Delta(t, \cdot) \in [m, M]$ [341], such as the second-order central-scheme by Kurganov and Tadmor [190].

In practice, none of the commonly adopted modern multidimensional HRSC schemes that we are going to present in this chapter are TVD or have been proven to be TV-stable, even though the numerical evidence seems to suggest that these schemes do converge to the entropic solution of conservation laws. Furthermore, in the case of general systems of non-linear conservation laws, no scheme has been proven to be stable or convergent⁶ even in the one-dimensional case [197], although there is numerical evidence suggesting that these schemes are convergent.

It is common belief that a proof of convergence for high order schemes in the scalar case will require the development of a more precise characterization of piecewise-regular entropic solutions of conservation laws [312]. Similarly, it is likely that a better understanding of the mathematical properties of systems of conservation laws is necessary in order to prove the convergence of their numerical approximation schemes. Meanwhile, the study of HRSC schemes for systems of conservation laws has mostly been done in a heuristic way, starting from the theory developed in the one-dimensional scalar case.

3.2 Finite-Volume Methods

Finite-volume (FV) methods are among the most popular methods in computational fluid-dynamics in general and in relativistic hydrodynamics in particular. The first FV scheme, the Godunov scheme [149], was one of the first examples of a *monotonicity-preserving scheme*⁷, *i.e.*, amongst the first schemes able to compute solutions with shocks without introducing spurious extrema in the solution. It is also the least-dissipative among the monotone methods. There are many different FV, or Godunov-type, methods for conservation laws currently available, see *e.g.*, [323] for a complete overview of the subject, but they are all built on top of the Godunov method. For this reason we here start with an extended treatment of the Godunov method. Next we discuss second-order Godunov-type schemes which constitute the current standard in relativistic hydrodynamics. Finally, we treat the extension of FV methods to even higher orders and discuss the difficulties involved in the process in the case of relativistic hydrodynamics.

3.2.1 The Godunov Method

The development of the theory behind all modern shock-capturing methods started with a very influential paper by Godunov [149]. First of all, Godunov studied the properties of numerical schemes in the case of the linear-advection equation. He proved a theorem stating that all monotonicity-preserving schemes for the advection equation are at most first-order accurate. In his proof Godunov tacitly assumed the discretization of the linear advection equation would be done using a linear scheme – a fact that was overlooked for many years until

⁶For generic initial data. Convergence has been shown in special cases, *e.g.*, the Riemann problem for the Euler equations in one dimension.

⁷Not to be confused with a monotone scheme!




Figure 3.1: Spacetime representation of the Godunov method. The solution at the time $n + 1$ is obtained by integrating in time the fluxes across the interfaces given by the solution of the local Riemann problems. In particular the spatial average of the solution in the green (cyan) regions is given by the integrated fluxes crossing the vertical green (cyan) lines. Note that the scheme works only under the assumption that the fluxes are not influenced by the results from nearby Riemann problems, but the Riemann problems themselves are allowed to interact. For this reason the Godunov scheme works with $\text{CFL} \leq 1/a$.

Boris [60] and van Leer [325] separately realized that higher-order, monotonicity-preserving schemes can be constructed, but they necessarily have to be non-linear, even for linear equations. More on this in the next sections.

In the same publication Godunov pointed out that the first-order upwind algorithm is in some sense the “best” algorithm for the advection equation: it is the most accurate (linear!) monotonicity-preserving scheme. He also went on to explain how the upwind method could be extended to the non-linear case, giving birth to the Godunov scheme.

In order to keep the notation simple, we consider first the case of a one-dimensional scalar hyperbolic equation

$$\partial_t u + \partial_x f(u) = 0. \quad (3.31)$$

We introduce a uniformly-spaced spatial and temporal grid

$$x_i = i\Delta^1, \quad i \in \mathbb{Z}, \quad t_n = n\Delta^0, \quad n \in \mathbb{N}, \quad (3.32)$$

we introduce also the so-called control volumes, $[x_{i-1/2}, x_{i+1/2}]$, and we define the averages of u over the control volumes to be

$$\bar{u}_i^n = \frac{1}{\Delta^1} \int_{x_{i-1/2}}^{x_{i+1/2}} u(x, t^n) dx. \quad (3.33)$$

If we average (3.31) over one control-volume and one time-step we obtain

$$\frac{\bar{u}_i^{n+1} - \bar{u}_i^n}{\Delta^0} = \frac{1}{\Delta^1} \int_{t_n}^{t_{n+1}} \{f[u(t, x_{i-1/2})] - f[u(t, x_{i+1/2})]\} dt. \quad (3.34)$$

It is important to point out that (3.34) is exact. No approximation has been done yet.

In a FV scheme one constructs a numerical scheme by approximating the integral formulation of the conservation law (3.34). At each time-step one starts with the approximate solution $\{U_i^n\}_{i \in \mathbb{Z}} \approx \{\bar{u}_i^n\}_{i \in \mathbb{Z}}$ and seeks a way to construct the solution at the next time-step, $\{U_i^{n+1}\}_{i \in \mathbb{Z}}$. Godunov's idea was the following: if we *assume* the solution to be piece-wise constant, *i.e.*, if we set

$$U^n(x) = \sum_{i \in \mathbb{Z}} U_i^n \chi_i(x), \quad (3.35)$$

where $\chi_i(x)$ is the characteristic function of $[x_{i-1/2}, x_{i+1/2}]$, *i.e.*,

$$\chi_i(x) = \begin{cases} 1, & \text{if } x \in [x_{i-1/2}, x_{i+1/2}], \\ 0, & \text{otherwise,} \end{cases} \quad (3.36)$$

then (3.34) can be solved exactly for small enough time-steps, Δ^0 . The reason why this is the case is that, in order to compute the integrals in the r.h.s. of (3.34), it is only necessary to solve a sequence of Riemann problems centered about the interfaces between the different control-volumes to obtain $u(t, x_{i-1/2})$ for all i and for all $t \in [t_n, t_{n+1}]$. As long as the time-step satisfies the CFL condition

$$\text{CFL} := \frac{\Delta^0}{\Delta^1} \leq \frac{1}{c}, \quad (3.37)$$

where c is the maximum propagation speed, the interface value of the solution of the various Riemann problems are independent from each other and can be computed exactly for most conservation laws. Finally the time-integration is trivially done because $u(t, x_{i-1/2})$ turns out to be constant in time, *e.g.*, [196].

The Godunov scheme can, then, be summarized as follows:

1. given $\{U_i^n\}_{i \in \mathbb{Z}}$ one constructs a piece-wise constant function $U^n(x)$, such that $U^n(x) = U_i^n$ if $x \in [x_{i-1/2}, x_{i+1/2}]$;
2. (3.34) is then evolved exactly for one time-step with initial data given by $U^n(x)$ to obtain $\{U_i^{n+1}\}_{i \in \mathbb{Z}}$.

Note that the only approximation made in the Godunov scheme is the assumption that the solution is piece-wise constant in each control-volume. A graphical representation of the Godunov scheme is given in Figure 3.1.

Since the solution $u(t, x_{i-1/2})$ depends only on U_{i-1}^n and U_i^n , *i.e.*, $u(t, x_{i-1/2}) =: u^*(U_{i-1}^n, U_i^n)$, we can define the *numerical flux* $F(U_{i-1}^n, U_i^n) := f[u^*(U_{i-1}^n, U_i^n)]$ and write the Godunov scheme in a more explicit way as

$$\frac{U_j^{n+1} - U_j^n}{\Delta^0} = \frac{1}{\Delta^1} [F(U_{i-1}^n, U_i^n) - F(U_i^n, U_{i+1}^n)]. \quad (3.38)$$

This immediately suggests the semi-discrete form of the FV method

$$\frac{dU_i}{dt} = \frac{1}{\Delta^1} [F(U_{i-1}, U_i) - F(U_i, U_{i+1})]. \quad (3.39)$$

Obviously if (3.39) is evolved in time with the Euler method this scheme reduces to (3.38). At the same time, the use of higher-order time-integrators is pointless since the time-update in the Godunov method is already exact. For these reasons the reader might wonder why we bothered to introduce the semi-discrete form of the Godunov scheme. The reason is two-fold. First of all higher-order FV schemes are easier to construct in the semi-discrete form, especially in the multi-dimensional case. Secondly, the semi-discrete form is very useful if one wants to couple the hydrodynamic equations with some other system of equations, *e.g.*, the spacetime evolution equations, that is not solved using a FV scheme. For these two reasons we do not discuss methods in their fully-discrete form and we refer the reader to *e.g.*, [323] for a more complete discussion of FV methods (see also [139] for a recent review concerning high-order one-step time discretization methods for FV and discontinuous Galerkin methods).

The FV method is easily extended to the multi-dimensional case and to general unstructured grids. Given a control volume Ω_α , the semi-discrete scheme reads

$$\frac{dU_\alpha}{dt} = \frac{1}{|\Omega_\alpha|} \int_{\partial\Omega} \mathbf{f} \cdot \mathbf{v} d\Sigma, \quad (3.40)$$

where $\mathbf{v}$ is the inwards pointing normal to Ω . In the case of three-dimensional Cartesian grids,

$$\mathbf{x}_{i,j,k} = (i\Delta^1, j\Delta^2, k\Delta^3), \quad i, j, k \in \mathbb{Z}, \quad (3.41)$$

this simplifies to

$$\begin{aligned} \frac{dU_{i,j,k}}{dt} = & \frac{1}{\Delta^1} [F^1(U_{i-1,j,k}, U_{i,j,k}) - F^1(U_{i,j,k}, U_{i+1,j,k})] \\ & + \frac{1}{\Delta^2} [F^2(U_{i,j-1,k}, U_{i,j,k}) - F^2(U_{i,j,k}, U_{i,j+1,k})] \\ & + \frac{1}{\Delta^3} [F^3(U_{i,j,k-1}, U_{i,j,k}) - F^3(U_{i,j,k}, U_{i,j,k+1})], \end{aligned} \quad (3.42)$$

where F^1, F^2 and F^3 are the numerical fluxes associated with f^1, f^2 and f^3 respectively.

Finally, we point out that the extension to systems of equations is straightforward as long as the solution of the Riemann problem can be constructed.

One important remark about FV schemes is that, while the full solution of the Riemann problem is computed at each interface, only a very small part of the information contained in this solution is retained, while the rest is destroyed in the averaging step, *i.e.*, when we set $U^n(x) = \sum_i U_i^n \chi_i(x)$. For this reason it should not be surprising that it is possible to construct monotone FV schemes even with the use of simplified or approximate Riemann solvers, at the price of increasing the diffusivity of the method [162]. This is particularly useful for systems, such as the relativistic Euler equations, where the solution of the Riemann problem is very complicated and requires an expensive iterative procedure [208, 257, 144]. Examples are the Roe Riemann solver [277], the so-called Marquina flux-formula [115] and the HLL solver [125]. A point to note here is that, although these flux-formulas are often called “approximate Riemann solvers”, not all of them are directly constructed as

$$F(U_L, U_R) = f(u^*(U_L, U_R)), \quad (3.43)$$

i.e., not all of them are really constructed starting from an approximate solution of the original Riemann problem. Instead, some of them approximate the flux-function *directly*. Harten et al. [162] showed that, as long as the flux formula satisfies

$$F(u, u) = f(u), \quad \forall u \in \mathbb{R}, \quad (3.44)$$

the scheme obtained by using the approximate numerical flux instead of the Godunov one is consistent and conservative. In particular the Lax-Wendroff theorem holds. As a consequence, if the resulting scheme is non-linearly stable⁸ and if the flux-formula is compatible with the entropy inequality⁹, then the resulting scheme will converge to the correct entropic solution of the conservation law. Finally, we point out that in the case of high order FV schemes, the benefits of using the Godunov flux instead of an approximate flux are reduced as higher-order schemes naturally have lower numerical dissipation.

3.2.2 TVD Finite-Volume Methods

The key behind the development of modern FV schemes is the observation that higher-order information on the solution can be *reconstructed* from the cell averages $\{U_i^n\}_{i \in \mathbb{Z}}$ without introducing spurious oscillations using a non-linear reconstruction procedure. Instead of setting $U(x) = U_i$ if $x \in [x_{i-1/2}, x_{i+1/2}]$, one constructs a second-order approximation of u in any control volume as

$$U_i(x) = U_i + \sigma_i(x - x_i), \quad x \in [x_{i-1/2}, x_{i+1/2}], \quad (3.45)$$

where σ_i is the reconstructed slope in the control volume $[x_{i-1/2}, x_{i+1/2}]$, which we specify later. We notice that the reconstructed profiles of U in the different control volumes might not agree at the interface, hence the introduction of the index in $U_i(x)$. If we denote as $U_{i-1/2}^\pm$ the two values $U_{i-1}(x_{i-1/2})$ and $U_i(x_{i-1/2})$, then a second order scheme can be obtained by setting

$$\frac{dU_i}{dt} = \frac{1}{\Delta t} [F(U_{i-1/2}^-, U_{i-1/2}^+) - F(U_{i+1/2}^-, U_{i+1/2}^+)]. \quad (3.46)$$

The reason why this works is that, at time t_n^+ , the solution at the interface between two cells depends only on the jump between the two states, *i.e.*, the instantaneous flux is given by $F(U_L, U_R)$, where U_L and U_R are the left and right states across the interface. Notice that in a fully-discrete scheme it would no longer be true that the state $u(t, x_{i-1/2})$ is time-independent and it would be necessary to either solve a generalized Riemann problem *i.e.*, with initial data having piecewise linear profile, or use a predictor-corrector approach to compute the fluxes with second-order accuracy in time. In the case of semi-discrete schemes it is sufficient to use a second order time-integrator to evolve (3.46) to obtain a scheme which is second order in time and space.

We now turn to the choice of σ_i . It can be proven that, if

$$\text{TV} \left[\sum_i U_i(x) \chi_i; \mathbb{R} \right] \leq \text{TV} \left[\sum_i U_i \chi_i(x); \mathbb{R} \right], \quad (3.47)$$

⁸This typically requires the approximate flux-formula to correctly upwind the solution.

⁹In the case of Riemann solver constructed from approximate solutions it is enough to choose an entropic solution.

then the scheme (3.46) is TVD [196]. This condition is satisfied with the use of special non-linear limiters, called *slope-limiters*. An example of the slope-limited method is given by the minmod limiter

$$\sigma_i = \frac{1}{\Delta^1} \minmod(U_{i+1} - U_i, U_i - U_{i-1}), \quad (3.48)$$

where

$$\minmod(z_1, \dots, z_n) = \begin{cases} \min_i z_i, & \text{if } z_i > 0 \forall i, \\ \max_i z_i, & \text{if } z_i < 0 \forall i, \\ 0, & \text{otherwise.} \end{cases} \quad (3.49)$$

Another popular example is given by the minmod2 reconstruction

$$\sigma_i = \frac{1}{\Delta^1} \minmod \left[2(U_{i+1} - U_i), \frac{U_{i+1} - U_{i-1}}{2}, 2(U_i - U_{i-1}) \right], \quad (3.50)$$

The idea behind all of these slope limiters is to reduce the scheme to first-order near shocks in order to avoid oscillations. We also point out that most of the slope-limiting techniques are only first-order accurate at extrema.

TVD FV schemes can be extended to the multi-dimensional case in a direction-by-direction fashion, for Cartesian grids, or by performing the reconstruction in barycentric coordinates in the unstructured case. The extension to systems of equations is also simple: the reconstruction can be performed component-by-component. Another popular choice, especially in the case of the relativistic Euler equation, is to perform a reconstruction of the primitive variables exploiting the fact that $\bar{u}_i = u(x_i)$ to second-order accuracy. Reconstructing in primitive variables reduces the possibility that the reconstructed values at the interface are non-physical and it is also computationally less expensive since the relation between primitive and conservative quantity does not need to be inverted.

3.2.3 Higher-Order Finite-Volume Methods

The procedure used to extend the Godunov scheme to second order suggest a strategy that can be applied to create even higher-order schemes. Instead of reconstructing only a second-order approximation of u as a piecewise linear function one could attempt to reconstruct it to higher order as a piecewise polynomial.

Reconstruction Operators

To explain the basic idea we start again from the one-dimensional case. We consider a generic function $v(x)$ and analyze the problem of reconstructing it at high order starting from the volume averages

$$\tilde{v}_i := \frac{1}{\Delta^1} \int_{x_{i-1/2}}^{x_{i+1/2}} v(x) dx. \quad (3.51)$$

A reconstruction operator, $\mathcal{R}$, is a non-linear operator yielding a high-order approximation of v at a given point x using its volume averages, $\tilde{v}_i$. Since $v(x)$ can be discontinuous, we distinguish two different reconstruction operators,

$\mathcal{R}^-$ and $\mathcal{R}^+$, called the left-biased and right-biased reconstruction operators, such that

$$[\mathcal{R}^-(\{\tilde{v}_i\})](x) = \lim_{y \rightarrow x^-} v(y) + O(\Delta^r), \quad (3.52)$$

$$[\mathcal{R}^+(\{\tilde{v}_i\})](x) = \lim_{y \rightarrow x^+} v(y) + O(\Delta^r), \quad (3.53)$$

where we have indicated with $\mathcal{R}^+(\{\tilde{v}_i\})$ the notion that $\mathcal{R}$ is an operator that acts on a set of averages $\tilde{v}_i$, and where r is the order of the reconstruction operator $\mathcal{R}$. Hereafter we use the notation $v_{i+1/2}^-$ and $v_{i+1/2}^+$ to denote the reconstructed values at $x_{i+1/2}$ using $\mathcal{R}^-$ and $\mathcal{R}^+$ respectively. Examples of such operators are the reconstruction operators given by the piecewise parabolic method (PPM) [96, 94], the piecewise hyperbolic method (PHM) [207], the essentially non-oscillatory (ENO) [163, 299, 298], weighted essentially non-oscillatory (WENO) [202, 176] and monotonicity-preserving (MP5) [310] algorithms.

The PPM reconstruction extended the slope-limiter approach to the case of a piecewise parabolic interpolation. It uses a series of limiters, that have been thoroughly studied numerically, to avoid the introduction of spurious oscillations. The PPM algorithm has not been proved to be TVD, but it has become the standard method of choice in many modern FV schemes and in numerical relativistic hydrodynamics [224, 33].

The ENO scheme is based on a standard Lagrange interpolation and uses a recursive procedure, which in principle can be extended to any order of accuracy, to determine the stencil on which the solution is smoothest in order to avoid Gibbs oscillations due to interpolation across discontinuities. The ENO interpolation is not TVD, but [163] showed that

$$\text{TV}[\mathcal{R}(\{\tilde{v}_i\}); \mathbb{R}] \leq \text{TV}[\{\tilde{v}_i\}; \mathbb{R}] + O((\Delta^1)^r), \quad (3.54)$$

where r is the order of the reconstruction. The second order ENO reconstruction is equivalent to the minmod approach [196].

The WENO scheme is a modified version of the original ENO approach. Instead of selecting only the stencil on which the solution is smoothest, the WENO reconstruction takes a weighted average of the reconstructed polynomial on each stencil. The weights are chosen so that they are very small for non-smooth stencils, while, at the same time, the order of accuracy is maximized in smooth regions. WENO schemes are faster than ENO schemes using the same stencil because the WENO algorithm is free from conditional statements and, at the same time, it achieves a formal order of accuracy of $2r - 1$, where r is the order of the ENO algorithm, by combining the results of multiple reconstructions. WENO schemes are theoretically available for any order, but schemes of order higher than 7 ($r > 4$) are empirically found to require the use of additional limiters [35] or order-reducing technique [317, 141] to be stable.

Many different versions of the WENO algorithm are available that differ in the way in which the weights are constructed. Examples are the mapped-WENO schemes of [166], that use a mapping procedure to decrease the amount of dissipation of the scheme, or the WENOZ scheme by [59] that employs improved non-linear weights to obtain results comparable with the fifth order mapped-WENO schemes at a lower computational cost. Gerolymos et al. [141] presented a comparative study of different WENO schemes and provide

tabulated coefficients and implementation details for methods of order up to 17.

Another interesting type of WENO scheme is given by the bandwidth-optimized schemes. In classical WENO schemes, the reconstruction is obtained from a weighted average of a set of lower-order reconstructions of the $\tilde{v}_i$'s on a number of overlapping stencils. The weights are computed using some non-linear smoothness indicators designed in such a way that the maximum possible order of accuracy is obtained in the case of smooth solutions. On the other hand, when discontinuities are detected, the order is automatically reduced to avoid spurious oscillations. The bandwidth-optimized WENO schemes differ from the classical ones because they use a more symmetric stencil and because their weights are not designed to yield the maximum possible formal order of accuracy for smooth solutions, but are instead tuned to minimise the attenuation of high-frequency modes. In other words, while in the classical WENO case the weights for a smooth function are chosen to match as many terms as possible in the Taylor expansion of the target function, in the bandwidth-optimized case the coefficients are chosen to yield the best possible approximation of the Fourier coefficients of the function to be reconstructed. The non-linear smoothness indicators are also modified to avoid "over-adaptation" of the scheme and minimise the amount of numerical dissipation, see [211, 316] for details.

Finally the MP5 scheme is basically a fifth-order extension of the PPM scheme. It uses fifth order reconstruction combined with a flattening procedure designed to avoid the creation of artificial extrema in the function to be reconstructed. The monotonicity-preserving reconstruction has the nice property that no spurious oscillations can be produced by the reconstruction, which is not guaranteed for the WENO schemes. On the other hand the limiting procedure employed by MP5 requires the use of a series of conditional statements in the code, that are not present in the WENO case.

Very-High-Order Finite-Volume Schemes

Once high-order, non-oscillatory reconstruction operators are available one can obtain a high-order version of the FV method in 1D by simply repeating the construction we have followed in deriving second-order schemes. In particular we can set

$$\frac{dU_i}{dt} = \frac{1}{\Delta^1} [F(U_{i-1/2}^-, U_{i-1/2}^+) - F(U_{i+1/2}^-, U_{i+1/2}^+)], \quad (3.55)$$

where U^- and U^+ are computed using the $\mathcal{R}^+$ and $\mathcal{R}^-$ operators. The scheme obtained in this way has formally the same order of accuracy as the reconstruction algorithm. The actual order of accuracy depends on the solution, since all stable reconstruction operators reduce to first order near discontinuities (or under-resolved, but smooth features of the solution).

The main problem with high-order FV methods is that they are much more complex to extend to the multidimensional case. The scheme can not be written in the form (3.42) as the fluxes have to be computed using suitable quadrature formulas, otherwise the accuracy is limited to second order independently on the order of the reconstruction.

In the case of systems of equations, experience seems to suggest that reconstruction should not be performed component-by-component, but on local characteristic variables, to avoid spurious oscillations in the numerical solution,

especially for schemes with order higher than two. More on this in Section 3.4. We also point out that the reconstruction can not be performed on primitive variables as in general it is not possible to convert the volume-averaged conserved variables to volume-averaged primitive variables with accuracy higher than second order.

Finally, in the general-relativistic case, high-order FV methods are particularly expensive because, on the one hand, they require the use of very high-order schemes to interpolate the metric at the quadrature points for the calculation of the fluxes and, on the other hand, they require high-order quadrature of the metric source terms.

Despite these difficulties, due to their superior accuracy with respect to second-order schemes, high-order FV schemes are available in Newtonian and relativistic hydrodynamics, *e.g.*, [317], and have been also extended to unstructured grids, using a generalized WENO algorithm *e.g.*, [122].

3.3 Central Methods

The first monotone scheme for conservation laws was introduced by Lax and Friedrichs (LxF) [195, 134]. This scheme has a significantly larger numerical diffusion with respect to the Godunov scheme, but has the advantage of not requiring any Riemann solver and it is for this reason called a *central* scheme, as opposed to upwind Godunov-like schemes, thus being much more simple to implement and computationally less expensive. Unfortunately the Lax-Friedrichs scheme is very inaccurate due to its high dissipation and for this reason is seldom employed. The interest in central schemes increased when Nessyahu and Tadmor (NT) [235] showed the first extension of the Lax-Friedrichs scheme to second order. This scheme was later improved by Kurganov and Tadmor (KT) [190] in an influential paper that set the basis for a whole new family of schemes: high-order central schemes.

The Lax-Friedrichs scheme is in some sense the dual of the Godunov scheme. Given the local averages of the numerical solution at time t_n , $\{U_i^n\}_{i \in \mathbb{Z}}$, one constructs a solution at time t_{n+1} by considering the local averages *on the dual grid*, $\{x_{i+1/2}\}_{i \in \mathbb{Z}}$, hence the name central schemes for the family of schemes based on the Lax-Friedrichs approach, *i.e.*, one obtains $\{U_{i+1/2}^{n+1}\}_{i \in \mathbb{Z}}$. The scheme then proceeds by evolving the solution by alternating between the primal and the dual grid.

In order to see how this works, suppose we have the solution at time t_n , $\{U_i^n\}_{i \in \mathbb{Z}}$ on the primal grid. In order to find $\{U_{i+1/2}^{n+1}\}_{i \in \mathbb{Z}}$, one proceeds as in the Godunov scheme and integrates (3.31) over the dual grid to obtain

$$\frac{U_{i+1/2}^{n+1} - U_{i+1/2}^n}{\Delta^0} = \frac{1}{\Delta^1} \int_{t_n}^{t_{n+1}} \{f[u(t, x_i)] - f[u(t, x_{i+1})]\} dt. \quad (3.56)$$

The main difference with respect to the Godunov scheme is that, as long as $\text{CFL} \leq 1/2a$, a being the maximum local-characteristic speed, the term in the r.h.s. of (3.56) can be computed without any need for a Riemann solver using a simple quadrature in time, because the solution at cell centers remains free of discontinuities for the duration of the whole time-step. In particular if we




Figure 3.2: Spacetime representation of a central scheme. The solution at the time $n + 1$ is obtained on the staggered grid by integrating the fluxes at the cell-centers. In particular the spatial average of the solution in the green region is given by the integrated fluxes crossing the green vertical lines. Note that the fluxes can be computed with high accuracy as long as $\text{CFL} \leq 1/2a$, because the solution in the vertical green bands remains smooth over this time scale.

assume the solution at time t_n to be piecewise constant as in the Godunov scheme,

$$U^n(x) = \sum_i U_i^n \chi_i(x), \quad (3.57)$$

then we have $u(t, x_i) = U_i^n$ for $t \in [t_n, t_{n+1}]$ and

$$U_{i+1/2}^n = \frac{U_i^n + U_{i+1}^n}{2} \quad (3.58)$$

so that (3.56) becomes

$$U_{i+1/2}^{n+1} = \frac{U_i^n + U_{i+1}^n}{2} + \frac{\Delta^0}{\Delta^1} [f(U_i^n) - f(U_{i+1}^n)]. \quad (3.59)$$

A graphical representation of the Lax-Friedrichs scheme is given in Figure 3.2.

The central scheme was extended to second order by Nessyahu and Tadmor [235] who used minmod to reconstruct a piecewise linear approximation of u^n , instead of the piecewise constant reconstruction, and a predictor-corrector approach to obtain second order in time flux quadratures.

One of the things to point out is that neither the LxF nor the NT scheme admit a semi-discrete form. Actually their dissipation is CFL-dependent, making these schemes extremely dissipative if a small time-step is used. To overcome this issue Kurganov and Tadmor [190] modified the NT scheme in such a way so as to limit the averaging on the dual grid to the region actually spanned by the “Riemann fan”¹⁰. In other words, in the KT scheme $U_{i+1/2}^{n+1}$ is only supported

¹⁰The Riemann fan is the region of the spacetime spanned by the largest characteristics of the Riemann problem.

on the region influenced by the results of the local Riemann problem and not on the whole dual mesh. In this way they managed to reduce the dissipation of the scheme and obtained a method admitting the semi-discrete form

$$\frac{dU_i}{dt} = \frac{1}{\Delta^1} [F(U_{i-1/2}^-, U_{i-1/2}^+) - F(U_{i+1/2}^-, U_{i+1/2}^+)], \quad (3.60)$$

with numerical flux

$$F(U_L, U_R) = \frac{f(U_L) + f(U_R)}{2} - \frac{a}{2}[U_R - U_L]. \quad (3.61)$$

The KT scheme is substantially equivalent to a Godunov method in which the fluxes are computed with the flux-formula (3.61), commonly called the local-Lax-Friedrichs or Rusanov flux [190]. For this reason the scheme is stable even with $1/2 < a \cdot \text{CFL} \leq 1$, even though the underlying interpretation of the scheme is lost in this case.

Finally a comment on the distinction between Godunov-type and central scheme: as the reader might have guessed this is a bit arbitrary if modern central schemes are considered, because they all can be interpreted as being FV schemes with the Rusanov “Riemann solver”. Nevertheless the construction and analysis of upwind and central schemes is significantly different. This is important when studying the properties of the schemes. For instance, the KT scheme is the only scheme for which a maximum principle has been proven in the multi-dimensional case even without the use of special maximum-principle-enforcing limiters.

3.4 Finite-Difference Methods

High-order upwind finite-difference (FD) schemes work by directly approximating the point-wise value of the solution, as opposed to its volume-averages. They were mainly introduced as a more efficient alternative to very high-order ENO and WENO schemes [299, 298, 176]. High-order FD schemes are equivalent to high-order FV methods in one spatial dimension, but can be several times more efficient in the multidimensional case [296, 300]. We can only provide here a minimal description of FD-HRSC methods and we refer the interested reader to [296] for a more detailed description of FD ENO/WENO HRSC schemes, and to [226] for a detailed description of the FD MP5 scheme.

In order to illustrate the differences between FV and FD schemes we consider directly the case of a system of hyperbolic balance-laws in the form

$$\partial_t \mathbf{F}^0(\mathbf{u}) + \partial_i \mathbf{F}^i(\mathbf{u}) = \mathbf{S}(\mathbf{u}). \quad (3.62)$$

Similarly to the case of FV methods we introduce a uniform Cartesian grid

$$\mathbf{x}_{i,j,k} = (i\Delta^1, j\Delta^2, k\Delta^3), \quad i, j, k \in \mathbb{Z}. \quad (3.63)$$

Also, since we do not have to distinguish between the exact solution, its volume averages, their approximation and the reconstructed solution, we use a simplified notation in which $u_{i,j,k}$ denotes the numerical approximation of the quantity u at the point $\mathbf{x}_{i,j,k}$.

The FD scheme for (3.62) is written in a form which is formally identical to (3.42):

$$\begin{aligned} \frac{d\mathbf{F}_{i,j,k}^0}{dt} = \mathbf{S}_{i,j,k} + & \frac{\mathbf{F}_{i-1/2,j,k}^1 - \mathbf{F}_{i+1/2,j,k}^1}{\Delta^1} \\ & + \frac{\mathbf{F}_{i,j-1/2,k}^2 - \mathbf{F}_{i,j+1/2,k}^2}{\Delta^2} + \frac{\mathbf{F}_{i,j,k-1/2}^3 - \mathbf{F}_{i,j,k+1/2}^3}{\Delta^3}. \end{aligned} \quad (3.64)$$

The important point is that here the terms like $(\mathbf{F}_{i-1/2,j,k}^1 - \mathbf{F}_{i+1/2,j,k}^1)/\Delta^1$ are not defined as integrals along the boundary of the control-volume, but represent directly an high-order, non-oscillatory, approximation of the point-wise value of $-\partial_1 \mathbf{F}^1$ at $\mathbf{x}_{i,j,k}$, which we explain in a moment.

It is already clear that, in the low-order case, FV and FD schemes are completely equivalent, as point-wise values and averages are identical at second order. In the high-order case and especially in multiple spatial dimensions, however, FD schemes are significantly less expensive than FV schemes. First of all they do not require any quadrature, Riemann solver or extra primitive recovery calls at region boundaries. Secondly the treatment of source terms, which is important in the general relativistic case also in terms of computational costs, is significantly less expensive as only the point-wise values are needed. This makes FD schemes very attractive in the quest for more accurate general relativistic hydrodynamics calculations.

To illustrate how to compute the discrete derivatives in the right-hand-side of (3.64) it is useful to take a step back and consider first a simpler scalar hyperbolic equation in one dimension (3.31). The reconstruction operators are the core ingredients of both FV and FD schemes. As we have seen, in a FV scheme they are used to compute the left and right state to be used in the (usually approximate) Riemann solver to compute the fluxes. In a FD scheme, instead, they are used to compute the above-mentioned non-oscillatory approximation of $\partial_x f$. Following [299] we introduce a function $h(x)$ such that

$$f[u(x_i)] = \frac{1}{\Delta} \int_{x_{i-1/2}}^{x_{i+1/2}} h(\xi) d\xi, \quad (3.65)$$

that is, the average of $h(x)$ between $x_{i-1/2}$ and $x_{i+1/2}$ corresponds to the value of f at x_i . Next, we note that

$$\left. \frac{\partial f}{\partial x} \right|_{x_i} = \frac{h(x_{i+1/2}) - h(x_{i-1/2})}{\Delta}, \quad (3.66)$$

where both (3.65) and (3.66) are exact expressions. Hence, by using the usual reconstruction operators $\mathcal{R}$ of order r to reconstruct $h_{i+1/2}$, one obtains a correspondingly accurate approximation of order r of the derivative $\partial f/\partial x$ at x_i . Note that h is never actually computed at any time during the calculation as we only need the values of f at the gridpoints, *i.e.*, $f[u(x_i)]$.

In order to ensure the stability of the resulting scheme, one has to take care to upwind the reconstruction appropriately. Let us first consider the case in which $f'(u) > 0$. If we set

$$\tilde{v}_i = f[u(x_i)] = \frac{1}{\Delta} \int_{x_{i-1/2}}^{x_{i+1/2}} h(\xi) d\xi \quad (3.67)$$

and

$$f_{i+1/2} := v_{i+1/2}^-, \quad f_{i-1/2} := v_{i-1/2}^-, \quad (3.68)$$

then

$$\frac{\partial f(u)}{\partial x} = \frac{f_{i+1/2} - f_{i-1/2}}{\Delta} + O(\Delta^r), \quad (3.69)$$

gives the wanted high-order approximation of $\partial_x f$ at x_i .

In the more general case, where the sign of $f'(u)$ is undetermined, in order to compute $f_{i+1/2}$, we have to split f in a right-going, f^+ , and a left-going, f^- , flux, $f = f^+ + f^-$, and use the appropriate upwind-biased reconstruction operators separately on both parts, in order to guarantee the stability of the method. We point out that the use of the flux-splitting procedure to make the scheme upwind is akin to the use of Riemann solvers in FV schemes. Indeed some of the flux-splitting algorithms can be associated to particular Riemann solvers in the FV case.

There are several different ways to perform such a split. The ones relevant for this thesis are two: the Roe flux-split, *i.e.*,

$$f = f^\pm, \quad \text{if } [f'(\bar{u})]_{x_{i+1/2}} \gtrless 0, \quad (3.70)$$

where $\bar{u}_{i+1/2} := \frac{1}{2}(u_i + u_{i+1})$, and the Lax-Friedrichs or Rusanov flux-split [295], *i.e.*,

$$f^\pm = f(u) \pm \alpha u, \quad \alpha = \max[f'(u)], \quad (3.71)$$

where the maximum is taken over the stencil of the reconstruction operator. The Roe flux-split is less dissipative and yields a computationally less-expensive scheme, since only one reconstruction is required instead of two, but its use can result in the creation of entropy-violating shocks in the presence of transonic rarefaction waves, see, *e.g.*, [196], and it is also susceptible to the carbuncle (or odd-even decoupling) phenomenon [262]. To avoid these drawbacks, in [264] we proposed to switch from the Roe to the Lax-Friedrichs flux split when u or f are not monotonic within the reconstruction stencil¹¹.

We note that the condition that we use to switch from the Roe to the Lax-Friedrichs flux split is weaker than the commonly employed condition on the sign of $f'(u)$, see *e.g.*, [196], in the sense that it results in a more frequent use of the Lax-Friedrichs split with respect to the usual one. Our experience shows that this prescription works very well: it is computationally less expensive to compute with respect to the standard one, since u and f are already evaluated on the grid, while $f'(u)$ is not, and it seems to be sufficient to avoid the carbuncle phenomenon in all the tests that we performed. All the results that we are going to present in this thesis have been obtained using this Roe-split with this “entropy fix”.

We now go back to the more general system of equations (3.62). The derivatives $\partial_a F_{i,j,k}^a$ can be computed following the procedure outlined above on a component-by-component basis. This approach is commonly adopted in the case of low-order schemes, but it often results in spurious numerical oscillations in the high-order (usually higher than second) case. To avoid this issue,

¹¹Notice that we are tacitly assuming f to be a convex function. This assumption is satisfied in the case of the equations of relativistic-hydrodynamics, but not, for instance, in the case of relativistic-magneto-hydrodynamics.

the reconstruction should be performed on the local characteristic variables of the systems. To avoid an excessively complex notation, let us concentrate on the fluxes in the x -direction; in this case, to reconstruct $\mathbf{F}_{i+1/2,j,k}^1$ we introduce the Jacobian matrices

$$\mathbf{A}^\alpha = \frac{\partial \mathbf{F}^\alpha}{\partial \mathbf{u}} \bigg|_{\bar{\mathbf{u}}}, \quad \alpha = 0, 1, \quad (3.72)$$

where

$$\bar{\mathbf{u}} := \frac{1}{2}(\mathbf{u}_{i,j,k} + \mathbf{u}_{i+1,j,k}), \quad (3.73)$$

is the average state at the point where the reconstruction is to be performed. We point out that the average (3.73), and appearing in (3.70) and (3.72), is much simpler than the average of $\mathbf{u}_{i,j,k}$ and $\mathbf{u}_{i+1,j,k}$ suggested by [277]. In [339] it was checked that the use of (3.73) in place of the average suggested by [277] does not influence significantly the quality of the solution in the case of finite-differencing schemes, even in the relativistic case.

Strong hyperbolicity of (3.62) implies that $\mathbf{A}^0$ is invertible and the generalised eigenvalue problem

$$[\mathbf{A}^1 - \lambda_{(l)} \mathbf{A}^0] \mathbf{r}_{(l)} = 0, \quad (3.74)$$

has only real eigenvalues, $\lambda_{(l)}$, and N independent, real right-eigenvectors, $\mathbf{r}_{(l)}$ [see, e.g., [18]]. We denote with $\mathbf{R}$ the matrix of right eigenvectors, i.e.,

$$\mathbf{R}_J^I = r_{(J)}^I, \quad (3.75)$$

and with $\mathbf{L}$ its inverse. We define the local characteristic variables

$$\mathbf{w} = \mathbf{L} \mathbf{u}, \quad \mathbf{Q} = \mathbf{L} \mathbf{F}^1, \quad (3.76)$$

and compute $\mathbf{Q}_{i+1/2,j,k}$ doing a component-wise reconstruction, where $\mathbf{w}$ is used in place of $\mathbf{u}$ and $\mathbf{Q}$ in place of $\mathbf{f}$ in the (3.71). Finally we set

$$\mathbf{F}_{i+1/2,j,k}^1 = \mathbf{R} \mathbf{Q}_{i+1/2,j,k}. \quad (3.77)$$

This procedure is repeated in the other directions and yields the wanted approximations of the $\partial_a \mathbf{F}^a$ terms in $\mathbf{x}_{i,j,k}$. The results that we present in this thesis have been obtained performing the reconstruction of the local characteristic variables.

The main limitation of finite-differencing schemes is that they are particularly suited for Cartesian grids, either uniformly spaced or with Berger-Oliger-style AMR [45]. Finite-difference schemes also reduce to second-order at the boundary of refinement levels if cell-centered AMR with refluxing [44] is used. Both of these limitations are, on the other hand, of secondary importance for the kind of applications that we present in this thesis.

3.5 Discontinuous Galerkin Methods

Discontinuous Galerkin (DG) methods for hyperbolic equations were first introduced by Reed and Hill [267] for a problem of neutron-transport and then further developed by Cockburn and Shu in a series of papers [92, 88, 89, 90, 86].




Figure 3.3: L^2 -norm of the error measured in the numerical solution of the linear advection equation $\partial_t u + \partial_x u = 0$ with sinusoidal initial data, $u_0 = \sin(2\pi x)$ at time $t = 10$ and for different numerical schemes. On the x-axis N is the number of DG-elements/FD-points. We point out that, although the schemes have comparable computational costs, the DG methods uses five times more degrees of freedom at a fixed N than the other two schemes.

At the same time DG methods were also developed for elliptic and parabolic problems, see [22] and references therein. Since then, these methods have become increasingly popular and have been applied to a number of classical hyperbolic, parabolic and elliptic problems [91, 75, 168]. They have been also successfully applied to the solution of the Einstein equation in vacuum by Zumbush [347] and Field et al. [128]. Finally, the first DG general-relativistic hydrodynamics code was proposed by us in [263].

The reason for the popularity of DG methods and for the growing interest in this family of schemes lies in their numerical properties.

First of all, even though they became mainstream much later than FV or FD methods, they have the largest number of proven mathematical properties. For instance, they can be proven to be intrinsically non-linearly stable, at all orders of accuracy, even though limiters and/or filtering is necessary to treat under-resolved features in the solution, such as shock-waves, to avoid the aliasing instability¹². DG methods also satisfy a cell-entropy inequality, meaning that the solution obtained with DG methods is always entropic, see *e.g.*, [87]. In addition we mention that Zhang and Shu [341] were recently able to construct maximum-principle-satisfying DG schemes with the use of a special limiting technique.

¹²The aliasing instability occurs when the solution is not well represented by the truncated expansion used in the scheme. In particular the unrepresented modes of the solution tend to be “aliased” into the evolved ones, meaning that their content is improperly transferred by the numerical scheme to lower-order modes, leading to a non-linearly unstable behaviour (see *e.g.*, [62]).

Secondly, they can deliver very high, spectral, accuracy in regions where the solution is smooth. Furthermore the numerical dissipation of DG methods can be proven to depend only on the truncation error [87], implying that the dissipation is automatically deactivated where the solution is smooth and sufficiently resolved. For instance, this is of key importance in transport problems where the numerical diffusivity of other schemes can easily become orders of magnitude larger than the physical diffusivity, thus yielding completely wrong solutions. More on this in Chapter 7.

DG schemes can also naturally accommodate general unstructured meshes and, thanks to their very compact stencil, they have nearly-optimal scalability, which makes them very promising for massively-parallel computing [50].

Finally DG methods, as well as finite-element methods (FEM) from which they derive, have the very interesting property of being covariant, because they do not require any a priori choice of a coordinate system, but can be expressed in terms of push-forwards and pull-backs from some reference element [220]. This obviously makes them very interesting for general relativistic calculations.

An illustration of the superior accuracy of DG methods is shown in Figure 3.3, where we compare the numerical errors in the solution of the advection equation obtained with the fifth order DG scheme with the ones obtained using the fifth order and a seventh order WENO FD schemes. As we can see, the DG methods out-performs both the WENO5 and the WENO7 schemes, even though the latter has a higher formal order of accuracy. We point out that the comparison is made by using the same number of DG-elements/FD-points, because the computational costs of the two schemes are similar in this case. On the other hand, the DG method uses five times more degrees-of-freedom (d.o.f.) than the FD method. A comparison in which the two schemes are used employing the same number of d.o.f. yields a similar result, with the DG method being slightly better than WENO7, but it is, in our opinion, misleading since the FD method is significantly more expensive than the DG one under these conditions. Since the computational cost of DG scales with the number of elements in a similar way as the cost of FV/FD methods scales with the number of cells/points, we always present comparisons between DG and FD/FV methods by fixing the number of elements in the DG scheme to be equal to the number of points in the FD/FV scheme. Clearly our choice is motivated by the fact that we want to compare the methods in a situation in which the total computational times are similar. On the other hand, in situations where the memory requirements are the bottleneck of the simulation, a comparison based on the degrees of freedom is more adequate.

DG methods also have their limitations. First of all they have larger memory requirements with respect to other schemes. Secondly they are still somewhat less robust than FV schemes as all the standard flattening techniques are either not sufficiently reliable in the removal of oscillations near shock-waves or tend to degrade the accuracy of the scheme also in regions where limiting would not be necessary. Finally and most importantly DG schemes require a more strict CFL condition, with respect to FD/FV methods, to be linearly stable: if the time-discretization is done with the use of strongly-stability preserving (SSP) Runge-Kutta (RK) schemes [154] of the same order as the spatial discretization, the stability condition is

$$\text{CFL} \leq \frac{1}{c} \frac{1}{2k+1}, \quad (3.78)$$




Figure 3.4: Mapping between the reference element and an element in the physical space in a DG method.

where $k + 1$ is the order of accuracy of the scheme [93]. These limitations are being addressed in the current research on DG methods with the development of hybrid DG-FV schemes such as the DG-WENO [260, 259] and $P_N P_M$ [123, 121] schemes, and with the advance in the development of local-spacetime schemes [139] (see also [168] for alternative approaches).

3.5.1 Runge-Kutta Discontinuous-Galerkin Methods

Next we present the basic ideas behind the construction of the so-called spectral discontinuous Galerkin method with numerical integration (SDGM-NI) or nodal-DG scheme, which is the most commonly adopted variant of the DG method. We refer the reader to [168] for a more complete treatment.

In order to explain the key ideas behind the DG method and to keep our notation simple, we consider the scalar problem

$$\partial_t u + \nabla \cdot \mathbf{f}(u) = 0, \quad (t, x) \in \mathbb{R}_+ \times \Omega, \quad (3.79)$$

where $\Omega \subset \mathbb{R}^d$ is a sufficiently regular, bounded, domain.

First of all we create a triangulation, $\mathcal{T}_N$, of Ω , N being the number of *elements* of the triangulation. This is represented by the union of the images of a reference element T , usually a triangle (or a square) in 2D and a tetrahedron (or a cube) in 3D, through a family of diffeomorphisms, $\varphi_j: T \rightarrow \Omega$, $\Omega_j = \varphi_j(T)$, such that

$$\bigcup_{j=1}^N \Omega_j = \Omega, \quad \Omega_i \cap \Omega_j = \emptyset, \quad \text{if } i \neq j. \quad (3.80)$$

An example of such a triangulation is shown in Figure 3.4. A key point here is that these mappings are not merely a tool used to do “book-keeping” of the grid, but are actually an active part of the scheme: with the use of these mappings one can “pull-back” the equations to the reference element, where all the discrete differential operators have been pre-computed. In the classical case, this means that derivatives, interpolation and integration operators can be easily implemented once and for all in a simple geometry, while the shape of

the elements in the physical space can be arbitrarily complex. For instance one can easily construct elements whose boundaries are smooth curves adapting to the boundary of the physical domain. In addition, in the relativistic case, one could exploit the covariance of the equations to formulate them directly in the coordinate system generated by the diffeomorphisms (by pushing forward the ones in which the reference element is defined) effectively making the scheme completely coordinate free (because no coordinate system would need to be defined on the physical space).

The second step is to derive a condition similar to the weak formulation of (3.79). In particular, in order to construct a semi-discrete scheme, the classical approach is to derive a form of the weak formulation in which u is assumed to be a smooth function of time and of bounded-variation in space¹³. In Chapter 6 we follow a slightly different approach when working with the equations of general relativistic hydrodynamics, because we also discuss the possibility of solving general relativistic problems using a spacetime approach which more naturally suits the nature of relativistic equations. We also relax the hypothesis that the solution is BV, which might be too restrictive for systems of equations. Here, instead, we focus on the classical approach and concentrate on the case of scalar equations. Let $v \in C_0^1(\Omega)$, multiply (3.79) by v and integrate over Ω , we then find

$$\sum_{j=1}^N \left[\int_{\Omega_j} \partial_t u v \, dx - \int_{\Omega_j} \mathbf{f}(u) \cdot \nabla v \, dx \right] = - \sum_{j=1}^N \langle \mathcal{F} \cdot \mathbf{v}, v \rangle_{\partial\Omega_j} \quad (3.81)$$

where $\mathcal{F} \cdot \mathbf{v}$ is the normal trace of $\mathbf{f}(u)$ on the boundary, *i.e.*, it is a distribution whose value on any test function, v , is given by

$$\langle \mathcal{F} \cdot \mathbf{v}, v \rangle_{\partial\Omega_j} := \langle \operatorname{div} \mathbf{f}, v \rangle_{\Omega_j} + \int_{\Omega_j} \mathbf{f}(u) \cdot \nabla v \, dx, \quad (3.82)$$

where $\operatorname{div} \mathbf{f}$ is the distributional divergence of $\mathbf{f}$. In the case in which $\mathbf{f}$ is smooth,

$$\langle \mathcal{F} \cdot \mathbf{v}, v \rangle_{\partial\Omega_j} = \int_{\partial\Omega_j} \mathbf{f} \cdot \mathbf{v} v \, dx, \quad (3.83)$$

where $\mathbf{v}$ is the out-going unit-normal to Ω_j . The weak formulation of (3.79) then consists of finding

$$u \in \operatorname{BV}(\Omega), \quad \text{s.t., (3.81) holds} \quad \forall v \in C_0^1(\Omega). \quad (3.84)$$

The DG method works by projecting (3.84) on a finite-dimensional subspace of $\operatorname{BV}(\Omega)$. In particular we consider piecewise polynomial functions on $\mathcal{T}_N$:

$$V_N = \left\{ v \in \operatorname{BV}(\Omega) : v \circ \varphi_j \in \mathbb{P}_D(T), \, j = 1, \dots, N \right\}, \quad (3.85)$$

where $\mathbb{P}_D$ is the space of polynomials of degree D . Notice that we do not require the continuity of functions in V_N between different elements, hence the reason why these schemes are called *discontinuous* Galerkin. The DG formulation then consists of finding

$$u(t) \in C^1(\mathbb{R}_+; V_N), \quad \text{s.t., (3.81) holds} \quad \forall v \in V_N. \quad (3.86)$$

¹³This ensures the existence of the normal trace of $\mathbf{f}$ (see below equation 3.82).

Notice that so far we restricted the test function to be in V_N , *i.e.*, the test function is in the same functional space as the numerical solution u . The family of schemes where the test function is in the same space as the solution are called *Galerkin methods*. Secondly, since both u and v now live in a finite-dimensional space, a finite number of conditions of the form (3.81) is sufficient to determine all the degrees of freedom of u , *i.e.*, by writing down the condition (3.81) for a sufficiently large number of linearly-independent test functions v , one obtains a set of ODEs for the d.o.f. of u .

Before going into details of how to explicitly construct the scheme, we need to address two important points concerning the weak formulation (3.86).

The first one is that (3.81) is not really well defined if the test function is not smooth. This means that one has to be careful in the interpretation of the DG formulation: for the boundary terms of (3.81) to be well defined one should think of the v appearing in the boundary term as a smooth, $C_0^1(\Omega)$, extension of a function in V_N , done in such a way as to match the one-sided limit of v at $\partial\Omega_j$ from the interior of Ω_j , *i.e.*, we need to set

$$\langle \mathcal{F} \cdot \mathbf{v}, v \rangle_{\partial\Omega_j} := \langle \mathcal{F} \cdot \mathbf{v}, v_j \rangle_{\partial\Omega_j}, \quad (3.87)$$

where $v_j \in C_0^1(\Omega)$ and $v_j|_{\Omega_j} = v$.

The second point is that the normal trace actually has a very simple interpretation. To see that, let us consider the case in which $v = \chi_i$, so that (3.81) simplifies to

$$\partial_t \int_{\Omega_i} u \, dx = -\langle \mathcal{F} \cdot \mathbf{v}, v \rangle_{\partial\Omega_i}, \quad (3.88)$$

which is easily recognizable as being the FV method. In other words we can take the normal trace to be simply the flux as computed using an approximated Riemann solver.

Finally we turn to the detailed construction of the DG scheme. In order to simplify the notation we work in the one-dimensional case, so that (3.81) reads

$$\begin{aligned} \sum_{j=1}^N \left[\int_{x_{j-1/2}}^{x_{j+1/2}} \partial_t u v \, dx - \int_{x_{j-1/2}}^{x_{j+1/2}} f(u) \partial_x v \, dx \right] = \\ \sum_{j=1}^N \left[F^{j-1/2} v(x_{j-1/2}) - F^{j+1/2} v(x_{j+1/2}) \right], \end{aligned} \quad (3.89)$$

where F is now the numerical flux. We can expand both $u(t, \cdot)$ and $f[u(t, \cdot)]$ on a polynomial basis as

$$u(t, x) = \sum_{i=0}^D u_i^j(t) l_i^j(x), \quad f(t, x) = \sum_{i=0}^D f_i^j(t) l_i^j(x), \quad x \in [x_{j-1/2}, x_{j+1/2}], \quad (3.90)$$

where $l_i^j(x)$ is some polynomial basis over $[x_{j-1/2}, x_{j+1/2}]$, typically, but not necessarily, chosen to be composed of orthonormal polynomials. If we now choose $v(x) = l_k^j(x) \chi_j(x)$, $k = 0, \dots, D$, we obtain a set of evolution equations for the

expansion coefficients $u_i^j(x)$:

$$\begin{aligned} \sum_{i=0}^D \left[\int_{x_{j-1/2}}^{x_{j+1/2}} l_i^j(x) l_k^j(x) dx \right] \frac{du_i^j(t)}{dt} \\ - \sum_{i=0}^D \left[\int_{x_{j-1/2}}^{x_{j+1/2}} l_i^j(x) \partial_x l_k^j(x) dx \right] f_i^j = \\ F^{j-1/2} l_k^j(x_{j-1/2}) - F^{j+1/2} l_k^j(x_{j-1/2}), \end{aligned} \quad (3.91)$$

which we can write in more compact form as

$$\mathbf{M}^j \frac{d\mathbf{u}^j}{dt} - \mathbf{D}^j \mathbf{f}^j = \mathbf{F}^{j-1/2} - \mathbf{F}^{j+1/2}, \quad \forall j = 1, \dots, N, \quad (3.92)$$

where we have collected all the expansion coefficients of u in the vectors $\mathbf{u}^j$ and introduced the mass-matrix

$$(\mathbf{M}^j)_{ki} := \int_{x_{j-1/2}}^{x_{j+1/2}} l_i^j(x) l_k^j(x) dx, \quad (3.93)$$

the discrete co-differential matrix

$$(\mathbf{D}^j)_{ki} := \int_{x_{j-1/2}}^{x_{j+1/2}} l_i^j(x) \partial_x l_k^j(x) dx \quad (3.94)$$

and the flux vectors

$$(\mathbf{F}^{j-1/2})_i := \sum_k \delta_{ik} F^{j-1/2} l_k^j(x_{j-1/2}). \quad (3.95)$$

Before concluding it is useful to make some remarks.

First of all, if an orthonormal basis is chosen and the integrals are computed with enough accuracy, then the mass-matrix is diagonal. Otherwise the mass matrix has to be inverted in order to evolve u in time. In any case it is important to notice that the mass matrix is local, *i.e.*, defined only on the single element. For this reason its inverse is relatively inexpensive and can also be pre-stored. As an alternative the mass-matrix can be “diagonalized” using the mass-lumping technique, *e.g.*, [75].

Secondly, $f(u)$ is in general a non-linear function. This means that, while expanding the flux on a polynomial basis, we might introduce an aliasing error in the scheme, unless sufficiently many coefficients are used in the expansion and an appropriate quadrature formula is used to compute $\mathbf{D}^j$. This aliasing error is often unavoidable in practical applications and can spoil the non-linear stability of the scheme, but it is easily suppressed with the use of weak filtering as long as the solution is smooth [168]. Obviously, even in the linear case, if u has a jump discontinuity (*i.e.*, a shock), then the aliasing error is unavoidable and in order to achieve the stability of the scheme some more effort might be required. In this case the classical approach is to use the minmod limiter as a non-linear filter to detect discontinuities and flatten the profile of the solution within the troubled elements [93]. More on this in Chapter 6.

Part II

High-Order Methods for Relativistic Hydrodynamics

Chapter 4

Finite-Differencing Methods: Flat Spacetimes

4.1 Introduction

The use of flux-conservative finite-difference HRSC schemes is probably the easiest way of increasing the (formal) order of accuracy of the current generation of numerical codes: finite-difference schemes are much cheaper than high-order finite-volume codes since they do not require the solution of multiple Riemann problems at the interface between different regions [296, 297] and they are free from the complicated averaging and de-averaging procedures of high-order finite-volume codes, see, *e.g.*, [318].

Here we present a new code, the Templated-Hydrodynamics Code (THC), developed using the Cactus framework [150], that follows this approach. THC employs a state-of-the-art flux-vector splitting scheme: it uses up to seventh-order reconstruction in characteristic fields and the Roe flux split with a novel entropy-fix prescription. The “templated” aspect reflects the fact that the code design is based on a modern C++ paradigm called template metaprogramming, in which part of the code is generated at compile time. Using this particular programming technique it is possible to construct object-oriented, highly modular, codes without the extra computational costs associated with classical polymorphism, because, in the templated case, polymorphism is resolved at compile time allowing the compiler to inline all the relevant function calls, see *e.g.*, Yang [335]. Among the different reconstruction schemes that we implemented are the classical monotonicity-preserving (MP) MP5 scheme [310, 226], the weighted essentially non oscillatory (WENO) schemes WENO5 and WENO7 [202, 176, 295] and two bandwidth-optimized WENO schemes: WENO3B and WENO4B [211, 316], designed for direct simulations of compressible turbulence (we recall that the number associated to the different methods indicates the putative order of accuracy).

In this chapter we give the details of the algorithms used in THC, presenting a systematic comparison between the results obtained using the above mentioned reconstruction schemes, with emphasis on the application of these schemes to direct simulations of relativistic turbulence. To our knowledge this is the first time that such a comparison has been done in the relativistic case.

The rest of this chapter is organised as follows. In Section 4.2 we present the THC code in more detail: we discuss the numerical algorithms it uses and recall the equations of Newtonian and special-relativistic hydrodynamics. The results obtained with our code in a representative number of test cases for the Newtonian and special relativistic hydrodynamics are presented in Section 4.3. In Section 4.4 we present the application of our code to the study of the linear and non-linear development of the relativistic Kelvin-Helmholtz instability (KHI) in three dimensions as a nontrivial application of our code and a stringent test of its accuracy. In Section 4.5 we present the application of our code to the study of the statistical properties of driven relativistic turbulence in a hot, optically thick, pair-plasma, as another example of the potential of our code. Finally Section 4.6 is dedicated to the summary and the conclusions.

4.2 The THCCode

We here briefly outline the numerical infrastructure adopted by our templated-hydrodynamics code and report the formulations of the equations of Newtonian and special-relativistic hydrodynamics we actually solve.

4.2.1 Newtonian Hydrodynamics

The equations of classical (*i.e.*, Newtonian) hydrodynamics describe the conservation of mass, momentum and energy for a perfect fluid. They can be written in the form (3.62) with primitive variables,

$$\mathbf{u} = [\rho, \vec{v}, \epsilon], \quad (4.1)$$

where ρ is the density, v^i the velocity and ϵ the specific internal energy. The conserved variables are

$$\mathbf{F}^0(\mathbf{u}) = [\rho, \rho\vec{v}, E] = \left[\rho, \rho\vec{v}, \rho\left(\frac{1}{2}v^2 + \epsilon\right) \right], \quad (4.2)$$

the sources are zero and the fluxes are

$$\mathbf{F}^i(\mathbf{u}) = \left[\rho v^i, \rho v^i \vec{v} + p \delta^i, v^i(E + p) \right], \quad (4.3)$$

where p is the pressure and $[\delta^i]^j = \delta^{ij}$ is the Kronecker symbol. The system of equations is then closed by an equation of state $p = p(\rho, \epsilon)$ and we adopt that of an ideal fluid (or Gamma law)

$$p = (\Gamma - 1)\rho\epsilon, \quad (4.4)$$

where Γ is the adiabatic index of the fluid.

The Jacobians and their spectral decomposition for the equations of Newtonian hydrodynamics and for a generic equation of state, can be found, for example, in [189].

4.2.2 Special-relativistic hydrodynamics

In the case of the relativistic-hydrodynamic equations it is convenient to work using a system of units in which $c = 1$ and we adopt the standard convention

for the summation over repeated indices. We consider a perfect fluid having 4-velocity $\vec{u} = (W, W\vec{v})$, with $W := (1 - v^i v_i)^{-1/2}$ being the Lorentz factor. Then the rest-mass current 4-vector is given by

$$\vec{J} = \rho \vec{u}, \quad (4.5)$$

where ρ is here the rest-mass density. The stress-energy tensor is given by

$$T = \rho h \vec{u} \otimes \underline{u} + p \underline{g}, \quad (4.6)$$

where $h = 1 + \epsilon + p/\rho$ is the specific enthalpy and $\underline{g}$ is the spacetime metric, which we take to be that of flat spacetime, *i.e.*, where the only nonzero components are the diagonal ones and given by $g_{\mu\nu} = (-1, 1, 1, 1)$. Because our main interest with THC is in determining the statistical properties of special-relativistic turbulence and of unveiling novel and non-classical features we will consider the fluid not to affect the spacetime geometry, which we will consider to be that of a flat spacetime at all times.

Conservation of rest-mass, momentum and energy are expressed by the vanishing of the 4-divergence of $\vec{J}$ and T

$$\nabla \cdot \vec{J} = 0, \quad \nabla \cdot T = 0, \quad (4.7)$$

where ∇ is the covariant derivative associated with $\underline{g}$. Also in this case, the relativistic-hydrodynamic equations (4.7) can be cast in the form (3.62) with primitive variables

$$\underline{u} = [\rho, \vec{v}, \epsilon]. \quad (4.8)$$

The conservative variables are

$$\underline{F}^0(\underline{u}) = [D, \underline{s}, \tau], \quad (4.9)$$

where

$$D := \rho W, \quad \underline{s} := \rho h W^2 \vec{v}, \quad \tau := \rho h W^2 - p - \rho W, \quad (4.10)$$

and the fluxes are given by

$$\underline{F}^i(\underline{u}) = [D v^i, \underline{s} v^i + p \delta^i, s^i - D v^i]. \quad (4.11)$$

Finally, because we are working in special relativity, sources in the system (4.7) are zero and the equations are closed by an equation of state, which we take again to be the ideal-fluid equation of state (4.4).

An important difference between the Newtonian and the (special-)relativistic hydrodynamic equations is that in the latter case there is no simple analytic expression for the inverse transformation $\underline{F}^0 \rightarrow \underline{u}$, leading to the primitive variables from the conserved ones. For this reason, we use a numerical root-finding procedure to recover the primitive variables from the conservatives. In particular we follow the strategy of [179, 180]. The primitive variables can be easily written as a function of the conservative variables once a value for the enthalpy, $\tilde{h}$, is assumed. At the same time the enthalpy can be expressed as a function of the rest-mass density and of the internal energy using the equation of state. Thus we can construct the function

$$g(\tilde{h}) = h_{\text{EOS}}(\rho(\tilde{h}), \epsilon(\tilde{h})) - \tilde{h}, \quad (4.12)$$




Figure 4.1: Density (left panels), velocity (middle panels) and pressure (right panels) for the Newtonian strong shock test. We show both the analytic solution (solid line) and the numerical solutions obtained with different numerical schemes (dots). The resolution is $\Delta^1 = 1/100$ and the timestep is $\Delta^0 = 0.002$ for all the runs. Note that the numerical solution is not down-sampled and corresponds to a rather coarse resolution.

and use a one-dimensional root-finding procedure to find a self-consistent value of the enthalpy, since $g(\tilde{h}) = 0$ if and only if $\tilde{h}$ is the physical enthalpy¹. In particular, the root-finding algorithm uses a combination of the Newton-Raphson method with the regula-falsi and the bisection schemes: we check the Newton-Raphson method for convergence and, in case of failure, we switch to the regula-falsi. The bisection scheme is used as a “fail-safe” root-finder in situations where the regula-falsi is converging too slowly: this is necessary only when the values of the conservative variables are close to an unphysical region. In the large majority of cases, the Newton-Raphson method usually converges to the required level of accuracy with only three iterations on average.

The Jacobians and their spectral decomposition for the equations of special-relativistic hydrodynamics and for a generic equation of state, can be found in [116].

4.3 Numerical tests

This Section is dedicated to the presentation of some of the results obtained with THC in a series of tests in Newtonian and special-relativistic hydrodynamics.

4.3.1 Newtonian hydrodynamics

We begin with a series of tests in classical hydrodynamics, before switching to the special-relativistic case.

Strong shock

The first test is a classical one-dimensional shock tube: the initial data describes two regions filled with a $\Gamma = 5/3$ ideal-fluid in equilibrium separated by a membrane. At $t = 0$ the membrane is removed and the two regions start to interact. The initial left and right states are

$$(\rho_L, v_L, p_L) = (10, 0, 10), \quad (\rho_R, v_R, p_R) = (1, 0, 10^{-5}), \quad (4.13)$$

for $t > 0$ the analytic solution consists in a right-going shock wave, followed by a right-going contact discontinuity and a transonic left-going rarefaction wave.

In Figure 4.1 we show with a solid line the analytic solution, while filled circles are used to represent the solution obtained at time $t = 0.2$ with the different numerical schemes. The grid resolution is $\Delta^1 = 1/100$ and the timestep $\Delta^0 = 0.002$ for all the runs. Note that the numerical solution is not down-sampled and hence it corresponds to a genuinely coarse resolution. We do not show the results obtained with the bandwidth-optimised schemes since they are basically indistinguishable from the ones obtained using their traditional counterpart, *i.e.*, the solution obtained with WENO3B is basically on top of the one obtained with WENO5 and the solution obtained with WENO4B is identical, at the plot scale, with the one obtained with WENO7. Since the results obtained with the bandwidth-optimized schemes are found to be very close to

¹We are tacitly assuming the enthalpy to be a convex function of ρ and ϵ . This is satisfied by the EOSs that we consider here, but does not hold for general EOSs. For this reason in the general-relativistic extension of THC, presented in Chapter 5, we make use of a different procedure which has been kindly provided us by Wolfgang Kastaun and Filippo Galeazzi [136].




Figure 4.2: The same as in Figure 4.1, but for the Newtonian blast-wave test. The resolution is $\Delta^1 = 1/100$ and the timestep is $\Delta^0 = 0.0001$ for all the runs.

the ones obtained with the standard WENO schemes in all the shock-tube tests that we performed, we will show only the numerical solutions obtained with WENO5, WENO7 and MP5 in all the cases.

Even at this fairly low resolution, all the schemes are able to capture well both the shock wave and the rarefaction wave, showing the good behaviour of our entropy fix. The shock wave is captured within three gridpoints. There are no appreciable post-shock oscillations in the solution obtained with the MP5 reconstruction, while small oscillations are present in the velocity field with WENO5 and, in particular, with WENO7.

We note that the contact discontinuity is resolved, but not without oscillations. However, we should bear in mind that, although the density contrast is small across the contact discontinuity, this test is a severe one due to the high Mach number of the shock wave, *i.e.*, $\mathcal{M}_s \approx 360$.

Blast wave

The second test is similar to the first one, but results in a much larger density contrast at the contact discontinuity. The initial data is given by

$$(\rho_L, v_L, p_L) = (10^{-3}, 0, 1), \quad (\rho_R, v_R, p_R) = (10^{-3}, 0, 10^{-5}), \quad (4.14)$$

and the adiabatic index is still $\Gamma = 5/3$. Also in this case, the analytic solution consists of a right-going shock wave, followed by a contact discontinuity and a left-going rarefaction wave.

In Figure 4.2 we show the exact solution at time $t = 0.015$ (solid line), as well as the numerical solution (filled circles) obtained using different numerical schemes. The grid resolution is $\Delta^1 = 0.01$ ($N = 100$) and the timestep is




Figure 4.3: The same as in Figure 4.1, but for the Newtonian rotated Sod test. The resolution is $\Delta^i = 1/200$ and the timestep is $\Delta^0 = 0.000625$ for all the runs.

$\Delta^0 = 0.0001$ for all the runs. The Mach number of the shock wave is $\mathcal{M}_s \approx 190$ and in this case all our schemes are free from major oscillations. The shock wave is resolved within two grid points, while the contact discontinuity is smeared over 5-6 grid points.

The MP5 scheme is able to properly capture the constant state between the shock wave and the contact discontinuity, while the WENO schemes result in more “rounded” solutions. In particular both WENO5 and WENO7 overestimated the density contrast.

Rotated Sod test

A genuinely three-dimensional shock-tube test in Newtonian hydrodynamics is offered by the classical rotated Sod test [305]. In this case, the adiabatic index is $\Gamma = 1.4$ and the right and left states are

$$(\rho_L, v_L, p_L) = (1, 0, 1), \quad (\rho_R, v_R, p_R) = (0.125, 0, 0.1), \quad (4.15)$$

the initial, 1D data, is rotated by 45° about the z and y axes to yield a shock wave that is diagonal to the principal axes of the grid. The analytic solution consists of a left-going rarefaction wave and a right-going shock wave separated by a right-going contact discontinuity.

In Figure 4.3 we show the analytic solution in the diagonal direction (solid line), as well as the numerical solutions (filled circles), at time $t = 0.2$. The spatial resolution is $\Delta^i = 1/200$ for $i = 1, 2, 3$ and the timestep is $\Delta^0 = 0.000625$ for all the runs. All the schemes are able to properly capture the main features of the solution: the discontinuities are captured within 1 or 2 gridpoints and both WENO5 and MP5 are able to capture the plateau in the velocity. The

solution obtained with the WENO7 scheme presents oscillations in the velocity after the shock wave and in the density and pressure fields downstream of the contact discontinuity. Similar oscillations are also observed with WENO5 and MP5 when the resolution is halved.

Overall, these tests demonstrate the accuracy of the dimensionally unsplit approach that we use to treat the multi-dimensional case.

Double Mach reflection test

A final test in Newtonian hydrodynamics is the double-Mach reflection test proposed by [334]. The initial data describes a right-going Mach 10 shock wave making a 60° angle with the computational grid and intersecting the x -axis at $x = 1/6$. A perfectly reflecting wall is placed along the x -axis in the $x > 1/6$ region, while the values along the other regions of the boundary are set to the pre and post-shocked values on the left/right side of the shock wave [see [334] for details on the boundary conditions and the initial data].

We considered the numerical solutions obtained with WENO5, WENO7, MP5, WENO3B and WENO4B in the computational domain $0 \leq x \leq 4$, $0 \leq y \leq 1$. For each of these schemes we performed five runs with resolutions 120×30 , 240×60 , 480×120 , 960×240 and 1920×480 . In Figure 4.4 we show isocontours (with an equal spacing of 0.15 between 0 and 3) of the rest-mass density obtained at the highest resolution at time $t = 0.2$ by the different numerical schemes. In order to ease the comparison with the results reported in [334], we show only the region $x < 3$.

Our code performs well in this test across all the reconstruction schemes that we tried. All the discontinuities, including the contact discontinuity ahead of the jet region, are well resolved within a few grid regions, even at the lowest resolution. As we increase the number of gridpoints we do not observe any sign of the carbuncle phenomenon and this seems to be an indication that our algorithm is able to introduce enough numerical dissipation to avoid the odd-even decoupling.

At high enough resolution it is possible to observe the development of instabilities upstream from the reflected shock wave, generating small scale structures along the contact discontinuity. The ability of the different codes to resolve these structures can be used as an indication of their numerical viscosity. In particular we can see how the bandwidth-optimized schemes gain with respect to their “standard” counterparts. For instance, WENO3B, which uses the same stencil as WENO5, yields a solution which is in qualitative agreement with the one obtained using WENO5 at twice the spatial resolution. The same is also true if we compare WENO4B and WENO7.

All things considered, we find that the best performance is given by the MP5 scheme. This algorithm has a computational cost which is comparable with the one of WENO5, as it uses the same stencil. Nevertheless the solution obtained with MP5 is very close to the one obtained with the optimized WENO4B scheme, which, in 3D, is almost twice as expensive as WENO5.

4.3.2 Special-relativistic hydrodynamics

In this Section we will present the results obtained in a series of tests in special-relativistic hydrodynamics.




Figure 4.4: Isocontours of the rest-mass density at time $t = 0.2$ for the Newtonian double Mach reflection test, obtained with different schemes. We show 20 contours levels equally spaced between 0 and 3. The resolution is $\Delta^i = 1/480$ and the timestep is $\Delta^0 = 1/40000$ for all the runs.




Figure 4.5: Analytic solution and numerical solution computed with WENO3B with 50 gridpoints for the case of a smooth wave in an adiabatic relativistic fluid. The figure shows both the initial data (top panel) and the solution at time $t = 0.8$ (bottom panel). The solid line represent the analytic solution, while the filled circles represent the numerical one.

Adiabatic smooth flow

The first test we present is designed to show the accuracy of the code in the case of smooth solutions and hence to measure rigorously the convergence order of THC for the different schemes implemented. This test is very similar to the one discussed by [339].

We consider a one-dimensional, large-amplitude, smooth, wave propagating in an isentropic fluid, with polytropic equation of state,

$$p = K\rho^\Gamma, \quad (4.16)$$

where $K = 100$ and $\Gamma = 5/3$. The rest-mass density at $t = 0$ is given by

$$\rho_0(x) = \begin{cases} 1 + \exp\left[-1/(1 - x^2/L^2)\right], & \text{if } |x| < L; \\ 1, & \text{elsewhere;} \end{cases} \quad (4.17)$$

where, differently from [339], the initial profile of the rest-mass density is chosen to be C^∞ but is not analytic. We have found this choice important to obtain the correct convergence order at very high resolutions. Indeed, when adopting the same profile as in [339], we found that the jump discontinuity in the fifth derivative of the initial data prevents the WENO7 scheme from achieving a convergence order larger than five. Besides this small difference, our initial data is basically identical to the one used by [339]. The initial velocity is set up assuming that one of the two Riemann invariants [18],

$$J_- = \frac{1}{2} \ln\left(\frac{1+v}{1-v}\right) - \frac{1}{\sqrt{\Gamma-1}} \ln\left(\frac{\sqrt{\Gamma-1} + c_s}{\sqrt{\Gamma-1} - c_s}\right), \quad (4.18)$$

where c_s is the sound speed, is constant in the whole region and so that $v = 0$ if $|x| \geq L$. The other Riemann invariant,

$$J_+ = \frac{1}{2} \ln\left(\frac{1+v}{1-v}\right) + \frac{1}{\sqrt{\Gamma}-1} \ln\left(\frac{\sqrt{\Gamma}-1+c_s}{\sqrt{\Gamma}-1-c_s}\right), \quad (4.19)$$

is not constant, so that the initial data describes a right-going wave.

The analytic solution can be easily found in Lagrangian coordinates, up to the caustic point, using the method of characteristics [18], while its calculation in Eulerian coordinates involves the solution of a transcendental equation. For the purposes of our calculation, a good-enough approximation of the exact solution was obtained by computing it on a very fine Lagrangian grid (we have used 100,000 gridpoints), and interpolated on the Eulerian grid using a cubic spline interpolation. This solution is then used as the reference solution against which the numerical solutions obtained with THC have been compared.

In our test we set $L = 0.3$ and use a computational grid in the region $-0.4 \leq x \leq 2$, evolving the initial data up to time $t = 1.6$, which is approximately the time when a caustic appears and the solution becomes discontinuous. We performed this test using different schemes and different resolutions and we measured the L^1 -norm of the error against the reference solution. Differently from all the other tests, instead of the third-order SSP-RK scheme, we adopt here a fourth-order RK time integrator. The reason for this choice is that, since this test involves a smooth solution, a SSP time-integrator is not required. Moreover, as we will show in the following, the use of a more accurate time integrator enabled us to measure a convergence order of the spatial discretization which is not spoiled by the order in time, the only exception being the highest resolution run done with WENO7. The CFL factor is set to be $C \approx 0.2$.

The initial rest-mass density profile, as well as the solution at time $t = 0.8$, which we take as a reference time for the measure of the error, is shown in Figure 4.5, together with the solution obtained with WENO3B using 50 gridpoints. As can be seen from the Figure, the solution at the considered time is still smooth and our scheme is able to properly capture it very well even at this very coarse resolution.

The outline of the obtained results is shown in Table 4.1. The L^1 -norm of the error, as measured at the reference time, is shown in Figure 4.6. The first thing one notices is that all our schemes approach the expected convergence order only asymptotically, at very high resolution. This is more evident by comparing Figure 4.6 with Figure 4.7 where we show the results obtained by our code in a similar test, but in the case of the scalar advection equation. The reason for this behaviour is in the “kinks” ahead and behind the pulse, where the numerical error is largest. These regions are “misinterpreted” as discontinuities by the shock-detection part of our schemes, unless they are resolved with enough gridpoints.

The best performing scheme in this test is the MP5 one: at low resolution it yields a smaller error than the seventh-order WENO scheme, which, in turn, is able to attain an higher convergence order only at a resolution which is unfeasible in any practical multidimensional application. The bandwidth optimized schemes present a somewhat errant behaviour in their convergence order, with WENO4B, not even showing a monotone trend in the L^1 -norm of the error as a function of the number of gridpoints. We do not presently have an explanation

Table 4.1: Adiabatic smooth flow test. For each numerical scheme we quote the L^1 norm of the error and the relative convergence order as measured against the analytic solution at time $t = 0.8$.

Scheme	N	L^1 -Error	Convergence
WENO5	50	0.00330632	–
	100	0.00123249	1.42
	200	0.000294511	2.06
	400	4.4764e-05	2.72
	800	5.57555e-06	3.01
	1600	3.90326e-07	3.84
	3200	2.09857e-08	4.21
	6400	7.68527e-10	4.77
WENO3B	50	0.00309225	–
	100	0.00118655	1.38
	200	0.000376258	1.66
	400	8.28337e-05	2.18
	800	2.86678e-05	1.53
	1600	5.53009e-06	2.37
	3200	3.15153e-08	7.46
	6400	3.5237e-09	3.16
WENO7	50	0.00194469	–
	100	0.000751681	1.37
	200	0.000123279	2.61
	400	9.62907e-06	3.68
	800	7.32092e-07	3.62
	1600	3.00016e-08	4.61
	3200	4.09176e-10	6.19
	6400	7.03027e-12	5.86
WENO4B	0	0.00191339	–
	100	0.00072531	1.40
	200	0.000188846	1.94
	400	1.98584e-05	3.25
	800	2.27458e-06	3.13
	1600	2.33667e-06	-0.04
	3200	5.38905e-09	8.76
	6400	3.55794e-10	3.92
MP5	50	0.00157656	–
	100	0.000554439	1.51
	200	9.89946e-05	2.49
	400	1.3236e-05	2.90
	800	9.15752e-07	3.85
	1600	6.27508e-08	3.87
	3200	2.74506e-09	4.51
	6400	9.15411e-11	4.91




Figure 4.6: L^1 -norm of the error for different resolutions and for different numerical schemes for the case of a smooth simple wave in an adiabatic relativistic fluid




Figure 4.7: L^1 -norm of the error for different resolutions and for different numerical schemes for the linear advection of a sinusoidal wave over ten characteristic times.




Figure 4.8: Order of convergence, as measured using the two highest resolution runs, as a function of the time to the caustic for different numerical schemes for the smooth simple wave test.

we find sufficiently convincing for the behaviour shown.

Finally it is interesting to study the convergence order of the different numerical schemes as a function of $t - t_c$, where t_c is the time when the caustic is formed. The convergence order is computed using the error with respect to the exact solution of the two highest resolution runs and is shown in Figure 4.8. At these very high resolutions, all the schemes appear to be converging at their nominal convergence order away from the caustic, apart from WENO7 that appears to have already reached saturation. Its order of convergence increases up to almost seven close to the caustic, at time $t - t_c \approx -0.4$, because there the error is dominated by the presence of a steeper gradient ahead of the pulse. Before that, the error from the spatial discretization is probably very close to the one from the time discretization (we recall that we use a fourth-order RK integrator), thus degrading the convergence order.

Indeed, as the time of shock formation approaches, the order of the schemes decreases slowly to the first-order expected in the case of discontinuities. The bandwidth optimized schemes and, in particular, WENO4B show, again, an erratic behaviour of their convergence order. The reasons for this behaviour are most probably the same ones behind the similar behaviour observed while looking at the error as a function of the number of gridpoints.

Blast wave

When contrasted with its Newtonian counterpart, one of the most striking features of relativistic hydrodynamics is that relativistic fluids can exhibit much stronger shock waves. For this reason, it is important to assess the capability of the code to handle very strong shocks. As a first example we consider a




Figure 4.9: The same as in Figure 4.1, but for the relativistic blast-wave test. The resolution is $\Delta^1 = 1/400$ and the timestep is $\Delta^0 = 0.0005$ for all the runs.

one-dimensional shock-tube where the initial data is given by,

$$(\rho_L, v_L, p_L) = (10^{-3}, 0, 1), \quad (\rho_R, v_R, p_R) = (10^{-3}, 0, 10^{-5}), \quad (4.20)$$

and the adiabatic index is still $\Gamma = 5/3$. This initial data is formally identical to the one used in Section 4.3.1 in the Newtonian case. The analytic solution consists again of a transonic left-going rarefaction wave and a right-going shock wave separated by a right going contact discontinuity. The shock wave has a relativistic Mach number $M_s \approx 50$ [85, 185].

The results obtained with the different numerical schemes are reported in Figure 4.9, where we show the analytic solution (solid line) as well as the numerical ones (filled circles) obtained with WENO5, WENO7 and MP5, at time $t = 0.4$. As in the Newtonian case, we do not show the results obtained with the bandwidth optimized schemes as they are basically identical to the ones of their traditional variant. The spatial resolution that we use is $\Delta^1 = 1/400$ and the CFL factor is $C = 1/5$ in all cases.

The CFL factor in this test is basically constrained by the MP5 scheme: while the WENO schemes appear to be robust and stable up to CFL factor of $C \approx 2/5$, the MP5 algorithm produces large oscillations and yields non-physical values in the conservative variables unless a smaller CFL factor is used. We point out that for this particular test the necessity of using small timesteps to avoid the creation of unphysical states has been observed also when using other numerical schemes. For instance, when using the finite-volume code *Whisky* [28, 29], with the HLLC approximate Riemann solver [see, *e.g.*, [323]] and the PPM reconstruction [96], the maximum allowed CFL factor was also found to be $C \approx 2/5$. It is also known that, for the monotonicity-preserving property to hold for the MP5 scheme, the timestep must be subject to an additional




Figure 4.10: The same as in Figure 4.1, but for the relativistic shock-heating test. The resolution is $\Delta^1 = 1/100$ and the timestep is $\Delta^0 = 0.001$ for all the runs.

constraint which is distinct from the standard CFL condition [310]. Yet, it is somewhat surprising to observe that MP5 requires a timestep which is smaller by a factor of order two with respect to the WENO schemes. Furthermore, this property does not seem to be a peculiarity of this specific problem, since we observed a similar behaviour also for the other tests that we performed.

[339] report that the use of MP5 reconstruction results in large numerical oscillations in their special relativistic code. In our code, we see a similar behaviour unless we use a timestep which is about half of the one considered “safe” for the WENO scheme. If the timestep is sufficiently small, on the other hand, the MP5 algorithm results in very accurate solutions, as in the Newtonian case. In particular in Figure 4.9 we can see the results obtained with this particular test. Our code, with MP5, is able to capture the shock wave within three gridpoints while the contact discontinuity is spread across six points. WENO7 yields a solution of similar quality, while WENO5 is slightly more diffusive.

The capability of a numerical code to capture the density contrast for this particular test is a classical benchmark for relativistic hydrodynamics codes. Again, using as reference the *Whisky* code, the use of the HLLE solver with PPM reconstruction and with artificial compression, leads to a maximum density which is 71% of the analytic solution. At the same resolution, our THC using a WENO5, WENO7 and MP5 reconstruction is able to attain a maximum density which is respectively 78%, 90% and 91% of the analytical value. These results are in very good agreement with the ones reported by [339], who measured a relative value of 72% and 79% for their implementation of PPM and WENO5 schemes.

Shock-heating

An even more striking example of how relativistic effects can enhance the density contrasts in shock waves is given by the classical shock-heating test. In this case, the initial data is given by

$$(\rho_L, v_L, p_L) = (10^{-3}, v, 10^3), \quad (\rho_R, v_R, p_R) = (10^{-3}, -v, 10^3), \quad (4.21)$$

the polytropic index is $\Gamma = 4/3$ and

$$v = \sqrt{1 - \frac{1}{W^2}} \simeq 0.99999949, \quad (4.22)$$

where the Lorentz factor is set to be $W = 1000$. In this case, the analytic solution is represented by two shocks whose collision compresses the fluid, converting its kinetic energy into thermal energy, that is, through “shock heating”.

In Newtonian hydrodynamics the maximum compression ratio can be computed as

$$\sigma_{\text{newt}} = \frac{\Gamma + 1}{\Gamma - 1} = 7, \quad (4.23)$$

for all values of v , while it is easy to show [209] that in the relativistic case the compression ratio is

$$\sigma = \frac{\Gamma + 1}{\Gamma - 1} + \frac{\Gamma}{\Gamma - 1}(W - 1) \simeq 4003, \quad (4.24)$$

thus about three orders of magnitude larger for the same adiabatic index and growing linearly with the Lorentz factor.

The exact solution at time $t = 0.4$, as well as the numerical solutions obtained with our code, are shown in Figure 4.10. As can be seen from the Figure, THC is able to handle very well this large compression ratio. The WENO5 and WENO7 solutions are affected by some small wall-heating effect [238, 276], resulting in a slight under-density near $x = 0$ and some numerical oscillations in the pressure. The MP5 scheme, on the other hand, yields a solution which is essentially free from oscillations in the pressure and much less affected by the wall-heating effect in the density variable, although at the cost of a smaller timestep, as discussed in the previous Section.

Transverse shock

Another peculiar difficulty of relativistic hydrodynamics and without a Newtonian counterpart, is that the equations for the momentum in the different directions are coupled together by the Lorentz factor: even in one-dimensional problems the application of a transverse velocity can change completely the solution. This feature was first pointed out by [257] and [272], and then used by [273] and [15] to discover a new physical effect [see also [224, 339] for a description of the numerical consequences of this property].

To explore the flow dynamics in this case, we consider the same initial data as for the blast-wave test [see Section 4.3.2], with the only difference being that we add a transverse velocity to the initial data, *i.e.*,

$$v_L^t = 0, \quad v_R^t = 0.99. \quad (4.25)$$




Figure 4.11: The same as in Figure 4.1, but for the relativistic blast-wave test. The resolution is $\Delta^1 = 1/100$ and the timestep is $\Delta^0 = 0.002$ for all the runs.

As discussed in [339] this is not a very challenging test, since the velocity is added only to the cold fluid and it does not interact with the contact discontinuity. Nevertheless it is a good test to evaluate the capability of the code to handle the presence of a transverse-velocity at a moderate resolution, while more extreme configurations require a resolution which is unreasonably high for multi-dimensional applications to obtain decent solutions. A way around this issue is the use of adaptive-mesh-refinement (AMR) [339], which however would not be useful for our use of THC to study of relativistic turbulence (it is in fact debatable whether the use of AMR is advantageous in this case).

In Figure 4.11 we show the analytic and numerical solutions for this problem at time $t = 0.4$. The presence of the transverse velocity widens the state between the shock wave and the contact discontinuity. The density contrast is also smaller with respect to the case where no tangential velocity is present. THC is able to capture the solution even at the low resolution, $\Delta^1 = 1/100$, shown in the Figure. The MP5 scheme overestimates slightly the density contrast, but all of the algorithms are able to capture the correct location of the shock wave.

Spherical explosion

As an example of a test involving non-grid-aligned shocks we consider the classical test of the spherical explosion in relativistic hydrodynamics. The initial data in this case is given by

$$(\rho, v, p) = \begin{cases} (1, 0, 1), & \text{if } r < 0.4; \\ (0.125, 0, 0.1), & \text{elsewhere.} \end{cases} \quad (4.26)$$




Figure 4.12: Rest-mass density (left panels) velocity (middle panels) and pressure (right panels) for the relativistic spherical explosion test. We show the results obtained with different schemes (dotted) as well as a reference solution (solid line) obtained with the 1D spherically symmetric code EDGES, [263], using 2000 elements of degree 3. The resolution is $\Delta^i = 1/100$ and the timestep is $\Delta^0 = 0.00125$ for all the runs.




Figure 4.13: RMS value of the y -component of the velocity during the linear-growth phase of the 2D relativistic Kelvin-Helmholtz instability, for different numerical schemes. The solid lines represent the results obtained at the highest resolution, 1024×2048 , while the dashed lines represent the results obtained at the lowest resolution, 128×256 . The timestep is chosen so that the CFL factor is $C \approx 0.25$ for all of the runs.

Since no analytic solution is known in this case, we use as reference solution the one computed using the one-dimensional spherically-symmetric discontinuous-Galerkin code EDGES [263] using 2000 elements of degree 3, solid lines, and compare it with the numerical solutions obtained in three-dimensions with THC in the diagonal direction. Both solutions at time $t = 0.25$ are shown in Figure 4.12, when using a resolution $\Delta^i = 1/100$ and a timestep $\Delta^0 = 0.00125$ for all the numerical schemes. As in the one-dimensional case, a small timestep is necessary in order to avoid numerical oscillations with the MP5 algorithm, while the other schemes appear to be stable even with a timestep which is twice as large.

Overall, all the schemes are able to properly capture the reference solution even at this very low resolution: the contact discontinuity is captured within two grid points and no oscillations are found in any of the physical quantities.

Kelvin-Helmholtz instability in 2D

The last test that we present is a classical benchmark for multidimensional hydrodynamics codes: the simulation of the KHI in two-dimensions, x, y . In order to ease the comparison with the existing literature, the initial conditions are chosen following [39]. The shear velocity is given by

$$v^x(y) = \begin{cases} V_{\text{shear}} \tanh[(y - 0.5)/a], & \text{if } y > 0; \\ V_{\text{shear}} \tanh[(y + 0.5)/a], & \text{if } y \leq 0; \end{cases} \quad (4.27)$$

where $a = 0.01$ is the characteristic size of the shear layer and $V_{\text{shear}} = 0.5$, corresponding to a relative Lorentz factor, *i.e.*, the Lorentz factor of a part of the




Figure 4.14: Rest-mass density for the 2D relativistic Kelvin-Helmholtz instability test at time $t = 3$ and for different numerical schemes. The resolution is 512×1024 for all the schemes and the CFL factor is $C \approx 0.25$. This figure should be compared with its equivalent Figure 4.15 at a resolution 1024×2048 .

fluid as seen by an observer comoving with the other one, of $W = 2.29$. The instability is seeded by adding a small perturbation in the transverse component of the velocity,

$$v^y(x, y) = \begin{cases} A_0 V_{\text{shear}} \sin(2\pi x) \exp\left[-(y - 0.5)^2 / \sigma\right], & \text{if } y > 0; \\ -A_0 V_{\text{shear}} \sin(2\pi x) \exp\left[-(y + 0.5)^2 / \sigma\right], & \text{if } y \leq 0; \end{cases} \quad (4.28)$$

where $A = 0.1$ is the perturbation amplitude and $\sigma = 0.1$ its characteristic lengthscale. The adiabatic constant is $\Gamma = 4/3$ and the initial pressure is uniform, $p = 1$. The rest-mass density distribution, which is uniform in the x -direction, is instead given by

$$\rho(y) = \begin{cases} \rho_0 + \rho_1 \tanh[(y - 0.5)/a], & \text{if } y > 0; \\ \rho_0 - \rho_1 \tanh[(y + 0.5)/a], & \text{if } y \leq 0; \end{cases} \quad (4.29)$$

where $\rho_0 = 0.505$ and $\rho_1 = 0.495$, so that $\rho = 1$ in the regions with $v^x = 0.5$ and $\rho = 0.1$ in the regions with $v^x = -0.5$. Finally the computational domain is $-0.5 \leq x \leq 0.5$, $-1 \leq y \leq 1$ and we use periodic boundary conditions in all the directions. Differently from [39] we do not add a random perturbation to the initial data and we do not take into account the effects of magnetic fields. Nevertheless, our results are in good agreement with the ones by [39] in the linear phase of the instability, since the magnetic field that they use is too weak to play a dynamical role in this phase.

We performed a series of simulations using six different numerical schemes: WENO5, WENO7, MP5, the bandwidth optimized WENO3B, WENO4B and using the WENO5 scheme, but with the Lax-Friedrichs flux-split, WENO5-LF. For each of these schemes, we ran with four different resolutions: 128×256 , 256×512 , 512×1024 and 1024×2048 . For all the runs, the CFL factor was taken to be $C \approx 0.25$.

The first quantity of interest to check is the growth rate of the transverse velocity during the linear-growth phase of the KHI, as computed with the different numerical schemes. This can be done by comparing the evolution of the root-mean-square (RMS) value of the transverse component of the velocity and defined as

$$\langle |v^y|^2 \rangle := \frac{1}{V} \int_V |v^y|^2 dV, \quad (4.30)$$

where V is the volume of the computational domain. This is shown in Figure 4.13, where, for each numerical scheme, we show the results taken at the lowest resolution (dashed lines) and at the highest one (continuous lines). First, we notice that with our setup the linear-growth phase of the instability lasts up to until $t \approx 3$. At that time, in fact, the transverse velocity reaches saturation and afterwards the fluid starts to become turbulent and the velocity shows large fluctuations.

[225] showed the importance of including the contact wave in the approximate Riemann solver in the case of a finite-volume code. In particular, they observed that while the growth rate of v^y was already accurate at low resolution when using their HLLD Riemann solver, in the cases where the more dissipative HLLC Riemann solver was employed the correct growth rate was recovered only at high resolution [similar results were also reported by [39]]. In analogy with what is observed with finite-volume codes, we also note the importance of avoiding excessive dissipation in the contact discontinuity by comparing the results obtained with the Lax-Friedrichs and the Roe flux-split when using WENO5. The Lax-Friedrichs flux-split underestimates the growth rate at low resolution and achieves a good accuracy in its measure only at resolution 256×512 (not shown in the figure). WENO5 and WENO7 with the Roe flux split already have a growth rate which is consistent with the highest resolution runs at the lowest resolution of 128×256 .

The behaviour of the MP5 scheme, as well as that of the bandwidth-optimized WENO schemes, is more surprising: all of these schemes overestimate the growth of the RMS transverse velocity at low resolution. This problem disappears as we increase the resolution and in the 256×512 case the growth rate is already consistent with the highest-resolution one for all the numerical schemes.

Some insight about the numerical viscosity can be gained by looking at the topology of the flow during the linear-growth phase of the KHI. In particular in Figure 4.14 we show a colour map of the rest-mass density obtained with the different schemes at time $t = 3$ using the 512×1024 resolution. [225] noticed that their solution obtained with the HLLD Riemann solver was showing a different structure, with the development of secondary small-scale instabilities along the shear layer, and that were not present when using the more diffusive HLLC solver. Similar differences were also observed in other works [see, e.g., [8, 308]], and the presence (or absence) of these secondary instabilities has been often

used as an indication of the numerical viscosity of the codes. [39] interpret these differences as an indication that HLLC is converging to a different weak solution from HLLD. Since both solvers are entropic (*i.e.*, with non-decreasing entropy), this would imply the non-uniqueness of the entropic solution for the Euler equations in this case.

In agreement with the conclusions of [39], we also note that these secondary instabilities, although only numerical artifacts (see below), appear only in schemes able to properly treat the initial contact discontinuity. As a result, the rest-mass density in Figure 4.14 obtained with WENO5 and WENO5-LF match very well the ones they obtained using HLLD and HLLC, respectively. However, our results do not support their conjecture that different schemes are converging to different weak solutions of the equations. The reason is that these secondary instabilities appear not to be genuine features of the solution and, rather, tend to disappear as the resolution is increased. For instance, the number of secondary vortices seems to change in a non-predictable fashion with the different numerical schemes and also with the resolution. This can be seen in Figure 4.14, which shows the great variability in the solutions obtained with different schemes. A similar variability is also present in results obtained at different resolutions with each scheme, WENO5-LF being the only exception. Finally, we point out that the size of the vortices is also shrinking as the resolution is increased.

More convincing evidence that these are indeed artifacts is shown in Figure 4.15, which is the same as Figure 4.14, but at the higher resolution of 1024×2048 . When comparing the two figures it appears clear that these secondary instabilities are much smaller. Interestingly, therefore, the more dissipative scheme WENO5 with the Lax-Friedrichs flux-split, WENO5-LF, seems to converge more rapidly to the correct solution (at least in these initial stages) than its less diffusive counterparts.

In essence, therefore, all of these considerations lead us to the conclusion that the secondary instabilities are triggered by the non-linear dissipation mechanism of the different schemes, emerge neatly when computed with numerical schemes that treat properly the initial contact discontinuity, but do not have a physical meaning.

A similar interpretation is also given by [219], who go one step further and suggest to add additional viscosity to numerical codes displaying these secondary instabilities in order to prevent their growth. While the addition of extra dissipation is going to smooth out small scales numerical perturbations, we argue that this issue can only be resolved with the inclusion of physical viscosity. Artificial viscosity would probably compete with the internal, non-linear, dissipation mechanism of the schemes yielding results that would be even more difficult to interpret (as the viscous scale will now depend both on the resolution and on the artificial viscosity strength).

A more quantitative way of estimating the numerical viscosity of the code in this test has been proposed by [39] and is based on the study of one-dimensional integrated power-spectra. Given a quantity $u(x, y)$ we define its integrated power-spectrum

$$P_u^1(k) = \int_{-1}^1 |\hat{u}(k, y)|^2 dy, \quad (4.31)$$




Figure 4.15: Rest-mass density for the 2D relativistic Kelvin-Helmholtz instability test at time $t = 3$ and for different numerical schemes. The resolution is 1024×2048 for all the schemes and the CFL factor is $C \approx 0.25$. This figure should be compared with its equivalent Figure 4.14 at a resolution 512×1024 .




Figure 4.16: One dimensional power spectrum of the rest-mass density for the 2D relativistic Kelvin-Helmholtz instability test at time $t = 3$ and for different numerical schemes. The resolution is 1024×2048 for all the schemes and the CFL factor is $C \approx 0.25$.

where k is the wavenumber and

$$\hat{u}(k, y) = \int_{-1/2}^{1/2} u(x, y) e^{-2\pi i k x} dx, \quad (4.32)$$

is the one-dimensional Fourier transform of u . To ease the comparison with the spectra reported by [39], the power spectrum is then normalized so that

$$\sum_{k=1}^{N/2} P_u^1(k) = 1, \quad (4.33)$$

where N is the number of gridpoints. The one-dimensional power spectrum can be used to quantify the typical scale of structures, such as the secondary vortices discussed above, stretched in the direction of the bulk shear flow.

In Figure 4.16 we show the spectra of the rest-mass density for the different schemes at time $t = 3$ and for the highest resolution. As expected the WENO5-LF scheme stands out as the scheme having the largest dissipation. More surprising is the behaviour of the bandwidth optimized schemes: they appear to be improving over their classical counterparts only at high wavenumbers, that is, at scales that are dominated by the (non-physical) dissipation of the scheme. Even more unexpected is the ability of the MP5 scheme to resolve small scales structures and that, on the basis of the argument about the development of the secondary instabilities, should be more dissipative than WENO4B, but which instead appears to yield more small-scale structures in the rest-mass density.

A similar conclusion can be obtained by studying the spectrum of the kinetic energy

$$E_k = \rho W(W - 1), \quad (4.34)$$

but which we do not report since its behaviour is very similar to the one shown for the rest-mass density.

4.4 The relativistic Kelvin-Helmholtz instability in 3D

As an example of a non-trivial application of THC and a perfect playground to evaluate the performances of the different numerical schemes for the simulation of turbulence, we present here a study of the relativistic turbulence generated by the KHI in three-dimensions. Our analysis is not meant to be a systematic assessment of the accuracy of these methods for direct numerical simulations of compressible relativistic turbulence, as done, for instance by [181], or [177] in the case of classical turbulence. Rather, our analysis is meant to assess how the different methods reproduce the same turbulent initial-value problem and to provide some insight on the spectral properties of the different schemes.

The relativistic KHI [see, *e.g.*, [51]] is of particular interest because of its relevance for the stability of relativistic jets [see, *e.g.*, [251, 250]], and because of its potential role in the amplification of magnetic fields in gamma-ray bursts [see, *e.g.*, [338]], and binary neutron-star mergers [25, 143, 240, 274].

We consider a triply-periodic box, $-0.5 \leq x \leq 0.5$, $-1 \leq y \leq 1$, $-0.5 \leq z \leq 0.5$. The initial conditions are the same ones employed in Section 4.3.2, and the




Figure 4.17: RMS value of the y -component of the velocity during the linear-growth phase of the 3D relativistic Kelvin-Helmholtz instability, for different numerical schemes. The thick-dashed line represent the results obtained at the highest resolution, 1024×2048 , in 2D with WENO5. The resolution is $256 \times 512 \times 256$ for all the schemes.

symmetry in the z -direction is broken with the application of random perturbations, uniformly distributed in the range $[0, 0.01]$, applied to the velocity in the z -direction v^z . We simulated this system up to the time $t = 30$ on a grid of $256 \times 512 \times 256$ cells using the WENO5, WENO7, MP5, WENO3B and WENO4B reconstructions. In addition, we have also performed one run at the same resolution using the second-order MINMOD reconstruction. Furthermore, because of the high computational costs involved to check the convergence of the code, we have managed to run only one model at twice the reference resolution, *i.e.*, $512 \times 1024 \times 512$, using the WENO5 reconstruction scheme.

As discussed in the previous Section, the MP5 scheme requires a timestep which is half that of the other schemes. On the other hand, there is no need to run the WENO schemes with such a small timestep and in any “real” application we will run the WENO scheme with the maximum possible timestep. For this reason, and since we are interested in the performance of the different schemes in their “real-world” configuration, we have not used the same CFL factor for all of them, with a considerable saving in computational costs. In particular, for all the runs we have used a CFL factor $C \approx 0.25$, with the exception of the one with the MP5 scheme, where we have used $C \approx 0.125$.

4.4.1 The linear evolution of the instability

First of all, we consider the evolution of the instability during its linear-growth phase. At this stage, the velocity perturbations in the direction perpendicular to the shearing one are growing exponentially and three-dimensional effects are still negligible.

The evolution of the RMS value of the y -component of the velocity is re-




Figure 4.18: Rest-mass density in the $z = 0$ plane for the 3D relativistic Kelvin-Helmholtz instability study at time $t = 3$ and for different numerical schemes. The resolution is $256 \times 512 \times 256$ for all the schemes.

ported in Figure 4.17, where we show the results obtained with the different schemes, as well as a reference solution computed in two dimensions (2D) using WENO5 on 1024×2048 grid points. As expected, all the numerical schemes, with the exception of MINMOD, are in very good agreement with the 2D solution up to the end of the linear-growth phase at time $t = 3$, when 3D effects become important and turbulence starts to play an important role in the dynamics. It is interesting to note that MINMOD, which is the most dissipative of the schemes we are using, is actually overestimating the growth of the KHI. This suggests that some care should be taken when interpreting the results from under-resolved simulations, since it is not always safe to assume that the high numerical viscosity of the low-resolution runs tends to suppress the instability, yielding a lower-bound on its growth.

The rest-mass density at time $t = 3$ is shown in Figure 4.18, where, in analogy with the 2D case, we find that the more dissipative schemes (*i.e.*, MINMOD), do not show any sign of secondary instabilities apart from the seeded ones, while the least dissipative ones (*i.e.*, WENO3B, WENO4B) show the emergence of secondary vortices.

At this point in time, the flow is still mostly two-dimensional, but it is interesting to notice the effects resulting from the small perturbations in the vertical velocity v^z . The perturbations have the same statistical properties in all the different models, but their effects are appreciably different for the different schemes, as can be seen from Figure 4.18. Although the perturbation in v^z is random and it does not preserve any of the symmetries of the problem, it is still symmetric, on average, with respect to the y -axis. For this reason, one does not expect to find a large symmetry violation until the time




Figure 4.19: Evolution of the concentration of the passive tracer in the $z = 0$ plane for the 3D Kelvin-Helmholtz instability obtained using WENO5 on $512 \times 1024 \times 512$ gridpoints.

when these perturbations have had enough time to grow and the dynamics has started to become dominated by three-dimensional effects. Yet, the optimized schemes WENO3B, WENO4B appear to be much more sensitive to the symmetry breaking than the others. The reason for this is probably that the bandwidth optimized-algorithms appear to trigger smaller-scale secondary instabilities, which, in turn, are more easily affected by the perturbations in the vertical velocity, since the perturbation does not average out at small scales.

4.4.2 The non-linear evolution of the instability

The linear-growth phase of the KHI instability ends when the primary vortices become unstable to secondary instabilities and the flow starts the transition to turbulence. At this point, three-dimensional effects start to become dominant and the dynamics is qualitatively different from the one observed in the 2D case.

In order to better track the evolution of the fluid in this phase, we monitor the concentration of a passive scalar field, ϕ , transported with the fluid via the advection equation

$$\frac{\partial \phi}{\partial t} + \sum_{i=1}^3 v^i \frac{\partial \phi}{\partial x^i} = 0. \quad (4.35)$$

Equation (4.35) is not a conservation law, so that it can not be written directly²

²Equation (4.35) can be written in conservation form at the price of introducing an additional, conserved variable, $\tilde{\phi} := \rho W \phi$. In our study this complication is not necessary, as we use the tracer

in the form (3.62), but it is nevertheless a hyperbolic equation of the Hamilton-Jacobi type. For this reason, in order to solve numerically (4.35) we use a class of numerical schemes designed for this family of equations and that can be built using the non-oscillatory reconstruction of the derivative introduced in Section 4.2 [see, *e.g.*, [301]]. In particular, the semi-discrete form of equation (4.35) becomes

$$\begin{aligned} \frac{d\phi_{i,j,k}}{dt} = & v_{i,j,k}^x \frac{\phi_{i-1/2,j,k} - \phi_{i+1/2,j,k}}{\Delta^1} + v_{i,j,k}^y \frac{\phi_{i,j-1/2,k} - \phi_{i,j+1/2,k}}{\Delta^2} + \\ & v_{i,j,k}^z \frac{\phi_{i,j,k-1/2} - \phi_{i,j,k+1/2}}{\Delta^3}, \end{aligned} \quad (4.36)$$

where $(\phi_{i-1/2,j,k} - \phi_{i+1/2,j,k})/\Delta^1$ is a high-order non-oscillatory approximation of $-\partial\phi/\partial x^1$ at $x_{i,j,k}$ obtained using an upwind-biased reconstruction, *i.e.*,

$$\phi_{i-1/2,j,k} = \phi_{i-1/2,j,k}^\pm \quad \text{if } v^x \lessgtr 0. \quad (4.37)$$

The “fluxes” in the other directions are obviously computed in an analogous way.

The time evolution of the tracer, as computed with WENO5 at the highest-resolution is shown in Figure 4.19. At the initial time, we set the scalar field ϕ to be equal to the rest-mass density, so that the initial configuration consists in two regions with opposite “colour”, separated by a thin transition layer. At the end of the linear regime, *i.e.*, at time $t \simeq 3$, the tracer profile is distorted by the presence of the primary vortices as well as of the secondary ones, but there is only a marginal mixing of the two “phases” of the fluid. Note that, since we do not include any explicit dissipation term in the advection equation (4.35), the mixing of the tracer happens only due to numerical dissipation. As the vortices start to become unstable, the fluid starts to concentrate the scalar field in thin structures and the two regions of the fluid start to mix around the shearing region. As the simulation proceeds, the width of the region affected by the mixing grows, till at time $t \simeq 24$, when there are only small patches of the fluid that still have a uniform tracer (colour). At the final time, $t = 30$, there are no regions in the flow that are unaffected by the mixing and the initial structure is mostly destroyed, even though perfect mixing has not been achieved yet.

We track the evolution of the variance of the scalar field, which we compute as

$$\text{Var}[\phi] := \langle |\phi - \langle \phi \rangle|^2 \rangle, \quad (4.38)$$

and which, for $t \geq 5$, we find to be very-well fitted by a simple exponential law of the type

$$\text{Var}[\phi] = K e^{-t/\tau}. \quad (4.39)$$

The values of the fitting constant for the WENO5 scheme at the highest resolution are $K = 0.28$ and $\tau = 14.75$. The mixing timescale, τ , exhibits only small changes between the runs at different resolutions, with the exception of the results obtained with MP5, where the timescale is $\tau = 17.5$. Hence, the total time of the simulation, $t = 30$, is roughly equivalent to two e-folding times for the mixing of the passive tracer.

only as a diagnostic quantity. On the other hand, in situations where, for instance, the tracers are used to model the chemical composition of a fluid in a reacting flow, it may be important to ensure the conservation of the different species and a conservative approach may be preferred [see, *e.g.*, [252] for an example of this approach].




Figure 4.20: Compensated power spectrum of the velocity at time $t = 30$ for the 3D Kelvin-Helmholtz instability test, using WENO5 at two different resolutions, $256 \times 512 \times 256$ and $512 \times 1024 \times 512$.

Fully-Developed turbulence

After the linear-growth phase, the flow quickly becomes turbulent. By the time we stop the simulation at $t = 30$, the turbulence is fully developed even though the flow is still not isotropic, but preserving some memory of the initial configuration. Because of the large computational costs involved in these tests (which, we recall, have been performed for six different methods) we have not carried out the simulations for longer times, assuming that the properties of the non-perfectly isotropic turbulence at the final time are sufficiently close to the fully isotropic turbulence. We point out that our simulation time is more than two times longer than the one reported in [338], where a setup similar to ours was used, but in the more challenging regime of relativistic MHD.

By far the most interesting quantity to study is the three-dimensional velocity power spectrum

$$P_v(k) = \frac{1}{2} \int_{|\mathbf{k}|=k} |\hat{v}(\mathbf{k})|^2 d\mathbf{k}, \quad (4.40)$$

where $\hat{v}(\mathbf{k})$ is the three-dimensional Fourier transform of $v(\mathbf{x})$,

$$\hat{v}(\mathbf{k}) = \int_V v(\mathbf{x}) e^{-2\pi i \mathbf{k} \cdot \mathbf{x}} d\mathbf{x}, \quad (4.41)$$

and V is the volume of the computational region. The integral in (4.40) is computed following [126] as

$$P_v(k) = \frac{1}{2} \frac{4\pi k^2}{N_k} \sum_{k-1/2 < |\mathbf{k}| \leq k+1/2} |\hat{v}(\mathbf{k})|^2, \quad (4.42)$$

where N_k is the number of discrete wave-numbers in the shell $k - 1/2 \leq |\mathbf{k}| < k + 1/2$. For simplicity, we study the data in the artificially-enlarged cubic




Figure 4.21: Compensated power spectrum of the velocity at time $t = 30$ for the 3D Kelvin-Helmholtz instability test, using different numerical schemes. In all the cases the resolution is $256 \times 512 \times 256$.

domain $[-1, 1]^3$ and we do that by duplicating the original data in the x and z -directions, exploiting the symmetry of the problem. This procedure avoids the complications of a computational domain which does not have the same extent in all directions and yields an “equivalent resolution” of 512^3 and 1024^3 points for the low and high-resolution runs respectively. Clearly the first few wave-numbers of the spectrum are influenced by this procedure, but all the higher wave-numbers are essentially unaffected.

At the final time, the flow is subsonic ($M_{\max} \lesssim 1$) and relativistically warm ($\epsilon \gtrsim 1$). Under these conditions, studies of Newtonian [258] and relativistic [338, 172, 346] transonic turbulence suggest that the velocity spectrum should be consistent with the Kolmogorov phenomenology [183]. In particular, the power spectrum should scale as

$$P_v(k) \sim k^{-5/3}, \quad (4.43)$$

in the so called *inertial range*, that is at those scales where the fluid motion is sufficiently independent from the large-scale dynamics and from the small-scale viscosity, so as to exhibit a self-similar universal behaviour.

In Figure 4.20 we show the compensated velocity spectrum, *i.e.*, $k^{5/3}P_v(k)$, at time $t = 30$ obtained from the data of the two WENO5 runs. More specifically, the low-resolution spectrum is shifted to larger wavenumbers by a factor two and scaled assuming a $k^{-5/3}$ scaling of the spectrum. The rationale behind this procedure is that we are interested in the (eventual) self-similar behaviour of the spectrum and it is therefore useful to consider the low-resolution run as a realisation of the same flow as the high-resolution one, but in a smaller volume.

In other words, we interpret the spectrum of the low-resolution run as being computed from a subsample of the data of the high-resolution one [see, *e.g.*, [52] for a more detailed discussion of the issue of convergence for direct simulations of turbulent flows].

The plot demonstrates the statistical convergence of the simulation, with the exception of the very low-wavenumber part of the spectrum, where the convergence is probably spoiled by the symmetry. At the same time, the plot also shows that only the high-resolution run seems to be able to cover a sufficiently wide part of the inertial range to give a clear indication of the Kolmogorov $-5/3$ scaling.

In Figure 4.21 we show the compensated velocity spectra at time $t = 30$ obtained with the different numerical schemes. For each scheme, we highlight with a grey shaded area the part of the spectrum that appears to be “compatible” with the $-5/3$ scaling inferred from the high-resolution WENO5 run. The width of this region, expressed in terms of wavenumbers, is also indicated on the figure. Clearly, this measure has to be taken with a bit of caution, since there is a certain degree of arbitrariness in the identification of the “inertial range”; furthermore, the judgment is also made more difficult by the fact that all of the results are obtained at a resolution which probably is not high enough and that a convergence study is only available for the WENO5 case. Notwithstanding these caveats, the difference between the various schemes is already sufficiently clear to deduce some of their properties despite the limited amount of data that we were able to generate.

A first conclusion to be drawn is about the importance of the use of high-order schemes, which is apparent if we compare the different spectra with the one obtained with the MINMOD scheme. This second-order algorithm, in fact, yields a spectrum which is completely dominated by the so-called “bottleneck-effect” [see, *e.g.*, [65, 287]], *i.e.*, by a region where the power-spectrum shows an excess due to viscous effects. No clear indication of an inertial range appears anywhere in the spectrum, with the only “flat” region being the middle of the bottleneck zone. This could be easily mistaken as the inertial range in the absence of a proper convergence analysis. For this reason, and as remarked many times already, care should be taken while interpreting the results obtained with low-order schemes.

Secondly, we observe that WENO4B has an effective resolution which is about 50% larger than the one from WENO5 and 20% larger than WENO7. Given that saving a factor 1.5 in resolution corresponds, in 3D to a decrease of the computation costs by a factor five³, we conclude that the use of WENO4B over WENO5 is well justified, since WENO4B is roughly twice as expensive as WENO5 in 3D. On the other hand, the spectral resolution of WENO4B does not appear to be better than the one of the MP5 scheme, which has similar computational costs (due to the stricter CFL limitation), but which is also expected to have better parallel scaling because of its more compact stencil. Overall, MP5 shows an “inertial-range” as large as WENO4B. We also note that WENO3B does not seem to yield any improvements over WENO5.

All things considered, the main differences between the bandwidth-optimized schemes and their traditional counterparts seem to lay in the bottleneck region:

³We note that when increasing the resolution, also the timestep is reduced via the CFL condition. Hence, the additional cost goes like the fourth power of the ratio in resolutions.

WENO3B and WENO4B have a much less pronounced bottleneck with respect to WENO5, WENO7 and MP5. This suggests that these schemes should be considered especially for under-resolved turbulent flows, where the spectrum is basically entirely dominated by the dissipation. MP5, on the other hand, can be particularly useful for those problems where one is interested in well-resolved quantities, such as in direct numerical simulations of turbulence, since the scales affected by the numerical dissipation are more easily identified by the pronounced bottleneck. MP5 should also be the scheme of choice in Newtonian hydrodynamics, since there its timestep constraint seems to be less severe.

4.5 Driven Relativistic Turbulence

4.5.1 Introduction

Turbulence is a ubiquitous phenomenon in nature as it plays a fundamental role in shaping the dynamics of systems ranging from the mixture of air and oil in a car engine, up to the rarefied hot plasma composing the intergalactic medium. Relativistic hydrodynamics is a fundamental ingredient in the modeling of a number of systems characterized by high Lorentz-factor flows, strong gravity or relativistic temperatures. Examples include the early Universe, relativistic jets, gamma-ray-bursts (GRBs), relativistic heavy-ion collisions and core-collapse supernovae [130].

Despite the importance of relativistic hydrodynamics and the reasonable expectation that turbulence is likely to play an important role in many of the systems mentioned above, extremely little is known about turbulence in a relativistic regime. For this reason, the study of relativistic turbulence may be of fundamental importance to develop a quantitative description of many astrophysical systems. To this aim, we have performed a series of high-order direct numerical simulations of driven relativistic turbulence of a hot plasma.

4.5.2 Model and method

We consider an idealized model of an ultrarelativistic fluid with four-velocity $u^\mu = W(1, v^i)$, where $W := (1 - v_i v^i)^{-1/2}$ is the Lorentz factor and v^i is the three-velocity in units where $c = 1$. The fluid is modeled as perfect and described by the stress-energy tensor

$$T_{\mu\nu} = (e + p)u_\mu u_\nu + p g_{\mu\nu}, \quad (4.44)$$

where e is the (local-rest-frame) energy density, p is the pressure, u_μ the four-velocity, and $g_{\mu\nu}$ is the spacetime metric, which we take to be the Minkowski one. We evolve the equations describing conservation of energy and momentum in the presence of an externally imposed Minkowskian force F^μ , *i.e.*, $\nabla_\nu T^{\mu\nu} = F^\mu$, where the forcing term is written as $F^\mu = \tilde{F}(0, f^i)$. More specifically, the spatial part of the force, f^i , is a zero-average, solenoidal, random, vector field with a spectral distribution which has compact support in the low wavenumber part of the Fourier spectrum. Moreover, f^i is kept fixed during the evolution and it is the same for all the models, while $\tilde{F}$ is either a constant or a simple function of time (see below for details).




Figure 4.22: Average Lorentz factor as a function of time for the different models considered. Note that a quasi-stationary state is reached before $t \sim 10$ for all values of the driving force.




Figure 4.23: Logarithm of the Lorentz factor on the (y, z) plane at the final time of model D. Note the large spatial variations of the Lorentz factor with front-like structures. The time-averaged PDFs are shown in the lower left corner for the different models considered.

The time component of the forcing term, F^0 , is set to be zero, so that the driving force is able to accelerate fluid elements without changing their total energy (in the Eulerian frame). Note that this is conceptually equivalent to the addition of a cooling term balancing the effect of the work done on the system by the driving force. On the other hand, we impose a minimum value for the energy density in the local-rest-frame, $e_{\min}$. This choice is motivated essentially by numerical reasons (the very large Lorentz factor produced can lead to unphysical point-wise values of e) and has the effect of slowly heating up the fluid. Furthermore, this floor does not affect the momentum of the fluid and only the temperature is increased. From a physical point of view, our approach mimics the fact that in the low-density regions, the constituents of the plasma are easily accelerated to very high Lorentz factors, hence emitting bremsstrahlung radiation heating up the surrounding regions. The net effect is that energy is subtracted from the driving force and converted into thermal energy of the fluid, heating it up. In general $e_{\min}$ is chosen to be two orders of magnitude smaller than the initial energy density, but we have verified that the results presented here are insensitive to the specific value chosen for $e_{\min}$ by performing simulations where the floor value is changed by up to two orders of magnitude without significant differences.

The set of relativistic-hydrodynamic equations is closed by the equation of state (EOS) $p = \frac{1}{3}e$, thus modelling a hot, optically-thick, radiation-pressure dominated plasma, such as the electron-positron plasma in a GRB fireball or the matter in the radiation-dominated era of the early Universe. The EOS used can be thought as the relativistic equivalent of the classical isothermal EOS in that the sound speed is a constant, *i.e.*, $c_s^2 = 1/3$. At the same time, an ultrarelativistic fluid is fundamentally different from a classical isothermal fluid. For instance, its “inertia” is entirely determined by the temperature and the notion of rest-mass density is lost since the latter is minute (or zero for a pure photon gas) when compared with the internal one. For these reasons, there is no direct classical counterpart of an ultrarelativistic fluid and a relativistic description is needed even for small velocities.

To solve the equations of relativistic hydrodynamics in 3D we use the THC code described in this chapter and published in [264]. In particular, here, we use the MP5 reconstruction in local characteristic variables [165].

4.5.3 Basic flow properties

Our analysis is based on the study of four different models, which we label as A, B, C and D, and which differ for the initial amplitude of the driving factor $\tilde{F} = 1, 2, 5$ for models A–C, and $\tilde{F}(t) = 10 + \frac{1}{2}t$ for the extreme model D. Each model was evolved using three different uniform resolutions of 128^3 , 256^3 and 512^3 grid-zones over the same unit lengthscale. As a result, model A is subsonic, model B is transonic and models C and D are instead supersonic. The spatial and time-averaged relativistic Mach numbers $\langle vW \rangle / (c_s W_s)$ are 0.362, 0.543, 1.003 and 1.759 for our models A, B, C and D, while the average Lorentz factors are 1.038, 1.085, 1.278 and 1.732 respectively.

The initial conditions are simple: a constant energy density and a zero-velocity field. The forcing term, which is enabled at time $t = 0$, quickly accelerates the fluid, which becomes turbulent. By the time when we start to sample the data, *i.e.*, at $t = 10$ (light-)crossing times, turbulence is fully developed and

the flow has reached a stationary state. The evolution is then carried out up to time $t = 40$, thus providing data for 15, equally-spaced timeslices over 30 crossing times. As a representative indicator of the dynamics of the system, we show in Fig. 4.22 the time evolution of the average Lorentz factor for the different models considered. Note that the Lorentz factor grows very rapidly during the first few crossing times and then settles to a quasi-stationary evolution. Furthermore, the average grows non-linearly with the increase of the driving term, going from $\langle W \rangle \simeq 1.04$ for the subsonic model A, up to $\langle W \rangle \simeq 1.73$ for the most supersonic model D.

Flow quantities such as the energy density, the Mach number or the Lorentz factor show large spatial variance, even in our subsonic model. Similar deviations from the average mass density have been reported also in classical turbulent flows of weakly compressible fluids [42], where it was noticed that compressible effects, leading to the formation of front-like structures in the density and entropy fields, cannot be neglected even at low Mach numbers. In the same way, relativistic effects in the kinematics of the fluid, such as those due to non-linear couplings via the Lorentz factor [272], have to be taken into account even when the average Lorentz factor is small. The probability distribution functions (PDFs) of the Lorentz factor are shown in Fig. 4.23 for the different models. Clearly, as the forcing is increased, the distribution widens, reaching Lorentz factors as large as $W \simeq 40$ (*i.e.*, to speeds $v \simeq 0.9997$). Even in the most “classical” case A, the flow shows patches of fluid moving at ultrarelativistic speeds. Also shown in Fig. 4.23 is the logarithm of the Lorentz factor on the (y, z) plane and at $t = 40$ for model D, highlighting the large spatial variations of W and the formation of front-like structures.

4.5.4 Universality

As customary in studies of turbulence, we have analyzed the power spectrum of the velocity field

$$E_v(k) := \frac{1}{2} \int_{|\mathbf{k}|=k} |\hat{v}(\mathbf{k})|^2 d\mathbf{k}, \quad (4.45)$$

where $\mathbf{k}$ is a wavenumber three-vector and

$$\hat{v}(\mathbf{k}) := \int_V \mathbf{v}(\mathbf{x}) e^{-2\pi i \mathbf{k} \cdot \mathbf{x}} d\mathbf{x}, \quad (4.46)$$

with V being the three-volume of our computational domain. A number of recent studies have analyzed the scaling of the velocity power spectrum in the inertial range, that is, in the range of wavenumbers between the lengthscale of the problem and the scale at which dissipation dominates. More specifically, Inoue et al. [172] has reported evidence of a Kolmogorov $k^{-5/3}$ scaling in freely-decaying MHD turbulence, but have not provided a systematic convergence study of the spectrum. Evidence for a $k^{-5/3}$ scaling was also found by Zhang et al. [338], in the case of the kinetic-energy spectrum, which coincides with the velocity power-spectrum in the incompressible case. Finally, Zrake and MacFadyen [346] have performed a significantly more systematic study for driven, transonic, MHD turbulence, but obtained only a very small (if any) coverage of the inertial range.




Figure 4.24: Power spectra of the velocity field. Different lines refer to the three resolutions used and to the different values of the driving force. The spectra are scaled assuming a $k^{-5/3}$ law.

The time-averaged velocity power spectra computed from our simulations are shown in Fig. 4.24. Different lines refer to the three different resolutions used, 128^3 (dash-dotted), 256^3 (dashed) and 512^3 (solid lines), and to the different values of the driving force. To highlight the presence and extension of the inertial range, the spectra are scaled assuming a $k^{-5/3}$ law, with curves at different resolutions shifted of a factor two or four, and nicely overlapping with the high-resolution one in the dissipation region. Clearly, simulations at higher resolutions would be needed to have power-spectra which are more accurate and with larger inertial ranges, but overall, Fig. 4.24 convincingly demonstrates the good statistical convergence of our code and gives strong support to the idea that the *key* prediction of the Kolmogorov model (K41) [183] carries over to the relativistic case. Indeed, not only does the velocity spectrum for our subsonic model A shows a region, of about a decade in length, compatible with a $k^{-5/3}$ scaling, but this continues to be the case even as we increase the forcing and enter the regime of relativistic supersonic turbulence with model D. In this transition, the velocity spectrum in the inertial range, the range of lengthscales where the flow is scale-invariant, is simply “shifted upwards” in a self-similar way, with a progressive flattening of the bottleneck region, the bump in the spectrum due to the non-linear dissipation introduced by our numerical scheme. Steeper or shallower scalings, such as the Burgers one, k^{-2} , or a $k^{-4/3}$ one, are also clearly incompatible with our data.

These results have been confirmed in a preliminary study where we pushed our resolution for model D, the most extreme one, to 1024^3 .

All in all, this is one of our main results: the velocity power spectrum in the inertial range is *universal*, that is, insensitive to relativistic effects, at least in the subsonic and mildly supersonic cases. Note that this does *not* mean that relativistic effects are absent or can be neglected when modelling relativistic turbulent flows.




Figure 4.25: Compensated, third-order, parallel structure function computed for the different models as functions of r/Δ . Note the very good match with the classical $S_3^{\parallel} \sim r$ behaviour.

4.5.5 Intermittency

Not all of the information about relativistic turbulent flows is contained in the velocity power spectrum. Particularly important in a relativistic context is the intermittency of the velocity field, that is, the local appearance of anomalous, short-lived flow features, which we have studied by looking at the parallel-structure functions of order p

$$S_p^{\parallel}(r) := \langle |\delta_r v|^p \rangle, \quad \delta_r v = \left[\mathbf{v}(\mathbf{x} + \mathbf{r}) - \mathbf{v}(\mathbf{x}) \right] \cdot \frac{\mathbf{r}}{r} \quad (4.47)$$

where $\mathbf{r}$ is a vector of length r and the average is over space and time.

Figure 4.25 reports the compensated, third-order, parallel structure function, $S_3^{\parallel}$, as functions of r/Δ , where Δ is the grid spacing. Within the inertial range, classical incompressible turbulence has a precise prediction: the Kolmogorov 4/5-law, for which $\langle (\delta_r v)^3 \rangle = \frac{4}{5} \epsilon r$, where ϵ is the kinetic-energy dissipation rate. This translates into $S_3^{\parallel} \sim \epsilon r$. As shown in the figure, the structure functions are somewhat noisy at small scales, but are consistent with the classical prediction over a wide range of lengthscales, with linear fits showing deviations of $\sim 5\%$, and an increase of ϵ with the driving force.

Although even in the classical compressible case, the 4/5-law is not strictly valid, we can use it to obtain a rough estimate of the turbulent velocity dissipation rate [258]. We find that ϵ , as measured from $S_3^{\parallel}$ or directly from $\langle (\delta_r v)^3 \rangle$, grows linearly with the Lorentz factor, in contrast with the classical theory, where it is known to be independent of the Reynolds number. This is consistent with the observations that in a relativistic regime the turbulent velocity shows an exponential decay in time [345, 172], as opposed to the power-law decay seen in classical compressible and incompressible turbulence.

The scaling exponents of the parallel structure functions, $\zeta_p^{\parallel}$ have been computed up to $p = 10$ using the extended-self-similarity (ESS) technique [41] and are summarized in Table 4.2. The errors are estimated by computing the ex-

Table 4.2: Scaling exponents of the parallel structure functions computed using the ESS technique and analytical predictions from the KS41, SL and Boldyrev models.

Order	K41	SL	Boldyrev	A512	B512	C512	D512
1	0.33	0.36	0.41	0.37 ± 0.01	0.36 ± 0.01	0.37 ± 0.01	0.38 ± 0.005
2	0.67	0.70	0.74	0.70 ± 0.02	0.70 ± 0.03	0.70 ± 0.02	0.71 ± 0.01
3	1	1	1	1 ± 0.02	1 ± 0.04	1 ± 0.03	1 ± 0.03
4	1.33	1.28	1.21	1.27 ± 0.03	1.27 ± 0.05	1.26 ± 0.04	1.25 ± 0.03
5	1.67	1.54	1.39	1.51 ± 0.02	1.50 ± 0.07	1.48 ± 0.05	1.46 ± 0.05
6	2	1.78	1.56	1.72 ± 0.03	1.70 ± 0.08	1.68 ± 0.07	1.64 ± 0.07
7	2.33	2.00	1.70	1.89 ± 0.04	1.86 ± 0.12	1.84 ± 0.09	1.79 ± 0.09
8	2.67	2.21	1.84	2.04 ± 0.04	1.99 ± 0.16	1.98 ± 0.11	1.92 ± 0.11
9	3	2.41	1.96	2.17 ± 0.03	2.10 ± 0.21	2.09 ± 0.13	2.04 ± 0.14
10	3.33	2.59	2.08	2.27 ± 0.02	2.18 ± 0.26	2.19 ± 0.16	2.14 ± 0.16

ponents without the ESS or using only the data at the final time. We also show the values as computed using the classical K41 theory, as well as using the estimates by She and Leveque (SL) [293] for incompressible turbulence, *i.e.*, $\zeta_p^{\parallel} = \frac{p}{9} + 2 - 2(\frac{2}{3})^{p/3}$, and those by Boldyrev [53] for Kolmogorov-Burgers supersonic turbulence, *i.e.*, $\zeta_p^{\parallel} = \frac{p}{9} + 1 - (\frac{1}{3})^{p/3}$.

Not surprisingly, as the flow becomes supersonic, the high-order exponents tend to flatten out and be compatible with the Boldyrev scaling, as the most singular velocity structures become two-dimensional shock waves. $\zeta_2^{\parallel}$, instead, is compatible with the She-Leveque model even in the supersonic case. This is consistent with the observed scaling of the velocity power spectrum, which presents only small intermittency corrections to the $k^{-5/3}$ scaling. Previous classical studies of weakly compressible [42] and supersonic turbulence [258] found the scaling exponents to be in very good agreement with the ones of the incompressible case and to be well described by the SL model. This is very different from what we observe even in our subsonic model A, in which the exponents are significantly flatter than in the SL model, suggesting a stronger intermittency correction. This deviation is another important result of our simulations.

One non-classical source of intermittency is the genuinely relativistic constraint that the velocity field cannot be Gaussian as the PDFs must have compact support in $(-1, 1)$. This is shown by the behaviour of the PDFs of v_z and plotted as solid lines in the shaded area of Fig. 4.26. Clearly, as the Lorentz factor increases, the PDFs become flatter and, as a consequence, the velocity field shows larger deviations from Gaussianity (dashed lines). Stated differently, relativistic turbulence is significantly more intermittent than its classical counterpart.

4.6 Conclusions

We have presented THC, a new multi-dimensional, finite-difference, high-resolution shock-capturing code for classical and special-relativistic hydrodynamics. THC




Figure 4.26: PDFs of the velocity v_z for the different models considered (solid lines). As the forcing is increased, the PDFs flatten, while constrained to be in $(-1, 1)$ (shaded area). Increasingly large deviations from Gaussianity (dashed lines) appear in the relativistic regime.

employs up to seventh-order accurate reconstruction of the fluxes in local characteristic variables and the Roe flux-vector-splitting algorithm with a novel entropy-fix prescription. The multi-dimensional case is treated in a dimensionally unsplit fashion and the time integration is done with a third-order strongly-stability-preserving Runge-Kutta scheme.

We have carried out a systematic comparison of the results obtained with our code using five different reconstruction operators: the classical WENO5, WENO7, MP5, as well as two specialised bandwidth-optimized WENO schemes: WENO3B and WENO4B. For all schemes, we have checked their ability to sharply capture grid-aligned, diagonal or spherical shock waves, and have carried out a rigorous assessment of their accuracy in the case of smooth solutions. Finally, we have contrasted the performance of the different methods in the resolution of small scale structures in turbulent flows. To the best of our knowledge, this is the first time that such a comparison has been carried over in the special-relativistic case.

Among the different tests studied, some are highly nontrivial, such those involving the linear and the non-linear phases of the development of the Kelvin-Helmholtz instability for a relativistic fluid in two and three dimensions. In particular, we have shown the importance of using schemes that are able to properly capture the initial contact discontinuity in order to obtain the correct growth rate of the RMS transverse velocity in the linear-growth phase of the instability at low resolution, confirming the findings by [225] and [39].

When studying the Kelvin-Helmholtz instability in two dimensions, we have investigated the nature of the secondary vortices that appear during the initial stages of the instability when using some of the numerical schemes considered. We have then clarified that these vortices are not genuine features of the solution of the equations, but rather numerical artefacts that converge away with resolution. When studying the Kelvin-Helmholtz instability in three dimensions, we have instead investigated the “mixing timescale” of the

instability and the subsequent turbulent flow, showing that we are able to obtain a converged measure of the velocity power spectrum, using the WENO5 scheme. Our data offers a clear indication that the Kolmogorov phenomenology [183] holds also for the turbulence in a subsonic relativistically-warm fluid. Using the Kolmogorov $-5/3$ scaling as a reference, we have estimated the effective inertial range of the different schemes, highlighting the importance of using high-order schemes to study turbulent flows.

Furthermore, using a series of high-order direct numerical simulations of driven relativistic turbulence in a hot plasma, we have explored the statistical properties of relativistic turbulent flows with average Mach numbers ranging from 0.4 to 1.7 and average Lorentz factors up to 1.7. We have found that relativistic effects enhance significantly the intermittency of the flow and affect the high-order statistics of the velocity field. Nevertheless, the low-order statistics appear to be universal, *i.e.*, independent from the Lorentz factor, and in good agreement with the classical Kolmogorov theory. In the future we plan to pursue a more systematic investigation of the properties of relativistic turbulent flows at higher resolution.

Finally, THC represents the first step towards the implementation of new and high-order methods for the accurate study of general relativistic problems, such as the inspiral and merger of binary neutron stars [26] and their relation with gamma-ray bursts [274]. We are in fact convinced that the transition towards higher-order methods is now a necessary and an inevitable step for a more realistic description of the complex phenomenology associated with these relativistic-astronomy processes.

Chapter 5

Finite-Differencing Methods: General Spacetimes

5.1 Introduction

The inspiral and merger of binary neutron stars (BNS) is one of the most promising sources of gravitational waves (GWs) for next-generation ground-based laser-interferometers detectors such as LIGO, Virgo or GEO600 [286]. GWs could be used to infer parameters such as the mass, the spin, the compactness and the tidal Love number of the neutron stars, providing very valuable informations concerning the EoS of matter at super-nuclear densities. In order to be able to 1) identify the GWs in the detector noise and, especially, 2) in order to do parameter estimations, it is however necessary to construct accurate analytical or semi-analytical GW templates. The validation and tuning of these models, in turn, must be done by matching them with the prediction of fully non-linear numerical relativity (NR) calculations, which provide the only means to completely describe the late inspiral of BNS [32, 26, 46, 169].

While very high-quality numerical-relativity waveforms are available for binary black-hole (BBH) mergers, *e.g.*, [24, 231], however see also [344], BNS simulations have been plagued by low convergence order and large phase uncertainties [31, 48]. This lead to the widespread perception, in the numerical relativity community, that matter simulations are inherently more difficult than vacuum simulations and that larger errors are to be expected. On the other hand there is no real reason why this should be the case: significant shocks are not expected to form during the inspiral in the bulk of the neutron stars and the spacetime of a BNS system is much more regular and smooth than the one of a BBH.

The goal of the work presented in this chapter is to show that it is indeed possible to obtain BNS waveforms of almost comparable quality with the ones present in the BBH literature by using higher-order numerical schemes with respect to the traditional 2nd order schemes commonly adopted for matter simulations in numerical relativity.

In particular, here, we describe our new high order, high-resolution shock-capturing, finite-differencing code: `WhiskyTHC`, which constitutes the extension to general relativity of the `THC` code. With respect to other high-order relativistic

hydrodynamics codes, such as `wham` [317] and `ECHO` [111, 71], this is the first higher than second order code that works in full general relativity, with evolved spacetime, and in three spatial dimensions.

We demonstrate the capabilities of our new code in a series of tests involving the evolution of isolated neutron stars. There we show that our code is able to yield stable and accurate evolutions of stable and unstable stars. We measure the accuracy of our code for the case of an unstable star collapsing to BH and we show that we are able to achieve third order convergence.

Finally we applied our new code to the inspiral and merger of binary neutron stars in quasi-circular orbits. We show, for the first time, higher than second order convergence for the phase and the amplitude of the gravitational wave of a binary neutron star merger. As a result, in a comparison with the old `Whisky` code, we show that our new code, when running at the same resolution and with similar computational costs as `Whisky`, is able to yield a decrease in the phase error of a factor 50. Furthermore we study the inspiral and merger of binary neutron stars initially at the “large separation” of 60 km. We consider a high-compactness model and we find only minor de-phasing, localized in the last few cycles of the inspiral, between our NR waveforms and the one from the point-particle post-Newtonian (PN) approximation Taylor-T4 [64] waveform.

The rest of this chapter is organized as follows. In Section 5.2 we give a quick summary of the numerical methods that we employed as well as a detailed description of the treatment of the fluid–vacuum interfaces, which was one of the main challenges in the application of higher-order numerical schemes to BNS simulations. In Section 5.3 we present the results obtained with our code in a series of representative tests involving the evolution of isolated neutron stars, with particular focus on the properties of the different vacuum treatments that we implemented in our code. In Section 5.4 we present some preliminary results obtained with our code in the case of BNS. Finally Section 5.5 is dedicated to the discussion of the our results as well as the conclusions.

5.2 WhiskyTHC

In this section we give an overview of `WhiskyTHC`. First of all we describe the numerical methods used in `WhiskyTHC` and secondly we give a detailed description of our treatment of fluid–vacuum interfaces which is one of the key problems one has to address in order to attain stable binary evolution, especially with high-order codes.

5.2.1 Numerical Methods

`WhiskyTHC` comes from the merger of two codes: `Whisky` [29] and `THC` [264], Chapter 4. It inherited from `THC` the use of high-order flux-vector splitting finite-differencing techniques and from `Whisky` the new module for the recovery of the primitive quantities as well as the new equation of state framework recently introduced in Galeazzi et al. [136]. This code can make use of tabulated, temperature and composition dependent equation of states, but here we are concerned only with gamma-law and polytropic evolutions. More specifically `WhiskyTHC` solves the equations of general-relativistic hydrodynamics in

conservation form (2.103) using a finite difference scheme. It employs flux reconstruction in local-characteristic variables using the MP5 scheme, for which it uses the explicit expression for the eigenvalues and left and right eigenvectors which can be found in, *e.g.*, [275].

For the spacetime we can either use the standard BSSNOK formulation (2.59), or the new CCZ4 formulation (2.64) introduced in Chapter 2, which we numerically evolve using a finite-difference code publicly available through the `Einstein Toolkit` [203, 4]. In this code all the derivatives, with the exception of the terms associated with the advection along the shift vector, for which we use a stencil upwinded by one grid point, are computed with a centered stencil. Typically all these terms are computed with a fourth order accurate scheme, but sixth and 8th order are also available. To ensure the non-linear stability of the scheme we add a fifth order Kreiss-Oliger style artificial dissipation [186]. More details on the code that we use can be found in [66] where its BSSNOK variant is described.

Our code can use a rather basic form of mesh refinement through the `Carpet` [289] AMR driver of the `Cactus` computational toolkit [150] on top of which our code is built. `Carpet` supports Berger-Oliger-style mesh refinement [45, 44] with subcycling in time and refluxing, but, at the moment, our code can only use static grid refinement, with subcycling in time, but with no refluxing. While we have plans to implement refluxing and test high-order prolongation operators in order to use dynamical grids, we also note that these features are not of fundamental importance for the study of gravitational waves from inspiraling binary neutron stars, which is the main aim of our code.

The coupling between the hydrodynamic and the spacetime solvers is done using the method of lines (MoL). For the simulations performed in this work we use either the optimal, strongly-stability preserving (SSP) third order Runge-Kutta (RK) scheme or the standard fourth order one.

5.2.2 Atmosphere Treatment

The treatment of interfaces between vacuum region and fluid regions is one of the most challenging problems in Eulerian (relativistic) hydrodynamics codes (see *e.g.*, [137, 179, 228]). Especially when studying near-equilibrium configuration, such as oscillating stars, having large density gradients close to the surface (as is the case for stiff EoS describing high-density matter) and over long timescales. The most commonly used approach to treat vacuum regions is to fill them with a low-density fluid, the “atmosphere”, such that if a fluid element is evolved to have a rest-mass density below a certain threshold, it is set to have a floor value and zero coordinate velocity [131, 29]. This approach works reasonably well for standard 2nd order codes and has been adopted by the vast majority of the relativistic-hydrodynamics codes, but it is problematic for higher-order codes [263]. The reason is that small numerical oscillations can easily result in the creation of vacuum and couple with the prescription for the floor and violate conservation, artificially creating mass, energy and momentum. As a result they are subsequently amplified, ultimately destabilizing the evolution. The situation is even more complicated for a code, such as ours, which relies on characteristic variables as they become degenerate in the low-density, low-temperature limits.

We notice that for many applications, such as the inspiral of binary neutron stars (at least up to contact), or the oscillation of single stars, the treatment of the stellar surface is one of the main challenges and the only reason why low-order, but robust shock capturing codes are commonly used. Indeed the problem of the vacuum treatment is one of the main reasons why higher order general-relativistic hydrodynamics codes have not been developed before. For this reason we think that it is instructive to address this problem in detail as we do in the following.

First of all we point out that the MP5 scheme is remarkably robust even in conjunction with the most basic atmosphere treatment that we implemented, *i.e.*, one in which no additional modification is made on the scheme at the interface between vacuum and fluid region, beside the imposition of a minimum density level (more on this below). In our preliminary tests, other schemes, such as WENO5, which do not enforce the monotonicity of the reconstruction, could not yield stable evolutions even for single stars in the Cowling approximation (*i.e.*, where the spacetime is kept fixed, while the matter is evolved).

Secondly most of the problems with the atmosphere appear in points where the surface of the star is aligned with the grid, because along these directions numerical dissipation is minimal. These artifacts, that we discuss in more detail in the next Section, are easily “fixed” with the use of extra numerical dissipation close to the surface of the star.

Keeping this in mind, we now give the details of the three different prescriptions that we developed for the treatment of the low-density regions.

Standard Atmosphere Treatment

The first prescription is what we call the “ordinary MP5” approach. It follows the lines of what is most commonly done to treat vacuum in general-relativistic hydrodynamics. First of all we choose a minimum rest-mass density ρ_{atmo} , which we take to be, typically, in the range $(10^{-7} - 10^{-9})\rho_{\text{ref}}$, ρ_{ref} being some reference density (for instance the initial maximum density). Secondly we choose a tolerance parameter, ϵ , typically 10^{-2} , chosen to avoid excessive oscillations of the fluid–vacuum interface so that points where the density falls below $(1 + \epsilon)\rho_{\text{atmo}}$, are set to atmosphere. In particular, their density is set to ρ_{atmo} , their velocity to zero and their internal energy is calculated assuming a polytropic EoS. In addition to this, we enforce a floor for the conserved energy density τ , $\tau_{\text{atmo}} \sim \rho_{\text{atmo}}\epsilon_{\text{atmo}}$.

This approach, as we show below, is already perfectly adequate for inspiraling binary neutron stars, but might have problems in the case of slowly moving vacuum–fluid interfaces aligned with the grid, especially in the case of isentropic evolutions where the surface remains sharper as no spurious heating can occur.

An Improved Atmosphere Treatment

In order to improve our atmosphere treatment, we introduced an alternative method in which we increase the level of dissipation of the scheme by switching to the component-wise Lax-Friedrichs flux split below a certain density. Typical values for this new threshold are chosen so that the first one or two grid points in the star’s interior are evolved using the Lax-Friedrichs flux split. The use

of component-wise reconstruction, as opposed to characteristic-wise, is done to avoid problems due to the degeneracy of the characteristic variables close to vacuum and to avoid polluting quantities, such as the momentum, with the numerical errors present in the internal energy (which is typically large in the immediate vicinity of the atmosphere). This is what we refer to as “MP5+LF” approach.

This latter approach is more robust, but can result in the creation of artifacts in the case in which low-density matter is ejected from the stellar surface. In this case, the fluid typically presents a rather smooth interface with vacuum, so that one would expect to be able to treat it with high accuracy. Unfortunately, as we show in the next section, if the density of the ejecta falls below the Lax-Friedrich threshold, the use of a component-wise reconstruction yields a rather oscillatory solution.

Positivity Preserving Limiter

Overall neither of the previous methods is completely satisfactory. For this reason we propose here a novel approach based on the use of the positivity preserving limiter recently proposed by [170], which is significantly simpler to implement with respect to the “classical” positivity-preserving limiters already proposed in the literature, *e.g.*, [340, 341, 342, 34].

For the sake of completeness, we give here a brief overview of the key ideas presented in [170] to which we refer to for a more complete presentation. To keep the notation simple we consider, at first, a scalar conservation law in one dimension

$$\frac{\partial u}{\partial t} + \frac{\partial f(u)}{\partial x} = 0. \quad (5.1)$$

We notice that any scheme able to guarantee the positivity of u over one first-order Euler timestep, will automatically guarantee positivity when used with any SSP time integrator as, in these schemes, the time update is always constructed as a convex combination¹ of Euler steps. For this reason we consider a discretization of (5.1) of the form

$$\frac{u_i^{n+1} - u_i^n}{\Delta^0} = \frac{f_{i-1/2} - f_{i+1/2}}{\Delta^1}. \quad (5.2)$$

If we let $\lambda := \Delta^0 / \Delta^1$, then the previous can be written as

$$\begin{aligned} u_i^{n+1} &= \frac{1}{2}(u_i^+ + u_i^-) \\ &= \frac{1}{2}[(u_i^n + 2\lambda f_{i-1/2}) + (u_i^n - 2\lambda f_{i+1/2})]. \end{aligned} \quad (5.3)$$

where $u_i^+ = u_i^n + 2\lambda f_{i-1/2}$ and $u_i^- = u_i^n - 2\lambda f_{i+1/2}$. Clearly, if u_i^+ and u_i^- are positive, so will be u_i^{n+1} . The key observation by [170] is that, if $f_{i+1/2}$ and $f_{i-1/2}$ are computed with the first-order Lax-Friedrichs scheme with $\lambda \leq \frac{1}{2a}$, a being the largest propagation speed, then $u_i^+, u_i^- \geq \min_i u_i^n$ [340].

¹A convex combination of a set of vectors, $\vec{x}_i$, is a combination of the form $\sum_i c_i \vec{x}_i$, where $0 \leq c_i \leq 1$ and $\sum_i c_i = 1$.

The idea is to modify (5.2) as

$$f_{i+1/2} = \theta f_{i+1/2}^{\text{HO}} + (1 - \theta) f_{i+1/2}^{\text{LF}}, \quad (5.4)$$

where $f_{i+1/2}^{\text{HO}}$ is the high-order flux of the original scheme, $f_{i+1/2}^{\text{LF}}$ is the flux associated with the first order Lax-Friedrichs scheme and $\theta \in [0, 1]$ is chosen to be the maximum value such that both u_i^- and u_{i+1}^+ are positive. In regions where the solution is far from vacuum $\theta = 1$ so that the high-order fluxes are used (and in particular the formal order of accuracy of the scheme remains unchanged). In regions close to vacuum it is always possible to find some $\theta \geq 0$ such that positivity is guaranteed, since for $\theta = 0$ the scheme reduces to the Lax-Friedrichs scheme which is known to be positivity preserving.

The multi-dimensional extension is done in a component-by-component fashion. For instance in three dimensions (5.3) becomes

$$\begin{aligned} u_{i,j,k}^{n+1} = & \frac{\alpha_x}{2} \left[\left(u_{i,j,k}^n + 2 \frac{\lambda_x}{\alpha_x} f_{i-1/2,j,k} \right) + \left(u_{i,j,k}^n - 2 \frac{\lambda_x}{\alpha_x} f_{i+1/2,j,k} \right) \right] + \\ & \frac{\alpha_y}{2} \left[\left(u_{i,j,k}^n + 2 \frac{\lambda_y}{\alpha_y} f_{i,j-1/2,k} \right) + \left(u_{i,j,k}^n - 2 \frac{\lambda_y}{\alpha_y} f_{i,j+1/2,k} \right) \right] + \\ & \frac{\alpha_z}{2} \left[\left(u_{i,j,k}^n + 2 \frac{\lambda_z}{\alpha_z} f_{i,j,k-1/2} \right) + \left(u_{i,j,k}^n - 2 \frac{\lambda_z}{\alpha_z} f_{i,j,k+1/2} \right) \right], \end{aligned} \quad (5.5)$$

where the α_x 's are positive constants such that $\alpha_x + \alpha_y + \alpha_z = 1$, typically chosen to be equal to $1/3$ (but see the remarks at the end of the section). The limiter at each interface is then chosen enforcing positivity of the terms in round brackets.

This approach can then be easily extended to systems of equations [170]. In particular [170] constructed a limiter able to guarantee positivity of density and pressure for the classical equations of gasdynamics, also in the case in which sources are present (but in this case a smaller timestep might be required, depending on the nature of the source terms).

In the general relativistic case it does not appear to be trivial to enforce the positivity of the pressure, especially for tabulated EoS, because of the presence of complex source terms in the energy equations. For this reason, as was the case for the atmosphere treatment, we need to enforce positivity of the pressure with the imposition of a floor on τ . On the other hand, the continuity equation

$$\partial_t \hat{D} + \partial_j [\hat{D} w^j] = 0, \quad (5.6)$$

where $\hat{D} = \sqrt{\gamma} \rho W$ and $w^j = \alpha v^j - \beta^j$, is formally equivalent to the Newtonian continuity equation with the identification

$$\hat{D} \longleftrightarrow \rho, \quad w^j \longleftrightarrow v^j, \quad (5.7)$$

thus one can construct a scheme ensuring the positivity of $\hat{D}$ by simply adopting the prescriptions used by [170] to guarantee the positivity of the density for the Newtonian Euler equations.

Some comments on the positivity preserving limiter. First of all the positivity preserving limiter is not directly a way to treat vacuum–fluid interfaces in a physically accurate way, for the simple reason that the fluid model is not

adequate to represent such transitions. A proper modeling of the stellar surface can only be done by treating it as a free boundary of the problem determined by the balance between inertial and gravitational forces on the fluid as done, for instance, in [179]. For this reason its use does not free us from having a low density fluid everywhere or from having to manually enforce $\hat{D} > \hat{D}_{\text{atmo}}$, because in some situations, for instance at the surface of a star, the high-order fluxes and the Lax-Friedrichs fluxes can differ by several orders of magnitude, so that small floating point errors can drift the conserved density, $\hat{D}$, below the minimum. This is done by simply resetting $\hat{D}$ to $\hat{D}_{\text{atmo}}$ whether $\hat{D} < \hat{D}_{\text{atmo}}$, without changing the other quantities.

What the positivity preserving limiter does, however, is to ensure local conservation of the solution up to floating point precision. In particular it prevents the scheme from extracting/losing mass from/to the atmosphere because of numerical oscillations. In particular it provides a way to setup a floor where the fluid density can be arbitrarily small and that does not require any hand tuning. In contrast, the classical atmosphere prescriptions usually work only in a limited range of ρ_{atmo} and ε as these coefficients must be tuned in order to achieve a balance between the amount of mass extracted from the atmosphere (which typically increases as ρ_{atmo} decreases) and the mass lost into it (which typically increases as ε increases). The reason is that, even in situations where we need to reset $\hat{D}$, we are guaranteed that this correction is of the order of the floating point precision (with respect to the typical densities we are actually interested in tracking).

In practice the way in which we use the positivity preserving limiter is rather simple: we fill the vacuum with a low density floor at the beginning of a simulations and we let it evolve freely, only relying on the positivity preserving limiters to ensure its well behaviour. This typically results in the creation of accretion flows onto our compact objects. However, given the low density of the floor, which we take to be $\sim 10^{-16} \rho_{\text{ref}}$ (*i.e.*, below floating point precision!) in production runs, the effects of this artificial accretion are completely negligible. Note that, to avoid issues with the decomposition of the Jacobian in eigenvalues and eigenvectors, we also switch to component-wise reconstruction below a certain density, typically $10^{-7} \rho_{\text{ref}}$, but this, in contrast to the prescription outlined in the previous section, has little dynamical effect as flows at those densities are, anyway, completely dominated by numerical effects. Moreover, as we show in the next section, even if the floor density is taken to be unnecessarily large, the use of positivity preserving limiters results in much smaller perturbations with respect to the use of a more traditional atmosphere treatment.

Finally a comment concerning the timestep limitation. For the scheme to ensure positivity in the multi-dimensional case, one must ensure $a\alpha_{\times} \frac{\Delta^0}{\Delta_{\times}} < 1$. Since $a \sim 1$ and $\alpha_{\times} = 1/D$, D being the number of dimensions, this results in a rather stringent CFL condition. In practice we find our scheme to be robust even for much larger timesteps, probably because the advection velocity in the low density regions is typically smaller than the maximum velocity and the Lax-Friedrichs scheme is actually positivity preserving even with CFL = 1 in 1D (CFL = $1/D$ in D dimensions) [342] (even though it is not possible to guarantee that u_i^+ and u_i^- in (5.3) are separately positive). In order to run with larger timestep we simply compute the value of the limiter θ assuming

$\alpha_x = 1$, as in the one-dimensional case (note this *does not mean* that we evolve using (5.5) with the α_x 's equal to one), and we set it to zero (*i.e.*, we use Lax-Friedrichs fluxes) when it is not possible to enforce the positivity of u_i^+ and u_{i+1}^- . In our numerical experiments we found this procedure to be enough to prevent negative densities from occurring to a sufficient extent and to be computationally much less expensive with respect to the approach in which $\alpha_x = 1/D$.

5.3 Single Neutron Stars

In this section we describe a series of representative tests that we performed with `WhiskyTHC` in the case of single, isolated, non-rotating neutron stars (TOVs). First of all we present the results obtained in the Cowling approximation, *i.e.*, without evolving the spacetime, for perturbed, oscillating stars. Then we proceed analyzing the results in the case of linear oscillations of stable stars in full-GR. Finally we show the results obtained for the evolution of unstable stars: both for the migration and for the collapse to BH cases. The focus of our discussion is mostly on the effects of the different prescriptions for the treatment of the atmosphere. We denote the basic treatment (*i.e.*, “ordinary MP5”), the enhanced treatment (*i.e.*, with extra dissipation on the surface) and the positivity preserving treatments as “MP5”, “MP5+LF” and “MP5+PP”, respectively.

5.3.1 Linear Oscillations: Cowling Approximation

The first test that we consider is the long-term evolution of a perturbed, isolated, non-rotating, neutron star in the Cowling approximation. The goal of this test is to assess the impact on the accuracy of the three different atmosphere treatments over long timescales. We consider a model described by the polytropic EoS with $K = 100$ and $\Gamma = 2$. The initial central density is $\rho_c(0) = 1.28 \times 10^{-3} M_\odot^{-2}$, yielding a model with an ADM mass of $1.4 M_\odot$. The initial velocity is perturbed with the injection of an exact eigenfunction (in the Cowling approximation). The maximum amplitude of the perturbation is $|v'| \simeq 0.024$ and the initial perturbation is ingoing.

We evolve this model for $10\,000 M_\odot$, *i.e.*, $\simeq 130$ dynamical timescales, using our different prescriptions for the atmosphere. Our fiducial resolution is $h = 0.2 M_\odot$ so that the radius of the star is covered with $\simeq 45$ grid points. In order to make a fair comparison, we use the same atmosphere threshold for all the methods, $\rho_{\text{atmo}} = 10^{-10} M_\odot^{-2}$, and we evolve all the models with the third order SSP-RK3 with CFL = 0.4. The gravitational source terms are computed using sixth order finite-differencing. Finally the evolution is computed only in the octant $x, y, z \geq 0$ and we assumed reflection symmetry across the xy, xz and yz planes.

The evolution of the central density, ρ_c , is shown in Figure 5.1. In particular, in order to highlight the secular trend of the data, we show a moving average of ρ_c defined as

$$\langle \rho_c \rangle(t) := \frac{1}{2T} \int_{t-T}^{t+T} \rho_c(s) ds, \quad T \leq t \leq T_{\text{final}} - T, \quad (5.8)$$




Figure 5.1: Normalized central density for the perturbed TOV in the Cowling approximation and for different atmosphere prescriptions.

where $T = 5/f_F$, f_F being the frequency of the F-mode from linear perturbation theory. All the different prescriptions yield very similar evolutions of the central density which presents a series of slowly damped oscillations. The pulsation frequency agrees, to within the nominal error of the discrete Fourier transform, with the one expected from linear perturbation theory in the Cowling approximation. The power-spectrum also shows small contributions from higher-order overtones (*i.e.*, more than a factor ten smaller than the F-mode) as well as an even smaller non-linear component at integer multiples of the F-mode frequency. In the case of the MP5+LF prescription, we verified that the non-linear component decreases with decreasing perturbation amplitudes and that it is not distinguishable from the background for perturbation amplitudes $\simeq 2 \times 10^{-3}$. The ordinary MP5 prescription also shows a small secular increase in the central density. Apart from this, all the schemes appear to be able to yield very clean oscillations.

The difference between the different schemes can be better appreciated by looking at Figure 5.2 where we show the evolution of the total rest mass,

$$M_b(t) = \int_{\Sigma_t} \rho \sqrt{\gamma} d^3x, \quad (5.9)$$

for the different models. Overall the mass conservation is at acceptable levels for all the methods, but the ordinary MP5 prescription is clearly the one with the largest error, as it shows larger variations with respect to the other schemes. This, in turn, is responsible for the secular drift mentioned earlier (an increase in the total mass leads to an increase in the central density). The Lax-Friedrichs flux-switch at the surface, instead, results in a steady loss of matter which is slowly diffused into the atmosphere, while the MP5+PP approach yields a steady increase in the mass because of the accretion of the low density floor which is continuously “injected” from the outer boundary (we simply fix the density in the outer ghost regions to its initial value).

The reason for the bad behaviour of the standard MP5 prescription is, as




Figure 5.2: Normalized total rest mass variations for the perturbed TOV in the Cowling approximation and for different atmosphere prescriptions.

anticipated in the previous section, the fact that it lacks a sufficient amount of numerical dissipation in the case of surfaces aligned with the grid and especially for polytropic evolutions, such as the ones we are showing here. This is clearly seen in Figure 5.3 where we show a two-dimensional cut of the density, in $\log_{10}$ scale, at a representative time during the evolution. One can clearly see the appearance of “jets” of low-density matter ($\rho \sim 10^{-5} M_{\odot}^{-2}$) aligned with the coordinate directions. These “jets” are launched at seemingly random times from the surface of the star, when the numerical errors “extract” from the atmosphere a large enough amount of mass. What happens is that the numerical oscillations create an imbalance at the surface of the star: the excess density coming from the atmosphere generates a pressure which is only balanced by the “potential barrier” at the surface of the star given by the double threshold on the density floor. As soon as the pressure is large enough, part of the matter is ejected in one of these streams. Counter-intuitively this process results in the increase of the total mass of the star because only part of this extra matter is actually lost from the outer boundary. In contrast, we can see that, with the addition of extra numerical dissipation at the surface of the star, these artifacts are completely suppressed, as shown for the MP5+LF case. This happens partly because dissipation prevents the scheme from extracting too much matter out of the atmosphere and partly because it diffuses the numerical errors back into the floor. Finally the positivity-preserving evolution does not show any kind of numerical ejecta out of the star’s surface because of its conservative nature. On the other hand it is affected by the accumulation of matter at the fluid-vacuum interface. As commented before, this accumulation can be greatly reduced by lowering the floor density and it is also somewhat less severe for the ideal-gas case, where the floor accretion is regulated by the thermal pressure.

The differences between the various methods are even more evident if we look at sensitive quantities such as the momentum in the radial direction. At the initial time it has a profile given by the eigenfunction of the F-mode in the Cowling approximation. In the linear regime one would expect the




Figure 5.3: Two dimensional cut of the $\log_{10}$ of the rest-mass density at time $t = 800 M_\odot$ for the perturbed TOV in the Cowling approximation and for different atmosphere prescriptions. The insets at the bottom of the plots show the one dimensional cuts along the x axis.




Figure 5.4: One dimensional cut along the x axis showing the x component of the conserved momentum, S_x , at the time $t = 4996 M_\odot$, for the perturbed TOV in the Cowling approximation and for different atmosphere prescriptions. The different symbols show the numerical solution as obtained with our code and with the different vacuum treatments, while the thick black line shows the exact eigenfunction from linear perturbation theory.

momentum to simply oscillate with the F-mode frequency. On the other hand, in a simulation, because of numerical errors, the profile of the eigenfunction is gradually lost. This is shown in Figure 5.4 where we plot the x component of the momentum, S_x , along the x axis at a representative time, $t = 4996 M_\odot$. At this particular time both the MP5 and the MP5+LF schemes have accumulated so much error that the profile of the eigenfunction is completely distorted. On the other hand the evolution using the positivity-preserving limiter still shows good agreement with the exact solution. Clearly the precision with which we recover the eigenfunction is resolution dependent and degrades over time also for the MP5+PP scheme. Nevertheless this figure clearly illustrates 1) how large the influence of the atmosphere is in this kind of simulation where nearly equilibrium configurations are evolved for long time, and 2) how small the perturbation due to the continuous, artificial, accretion is when we use our positivity preserving prescription, even when the floor density is rather high.

5.3.2 Linear Oscillations: Full-GR

The second test that we present is the evolution of a stable, non-rotating, star in full-GR. The goal of this test is to check the stability of the three different floor prescriptions in a fully general-relativistic setting. The model that we consider here is the same as the one described in Section 5.3.1, with the difference that we do not apply any perturbation to the initial data and we let it evolve under the sole effects of the numerical truncation error.

This test is performed using a grid covering $0 \leq x, y, z \leq 80 M_\odot$ and employing three refinement levels, with the finest one covering the star entirely and




Figure 5.5: *Top panel*: evolution of the normalized central density for the oscillating TOV in full-GR and for different atmosphere prescriptions. *Bottom panel*: power spectral density of the central density, normalized to have maximum value 1. In the calculation of the PSD we exclude the first $300 M_\odot$ of the evolution, to avoid contamination from the initial spike.

having a resolution $h = 0.2 M_\odot$. The spacetime is evolved using fourth order finite differencing and the CCZ4 formulation of the Einstein equations. Finally we assume reflection symmetry across the xy , xz and yz planes. We evolve the model with different atmosphere prescriptions and time integrators and we choose, for each of them, the values of ρ_{atmo} and ε giving the best results in order to showcase the capabilities of each method. We note, however, that, due to the high computational costs, we did not perform an extensive tuning of these parameters and we cannot exclude that another combination of parameters would give better results. The parameters that we use are summarized in Table 5.1.

The evolution of the central density for the different methods is shown in the top panel of Figure 5.5. First of all we notice that the ordinary MP5 prescription shows violent oscillations and a large secular growth. We evolve this model up to time $t \approx 1100 M_\odot$ where it has deviations from the initial

Table 5.1: Numerical parameters used for the oscillating TOV test in full-GR.

Model	Time integrator	CFL	$\rho_{\text{atmo}} [M_\odot^{-2}]$	ε
MP5	RK4	0.2	10^{-10}	1
MP5+LF	RK4	0.2	10^{-10}	0.01
MP5+PP	RK3	0.2	10^{-19}	—

density of the order of 1.5 %. We point out that we also tried to evolve it using the same prescription as the one used for the test in the Cowling approximation, but, in that case we obtained even larger oscillations and a more pronounced secular growth leading to an increase of about 4 % in the central density at time $t = 1000 M_{\odot}$.

As with the previous test, the simple addition of extra numerical dissipation at the star's surface seems to cure the most severe problems with the MP5 evolution. Indeed the MP5+LF scheme shows much smaller oscillations and only a weak secular trend (see also Figure 5.1).

The MP5+PP scheme yields very small oscillations and an almost zero trend in the central density. We point out, however, that in our preliminary tests the MP5+PP scheme showed a sudden increase in the oscillation amplitude and in the secular drift after time $t \gtrsim 3000 M_{\odot}$, which were at levels comparable to the MP5+LF ones. The reason for this behaviour is to be found in the prolongation operators used in our AMR setup as well as in the lack of refluxing in our code, which were resulting in spurious violations in the mass conservation at the mesh-refinement boundaries. In order to avoid this problem we have disabled the prolongation of the hydrodynamic variables, thus partially “decoupling” the various refinement levels. In the long term we plan to improve the AMR capabilities of our code to avoid these pathologies, but we point out that these are not of primary concern for the main purpose of our code, which is to compute gravitational waveforms from compact binaries.

The PSD of the central density is shown in the bottom panel of Figure 5.5. There we show the PSD normalized to have maximum amplitude of 1 for the different numerical schemes. The spectra are computed using the central density from the time $t \geq 300$, to remove the dependence from the relaxation of the initial data. In order to compute the spectra we also de-trend the data by removing its linear fit. Clearly the ordinary MP5 scheme has a more noisy spectrum, partly because of the shorter integration time. Apart from that all the three methods show spectra which are peaked at the frequencies corresponding to the F-mode and to the first overtone, H1, as computed from linear perturbation theory.

5.3.3 Non-linear Oscillations: the Migration Test

The third test that we discuss is the study of the large, non-linear, oscillations of a TOV migrating from the unstable branch of solutions to the stable one. This is a commonly adopted test for numerical-relativity codes, *e.g.*, [131, 28, 29, 98, 320], and has been studied in detail by [198, 265]. Here we consider a model initially described by a polytropic EoS with $\Gamma = 2$ and $K = 100$ and with central density $\rho_c = 0.007 M_{\odot}^{-1}$, yielding an ADM mass of $\approx 1.49 M_{\odot}$. The migration is triggered with the use of an outgoing velocity perturbation of the form $v^r = A r$, r being the areal radius, where A is chosen so that the maximum perturbation velocity is 0.01. The evolution is performed with a gamma-law equation of state to allow for shock heating. We point out that we do not solve the constraints equations after the application of the initial perturbation, but we rely on the constraint damping nature of CCZ4 to bring the evolution back to the constraint “hypersurface” as done in [178].

The grid setup is identical to the one described in Section 5.3.2, with the only difference that we enlarge the area covered by the finest grid to ensure that the




Figure 5.6: Evolution of the central rest-mass density, normalized to its initial value, for the TOV migration test and for different numerical schemes. The vertical dashed line marks the point in time shown in Figure 5.7. The horizontal dotted line marks the value of the central density for the equilibrium model on the stable branch corresponding to the unstable model evolved in this test and to which the solution is expected to relax.

star is completely covered by it even when it is fully expanded. Here, again, we use the same atmosphere prescription for all the schemes with $\rho_{\text{atmo}} = 10^{-12} M_\odot^{-2}$ and we evolve all the models using the SSP RK3 scheme with a CFL of 0.1. Finally the spacetime is evolved using sixth order finite-differencing and with the addition of a fifth order Kreiss-Oliger dissipation.

The evolution of the system is summarized by Figure 5.6 where we show the evolution of the central rest-mass density, normalized to its initial value, for our three different schemes. As can be seen from the figure, the star undergoes a sequence of violent expansion, contraction cycles after it has migrated to a model on the stable branch of equilibria having the same baryonic mass. During the contraction phase, shocks are formed and part of the shock heated matter is ejected with large velocities from the central object. All of the methods are perfectly adequate for this test and only minimal differences appear between the MP5+LF scheme and the other two in the amplitude of the first peak.

The difference between the various atmosphere prescriptions is better appreciated by looking at Figure 5.7, where we show a two-dimensional cut of the $\log_{10}$ of the density at the time when the matter ejected at the first bounce reaches the grid boundaries (this is indicated with a vertical line in Figure 5.6). As can be seen from the figure, both MP5 and MP5+PP are able to capture the dynamics of the low-density ejecta without introducing large numerical oscillations or excessive deviations from spherical symmetry. The front of the ejecta is reasonably well captured even if it has crossed two mesh-refinement boundaries, where our code cannot currently ensure mass-conservation and hence the right propagation speed for shocks. On the other hand, the MP5+LF prescrip-




Figure 5.7: Two-dimensional cuts of the $\log_{10}$ of the rest-mass density for the TOV migration test and for different numerical schemes at time $t = 268.8 M_{\odot}$. The insets show the one-dimensional cuts along the x axis.




Figure 5.8: *Top panel:* evolution of the normalized central lapse for the TOV collapse test and for different resolutions. *Bottom panel:* evolution of the normalized central density for the TOV collapse test and for different resolutions.

tion seems to be completely inadequate to treat the ejecta as it exhibits large numerical oscillations/fragmentation. The reason is probably due to our choice of avoiding the reconstruction in characteristic variables at low densities, when it is well known that component-by-component reconstruction typically results in oscillatory solutions when used in conjunction with high-order schemes.

5.3.4 Gravitational Collapse to Black-Hole

The final test that we are going to describe involving the evolution of isolated neutron star is the gravitational collapse of a TOV to a black hole. This is another commonly adopted benchmark for general-relativistic hydrodynamics codes and has been studied in great detail in [27, 319]. The model that we consider here is initially described by a polytropic EoS with $\Gamma = 2$ and $K = 100$ and has an initial central density of $0.008 M_\odot^{-2}$, yielding an ADM mass of $\approx 1.43 M_\odot$. The collapse is triggered with the addition of a velocity perturbation with the same characteristics of the one used in the migration test, but with opposite sign, *i.e.*, an ingoing perturbation. Finally the model is evolved using an ideal-gas EoS to allow for thermal effects.

In the case of the collapse the influence of the atmosphere is negligible, so here we consider only evolutions performed with the MP5+LF prescription. As in the migration test, our computational domain covers $0 \leq x, y, z \leq 80 M_\odot$ and we assume reflection symmetry across the xy , xz and yz planes. We employ three different refinement levels with the finest one covering the star entirely and we study the convergence of the code as we vary the resolution. In particular we considered six different resolutions having grid spacing (in the finest refinement level) of $h = 0.2, 0.16, 0.13333, 0.11429, 0.1$ respectively. We also perform a higher-resolution run, with $h = 0.08$, which we evolve only up to time $\approx 45 M_\odot$ and that we use as a reference solution to measure the self-convergence




Figure 5.9: Estimated L^1 -norm of the error of the lapse function on the xy plane on the finest refinement level for the TOV collapse test and for different resolutions. The errorbars show the excursion between the maximum and minimum normalized error in the time interval. The solid black line shows the curve for third order of convergence.

of the code. We adopt sixth order finite-differencing for the spacetime, which is evolved using the CCZ4 formulation, with fifth order Kreiss-Oliger artificial dissipation on the metric variables. We use the fourth order Runge-Kutta scheme as time integrator, so that our scheme is formally fourth order (fifth order in space and fourth in time). Finally, in order to avoid excessive oscillations in the matter fields inside the forming black hole, we artificially evacuate the regions where $\alpha < 0.1$ by adding an artificial damping term in the sources of the hydrodynamic variables as in [13]. We also switch to the component-wise reconstruction with Lax-Friedrichs split in regions where $\alpha < 0.2$.

The evolution of the lapse and the rest-mass density at the center and for the different resolutions are shown in Figure 5.8. As the star collapses the central density rapidly increases and the lapse function approaches zero. At time $t \approx 40 M_\odot$ the lapse at the center becomes smaller than 0.1 and the density starts to be dissipated. Finally, after a small re-bounce, the lapse settles and the evolution reaches quasi-stationarity. The exact behaviour of the lapse is determined by the way in which we evacuate the fluid as, in our preliminary tests, we found that, in simulations without matter damping, the lapse showed a smaller bounce, at the moment of the collapse, but a more irregular evolution at intermediate times.

In order to estimate the convergence rate of our code we use the highest resolution simulation, $h = 0.08$, as a reference solution and we compute an estimate of the error of a given physical quantity, ϕ , at the time t_n as

$$E_h^n = \frac{1}{N} \sum_x |\phi_h(x, t_n) - \phi_{h=0.08}(x, t_n)|, \quad (5.10)$$

where the sum is taken over the common grid points between the resolution

h and the highest resolution run on the xy plane. E_h^n is then computed using 14 equally spaced data dumps in the interval $[3.2, 44.8] M_\odot$ (including the first and the last time). In order to have an absolute measure of the relative errors between the different resolutions over the whole time-interval, we normalize the error estimates with respect to the deviations as measured between the lowest resolution simulation and the reference one, *i.e.*, $E_h^n/E_{0.2}^n$. Finally we take as relative error the average in time of the normalized error estimates and we use the maximum and minimum relative errors (between the different times) as a measure of the uncertainty of this procedure.

The results obtained for the lapse function are shown in Figure 5.9. Also shown, as a solid black line, is the curve for third order convergence. As can be seen from our figure our data is consistent with third order convergence for $h \gtrsim 0.16$. As commented before, our code is formally fourth order convergent in time and fifth order in space, on the other hand, based on our previous experience with MP5 [264], we argue that the observed third order convergence is most probably related to the fact that high-order shock-capturing codes are able to converge at their nominal order only at extremely high resolutions because their accuracy is typically spoiled by the activation of the flattening procedure close to under resolved features of the solution [264].

5.4 Binary Neutron Stars

In this section we present the results obtained for the inspiral and merger of binary neutron stars in quasi circular orbit. We consider two different models: one having an initial small separation of 45 km, which we discuss in Section 5.4.1, and one having a larger separation of 60 km, which is discussed in Section 5.4.2. We use the first binary, which we can run with relatively small computational costs, to explore the different atmosphere prescriptions and make a detailed comparison between the results obtained with our code and the ones obtained for the original Whisky code. The second binary, instead, is used to assess the accuracy of our code for the production of long waveforms.

We recall here that the Whisky code is a second-order finite-volume code with high-order primitive reconstruction and implements several different approximate Riemann solvers. For the runs presented here we make use of the PPM reconstruction [95] and of the HLLC Riemann solver [162, 125].

The initial data we consider describes two neutron stars in quasi-circular orbit. It is computed in the conformally-flat approximation using the LORENE pseudo-spectral code [157] and has been made publicly available by the Meudon group [3]. The EoS assumed for the initial data is polytropic with $K = 123.56$ and $\Gamma = 2$, while the evolution is performed using the ideal-gas EoS to allow for thermal effects in the merger phase. The details of the models we consider are listed in Table 5.2. Here we point out that these binaries consist of stars with a rather high baryonic mass, $M_b \simeq 1.9 M_\odot$, close to the maximum mass allowed by the EoS for non-rotating models, $M_{b,\max} = 2 M_\odot$, and having high compactness, $C = M_\infty/R_\infty = 0.18002$, M_∞ being the gravitational mass of each of the two stars when at infinite separation and R_∞ the corresponding areal radius. We remark that binaries with a similar compactness have been already considered by Hotokezaka et al. [169] where it was found that high-compactness binaries are much more challenging to evolve accurately with respect to low-compactness

ones.

5.4.1 Small Separation

Here we discuss the results obtained from eight different evolutions of the model A described in Table 5.2. All of these runs are performed on a grid covering $0 < x, z \leq 512 M_\odot$, $-512 M_\odot \leq y \leq 512 M_\odot$, where we assume reflection symmetry across the xy plane and π symmetry across the yz plane. The grid employs several refinement levels, 5 or 6 depending on the run, with the finest refinement levels covering both stars, *i.e.*, our grid is static and we have no moving boxes. A summary of the main numerical parameters can be found in Table 5.3.

Finally we evolve this model using the CCZ4 formulation with damping constants $\kappa_1 = 0.036$ and $\kappa_2 = 0$, with $\kappa_3 = 1/2$ and with beta-driver $\eta = 0.71$. The spacetime is evolved using fourth order finite-differencing and with fifth order Kreiss-Oliger artificial dissipation. The evolutions are performed *without* resetting the shift to zero at the beginning of the simulation, which is known to yield a more oscillatory behaviour in the coordinates [25]. There is no particular reason for this choice: the gauges are only chosen so as to be able to leverage, in the debugging stage, on the comparison with previously existing Whisky simulations that were performed, with a different grid setup, by [13].

Since here our focus is mostly on the accuracy of the methods for the purpose of studying the gravitational radiation from compact binaries, we consider the accuracy of the code by mainly looking at the $\ell = 2, m = 2$ mode of the Weyl scalar Ψ_4 extracted at the fixed coordinate radius of $r = 450 M_\odot$ ($\approx 130 M_{\text{ADM}}$). We do not attempt to extrapolate Ψ_4 in radius or compute the strain as this involves other uncertainties [63, 269, 270, 268].

The dynamics of the inspiral and merger of BNS has been described many times and in great detail in the literature, *e.g.*, [25], for this reason we do not give a very in-depth discussion of it here. We only mention that the two neutron stars inspiral for about 2.5 orbits, touch and quickly merge into a single black-hole. For this particular model no significant disk is left behind. The gravitational-wave signal consists of about 6 cycles up to merger, followed by the black-hole ring-down.

The GW signal is shown in Figure 5.10, where we plot the 22 mode of Ψ_4 as extracted at $r = 450 M_\odot$ and as a function of the retarded time $t - r_*$, where $r_* = r + 2M_{\text{ADM}} \log(r/(2M_{\text{ADM}}) - 1)$. In particular we show the results

Table 5.2: Summary of the considered BNS models. For each model we report the total baryonic mass, M_b , the ADM mass, M_{ADM} , the initial separation, r and the initial orbital frequency f_{orb} , the gravitational mass of each of the two stars when at infinite separation, M_∞ , as well as the compactness, $C = M_\infty/R_\infty$, where R_∞ is the areal radius of the two stars when at infinite separation.

Model	$M_b [M_\odot]$	$M_{\text{ADM}} [M_\odot]$	$r [\text{km}]$	$f_{\text{orb}} [\text{Hz}]$	$M_\infty [M_\odot]$	C
A	3.8017	3.44537	45	309.702	1.7428	0.18002
B	3.8017	3.45366	60	208.431	1.7428	0.18002

Table 5.3: Summary of the main numerical parameters used in the numerical simulations presented here. For each run we give the name of the code used to perform it, *WhiskyTHC* or *Whisky*, the numerical method employed, the time integrator used for the method of lines, MoL, the CFL, the number of refinement levels of the grid, N_{refl} , and the grid spacing in the finest refinement level, h .

Run	Code	Method	MoL	CFL	N_{refl}	h
A.MP5.H1	WhiskyTHC	MP5	RK4	0.30	5	0.40000
A.MP5.H2	WhiskyTHC	MP5	RK4	0.30	6	0.20000
A.MP5.H4	WhiskyTHC	MP5	RK4	0.30	6	0.12800
A.MP5+LF.H2	WhiskyTHC	MP5+LF	RK4	0.30	6	0.20000
A.MP5+PP.H2	WhiskyTHC	MP5+PP	RK3	0.15	6	0.20000
A.PPM.H2	Whisky	PPM	RK4	0.30	6	0.20000
A.PPM.H3	Whisky	PPM	RK4	0.30	6	0.13333
A.PPM.H5	Whisky	PPM	RK4	0.30	6	0.10000
B.MP5.H1	WhiskyTHC	MP5	RK4	0.30	6	0.25000
B.MP5.H2	WhiskyTHC	MP5	RK4	0.30	6	0.20000
B.MP5.H3	WhiskyTHC	MP5	RK4	0.30	6	0.14545
B.PPM.H2	Whisky	PPM	RK4	0.30	6	0.20000




Figure 5.10: Real part of the 22 mode of Ψ_4 extracted at $r = 450 M_\odot$ for model A and for different atmosphere prescriptions using *WhiskyTHC*.

obtained for runs A.MP5.H2, A.MP5+LF.H2 and A.MP5+PP.H2. As can be seen from the plot, all our three different atmosphere prescriptions give identical results during the inspiral and yield very marginal differences in the merger phase. This provides an important result and suggests that the treatment of the neutron star surface is not a leading source of error in binary neutron star simulations, as far as the inspiral GW signal is concerned.

The particular choice of time-integrator, between SSP-RK3 and the standard RK4, also seems not to be of fundamental importance here, with the error most likely dominated by the spatial discretization. On the other hand we point out that the timestep used for A.MP5+PP.H2 is half of the one used in the other runs in order to use the proper positivity-preserving limiter, thus introducing a possible systematic difference.

The GW signal for the other runs of model A are shown in Figure 5.11. In the top panel we show the results obtained with the standard `Whisky` code and in the bottom one the ones obtained with `WhiskyTHC`. The first aspect to notice when comparing the two panels is that, when using a second order code, the phase difference between the GWs at different resolutions is significant. We can observe a difference between the low and the high resolution of about $\simeq 2$ radians at $t - r_* = 1350 M_\odot$, r being the extraction radius, $450 M_\odot$. In contrast the waveforms obtained with `WhiskyTHC` show a significantly smaller de-phasing: the difference between the low and the high resolution is about $\simeq 0.6$ radians at $t - r_* = 1350 M_\odot$, which is a factor four smaller than the one shown by `Whisky`, even though the `WhiskyTHC` runs span a wider range of resolutions. The difference in phase between the high and the medium resolution of `WhiskyTHC` at $t - r_* = 1350 M_\odot$ is as small as $\simeq 0.06$ radians.

The second interesting aspect is that, for this particular model and with `Whisky`, the merger happens *earlier* as we increase the resolution. This is the opposite of what it is observed in other, less compact binaries, *e.g.*, [46], where tidal effects have been found to be amplified at lower resolution.

We study the convergence of the waveforms by looking at the de-phasing between different resolutions. For each run we compute the phase, ϕ , of the 22 mode of Ψ_4 from its definition,

$$(\Psi_4)_{22} = A e^{i\phi}, \quad (5.11)$$

over the time interval $t - r_* \in [200, 1400] M_\odot$. Notice that we *do not* align the waveforms in any way, but we exclude from the calculation the first burst of gravitational radiation, due to the initial “junk” radiation present in the initial data. We also exclude the last part of the merger phase (where we expect large errors due to the presence of shock waves) and ringdown since, here, we are only concerned with the inspiral.

The results are shown in Figure 5.12 where we plot the convergence for both `Whisky` and `WhiskyTHC`. For `Whisky` we find a convergence order of $\simeq 1.7$, close to the 1.8 measured by [31], up to time $t - r_* \simeq 1200 M_\odot$, *i.e.*, about at the contact time $t \simeq 1200 M_\odot$. For `WhiskyTHC`, instead, we find a convergence order of $\simeq 2.8$ up to time $t - r_* \simeq 1300 M_\odot$. We point out that this is the first time that higher-than-second order of convergence has been shown for binary neutron star mergers.

As a consequence of having a higher convergence order, `WhiskyTHC` is also significantly more accurate. This is shown in Figure 5.13 where we compare the




Figure 5.11: Real part of the 22 mode of Ψ_4 extracted at $r = 450 M_\odot$ for model A at different resolutions and using two different codes: the original Whisky code and the new WhiskyTHC code.




Figure 5.12: Accumulated de-phasing between different resolutions for model A as evolved with Whisky and WhiskyTHC. In both plots we show the de-phasing between the medium and low resolution (blue lines), between the high and medium resolution (green lines) as well as the rescaled de-phasing between high and medium resolution (red lines) computed assuming convergence order of 1.7 and 2.8 for Whisky and WhiskyTHC respectively.




Figure 5.13: Estimated phase uncertainty due to finite resolution for `Whisky` and `WhiskyTHC` in the case of model A at the common resolution of $h = 0.2 M_{\odot}$. The blue region shows the phase differences between `Whisky` and its Richardson extrapolated waveforms computed assuming convergence order between 1.5 and 1.9. The green line region shows the phase differences between `WhiskyTHC` and its Richardson extrapolated waveforms computed assuming convergence order between 2.6 and 3.0. Finally the red region shows the same data as the green one, but rescaled by a factor 50.




Figure 5.14: One-dimensional cuts along the x axis of rest-mass density (dashed) and y component of the velocity, v^y , (solid) lines for **Whisky** (blue) and **WhiskyTHC** (red) for the model A. The velocity scale is shown on the left, while the density scale is shown on the right. The data is taken after ≈ 1 orbit, at the approximate moment when the center of one of the two stars crosses the x axis. Notice that we correct for the de-phasing accumulated between **Whisky** and **WhiskyTHC** by taking the data at slightly different times: $t = 625.92 M_\odot$ and $t = 622.08 M_\odot$ respectively.

estimated phase error between the runs A.MP5.H2, A.PPM.H2 and the Richardson extrapolated phase from **WhiskyTHC** and **Whisky** respectively. We give a rough estimate of the uncertainty in this procedure by performing two different extrapolations for each code, varying the convergence order by ± 0.2 with respect to the estimated one. The resulting range of phase errors are shown as shaded regions in the figure. A.PPM.H2 has an uncertainty in phase which is almost equivalent to one GW cycle, *i.e.*, of the order of $\sim 7\%$ of the entire accumulated phase. At the same time, **WhiskyTHC** has an error which, at the same resolution and for comparable computational costs, is ~ 50 times smaller than **Whisky**. We also point out that we tried computing the estimated phase error for **WhiskyTHC** using **Whisky** Richardson extrapolated data, constructed assuming convergence order of 1.7. In this case we found an even smaller estimated phase error, but with an uncertainty, measured by varying the order in the extrapolation by ± 0.1 , of more than 100 %.

We should stress that this error estimate only reflects the numerical truncation error. Other systematic errors and, in particular, finite extraction radius effects and inaccuracies in the initial data, are also present and might be relevant (especially for **WhiskyTHC**). On the other hand, here we are interested only in evaluating the accuracy of the two numerical methods.

A physical intuition of why **Whisky** has such a low accuracy with respect to **WhiskyTHC** can be gained by looking at Figure 5.14. There, we show one-dimensional cuts of the rest-mass density (solid lines) and of the y component of the velocity, v^y , (dashed lines) for runs A.MP5.H2 and A.PPM.H2 along the x axis.

The data is taken at the approximate time when the first orbit is completed and the centers of the two stars are aligned along the axis. Note that, since this happens at two different times for A.MP5.H2 and A.PPM.H2, the lines in the figure show data taken at two different (but equivalent) coordinate times. One should focus on the velocity profile inside the stars, as the velocity in the low density atmosphere around the star is expected to be dominated by unphysical numerical effects as neither of the two codes is able to provide a reasonable description of the stellar surface. The velocity in the interior, instead, is directly related to the orbital motion of the two stars. As can be seen from the figure, *Whisky* is not able to transport the velocity profile of the star correctly. The reason is that its dissipation is slowly flattening both the density and momentum profile and, since this flattening proceeds at different rates for the different fields, it results in a distortion of the profile of the physical quantities. This distortion, in turn, results in a small but artificial deformation of the stars and it is likely to be the leading source of error in the phase.

5.4.2 Large Separation

In this section we present the application of our code to the simulation of the inspiral and merger of binary neutron stars from large separation, *i.e.*, model B in Table 5.2.

The grid setup that we use here is the same as the one used in the previous section, with the difference that we extend the finest grid so that it covers both of the stars during the whole evolution. One important difference is that we evolve these binaries using the BSSNOK formulation of the Einstein equations, instead of CCZ4. The choice is motivated by the fact that we want to show the potential of our code using a setup as similar as possible to the one most commonly adopted in the literature. Furthermore the beta-driver coefficient, η , is set to 0.3 and we reset the shift to zero in the initial data as suggested by [25].

An overview of the dynamics can be seen in Figure 5.15 where we show six representative snapshots of the $\log_{10}$ of the rest-mass density for the run B.MP5.H2. The inspiral phase lasts for about $\simeq 13$ GW cycles, $\simeq 6.5$ orbits, and the two stars arrive at contact around time $t \sim 5000 M_{\odot}$ ($M\omega \simeq 0.11$). The total gravitational signal up to the merger, defined as the time where the gravitational waves amplitude has a first maximum, is of 16 cycles. Notice that contact happens before the bare contact angular frequency [102],

$$0.15276 = M\omega_{\text{contact}} := 2 C^{3/2}, \quad \omega := \dot{\phi}, \quad (5.12)$$

is reached. This is in any case expected because this approximation of the contact frequency does not take tidal deformations into account. The $M\omega_{\text{contact}}$ angular frequency is reached, in our highest resolution run, B.MP5.H3, at time $t \simeq 5200 M_{\odot}$. The merged object survives for about ~ 3 GW cycles and it quickly collapses to a BH. As in the small separation case, no significant accretion disk is left and most of the matter is quickly accreted into the BH. As can be seen from the figure, the stars are surrounded by a low density atmosphere which is the byproduct of the way in which we treat fluid–vacuum interfaces (and that we discussed in detail in the previous sections), nevertheless the profile of the two stars appear to be still sharply preserved. This should be contrasted with Figure 5.16 where we show the evolution obtained with the *Whisky* code at the




Figure 5.15: Two-dimensional visualization of the $\log_{10}$ of the rest-mass density for the run B.MP5.H2.




Figure 5.16: Two-dimensional visualization of the $\log_{10}$ of the rest-mass density for the run B.PPM.H2.




Figure 5.17: Real part of the 22 mode of Ψ_4 extracted at radius $r = 450 M_\odot$ for model *B*. We show a comparison between the highest resolution run, B.MP5.H3, and the predictions from post-Newtonian theory.

same resolution. The second order *Whisky* code yields an evolution where the two stars are significantly more smeared and their boundaries, which we take to be the yellow region where $\rho \approx 10^{-5} M_\odot^{-1}$, much less defined. Furthermore the merger time is significantly ($\approx 20\%$) delayed with respect to all of the *WhiskyTHC* runs.

The gravitational waveform is shown in Figure 5.17. There we plot the 22 mode of Ψ_4 as extracted at radius $r = 450 M_\odot$ for the run B.MP5.H3. As a reference we also show the analytical predictions from the PN model discussed in [64] and kindly provided by Francesco Pannarale. We align the two waveforms by fixing the position of the third maxima in the gravitational waveform, *i.e.*, after ~ 1 orbit, when the initial junk radiation has been propagated away and the two stars have settled to a quasi-circular inspiral. As can be seen from the figure, the two waveforms show good agreement in the earliest part of the late-inspiral, which is also magnified in the inset. Afterwards they start de-phasing as the binaries approach the contact point and the subsequent merger.

Before going into more details concerning the comparison between the numerical results and the analytical ones, we discuss the quality of the numerical relativity data. The first thing to point out is that the merger time, defined as the time where the gravitational wave amplitude has a first maximum, is very close between the different runs. This can be seen from Figure 5.18, where we plot the amplitude of the 22 mode of Ψ_4 , as extracted at radius $r = 450 M_\odot$, and as a function of the retarded time $t - r_*$. As we change the resolution, from low to high, by a factor 1.7, the differences in the merger time are only of the order of $\approx 2.5\%$. In comparison the results reported by [169] show, for a model with the same compactness, changes of the order of $\approx 20\%$ when changing the resolution by a factor 1.4, even though their highest resolution is about 35% higher than our highest one (this roughly corresponds to a factor 3 increase




Figure 5.18: Amplitude of the 22 mode of Ψ_4 extracted at radius $r = 450 M_\odot$ for model B and for three different resolutions.

in the computational costs). Having such small differences in the merger time allows us to perform a much simpler and cleaner analysis with respect to the one presented in [169]. In particular we do not need to perform any alignment / time scaling of the numerical waveforms when measuring their convergence or performing the Richardson extrapolation.

A more quantitative analysis of the accuracy of our code is shown in Figure 5.19. There we show the convergence of our code both in amplitude and in the phase. For both quantities we find clean convergence, with order 3.2, almost up to the contact time $t \sim 5000 M_\odot$. As in the case of model A and as remarked before, we stress again that we do not align the waveforms in any way and that this is the highest convergence order ever shown for binary neutron star simulations in full general-relativity. The relative de-phasing between the high and the medium resolution at the contact time is about 0.7 radians. Notice that the de-phasing is smaller than the one we observed for *Whisky*, at higher resolution, and over only the ≈ 2.5 orbits of the small separation binaries.

Finally in Figure 5.20, we show the comparison between our code and the predictions from post-Newtonian (PN) calculations, using either the point-particle Taylor-T4 formula (PP) [283] or with the inclusion of tidal effects up to the next-leading order (1PN) [329, 247, 212]. In particular we take as reference the Richardson extrapolated phase evolution, $\phi_{h=0}$, computed assuming convergence order 3.2 and we plot the de-phasing of the different models with respect to it. The alignment in phase of the PN waveforms is performed by time shifting the waveforms in such a way so that they match the location of the third maximum in the real part of the 22 mode of Ψ_4 with the one found in the highest resolution run.

The de-phasing between the highest resolution run and the extrapolated result is of ≈ 0.4 radians at NR contact point, $t - r_* = 5000 M_\odot$ (which is about 13.5 GW cycles), and of ≈ 1.4 radians over ~ 15 GW cycles at the bare contact frequency. As a comparison, we recall that [169] found, for a model with the same compactness, a de-phasing, between the highest resolution simulation




Figure 5.19: Accumulated de-phasing and amplitude differences for model B as evolved with `WhiskyTHC`. In both plots we show the difference between low and medium resolution (blue lines), between high and medium resolution (green lines), as well as the rescaled differences between high and medium resolution (red lines) computed assuming convergence order of 3.2.




Figure 5.20: Accumulated de-phasing for model B with respect to the Richardson extrapolated numerical-relativity waveform. In particular we show the de-phasing accumulated by the three runs B.MP5.H1, B.MP5.H2 and B.MP5.H3, as well as by the waveform predicted by the PN model. The resolution extrapolation was performed assuming convergence order of 3.2.

(which is 35 % higher than the one employed by us) and the extrapolated result, of $\gtrsim 5$ radians over 15 GW cycles at the bare contact frequency. Notice, however, that the resolution extrapolation after the NR contact point has to be considered with caution since convergence is lost after contact.

The de-phasing between the point-particle PN waveform and the extrapolated one is only of ≈ 0.6 radians at $t - r_* = 5000 M_\odot$. With the inclusion of tidal effects the de-phasing is reduced to the point that it can no longer be measured as it is smaller than the uncertainty due to the eccentricity in the initial data (more on this below). The tidally corrected PN waveform appears to be on top of the extrapolated numerical data up to NR contact point, $t - r_* = 5000 M_\odot$. The accumulated de-phasing at the bare contact frequency is of only 0.7 radians. This result clearly rules out, for this model, the presence of any significant tidal amplifications effects from next-next-leading order terms in the PN expansion.

Clearly, since we are talking about extrapolated results, there is a degree of uncertainty in the exact de-phasing, which we estimate, following [169], on the basis extrapolations performed with order varied by ± 0.2 , to be of ± 0.05 radians at $t = 5000 M_\odot$. Notice, however, that a larger uncertainty applies to the PN waveforms since the exact de-phasing depends on the alignment procedure. The reason is that their de-phasing with respect to the *WhiskyTHC* waveform oscillates by ± 0.2 radians during the initial part of the inspiral as a consequence of the residual eccentricity in the initial data in the numerical relativity simulations. Finally we point out that finite-radius extraction effects could also, in principle, be relevant in the amplitude of the de-phasing with PN. As discussed before, in order to keep our analysis simple, we do not extrapolate the waves in radius, but we point out that Baiotti et al. [26] estimated, for a grid setup similar to ours (but for a lower compactness and smaller total mass

model) a phase uncertainty of ± 0.05 radians, which is negligible with respect to the uncertainty due to the eccentricity in the initial data.

5.5 Conclusions

In this chapter we presented a new multi-dimensional, general-relativistic, hydrodynamics code, *WhiskyTHC*, born from the merger of the *Whisky* and the *THC* codes. This code inherited from *Whisky* the primitive recovery routine as well as a new EoS framework with support for composition and energy dependent realistic equation of state [136] and from *THC* the use of high-order flux-vector splitting finite-differencing schemes [264]. This is the first genuinely higher than second order fully general relativistic code.

Amongst the new techniques introduced with *WhiskyTHC* is the use of positivity-preserving limiters [170] as a way to treat low density regions alternative to the traditional “atmosphere” prescriptions. We showed that this treatment is able to significantly improve the quality of simulations involving isolated, stable, neutron stars. The introduction of positivity preserving limiters in any hydrodynamics code is rather straightforward and we think that it should become the method of choice for the evolution of isolated stars.

We demonstrated the accuracy of our code in a series of classical tests involving the linear and non-linear evolution of isolated stars. In particular we showed that our code is able to stably evolve isolated stars for a long time and can attain high order (third) of convergence in the simulation of the gravitational collapse of non-rotating stars to black hole.

We applied our code to the simulation of the late-inspiral and merger of two neutron stars in quasi-circular orbits. We used small separation models to test the dependence of the results on the atmosphere treatment and we showed that our results are completely independent on it. Furthermore we demonstrated the high order of convergence and accuracy of our new code with respect to the old *Whisky* code. In particular we found higher than second order convergence in the phase and an overall phase error which is, at the same resolution and with similar computational costs, estimated to be ≈ 50 times smaller than the one estimated for *Whisky*.

Finally we studied the tidal effects in the late-inspiral and merger of binary neutron stars from the initial “large separation” of 60 km. We showed that *WhiskyTHC* is able to accurately estimate the small tidal effects presents in the inspiral of highly compact binaries, with $C = 0.18$, at a much lower resolution and at a fraction of the cost of the studies published in the literature [169]. In particular we found a convergence order of 3.2 in both the amplitude and the phase up to the contact point in the numerical simulations, which is the highest convergence order ever shown for binary neutron star inspiral in full general relativity. We compared the numerical waveform with the analytic post-Newtonian predictions and we found remarkable agreement between the numerical and the analytic predictions, especially when tidal corrections are included in the analytic model. Our findings show that, for the model we considered, the tidally corrected Taylor-T4 waveform agrees with the numerical relativity one up to contact. In particular our results rule out, for this particular model, any significant amplification of tidal effects by next-next-leading order terms.

In the future we plan to exploit the efficiency of the high-order methods in `WhiskyTHC` to do a more systematic investigation of tidal effects in binary neutron stars mergers, as well as in black-hole neutron star binaries, using realistic equation of states and compactness parameters. We also plan to carefully assess the detectability of such effects by near future gravitational waves detectors.

Chapter 6

Discontinuous Galerkin Methods

6.1 Introduction

Even though FV and FD schemes have been particularly successful and are indeed the standard choice for modern numerical codes in relativistic hydrodynamics and MHD, they also suffer from some limitations, such as the difficulty of handling complicated grid structures and boundary conditions, or those associated with achieving high orders of accuracy. These are mainly due to the fact that high-order accuracy is generally attained with the use of large reconstruction stencils and expensive non-linear limiting operators, which quickly become cumbersome to handle when the grid is not structured and/or quadratures are required in the computation of the fluxes, *i.e.*, for higher than third order FV schemes. Large reconstruction stencils also come with large ghost regions when doing parallel calculations, leading to poor scalability results. Finally these schemes are also often overly dissipative in situations in which shock waves are not the dominant part of the dynamics and may fail to properly resolve fine structures of the flow [177]. This has important consequences for the accuracy of general relativistic hydrodynamics codes [31].

For these reasons, alternative approaches to general relativistic hydrodynamics such as finite-element methods [220], or spectral methods [155, 159] are worth consideration. This latter approach is particularly interesting because spectral methods are able to attain very high accuracy, but it is also limited by the well known fact that these methods fail spectacularly when the solution develops large gradients or discontinuities. For this reason, spectral methods for relativistic hydrodynamics have been limited to the generation of initial data [58, 19] or to situations in which strong gradients could be treated with shock-tracking techniques within a multi-domain framework [155].

More recently, however, a novel method was suggested by Dumbser and Zanotti [121], who presented a hybrid FV/discontinuous Galerkin (DG) approach for special relativistic resistive magnetohydrodynamics (MHD). In this approach a local spacetime DG method was used as an implicit predictor step in the context of a high-order FV scheme, in order to treat the stiff source term of resistive MHD.

Indeed, Spectral Discontinuous Galerkin Methods (SDGM) and their variant employing Gaussian numerical integration (SDGM-NI) were developed to overcome some of the above limitations of FV and spectral or pseudospectral methods respectively [75]. These methods work essentially by combining the classical Runge-Kutta discontinuous Galerkin approach by Cockburn [93] with the spectral element method (SEM) of Patera [249]. For this reason they are also often referred to as DG-SEM or DG-SEM-NI. These methods are particularly well suited for the solution of conservation laws and have been successfully applied to a number of classical hyperbolic, parabolic and elliptic problems (see, *e.g.*, [91, 75]). Finally they have been also successfully applied to the solution of the Einstein equation in vacuum by Zumbush [347] and Field et al. [128].

We develop here the necessary formalism for the application of fully explicit DG methods to relativistic hydrodynamics on curved spacetimes. As an application we present a prototype code employing SDGM-NI for general relativistic hydrodynamics in spherical symmetry. We show that the proposed scheme is able to properly resolve strong shocks and achieve high-order, spectral accuracy for smooth solutions. While we will not discuss explicitly the coupling of the solution of the hydrodynamics equations with that of the Einstein equations, it is clear that a natural choice would be to use discontinuous Galerkin methods, such as the ones recently proposed by [347] or [128], or finite-element-methods such as the ones introduced by [307, 306], also for the metric evolution equations. This approach would have the advantage, with respect to the solution proposed with the “Mariage des Maillages” [113, 119], that the fluid and the spacetime variables would share the same grid and no expensive interpolations would therefore be needed.

The chapter is organised as follows. In Sect. 6.2 we derive the general theory for the application of discontinuous Galerkin methods to relativistic hydrodynamics in curved spacetimes and we specialize it to the spherically-symmetric case. In Sect. 6.3 we present our prototype numerical code, EDGES (Extensible Discontinuous GalErkin Spectral library), which was used to test DG methods for general relativistic hydrodynamics in one-dimension (1D) and spherical symmetry. The results obtained on a representative number of test cases are then presented in Sect. 6.4. Finally Sect. 6.5 is dedicated to the summary and conclusions.

6.2 Discontinuous Galerkin methods for general relativistic hydrodynamics

Broadly speaking, Galerkin methods are projection methods for the weak formulation of the equations. In the case of the general relativistic hydrodynamics equations, such a formulation could be obtained in two different ways. The first one consists in starting from the relativistic hydrodynamics equations written in a conservative form in a chosen coordinate system, *e.g.*, the Valencia formulation [36], and then integrating them against a test function. The numerical scheme obtained with the Galerkin projection would then be a direct generalization of the standard HRSC schemes used in general relativistic hydrodynamics. Indeed, when considered at first-order only, DG schemes reduce to FV ones and it is for this reason that they are often interpreted as an alternative way to

attain high-order FV methods

While this approach is certainly possible and would seem to be quite natural, it has the limitation that since we start from the equations in their coordinate form, we also have to choose a metric with respect to which the volume integrals are performed. The choice of such a metric is effectively arbitrary, but any choice different from that of a flat metric corresponds to the absorption of a multiplicative factor into the definition of the test function and thus it is equivalent to a modification of the scheme in the higher-than-first-order case. As a result, the choice of the metric is unimportant only in the FV limit of Galerkin methods, but it plays a central role in higher-order Galerkin schemes.

A second way to obtain the weak formulation of the equations and the one actually outlined in this chapter, is to follow an approach which is instead manifestly covariant and thus does not require any assumed background. After the formulation is obtained, it can then be decomposed in the standard $3 + 1$ split of general relativity. This choice has the advantage of producing the most natural extension of the commonly used HRSC frameworks to the DG case. The resulting schemes will be naturally covariant, suited for use with standard spacelike or null foliations or even independent of any foliation or coordinate system. The reason why this is possible lies, as pointed out by Meier [220], in the covariant nature of finite-element methods and, by extension, of DG methods. In these methods, in fact, the equations are formulated on reference elements mapped into the physical space via diffeomorphisms, thus removing any need for a (preferred) coordinate system. The important difference between our approach and the one by Meier [220] is in the use of non-conforming, discontinuous, Galerkin methods. This gives us the possibility of reducing the coupling of the numerical solution across the elements to flux terms, thus enabling the construction of globally explicit, local schemes, in contrast to the need for solving implicit, global, non-linear problems.

6.2.1 Weak formulation of the equations of relativistic hydrodynamics

Let $(\mathcal{M}, g_{\alpha\beta})$ be a strongly hyperbolic, C^2 , spacetime with metric $g_{\alpha\beta}$ and let ∇ be the covariant derivative associated with $g_{\alpha\beta}$. We consider a perfect fluid described by a rest-mass-density 4-vector J^α and a stress energy tensor $T^{\alpha\beta}$ defined by

$$J^\alpha := \rho u^\alpha, \quad T^{\alpha\beta} := \rho h u^\alpha u^\beta + p g^{\alpha\beta}, \quad (6.1)$$

where ρ is the rest-mass density, u^α is the fluid 4-velocity, p is the pressure, ε is the specific internal energy and $h := 1 + \varepsilon + p/\rho$ is the specific enthalpy.

If we assume baryon-number conservation and a generic equation of state (EOS) of the form $p = p(\rho, \varepsilon)$, then the equations of motion for the fluid on $\mathcal{M}$ read

$$\nabla_\alpha J^\alpha = 0, \quad \nabla_\beta T^{\alpha\beta} = 0, \quad p = p(\rho, \varepsilon). \quad (6.2)$$

In general these equations are to be intended in the sense of distributions, since we expect the solution to develop singularities in the form of shock waves.

In general we are interested in solving (6.2) on an open, regular¹, finite

¹See [83] for a detailed discussion of the regularity requirements. Broadly speaking this amounts to having a domain which has a normal defined everywhere except for at most a discrete set of points (vertexes); *i.e.*, a cubic box is a regular domain.

domain $\Omega \subset \mathcal{M}$, with suitable initial/boundary data. A precise mathematical formulation of this problem can be done within the context of bounded divergence-measure vector fields using the theory developed in [82, 83]. In particular, we will look for solutions in the functional space, $\mathcal{V}$, of all the L^∞ , *i.e.*, “bounded”, vector fields over Ω , whose divergence, in the sense of the distributions, are Radon measures².

As a first step we introduce a triangulation of N “elements” of Ω , $\{\Omega_j\}_{j=1}^N$, by selecting a family of diffeomorphisms $\varphi_j: K \subset \mathbb{R}^4 \rightarrow \Omega$, $\Omega_j = \varphi_j(K)$ such that

$$\bigcup_{j=1}^N \Omega_j = \Omega, \quad \mathring{\Omega}_i \cap \mathring{\Omega}_j = \emptyset, \quad \forall i \neq j, \quad (6.3)$$

where K is the, so-called, “reference element”, usually an hypercube or a 4D simplex and $\mathring{\Omega}_j$ denotes the interior of Ω_j . We also arrange the local coordinate system, $\{x^\mu\}_j$, induced by φ_j , so that ∂_0 is timelike or null.

If we now look for solutions $J^\alpha \in \mathcal{V}$, the first of the equations (6.2) is equivalent, in the sense of distributions, to

$$\sum_{j=1}^N \left[\int_{\Omega_j} J^\alpha \nabla_\alpha \phi \epsilon - \int_{\partial\Omega_j} \phi \mathcal{J}^\alpha \epsilon_{\alpha\beta\gamma\delta} dx^\beta dx^\gamma dx^\delta \right] = 0, \quad (6.4)$$

for all $\phi \in C_0^1(\mathcal{M})^3$. Note that in expression (6.4) the symbol ϵ refers to the proper volume form of the spacetime, *i.e.*, in any local chart, $\{x^\mu\}$, $\epsilon = \sqrt{-g} e^0 \wedge e^1 \wedge e^2 \wedge e^3 = \epsilon_{\alpha\beta\gamma\delta} dx^\alpha dx^\beta dx^\gamma dx^\delta$, and $\mathcal{J}^\alpha$ is the internal normal trace of J^α . This object reduces simply to J^α , when J^α and Ω_k are regular, but in the general case the second integral has to be intended as the action of a measure, $\mathcal{J}^\alpha \epsilon_{\alpha\beta\gamma\delta} dx^\beta dx^\gamma dx^\delta$, on ϕ [83].

In the same way, if we look for solutions $T^{\alpha\beta} \in \mathcal{V} \otimes \mathcal{V}$, the second of the equations (6.2) is equivalent, in the sense of distributions, to

$$\sum_{j=1}^N \int_{\Omega_j} T^{\alpha\beta} \nabla_\beta \phi_\alpha \epsilon = \sum_{j=1}^N \int_{\partial\Omega_j} \phi_\alpha \mathcal{T}^{\alpha\beta} \epsilon_{\beta\gamma\delta\mu} dx^\gamma dx^\delta dx^\mu \quad (6.5)$$

for all the one-forms $\phi_\alpha \in C_0^1(\mathcal{M}; T^*\mathcal{M})$, $T^*\mathcal{M}$ being the co-tangent bundle of $\mathcal{M}$. Again, $\mathcal{T}^{\alpha\beta}$ is a generalization of $T^{\alpha\beta}$ and the integral has to be interpreted as the action of $\mathcal{T}^{\alpha\beta} \epsilon_{\beta\gamma\delta\mu} dx^\gamma dx^\delta dx^\mu$ on ϕ_α in the non-smooth case.

The solution of the relativistic hydrodynamics equations consists then in finding

$$J^\alpha \in \mathcal{V} \quad \text{s.t., (6.4) holds} \quad \forall \phi \in C_0^1(\mathcal{M}), \quad (6.6a)$$

$$T^{\alpha\beta} \in \mathcal{V} \otimes \mathcal{V} \quad \text{s.t., (6.5) holds} \quad \forall \phi_\alpha \in C_0^1(\mathcal{M}; T^*\mathcal{M}), \quad (6.6b)$$

together with an EOS and proper boundary-initial data, to be specified through $\mathcal{J}^\alpha$ and $\mathcal{T}^{\alpha\beta}$ on $\partial\Omega$. We remark, again, that (6.6) is perfectly equivalent, in

²For a precise definition see [82]. Broadly speaking this condition means that we restrict ourselves to cases in which the solution presents at most mild singularities, such as jump-discontinuities.

³We recall that a function of class $C_0^n(\Omega)$ is a function of class $C^n(\Omega)$ and, in addition, with compact support in Ω .

the sense of distributions, to (6.2) and that the triangulation $\{\Omega_j\}_{j=1}^N$ has been introduced mainly for later convenience.

6.2.2 Spacetime discontinuous Galerkin formulation

As mentioned above, within the Galerkin approach a numerical scheme is obtained by projecting (6.6) on a finite dimensional subspace $V \subset \mathcal{V}$. In general this space is constructed starting from the space of piecewise polynomials, in particular we define

$$X = \{u \in L^\infty(\Omega) : u \circ \varphi_j \in \mathbb{P}_D(K), j = 1, \dots, N\}, \quad (6.7)$$

where $\mathbb{P}_D(K)$ is the space of polynomials with at most degree D on K . Notice that the functions in X are allowed to be discontinuous at the edges of the elements, hence the name “discontinuous Galerkin” for the resulting numerical scheme. The space V is taken as the space of all the vector fields “whose components are elements of X ”, more precisely

$$V = \{u^\alpha \in \mathcal{V} : [\varphi_j]_* u^\alpha \in [\mathbb{P}_D(K)]^4, j = 1, \dots, N\}, \quad (6.8)$$

where $[\mathbb{P}_D(K)]^4$ is the space of 4-tuples of polynomials with at most degree D on K and $[\varphi_j]_*$ is the pull-back associated with φ_j . The Galerkin method is then simply the restriction of (6.6) to V so that it consists of finding $J^\alpha \in V$ and $T^{\alpha\beta} \in V \otimes V$ such that (6.4) and (6.5) hold for suitable choices of the test functions, ϕ and ϕ_α , that we are going to discuss in the following. The resulting equations can be solved numerically, because a finite number of conditions suffice to fully determine J^α and $T^{\alpha\beta}$ as long as we have a way to evaluate the fluxes $\mathcal{J}^\alpha$ and $\mathcal{T}^{\alpha\beta}$ on $\partial\Omega_j$.

As discussed above, in the continuous case these fluxes are simply the restriction of J^α and $T^{\alpha\beta}$ to $\partial\Omega_j$. More explicitly if we write symbolically $\mathcal{P} = \{\rho, u^1, u^2, u^3, \epsilon\}$ for the primitives variables and consider J^α and $T^{\alpha\beta}$ as functions of $\mathcal{P}$, i.e., $J^\alpha = J^\alpha(\mathcal{P})$ and $T^{\alpha\beta} = T^{\alpha\beta}(\mathcal{P})$, then $\mathcal{J}^\alpha = J^\alpha(\mathcal{P}^*)$ and $\mathcal{T}^{\alpha\beta} = T^{\alpha\beta}(\mathcal{P}^*)$, $\mathcal{P}^*$ being the restriction of $\mathcal{P}$ on $\partial\Omega_j$, as the fluxes can only depend on the location in the spacetime through $\mathcal{P}$. Stated differently, $\mathcal{J}^\alpha$ and $\mathcal{T}^{\alpha\beta}$ are simply the Godunov fluxes for the conservation law. In the general case $\mathcal{J}^\alpha$ and $\mathcal{T}^{\alpha\beta}$ can be determined with causality considerations on spacelike boundaries⁴, or as solutions of generalized Riemann problems⁵, on timelike and null-like boundaries, as they are known as soon as $\mathcal{P}^*$ is known on those boundaries. This is basically analogous to the Newtonian case when using spacetime DG methods to discretize balance laws. For this reason we refer to [246] for a more in-depth discussion. We limit ourselves to note that, in the context of a numerical scheme, the computation of the fluxes can be greatly simplified with the use of approximate Riemann solvers, such as the HLLE scheme. In that case

⁴Where the solution has different limits, $\mathcal{P}_1$ and $\mathcal{P}_2$ from different sides of $\Gamma \subset \partial\Omega_j$ and Γ is spacelike, we proceed as in the Godunov method and we set $\mathcal{J}^\alpha = J^\alpha[\mathcal{P}^*(\mathcal{P}_1, \mathcal{P}_2)]$. Causality requires that $\mathcal{P}^*$ must depend only on the past limit of $\mathcal{P}$ at Γ , say $\mathcal{P}_1$, this implies $\mathcal{P}^* = \mathcal{P}_1$, thus $\mathcal{J}^\alpha = [J^\alpha]_1$. The same argument can be applied to evaluate $\mathcal{T}^{\alpha\beta}$.

⁵Here we are assuming that a solution to the generalized Riemann problem exists under the conditions stated above. As discussed in Chapter 3 this is still an open problem for the equations of general-relativistic hydrodynamics.

$\mathcal{J}^\alpha$ and $\mathcal{T}^{\alpha\beta}$ are approximated directly without the need to explicitly compute the solution of the Riemann problem, $\mathcal{P}^*$, at the interface.

Once we have a way to compute the fluxes, the fully discrete equations are readily obtained by testing (6.4) and (6.5) on a set of linearly-independent test functions, ϕ and ϕ_α . The resulting finite set of equations can be cast in a set of non-linear equations for the spectral coefficients of the numerical solutions, when these are expanded over a linear basis of V , or $V \otimes V$, for J^α and $T^{\alpha\beta}$ respectively, and by following a standard finite-element method procedure, see e.g., [261]. In particular, it is sufficient to consider test functions $\phi \in X$ and $\phi^\alpha \in V$, so that the final numerical scheme consists in finding

$$J^\alpha \in V \quad \text{s.t., (6.4) holds} \quad \forall \phi \in X, \quad (6.9a)$$

$$T^{\alpha\beta} \in V \otimes V \quad \text{s.t., (6.5) holds} \quad \forall \phi^\alpha \in V, \quad (6.9b)$$

where the test functions on $\partial\Omega_j$ in the boundary integrals appearing in (6.9) have to be interpreted as being a C_0^1 extension to $\mathcal{M}$ of the original test function, created in such a way as to smoothly match the one-sided limit, from the interior of Ω_j , of the original test function.

To explicitly write down the method, in every finite element, Ω_j , we select a set of conserved quantities, $C = \{J^0, T^{0\mu}\}^6$, for which there exists a one-to-one relation, involving the EOS, with the set of primitive variables [248], $\mathcal{P}$, so that we can formally write $J^i = J^i(C)$ and $T^{i\mu} = T^{i\mu}(C)$. We can then obtain a set of non-linear equations for $C \in X^5$ simply expanding the Galerkin conditions (6.4) and (6.5)

$$\sum_{j=1}^N \left[\int_{\Omega_j} J^0 \partial_0 \phi \epsilon + \int_{\Omega_j} J^j(C) \partial_j \phi \epsilon \right] = \sum_{j=1}^N \int_{\partial\Omega_j} \mathcal{J}^\mu \phi \epsilon_{\mu\alpha\beta\gamma} dx^\alpha dx^\beta dx^\gamma \quad (6.10)$$

and, setting $\phi_\alpha = \phi \delta^\mu_{\alpha}$,

$$\begin{aligned} \sum_{j=1}^N \left[\int_{\Omega_j} T^{0\mu} \partial_0 \phi \epsilon + \int_{\Omega_j} T^{i\mu}(C) \partial_i \phi \epsilon \right] = \\ \sum_{j=1}^N \left[\int_{\partial\Omega_j} \mathcal{T}^{\nu\mu} \phi \epsilon_{\nu\alpha\beta\gamma} dx^\alpha dx^\beta dx^\gamma + \int_{\Omega_j} T^{\nu\lambda}(C) \Gamma^\mu_{\lambda\nu} \phi \epsilon \right], \end{aligned} \quad (6.11)$$

where $\Gamma^\alpha_{\beta\gamma}$ are the Christoffel symbols and $\phi \in X$.

The key point here is that, as the functions are discontinuous across the $\partial\Omega_j$'s, these equations are local equations for the spectral coefficients within the Ω_j 's coupled only through the fluxes. In particular this implies that, if the computational grid “follows the causal structure of the spacetime”, in the sense that it can be traversed with a succession of “causal slices” satisfying some sort of generalized Courant-Friedrichs-Lewy (CFL) condition ensuring the causal disconnection between the timelike boundaries of the elements, then the discontinuous Galerkin method becomes globally explicit. Under these conditions, in fact, the fluxes between the elements of the grid slice depend

⁶Other choices are possible, for example in the context of a 3 + 1 split we could use the same conserved quantities as the ones used in the Valencia formulation.




Figure 6.1: Schematic representation of a causal slice (I, S) containing the element Ω_j . The future and past boundaries $\partial^\pm \Omega_j$, as well as the time-like or null-like boundaries $\partial^\times \Omega_j$ of Ω_j are indicated with blue and red lines, respectively. Finally, the shaded region represent the causal slice.

only on the data on the previous slice. Once these are computed, we are left with a set of formally decoupled equations involving the spectral coefficients of the numerical solution in the different elements.

We now give the precise mathematical definitions of “causal slices” and of “grid that follows the causal structure of the spacetime”.

For every *open set* A , we introduce the notation $\partial^+ A$, to indicate the *future boundary* of A , that is the set of all the points, $p \in \partial A$, such that

$$J^+(p) \cap A = \emptyset, \quad (6.12)$$

$J^+(p)$ being the causal future of p . Analogously we call *past boundary* of A , $\partial^- A$, the set of all the points $p \in \partial A$ for which there exist an open neighborhood, U , such that

$$J^+(p) \cap U \subset \bar{A}, \quad (6.13)$$

where we have indicated with $\bar{A}$ the closure of A . Finally we will write $\partial^\times A$ for $\partial A \setminus [\partial^+ A \cup \partial^- A]$.

For a *generic set* E , instead, the above definitions are modified by considering the interior of the set, $\mathring{E}$. In the case of an element of the grid, Ω_j , these definitions are useful to identify the regions of the boundary of Ω_j where the characteristics are always ingoing, $\partial^- \Omega_j$, always outgoing, $\partial^+ \Omega_j$, or of mixed nature, $\partial^\times \Omega_j$. These different parts of the causal slice (I, S) containing the element Ω_j are shown in Fig. 6.1.

We also define as a *slice of the triangulation*, or simply *slice*, any tuple (I, S) where $I \subset \{1, 2, \dots, N\}$ and $S = \bigcup_{j \in I} \Omega_j$ is connected. We will say that a slice (I, S) is a *causal slice* if

$$\partial^\times S \subset \partial \Omega. \quad (6.14)$$

An example of such a slice of the triangulation is given by the shaded region in Fig. 6.1: a causal slice is basically a slice whose timelike spatial boundaries ∂^+ and ∂^- are parts of a Cauchy foliation of the spacetime. A causal slice (I, S) is said to be a *minimal causal slice* if it also satisfies the CFL-like condition

$$\left[\bigcup_{j \in I} (\partial^\times \Omega_j \cup \partial^- \Omega_j) \setminus \partial \Omega_k \right] \cap J^+(\Omega_k) = \emptyset, \quad \forall k \in I, \quad (6.15)$$

that is, if the characteristics originating in each element, Ω_k , only intersect the boundaries of the element itself, $\partial\Omega$, or the future boundaries of the other elements in the slice.

Finally we say that a triangulation *follows the causal structure of the spacetime* if it can be written as union of minimal causal slices. Intuitively a triangulation that follows the causal structure of the spacetime is one that can be sliced in minimal causal slices, which are grids associated with a Cauchy foliation of the spacetime and which satisfy a CFL condition on each hyper-surface of this foliation.

If the triangulation $\{\Omega_j\}_{j=1}^N$ follows the causal structure of the spacetime, in the sense defined above, then we can introduce a family of successive slices $\{(I_k, S_k)\}_{k=0}^Q$, parametrized by a discrete time k , such that $\partial^- S_{k+1} \subset \partial^+ S_k \cup \partial\Omega$ and $\Omega = \bigcup_k S_k$. We can then construct a globally explicit, locally implicit, scheme by solving (6.9) on the successive slices. The reason is that, for all $j \in I_{k+1}$, the fluxes across $\partial^- \Omega_j$ and $\partial^\times \Omega_j$ are completely determined by the data on S_k and the boundary conditions, while the fluxes across $\partial^+ \Omega_j$ do not couple Ω_j with other elements within S_{k+1} .

6.2.3 Discontinuous Galerkin formulation in the 3 + 1 split

Clearly the strategy outlined above could be useful in situations in which stiff sources are present, such as in the resistive MHD case, where implicit DG methods have already been shown to be suitable as predictors to treat stiff sources [121], but in the unmagnetized case an implicit time stepping is unnecessarily expensive in most situations. The generation of a triangulation which follows the causal structure of the spacetime could also be highly non-trivial, especially if the spacetime is evolved dynamically. For these reasons, instead of directly solving (6.10) and (6.11), one can use them in order to derive a fully explicit scheme. This can be accomplished by performing a 3 + 1 split directly at the level of the discrete problem (6.9).

As customary, we foliate the spacetime along $t = \text{const.}$ hypersurfaces, Σ_t , and consider a vector t^α such that $t^\alpha \nabla_\alpha t = 1$. Using this vector we define the three-volume form $\eta_{\alpha\beta\gamma} = \epsilon_{\delta\alpha\beta\gamma} t^\delta$. Also, as usually done in this context, we can use the integral lines of t^α to identify points on Σ_t with points on $\Sigma := \Sigma_0$ and interpret the variation of the fields across the Σ_t 's as being the result of dynamics on a three-manifold, Σ . We are then interested in studying (6.2) in a world-tube $S \times (0, t)$, $S \subset \Sigma$ being an open, bounded, regular domain in Σ , together with proper boundary-initial conditions. Clearly this is a particular case of the general problem studied above.

In order to apply the discontinuous Galerkin formulation, we consider a triangulation $\{S_j\}_{j=1}^N$ of S , by selecting a family of diffeomorphisms $\Phi_j: T \subset \mathbb{R}^3 \rightarrow \Sigma$, $S_j = \Phi_j(T)$ such that

$$\bigcup_{j=1}^N S_j = S, \quad \mathring{S}_i \cap \mathring{S}_j = \emptyset, \quad \forall i \neq j, \quad (6.16)$$

where T is now a three-dimensional reference element, a cube or a tetrahedron.

This induces a triangulation of Ω

$$\{\Omega_{j,n}\}_{j=1,n=1}^{N,Q} = \left\{ S_j \times \left(n\Delta t, (n+1)\Delta t \right) \right\}_{j=1,n=1}^{N,Q}, \quad (6.17)$$

that follows the causal structure of the spacetime on the n -th thin-sandwich $\Omega_n = S \times (t_n, t_n + \Delta t)$, at least for small Δt .

As we intend to use the method of lines in order to integrate the equations in time, we can factor out the time dependence from the test functions by considering the functional space

$$Y = \left\{ u \in L^\infty(S) : u \circ \Phi_j \in \mathbb{P}_D(T), j = 1, \dots, N \right\}, \quad (6.18)$$

and vector functions “whose components are elements of Y ”:

$$W = \left\{ u^\alpha \in \mathcal{V} : [\Phi_j]_* u^\alpha \in [\mathbb{P}_D(T)]^4, j = 1, \dots, N \right\}. \quad (6.19)$$

Given a function $u \in Y$, we can consider it as a function, $\tilde{u}$, over $S \times (0, t)$ with the identification $\tilde{u}(x^i, s) := u(x^i)$ so that $Y \hookrightarrow \tilde{Y} \subset X$ and, with a slight abuse of the notation, we can consider Y as being a subspace of X . In a similar way we can consider W to be a subspace of V . For these reasons we can choose $\phi \in Y$ in (6.10) and (6.11). We can thus obtain a fully explicit method by re-projecting (6.9) onto a new couple of subspaces or, equivalently, by finding

$$J^\alpha \in C^1((0, t); W) \quad \text{s.t., (6.4) holds } \forall \phi \in Y, \quad (6.20a)$$

$$T^{\alpha\beta} \in C^1((0, t); W \otimes W) \quad \text{s.t., (6.5) holds } \forall \phi^\alpha \in W. \quad (6.20b)$$

In the expressions above, and as customary when dealing with evolution problems, we have used the notation $u(\cdot) \in C^k((a, b); X)$, where $a, b \in \mathbb{R}$ and X is a Banach space, to indicate that, when u is regarded as a function only of time, it describes a regular, C^k , curve in X . In other words, when $u(t, x_1^i)$ is interpreted as a function of time only, it is of class C^k , while when $u(t_1, x^i)$ is interpreted as a function of x^i only, it is an element of the function space X . Note also that in this new formulation we do not allow the numerical solution to be discontinuous in time, so that the time integration can be performed with a standard solver for ordinary differential equations (ODEs).

We can derive a more explicit form for the (6.20) by projecting (6.10) with $\phi \in Y$, dividing both terms by Δt and by letting $\Delta t \rightarrow 0$, to obtain

$$\sum_{j=1}^N \partial_t \int_{S_j} J^0 \phi \boldsymbol{\eta} = \sum_{j=1}^N \left[\int_{S_j} J^i(C) \partial_i \phi \boldsymbol{\eta} - \int_{\partial S_j} \mathcal{J}^i \phi \eta_{i\alpha\beta} dx^\alpha dx^\beta \right]. \quad (6.21)$$

Reasoning along the same lines, we can derive an explicit discretization of the second of the (6.2), starting from the (6.11), to obtain

$$\begin{aligned} \sum_{j=1}^N \partial_t \int_{S_j} T^{0\mu} \phi \boldsymbol{\eta} = \sum_{j=1}^N \left[\int_{S_j} T^{i\mu}(C) \partial_i \phi \boldsymbol{\eta} - \right. \\ \left. \int_{\partial S_j} \mathcal{T}^{i\mu} \phi \eta_{i\alpha\beta} dx^\alpha dx^\beta - \int_{S_j} T^{\alpha\beta}(C) \Gamma_{\beta\alpha}^\mu \phi \boldsymbol{\eta} \right]. \end{aligned} \quad (6.22)$$

Finally the 3 + 1 discontinuous Galerkin formulation can be summarized in finding:

$$J^\alpha \in C^1((0, t); W) \quad \text{s.t., (6.21) holds } \forall \phi \in Y, \quad (6.23a)$$

$$T^{\alpha\beta} \in C^1((0, t); W \otimes W) \quad \text{s.t., (6.22) holds } \forall \phi^\alpha \in W. \quad (6.23b)$$

This scheme can be interpreted as a higher-order generalization of a FV discretization of the manifestly covariant formulation of relativistic hydrodynamics proposed by [248]. As a consequence, our scheme inherits properties such as hyperbolicity and the flexibility to work with spacelike or null-like foliations directly from [248]. This can be seen considering the case in which $D = 0$, that is, when looking for solutions that are constant over each element, S_j . Then a sufficient number of Galerkin conditions can be obtained by simply choosing $\phi = \chi_{S_j}$ for $j = 1, 2, \dots, N$, where χ_E is the indicator function of the set E , *i.e.*, a function which is equal to one in E and identically zero elsewhere. With this choice we obtain the set of equations

$$\partial_t \int_{S_j} J^0 \boldsymbol{\eta} + \int_{\partial S_j} J^i(C) \eta_{i\alpha\beta} dx^\alpha dx^\beta = 0, \quad (6.24)$$

and

$$\partial_t \int_{S_j} T^{0\mu} \boldsymbol{\eta} + \int_{\partial S_j} \mathcal{T}^{i\mu} \eta_{i\alpha\beta} dx^\alpha dx^\beta = - \int_{S_j} T^{\alpha\beta}(C) \Gamma^\mu_{\beta\alpha} \boldsymbol{\eta}, \quad (6.25)$$

for all $j = 1, 2, \dots, N$. These can easily be recognised as being the FV discretization of the formulation by [248].

6.2.4 Discontinuous Galerkin formulation in spherical symmetry

As a particular case of the formalism outlined above we consider the case in which the spacetime is spherically symmetric. This case is particularly interesting because the equations become 1 + 1 dimensional and are therefore well-suited for rapid prototyping and testing of new methods and techniques⁷.

In particular we consider a spherically symmetric spacetime in a radial-polar gauge

$$ds^2 = -\alpha^2 dt^2 + A^2 dr^2 + r^2(d\theta^2 + \sin^2 \theta d\phi^2), \quad (6.26)$$

where α and A are functions of t and r only. We next introduce the Bondi mass function, m , and the metric potential, v , by

$$A(t, r) = \left(1 - \frac{2m(t, r)}{r}\right)^{-1/2}, \quad \alpha(t, r) = e^{v(t, r)}. \quad (6.27)$$

Following [278], we define the physical velocity, v , by $v := Au^r/\alpha u^t$, where $W = \alpha u^t = (1 - v^2)^{-1/2}$ is the Lorentz factor. Furthermore we introduce the

⁷Of course in 1+1 dimensions a Lagrangian approach such as the one presented in [271, 232] is by far superior as it allows for natural spatial adaptivity and conservation properties. However, both of these advantages disappear in more than one spatial dimension. On the other hand Lagrangian approaches to multidimensional relativistic hydrodynamics have been recently proposed within the context of smoothed-particle-hydrodynamics schemes in special [279] and general relativity [302].

“conserved” quantities

$$\mathcal{D} := \alpha A J^t = \rho A W, \quad (6.28a)$$

$$S := \alpha T^{tr} = \rho h W^2 v, \quad (6.28b)$$

$$E := \alpha^2 T^{tt} = \rho h W^2 - p, \quad (6.28c)$$

$$\tau := E - \mathcal{D}. \quad (6.28d)$$

With these definitions, the Einstein equations reduce to the Hamiltonian constraint

$$\partial_r m = 4\pi r^2 E, \quad (6.29)$$

and the slicing condition $\partial_t K_{\theta\theta} = K_{\theta\theta} = 0$

$$\partial_r v = A^2 \left[\frac{m}{r^2} + 4\pi r(p + Sv) \right], \quad (6.30)$$

where K_{ij} is the extrinsic curvature.

These equations have to be integrated with the boundary conditions given by $m(0) = 0$ and by the requirement that v matches the Schwarzschild solution at the outer boundary of the computational domain (see *e.g.*, [237] for a detailed derivation and discussion of this equation).

The equations for the hydrodynamics are simply (6.21) and (6.22), where the elements, S_j , are taken to be the spherical shells $r_j < r < r_{j+1}$. These can be written in terms of the conserved quantities (6.28) and specialized for the metric (6.26) by substituting the explicit expression for the Christoffel symbols and the determinant of the metric, to obtain

$$\begin{aligned} \sum_{j=1}^N \int_{r_j}^{r_{j+1}} \partial_t \mathbf{F}^t(C) \phi r^2 dr &= \sum_{j=1}^N \left\{ \int_{r_j}^{r_{j+1}} X \mathbf{F}^r(C) \partial_r \phi r^2 dr - \right. \\ &\quad \left. \left[r^2 X \mathcal{F}^r \phi \right]_{r_j}^{r_{j+1}} + \int_{r_j}^{r_{j+1}} \mathbf{s}(C) \phi r^2 dr \right\}, \end{aligned} \quad (6.31)$$

where we define $X := \alpha/A$. Here, the fluxes are given by

$$\mathbf{F}^t = \{\mathcal{D}, S, \tau\}, \quad \mathbf{F}^r = \{\mathcal{D}v, Sv + p, S - \mathcal{D}v\}, \quad (6.32)$$

while the source term is

$$\mathbf{s} = \left\{ 0, (Sv - \tau - \mathcal{D})(8\alpha A \pi r p + \alpha A \frac{m}{r^2}) + \alpha A p \frac{m}{r^2} + 2 \frac{\alpha p}{Ar}, 0 \right\}. \quad (6.33)$$

In the derivation of (6.31) the momentum constraint was used to substitute the derivatives of the metric in the source term and a factor X was absorbed into the test function in the derivation of the equation for τ .

A close examination of the (6.31) reveals that, again, this formulation of the equations can be interpreted as an higher-order generalization of a classical FV discretization of the equations of relativistic hydrodynamics. In particular it can be seen that, in the case in which $\mathcal{D}$, S and τ are constant over each element, (6.31) reduces to the FV method discussed in [278].

6.3 The EDGES code

In order to test discontinuous Galerkin methods for relativistic hydrodynamics and their reduction to spherical symmetry, we have developed a new 1D code, EDGES. This consists of a general discontinuous Galerkin library which is then used by a code solving the general relativistic hydrodynamics equations in spherical symmetry. EDGES makes extensive use of advanced generic programming techniques such as static polymorphism via recursive templates, and expression templates see *e.g.*, [335]. The code employs the `Blitz++` high-performance array library [328, 1] and makes use of the `UMFPACK` multi-frontal sparse factorization method [104, 105, 106, 107, 103] for linear systems inversion. In what follows we describe in more detail the implementation of the discontinuous Galerkin approach within EDGES.

6.3.1 The DG equations in a fully discrete form

We consider a spherical shell $S = [0, R]$ in the spacetime (6.26), containing a fluid described by (6.1) and (6.2). Furthermore we consider a “triangulation” of $S := \bigcup_{j=1}^N S_j$ where (see Fig. 6.2 for a scheme of the triangulation)

$$S_i \cap S_j = \emptyset, \quad \forall i \neq j, \quad S_j = \varphi_j([-1, 1]). \quad (6.34)$$

The functional space that we consider in EDGES is,

$$Z = \{u \in L^\infty(S) : u \circ \varphi_j \in \mathbb{P}_D[-1, 1]\}, \quad (6.35)$$

the set of all the functions that are polynomials of degree D over each S_j .

If we denote by $L_D(x)$ the D -th *Legendre* polynomial on $[-1, 1]$, a Gaussian quadrature of order $2D - 1$ can be obtained with the formula

$$\int_{-1}^1 f(x) dx \approx \sum_{i=0}^D w_i f(x_i), \quad (6.36)$$

where $\{x_i\}_{i=0}^D$ are the zeros of $(1 - x^2)dL_D(x)/dx$, w_i are a set of weights given by

$$w_i = \int_{-1}^1 l_i(x) dx, \quad i = 0, 1, \dots, D \quad (6.37)$$

and $\{l_i(x)\}_{i=0}^D$ are the *Lagrange* polynomials associated with the nodes $\{x_i\}_{i=0}^D$, *i.e.*, a set of polynomials of degree D such that $l_i(x_k) = \delta_{ik}$ for $i, k = 0, 1, \dots, D$. Given two regular functions f and g in $r \in (0, R)$ we now define their *continuous* scalar product as

$$(f, g) := \int_0^R f(r) g(r) r^2 dr, \quad (6.38)$$

and use the quadrature formula (6.36) to introduce the *discrete* scalar product

$$(f, g)_D := \sum_{j=1}^N (f, g)_{j,D}, \quad (6.39)$$




Figure 6.2: Scheme of the spacetime grid structure in EDGES. The collocation nodes (filled squares) are generated and stored on the reference element, T , and then mapped onto the finite element S_j with an affine transformation, φ_j . Note that each point on the boundary of an element is associated with two distinct degrees of freedom, thus allowing the functions to have two different one-sided limits. Because the number of collocation points needed is $D + 1$, where D is the order of the polynomial representation, the figure refers to a polynomial of order three.

where

$$(f, g)_{j,D} := \sum_{i=0}^D w_i |\varphi'_j| f[\varphi_j(x_i)] g[\varphi_j(x_i)] \quad (6.40)$$

and $|\varphi'_j|$ is the Jacobian of the affine transformation $\varphi_j: [-1, 1] \rightarrow S_j$.

With these definitions in place we can construct a fully discrete system by looking for solutions $\mathcal{D}, S, \tau \in \mathbb{Z}$ and computing the integrals in (6.31) using the Gaussian quadrature (6.36) over each element. In particular, using the notation (6.39), we obtain

$$\left(r^2 \partial_t \mathbf{F}^t, \phi \right)_D = \left(r^2 X \mathbf{F}^r, \partial_r \phi \right)_D - \sum_{j=1}^N \left[r^2 X \mathcal{F}^r \phi \right]_{r_j}^{r_{j+1}} + \left(r^2 \mathbf{s}, \phi \right)_D, \quad (6.41)$$

where r_j is the left vertex of the j -th element in the discretization of the radial coordinate in the line element (6.26).

Considering now test functions with support contained within a given element, S_j , say

$$\phi(r) = l_i(\varphi_j^{-1}(r)) \chi_{[r_j, r_{j+1}]}(r) \quad (6.42)$$

where $l_i(\varphi_j^{-1}(r))$ is the Lagrange polynomial of order i evaluated at the position $x = \varphi_j^{-1}(r)$ of the reference element which is mapped into r , while $\chi_{[r_j, r_{j+1}]}$ is the indicator function and thus equal to one in the interval $[r_j, r_{j+1}]$ and zero elsewhere. Expanding $\mathbf{F}^t$ over the Lagrange basis of S_j

$$[\mathbf{F}^t \circ \varphi_j](x) = \sum_{k=0}^D \mathbf{F}_{jk}^t l_k(x), \quad (6.43)$$

we obtain a set of coupled ordinary differential equations for the coefficients $\mathbf{F}_{ji}^t$

$$r_{ji}^2 \partial_t \mathbf{F}_{ji}^t = \left(r^2 \mathbf{X} \mathbf{F}^r, \partial_r l_i \right)_{j,D} - \frac{1}{w_i |\varphi'_j|} \left[r^2 \mathbf{X} \mathcal{F}^r l_i \right]_{r_j}^{r_{j+1}} + r_{ji}^2 \mathbf{s}_{ji}, \quad (6.44)$$

where $r_{ji} = \varphi_j(x_i)$ and we used the fact that $(l_i, l_k)_{j,D} = \delta_{ik}$. Hereafter we will consider only affine maps φ_j , so that the transformation between the reference element T and the finite element S is given by

$$\varphi_j(x) = \left(\frac{x+1}{2} \right) r_{j+1} - \left(\frac{x-1}{2} \right) r_j, \quad (6.45)$$

and $|\varphi'_j| = |r_{j+1} - r_j|/2$.

Figure 6.2 offers a schematic representation of the spacetime grid structure in EDGES. The $D+1$ collocation nodes (*i.e.*, the four filled squares for a polynomial of degree three) are generated and stored on the reference element, T , and then mapped onto the finite element S_j with an affine transformation, φ_j . Each point on the boundary of an element is associated with two distinct degrees of freedom, so that the corresponding values of the functions can be different there.

This system of equations is coupled with the equations for the evolution of the metric quantities, described in Sect. 6.3.2, and closed with an EOS. In particular we use an ideal-fluid (or Γ -law) EOS

$$p = (\Gamma - 1) \rho \varepsilon. \quad (6.46)$$

Expression (6.44) is clearly a collocation scheme for $\mathcal{D}$, S and τ on the grid illustrated in Fig. 6.2, since we can interpret the expansion coefficients as collocated values, *i.e.*, $f_{ji} = f(r_{ji})$ for any function f . In our code we evaluate the fluxes, $\mathcal{F}^r$, using the relativistic HLL approximate Riemann solver and evolve numerically the resulting set of equations using a second-order strongly stability preserving (SSP) Runge-Kutta method [154]. We note that, as expected, if $D = 0$, then $l_j = \text{const.}$, the first term on the right-hand side vanishes and we are left with a standard FV scheme.

A final note concerns the CFL condition needed to ensure the linear stability of the scheme. The linear stability of Legendre pseudospectral methods for hyperbolic equations has been studied in [153], where it has been shown that L^2 -stability can be obtained if $\Delta t \sim D^{-2}$. As a result, in our code, we use a timestep given by

$$\Delta t = \frac{C_{\text{CFL}}}{(D+1)^2} \Delta x, \quad (6.47)$$

where Δx is the size of the smallest element, the C_{CFL} is a coefficient that is reminiscent (but distinct from) the traditional “CFL factor” and has to be determined empirically. This coefficient is usually taken to be $C_{\text{CFL}} \sim 0.2 - 0.3$, but our numerical experience (at least in 1D) seems to suggest that this condition is overly restrictive as stable evolutions can be obtained with $C_{\text{CFL}} \sim 1$ in most situations.

6.3.2 Coupling with the spacetime

As discussed in Sect. 6.2.4, in spherical symmetry and with the gauge chosen, the Einstein equations are simple constraints on each spatial slice and thus

ODEs of the form

$$\partial_r u(t, r) = f(t, r), \quad (6.48)$$

which could be easily integrated to very high accuracy. However, instead of using a standard Runge-Kutta method for their integration, we found that a more efficient and accurate approach to solve equations (6.29) and (6.30) consists of using an implicit discontinuous spectral ODE solver which makes use of the same grid as the hydrodynamical variables. Such an approach has the advantage of obtaining a numerical solution with a degree of accuracy which is of the same order or higher with respect to the one attained for the hydrodynamical variables, without requiring a very small step, as would have been the case for a Runge-Kutta method.

In order to implement this approach it is necessary to allow the solution to be discontinuous across the elements, while imposing the continuity using an interior penalty technique [21, 22]. In particular the discontinuous Galerkin formulation that we use reads

$$\sum_{j=1}^N \left[\int_{r_j}^{r_{j+1}} \partial_r u \phi \, dr - \mu_j \llbracket u \rrbracket_j \phi_j \right] = \sum_{j=1}^N \int_{r_j}^{r_{j+1}} f \phi \, dr, \quad (6.49)$$

where $\phi_j := \phi(r_j^+)$, $\llbracket u \rrbracket_j := u(t, r_j^-) - u(t, r_j^+)$ is the jump term, and $\mu_j \sim 1/\Delta r_j$ is the penalization coefficient. It is straightforward to see that this penalization term, which we refer to as “upwind penalization” in contrast to the usual symmetric penalization term $\sum_{j=1}^N \mu_j \llbracket u \rrbracket_j \llbracket \phi \rrbracket_j$, has the effect of enforcing $u(t, r_j^+) = u(t, r_j^-)$ without affecting the equation for the collocation value in r_j^- .

As customary in the context of DG-SEM-NI methods, we approximate (6.49) as

$$(\partial_r u, \phi)_D - \sum_{j=1}^N \mu_j \llbracket u \rrbracket_j \phi_j = (f, \phi)_D, \quad (6.50)$$

so that if we arrange the values of $u(t, r)$ and $f(t, r)$ on the collocation points in two arrays $\mathbf{u}(t)$ and $\mathbf{f}(t)$, the previous can be written as

$$\mathbf{A} \mathbf{u}(t) = \mathbf{f}(t), \quad (6.51)$$

where $\mathbf{A}$ is a large-sparse matrix. As $\mathbf{A}$ does not depend on t , in EDGES this matrix is pre-computed, stored and pre-factorized using UMFPACK (see e.g., [282]). At each step, then, we have simply to compute $\mathbf{f}(t)$ and use the factorized version of $\mathbf{A}$ to efficiently compute $\mathbf{u}(t) = \mathbf{A}^{-1} \mathbf{f}(t)$.

6.3.3 Limiters, spectral viscosity and spectral filtering

It is well known that high-order numerical methods suffer from numerical oscillations in the presence of discontinuities (Gibbs phenomenon). If these oscillations are not suitably handled, they tend to grow out of control and destabilize the method. To overcome this difficulty several different “flattening techniques” have been developed in the context of FD, FV and spectral methods to artificially lower the order of the methods in the presence of shock waves. Some examples are artificial viscosity, flux limiting, PPM or ENO/WENO reconstruction. Many of these techniques are implemented in EDGES and can be activated during the evolution.

Complications of spherical symmetry

While discontinuous Galerkin methods can in general be used in combination with the large majority of flattening techniques, this is not the case in spherical symmetry. The reason for this is that if we consider a function u and interpret any of the flattening methods above as the application of an operator $\mathcal{L}$ to u , it can be shown that, for the above mentioned flattening techniques,

$$\int_{\Omega_j} u \, dx = \int_{\Omega_j} \mathcal{L}(u) \, dx, \quad (6.52)$$

which is the desired behaviour for a conservative scheme in 1D. However, in the case of spherical symmetry, the conservation property that we should satisfy reads

$$\int_{\Omega_j} u r^2 \, dr = \int_{\Omega_j} \mathcal{L}(u) r^2 \, dr. \quad (6.53)$$

Unfortunately, the property (6.53) does not hold for almost all of the stabilization techniques which will be discussed in Sect. 6.3.3 and which are generally coordinate dependent (the only exception being given by the spectral viscosity method which is inherently a differential operator). In the case of FV codes this is not a problem because the reconstruction is only used in the computation of the fluxes and the volume averages are evolved using the correct cell volume, but in a scheme that works with the actual point-wise value of the solution, this leads to unacceptable variations of the volume integral of u (e.g., of the total “mass”) in the elements with $r \ll 1$ and $r \gg 1$, which in turn leads to the development of large numerical errors and/or instabilities. For this reason it is necessary to introduce a correction to the flattening procedure in the spherically symmetric case.

A naive way to obtain the desired result would be to modify the operator $\mathcal{L}$ as

$$\tilde{\mathcal{L}}(u) = \frac{1}{r^2} \mathcal{L}(ur^2), \quad (6.54)$$

so that $\tilde{\mathcal{L}}$ now satisfies (6.53). In practice, however, this strategy results, as can be easily foreseen, in large numerical errors near $r = 0$ and is thus of little use. For this reason we had to adopt a more radical approach and add a “correction step” after the application of $\mathcal{L}$ to enforce (6.53). In particular we implemented three different strategies which we discuss below.

The first one, which we refer to as “dummy” correction, consists of simply adding to $\mathcal{L}(u)$ a function C , constant on every element, such that (6.53) is satisfied. This is a very simple approach and has the advantage of not increasing the total variation of $\mathcal{L}(u)$ inside the single elements, as conservation is basically obtained at the expense of the amplification of the jumps of $\mathcal{L}(u)$ across the elements. In this way the additional total variation generated by the corrective procedure is concentrated over the grid points which constitute the “finite volume part” of the method.

The second one, which we refer to as “bubble” correction, consists of adding to $\mathcal{L}(u)$ a function b , which is a bubble function over each element, i.e., a function which is zero at the edges of the element, such as

$$[b \circ \varphi_j](x) = K_j(1 - x^2), \quad (6.55)$$

where the K_j 's are chosen in each element so that (6.53) is satisfied. This approach has the advantage of avoiding the creation of artificial discontinuities at the boundary of the elements, at the price of a small increase in the total variation of $\mathcal{L}(u)$.

Finally the third one, which we refer to “intrinsic” correction, consists of modifying the action of $\mathcal{L}$ so that $L_q = \mathcal{L}(L_q)$ for $q \leq 2$. In this way both (6.52) and (6.53) are satisfied. In practice this is obtained by overwriting the low-order coefficients of the discrete Legendre transform (DLT) of $\mathcal{L}(u)$ with the ones of u . This method has the advantage of not introducing any unwanted extra total variation and, for this reason, is the one that has the best mathematical properties. The only limitation is that it effectively weakens $\mathcal{L}$ and could thus fail to completely removing the Gibbs oscillations.

All these three techniques can be basically used in combination with any of the stabilization methods which are described in the following Section.

Stabilization techniques

The most commonly used stabilization technique in discontinuous Galerkin methods is based on slope limiting. These methods were originally developed for FV schemes [196], but are easily modified to work with discontinuous Galerkin schemes as done by Cockburn and Shu [93], who introduced the “ $\Lambda\Pi_1$ ” limiter based on a generalized version of the “minmod limiter”. In our code we implemented a refined version of this limiter originally proposed by [50] and subsequently improved and extended by [187]. This essentially works by recursively limiting the coefficients of the spectral representation of the solution on the various elements. The main advantage of this technique is that it does not require tuning and it is usually very reliable, while its main limitation is that its use often results in excessive flattening of the solution in the presence of discontinuities. Moreover we found that all these methods perform rather poorly in conjunction with the correction techniques outlined above and for this reason we have rarely used them.

Another method implemented in EDGES is the “spectral viscosity” method first proposed by Maday et al. [205] in the context of Legendre pseudospectral methods for conservation laws. This method consists of the addition to the right hand side of (6.44) of a term

$$- \varepsilon_D (Q \partial_x \mathbf{F}^t, \partial_x \phi)_D, \quad (6.56)$$

where ε_D is a coefficient depending on the number of grid points and Q is a viscosity kernel whose action on a scalar function u reads, in every element,

$$[Qu](x) = \sum_{k=0}^D \hat{Q}_k \hat{u}_k L_k(x). \quad (6.57)$$

Here, $\hat{u}_k$ are the coefficients of the DLT of u and $\hat{Q}_k$ are real numbers such that

$$\hat{Q}_k = 0, \quad \text{for } k \leq m_D, \quad (6.58a)$$

$$1 - \left(\frac{m_D}{k}\right)^4 \leq \hat{Q}_k \leq 1, \quad \text{for } k > m_D. \quad (6.58b)$$

Note that m_D effectively plays the role of the cut-off frequency of the filter. Meday et al. [205], were able to show that, in the context of scalar conservation laws in 1D and using only one domain, this method is able to stabilize the scheme, without spoiling its spectral accuracy if

$$\varepsilon_D \sim D^{-\theta}, \quad \text{for } D \gg 1, \quad (6.59)$$

where θ is just an exponent such that $0 < \theta \leq 1$ and

$$m_D \sim D^{q/4}, \quad \text{for } 0 < q < \theta \leq 1. \quad (6.60)$$

In our code we have set

$$\varepsilon_D = \mu \frac{\Delta t}{N} \frac{1}{D} \quad (6.61)$$

and

$$\hat{Q}_k = f, \quad \text{for } k \leq m_D, \quad (6.62a)$$

$$\hat{Q}_k = f + 1 - \left(\frac{m_D}{k}\right)^4, \quad \text{for } k > m_D, \quad (6.62b)$$

where μ , f and m_D are set by the user (standard reference values are $\mu = 1$, $f = 0$ and $m_D = 0, 1$).

Besides being particularly dissipative, this method has the important advantage that it can be adapted to spherical symmetry by modifying (6.56) as

$$- \varepsilon_D (r^2 Q \partial_r \mathbf{F}^t, \partial_r \phi)_D. \quad (6.63)$$

In this way, the condition (6.53) is satisfied without the need for adding corrective terms. In EDGES, the term (6.63) is added using an operator-splitting technique, so that its use corresponds to the application of a linear filter. This is discretized with two different techniques: locally on each element, or globally, using an interior penalty technique. We call the latter approach “interior-penalty spectral-viscosity” (IPSV) stabilization. While the first approach is completely local, the second one is able to relax also the jump terms between the elements. The main drawback of this method, however, is that we have found that the quality of the results can be quite sensitive to the tuning of μ and m_D , which are necessarily problem dependent.

The third method implemented in EDGES is usually referred to as “spectral filtering” (see e.g., [76]). The idea behind this technique is to filter the numerical solution, or its time derivatives, with a low-pass filter in order to remove high-frequency components and keep the Gibbs oscillations under control. In the context of Legendre pseudospectral methods, this filtering, which we indicate with $\mathcal{F}_D$, reads

$$[\mathcal{F}_D u](x) = \sum_{k=0}^D \sigma\left(\frac{k}{D}\right) \hat{u}_k L_k(x), \quad (6.64)$$

where $\hat{u}_k$ are the coefficients of the Legendre expansion of u and $\sigma(\eta)$ is a *filter function* of order p in the Vandeven’s sense, that is a function $\sigma \in C^p(\mathbb{R}^+; [0, 1])$

such that

$$\cdot \sigma(0) = 1; \quad (6.65)$$

$$\cdot \sigma^{(k)}(0) = 0, \quad \text{for } k = 1, 2, \dots, p-1; \quad (6.66)$$

$$\cdot \sigma(\eta) = 0, \quad \text{for } \eta \geq 1; \quad (6.67)$$

$$\cdot \sigma^{(k)}(1) = 0, \quad \text{for } k = 1, 2, \dots, p-1; \quad (6.68)$$

where $f^{(p)}(x)$ denotes the p -th derivative of f in x . These schemes were studied by Hesthaven and Kirby [167], who showed that if $\sigma(\cdot)$ is a filter of order p , and $u \in C^p(-1, 1)$, then

$$|u(x) - \mathcal{F}_D u(x)| \leq D^{1-p} \|u^{(p)}\|_{L^2(-1,1)}, \quad (6.69)$$

so that, if $p \sim D$, the filter does not spoil the spectral accuracy of the method.

The effects of filtering on the stability of the numerical method can be intuitively understood from the fact that some filters, in particular the exponential filter of order p ,

$$\sigma(\eta) = \exp \left[-\mu(D+1) \frac{\Delta t}{\Delta x} \eta^p \right], \quad (6.70)$$

can be interpreted as equivalent to the use of numerical diffusion of order p and strength μ [221]. At present, however, a mathematical understanding of the impact of filtering on the accuracy and the stability of Legendre pseudospectral methods is still lacking [167].

EDGES provides a number of different filter functions, but the most used are the exponential filter (6.70), which is only an approximate filter, but which is very popular due to its flexibility [76], and the “Erfc-Log filter” proposed by Boyd [61],

$$\sigma(\eta) = \frac{1}{2} \operatorname{erfc} \left\{ 2p^{1/2} \left(|\eta| - \frac{1}{2} \right) \sqrt{\frac{-\log[1 - 4(\eta - 1/2)^2]}{4(\eta - 1/2)^2}} \right\}, \quad (6.71)$$

which is a semi-analytic approximation of the Vandeven’s filter,

$$\sigma(\eta) = 1 - \frac{\Gamma(2p)}{\Gamma(p)^2} \int_0^\eta [t(1-t)]^{p-1} dt, \quad (6.72)$$

where Γ is the Euler Gamma function. Note that it is customary to use filters of order $p \sim 2D$ or larger when expecting regular solutions and $2D+1$ is a reasonable value for a strong filter such as Erfc-Log.

The main limitation of spectral filtering is that, being a high-order linear method, it cannot completely remove the Gibbs oscillations. This is a consequence of the well-known Godunov theorem stating that no linear monotonicity-preserving method exists of second-order or higher (see *e.g.*, [196]). The idea is, instead, to allow for small oscillations of the numerical solution at the location of shocks, while preventing them from growing out of control. In our numerical experience we found that spectral filtering is a robust and viable alternative to slope limiting in the shock-tube case, where its use often results in a sharper resolution of the discontinuities, and, at the same time, it is well-behaved in the spherically symmetric case. Spectral filtering seems also to be much less




Figure 6.3: Solution for the rest-mass density of a relativistic shock tube with different stabilization techniques. The solid black line represents the exact, analytic solution, while the red crosses and the blue squares refer to the use of the two different DLT filters, *i.e.*, the Erfc-Log and the exponential filter, respectively. The blue triangles and the violet circles refer instead to the use of the IPSV filter and of high-order slope limiters, respectively. To avoid excessive cluttering, only a data point every five is shown. The results were obtained on a grid of $N = 200$ elements with a polynomial of degree $D = 3$.

sensitive than the spectral viscosity methods to the tuning of the parameters of the filter functions. Finally, we point out that, in contrast to IPSV and slope limiters, this technique is completely local in the sense that its application in a given element does not require the knowledge of the solution elsewhere. For these reasons, most of the results that we will show in the next Sections were obtained with the use of spectral filtering.

Before doing that, however, and to illustrate the difference between the various stabilization techniques, we show in Fig. 6.3 a comparison of the results obtained using different filters, while keeping all the other discretization parameters fixed. In particular, we show the results obtained in the case of a relativistic shock tube in flat spacetime. The details of the physical setup will be given in Sect. 6.4.1, here instead we focus on the effects of the different stabilization techniques. As anticipated we find that slope-limiting, even in their high-order variant, results in a much larger smearing of the shock front than all the other methods. The 9th order Erfc-Log filter yields a much sharper resolution of the shock and a stable evolution, even though, as discussed before, filtering does not have such a strong mathematical basis, while slope-limiting is known to yield a total-variation diminishing in mean (TVDM) scheme [93]. The best overall results are the ones obtained with the second-order exponential filter, for which we used a small strength factor, $\mu = 1$, and the IPSV technique with $\mu = 1$ and $\hat{Q}_k = 1$. We remark that the results obtained with these last two

methods are very similar since, as discussed before, the action of an exponential filter is roughly equivalent to that of a spectral viscosity. The results in Fig. 6.3 obtained on a grid of $N = 200$ elements with a polynomial of degree $D = 3$.

6.3.4 Treatment of low-density regions

The treatment of interfaces between vacuum regions and fluid regions is one of the most challenging problems in Eulerian (relativistic) hydrodynamics codes (see *e.g.*, [137, 179, 228]). The most commonly used approach to treat vacuum regions is to fill them with a low-density fluid, such that if a fluid element is evolved to have a rest-mass density below a certain threshold, it is set to have a floor value and zero coordinate velocity [131, 29]. This approach works reasonably well and has been adopted by the vast majority of the relativistic-hydrodynamics codes. Nevertheless, the introduction of this “atmosphere” creates several numerical difficulties. First of all, this approach is gauge dependent and, for example, a star can accrete or lose mass from/to the atmosphere due to oscillations in the coordinate variables. Secondly, this approach often results in the introduction of errors at the surface of the star, which, on the one hand, usually have little influence on the overall dynamics because of the small momenta involved, but which, on the other hand are relatively large if compared with the magnitude of the involved quantities at the surface. Finally, and most importantly, it is possible that the algorithm governing the evolution of the floor can couple with the Gibbs oscillations, leading quickly to their amplification and destabilizing the scheme (After-all, the introduction of an atmosphere treatment is de-facto equivalent to the use of a boundary condition and this can very well lead to an unstable algorithm). For this reason this procedure could not be used in a straightforward way in our code, but required particular care.

Other solutions, such as the use of the equations in Lagrangian form, *e.g.*, [135], the use of moving grids techniques *e.g.*, [155], or the use of suitable limiters at the surface [137], do not suffer from these issues, but are restricted to the spherically symmetric case. As our code is meant to be a prototype code to study the viability of DG methods for relativistic hydrodynamics, the use of these techniques would have defeated our purpose.

As a result, several different approaches were implemented and tested in EDGES to overcome the difficulties with the atmosphere discussed above. The main idea behind all these approaches was to use stronger stabilization techniques at the interfaces between fluid and vacuum regions. Unfortunately all these techniques performed quite poorly because they resulted in an unacceptable lowering of the resolution at the surface and thus in large numerical errors. Moreover, in the spherically symmetric case, the situation is greatly worsened by the fact that these errors tend to build up coherently while traveling towards the center. For example, in the case of a neutron star, they tend to be amplified by a factor $R/r \sim 10^3$, $R \sim 10$ being the radius of the star and $r \sim 0.01$ being the location of the closest grid point to the center.

Within EDGES a solution to this problem was eventually found with the use of a method that is, at the same time, oscillation-free and capable of obtaining high-enough resolution at the surface. In particular, we derived a “sub-element method” approach that we later discovered to be very similar to the spectral volumes (SV) method already suggested by Wang [331]. In this method, we




Figure 6.4: Division of an element into control volumes for the spectral-volume method applied at the fluid-vacuum interface. The values of the solution u on the collocation points (filled squares) are interpreted as volume-averaged values of the solution in appropriate control volumes (hollow diamonds denote the boundaries). These values are then used to generate a reconstructed solution $\mathcal{R}u$ using an adapted slope-limiter method.

first flag those elements in which the rest mass density falls below a certain threshold. Secondly we interpret the value of the solution on the collocation points of those elements as being the volume-averaged value of the solution on appropriated *control volumes*. Finally these values are evolved as in a FV scheme: a linear reconstruction with slope limiting is used to compute the value of the primitive variables at the interface of the control volumes and the HLLC approximate Riemann solver is used to compute the fluxes.

The structure of the grid used for this spectral volume method is shown in Fig. 6.4. The control volumes are simply obtained by considering the points of the dual mesh as the cell interfaces, while their volume has to be corrected to take into account the weights of the Gaussian quadrature associated with the primal grid. The linear reconstruction can be performed only for the interior control volumes, while for the control volumes at the boundary of each element we already have the value of the solution at one side of the control volume, thus we are forced to reconstruct the solution there as a constant. In our tests we found that the “superbee” [130] limiter is the one which guarantees the best results among the ones that we tested. For this reason, all the results shown in the following make use of this limiter.

The important advantage of this approach, with respect to a more traditional approach in which the troublesome elements are split into equal-size cells and a classical FV method is used, see *e.g.*, [324], is that no interpolation is required in the switching from/to the discontinuous Galerkin method. The coupling between the two methods is also very natural and is done through the Riemann solvers between the elements. Again no special treatment is required to handle different type of elements: the DG ones and the SV ones. The main limitation of this approach is that, as we do not increase the number of degrees of freedom in the flagged elements, it has the effect of reducing the accuracy to second




Figure 6.5: First shock-tube test. Shown with solid lines are the exact solution for the rest-mass density (black line), the velocity (blue line), and the specific internal energy (red line). The corresponding numerical solution is represented by the black crosses, the blue triangles, and the red squares, respectively. To avoid excessive cluttering we show only one point every 15 of the numerical solution. The results were obtained on a grid of $N = 500$ elements with a polynomial of degree $D = 3$.

order where it is used. As will be shown in Sect. 6.4.4, this seems not to be limiting the accuracy of our code, probably because the use of the SV method is confined to small regions containing low-density fluid. In the case of a neutron star, for example, the use of the SV method is typically limited to one element containing the surface of the star.

As a final remark we note that, as we are working with a nonuniform grid, the slope limiting method in its standard form, also used by us, is not guaranteed to be TVD and/or second order [43]. Again this seems not to be a problem in practice.

6.4 Numerical tests

In what follows we present the results obtained from some of the tests performed with EDGES. These tests have been chosen because they highlight the capabilities of our code in idealized settings, such as in the simulation of shock tubes and shock waves, and in more “astrophysically motivated” settings, such as in the simulation of the dynamics of spherical stars and of the gravitational collapse to black hole. In all cases considered, the evolution has been made employing an ideal-fluid EOS to take account of non-isentropic transformations, such as shock-heating effects.

6.4.1 Shock tubes

The first tests performed are shock-tube tests and, more specifically, we first present the results obtained in the case of two standard benchmarks for relativistic-hydrodynamics codes described in [209]. These are referred to as “blast-wave” problems 1 and 2 in [209], and are essentially one-dimensional flat-space Riemann problems.

The first problem describes the propagation of a relativistic blast wave through a low-density medium. The solution obtained with EDGES, as well as the analytic solution, are shown in Fig. 6.5, where the solution is computed using polynomials of degree three (*i.e.*, $D = 3$) and 500 elements (*i.e.*, $N = 500$). The stabilization is obtained with the use of an Erfc-Log filter of 9th order, which has been applied directly to the conserved variables in the regions where they fail to be monotone at every timestep. As it can be seen from the figure, the stabilization is strong enough to eliminate any oscillations, thus reducing the order of the method to the expected first-order near the discontinuity, yet with a very small smearing of the shock.

The second shock-tube problem is similar to the first one, but is much more extreme. This benchmark is considered to be a challenging test also for modern HRSC codes [209]. The main difficulty is to handle the very high compression ratios, typical of relativistic hydrodynamics, produced in this case. In Fig. 6.6 we show a comparison of the results obtained with our code when using the exponential filter (6.70) at second-order (*i.e.*, $p = 2$) with $\mu = 1.0$, which is employed in the same way as the Erfc-Log filter in the previous test with the ones obtained with the high-order THC code introduced in Chapters 4 and 5 using the MP5 reconstruction scheme in local characteristic variables. The comparison is made when using the same number, 1000, of cells (for THC) and of elements (for EDGES); for both codes, in fact, we expect a convergence order (on




Figure 6.6: Rest-mass density profile for the second shock-tube test. Shown with different lines are the exact solution (solid black), the solution obtained with the THC code [111] (dotted blue). Both codes use 1000 cells/elements and algorithms at fifth order; EDGES also uses a exponential filter with $\mu = 1, p = 2$.




Figure 6.7: The same as in Fig. 6.6, but with a magnification of the shock and a comparison of the results at different orders D of the polynomial representation of the solution. Note that the accuracy with which the shock location is found increases with D .

smooth solutions) of the type Δx^{-p} , with Δx being the width of each cell/element and p being the order of the scheme. Note, however, that, while in THC high order is obtained with the use of data across multiple cells, in EDGES this is attained with the use of a polynomial representation of the solution in each element. Thus the EDGES code actually uses around six times more grid points than the THC code. Nevertheless, we argue that this is the correct comparison because, as mentioned before, the convergence properties of the two codes once the reconstruction procedure and the order of the polynomial representation is fixed, scale with the width of the cells/elements. This kind of comparison, where results obtained with discontinuous Galerkin and finite volume methods are obtained using the same number of elements/cells, but different number of degrees of freedom has already been discussed in the literature, see *e.g.*, [343].

As can be seen from Fig. 6.6 the quality of the solution obtained by the two codes is comparable over most of the domain, even though the solution obtained by the EDGES code shows signs of Gibbs oscillations ahead of the shock, while the THC yields a more sharp and non-oscillatory solution profile. Overall, the exponential filter can prevent the oscillations from growing and yields the best results when compared to the other flattening techniques discussed above. A more quantitative estimate of the quality of the numerical solution can be obtained by looking at the ratio $\sigma/\sigma_{\text{exact}}$, between the observed compression ratio and the analytic one [209]. We obtain a value of 0.96 with EDGES. At lower resolution, *i.e.*, using only 400 elements, we obtain, for EDGES, a value of 0.60 which is close to the ones reported for PPM at the same resolution: *i.e.*, $\sigma/\sigma_{\text{exact}} \approx 0.54 - 0.69$ [209]. We conclude therefore that in the case of discontinuous solutions, DG methods behave similarly to FV methods with the same number of elements/regions.

Finally Fig. 6.7 shows a magnification of the blast wave and in particular it shows a comparison between the exact solution, the solution obtained with THC and the solutions obtained with EDGES using different polynomial orders. As we increase the order of the polynomials used in the representation of the solution in EDGES we get a slow improvement in the shock location in the numerical solution, even though the maximum of the density seems to be better captured for $D = 5$ (for which we tuned the filter strength). Overall, however, the quality of the solution that we obtain depends only marginally on D . On the one hand, this is expected since the code should converge only to first order in Δx in the presence of shock waves independently from D . On the other hand, this is an important confirmation of the robustness of filtering as it shows that the filtering strength does not need to depend on D .

6.4.2 Spherical shock reflection

Another classical benchmark for relativistic-hydrodynamics codes is the spherical shock-reflection test. This is the relativistic version of the classical Noh test and its setup is described in detail in [209]. The initial data consists of a cold fluid converging at the center of the domain. As the fluid flows towards the center, a hot dense shell of matter is formed and a shock wave is propagated outwards. In this situation it is possible to generate arbitrarily high compression ratios by increasing the Lorentz factor of the ingoing flow. This is a peculiarity of relativistic hydrodynamics and, again, the key point of the test is to assess the capability of the numerical methods to handle such strong




Figure 6.8: The same as Fig. 6.5, but for the spherical shock-reflection test with a Lorentz factor of $W = 1000$. To avoid excessive cluttering in the main frame of the picture we plot only a point every five of the numerical solution, while the inset shows the solution on all the collocation points.

density contrasts.

This problem is also interesting to assess the quality of the correction procedure for the filtering in spherical symmetry. In particular, the quality of the results depends on the used correction procedure: they are of comparable quality when using the “dummy” and the “bubble correction” procedures (the former being slightly better), while the “intrinsic” correction yields a filtering which is too weak to guarantee stability. This is not a surprise since with the intrinsic procedure we are not able to filter the coefficients of order less than two in the Legendre expansion, so that we cannot reduce the order of the method below second even with a strong targeted filter.

In Fig. 6.8 we present the results obtained for a flow with $W = 1000$ solved using polynomials of degree three, on 200 elements and with an Erfc-Log filter with $p = 4$ and a dummy correction. As shown in the figure, the filtering procedure is able to suppress any spurious oscillation and the shock front is captured within two elements, *i.e.*, eight collocation points. The average relative error on the compression level is of 1.4 %, comparable with the 2.2 % reported by Romero et al. [278] using FV with a minmod reconstruction on a grid with the same number of cells. Furthermore, as in Romero et al. [278], the error in ϵ and ρ increases in the inner part of the solution. This is due to the well known wall-heating effect and is due to the small but nonzero dissipative and dispersive features of the numerical methods (see [238, 276] for an extended discussion of this effect).




Figure 6.9: L^2 -error in the spherical accretion test as a function of the number of elements N . Crosses and squares refer to the errors on C_1 and C_2 , respectively, while the black, blue and red colors denote refer to the simulations using third, fourth, and fifth-order polynomials. Finally the dashed lines show the slopes associated with third, fourth, and fifth-order convergence.

6.4.3 Spherical accretion onto a Schwarzschild black hole

Having verified the capability of our code to handle shock waves, we next present some results concerning its accuracy in the case of regular solutions. In particular, we consider the case of spherical accretion onto a Schwarzschild black hole in the Cowling approximation (*i.e.*, with a fixed spacetime). An analytic solution exists in the case of stationary flows and was first presented by Michel [223] and later used as a numerical test for other numerical codes, for example by [278]. This solution can be described in term of two constants:

$$C_1 := \sqrt{-g} J^r, \quad (6.73)$$

and

$$C_2 := \sqrt{-g} T^r_t. \quad (6.74)$$

In our simulation we consider a spherical shell with extent $3 - N/40 \leq r \leq 20$ in the spacetime metric of a Schwarzschild black hole of unit mass, while the initial conditions for the hydrodynamical variables describe a low-density fluid at rest with respect to an observer at infinity. At $t = 0$ we start injecting higher-density fluid from the outer boundary and after a short transient the solution reaches stationarity, allowing us to measure C_1 and C_2 and to compare them with the analytic values fixed by the outer boundary condition.

In Fig. 6.9 we show the L^2 -norm of the error for C_1 and C_2 as a function of the number of elements N , and for different orders D of the polynomials used for the representation of the numerical solution over the single elements. The stabilization was obtained using an exponential DLT filter of order $2D$ and strength




Figure 6.10: The same as in Fig. 6.9, but when the error is shown as a function of the polynomial degree when the number of elements N is kept fixed (black for $N = 20$ and red for $N = 30$). Note that an exponential convergence is measured.

$\mu = 1.0$, corrected using the “dummy” strategy. Clearly, the errors computed from C_1 and from C_2 are basically identical. More importantly the convergence order is the one that one would expect from the theory: third, fourth, and fifth-order using third, fourth, and fifth-order polynomials, respectively.

A complementary measure of the convergence properties of the code is shown in Fig. 6.10, where we report the L^2 -norm of the error obtained after keeping fixed the number of elements, $N = 20, 30$, but changing the order of the polynomials used for the representation of the solution. As expected, SDGM schemes behave in this case as multi-domain spectral methods [75], with the error decreasing exponentially with the polynomial degree. This is indeed what we observe in Fig. 6.10, where the errors decrease exponentially by almost three orders of magnitude for both C_1 and C_2 , and at both resolutions. No sign of saturation appears in the error and we find it remarkable that the same method that was able to capture so sharply the discontinuous solutions in the previous tests, is also able to attain exponential convergence in smooth flows.

6.4.4 Linear oscillations of spherical stars

The results presented so far refer to situations that are somehow idealized and are meant mainly as a way to highlight the code’s properties with respect to shock capturing, filtering correction and accuracy on smooth solutions. We next present the results obtained in tests describing systems of more direct physical interest. Furthermore, in contrast to the previous tests, the matter will not be considered as a test-fluid and the spacetime will be properly evolved.

A first interesting test is the study of linear oscillations of spherical stars. The initial data considered describes the equilibrium configuration of a self-




Figure 6.11: Evolution of the central rest-mass density normalized to its initial value for a stable TOV and evolved with a polynomial of degree five. Different lines refer to different resolutions: solid black, dashed blue and dotted red for $N = 75, 100$ and 125 , respectively. Shown in the insets are the evolution at the initial time (*i.e.*, $t \lesssim 2$ ms) and when the error at the surface is released triggering new small-scale oscillations (*i.e.*, $t \sim 45$ ms).

gravitating fluid sphere, described by a polytropic equation of state, $p = K\rho^\Gamma$, and is obtained by integrating the Tolman-Oppenheimer-Volkoff (TOV) equations, as described by [322], and interpolating the result on the computational grid used by EDGES. The system is then evolved under the sole effects of numerical perturbations, mainly due to the interpolation errors in the generation of the initial data and to the interaction between the star and the atmosphere. Note that in order to avoid the presence of collocation points at $r = 0$, we stagger our numerical grid so that the $r = 0$ point falls at the centre of the first element and we use only polynomials of odd degree.

More specifically, we have considered a TOV constructed with a polytropic EOS having $K = 100$ and $\Gamma = 2$ and whose initial central density is $\rho_c = 1.28 \times 10^{-3}$. This could be taken to represent a stable, nonrotating, neutron star with gravitational mass $M = 1.4 M_\odot$ and areal radius $R = 9.6 M_\odot \simeq 14.2$ km. The degree of the polynomial basis is five and an Erfc-Log filter of order eleven is used to ensure a stable evolution. This filter is corrected with the “intrinsic” procedure outlined in Sect. 6.3.3 and is applied to the right-hand-side of the time stepping scheme. This results in a very weak stabilization algorithm whose main effect is to diffuse back into the atmosphere the numerical errors that would otherwise accumulate at the surface of the star and destabilize high-resolution runs. The latter, in fact, have very low intrinsic viscosity and run for many more timesteps on timescales of 100 ms or longer (we recall that the dynamical timescale for this star is of the order of 0.7 ms).

In Fig. 6.11 we show the evolution of the central density for three of these




Figure 6.12: Power spectral density computed from the first $5000 M_{\odot}$ of the evolution of the central rest-mass density shown in Fig. 6.11. Different line types and colors refer to the different resolutions used (see legend). Shown as vertical dotted lines are the eigenfrequencies computed from linear perturbation theory.

runs using different resolutions. They span a time frame of $20\,000 M_{\odot} \sim 100$ ms, with the grid extending over the range $0 \leq r \leq 15$ and being composed of $N = 50, 75, 100, 125$ and 150 elements.

The dynamics is similar among the three runs, with the central density first exhibiting a “burst” with a variation of the order of 0.12% , which is due to the fact that the star is “cut” by our interpolation algorithm of the initial data there where the density falls below the atmosphere threshold. This phase is magnified by the inset in the top of the figure. In a second stage the “burst” is quickly damped and the star starts to vibrate radially with its characteristic eigenfrequencies and eigenmodes under the effect of numerical perturbations, mainly at the surface of the star. In a third phase these oscillations are damped by the numerical viscosity, which depends on the resolution, and by the crude treatment used to represent the surface of the star. High-frequency modes are damped quickly while low-frequency modes have longer damping times. As a result, towards the end of the simulation we are left only with slowly-damped sinusoidal oscillations associated with the *F*-mode. This is particularly evident in the $N = 75$ run, which has the largest numerical viscosity.

The evolution is stable for all the resolutions that we have considered, but a careful examination of the behaviour of the central density reveals the existence of a fourth stage of the dynamics in some of the runs with lower resolution. More specifically, it is possible to note that during the third phase of the dynamics numerical errors accumulate at the surface of the star, producing small variations in the density profile near the atmosphere. When these variations are large enough so that they cannot be controlled by the employed flattening methods, they are “released” and new energy is pumped into the high-frequency




Figure 6.13: Absolute relative difference between the estimated eigenfrequencies and the ones computed from linear perturbation theory, shown for different modes and as a function of the resolution. The lower-order modes are not considered because their error is smaller than the nominal one of the power spectral density and we report only the values for which a reliable measure of the error was possible. Finally, indicated as black dashed lines are the slopes associated with fifth and sixth-order convergence.

modes, starting a new phase in the dynamics. This happens at around 45 ms for the $N = 75$ run in Fig. 6.11, when a small high frequency component modulated by the F -mode suddenly appears, as shown by the inset at the bottom of the figure. The same happens around 80 ms for the $N = 100$ run, while this phenomenon is not observed within the simulation time for higher resolution runs. Although the onset of this “energy release” at the surface can be delayed with the use of stronger filtering, it occurs only on a secular timescale and it does not affect the stability of the star. As a result, the central density keeps oscillating with a very small amplitude, of the order of $10^{-3}\%$, even if, as discussed in Sect. 6.3.3, perturbations at the surface tend to be greatly amplified when they reach the center as an artifact of the spherical symmetry. For this reason we believe that the energy release does not impact the quality of our simulations.

A traditional measure of the accuracy of general relativistic hydrodynamics codes is the comparison of the power spectrum of the oscillations of TOVs against the values provided by linear perturbation theory. In Fig. 6.12 we make this comparison by showing the power spectrum of the first $5000 M_{\odot}$ of the evolution of the central density for different resolutions. The vertical dotted lines represent the eigenfrequencies computed from perturbation theory, which were kindly provided us by Kentaro Takami and computed using the method described in [336]. We note that, even at the lowest resolution, EDGES shows a perfect agreement between the observed proper frequencies and the perturbative ones for the F -mode and the first four overtones, $H1, H2, H3$ and

H4. Furthermore, as the resolution increases we are able to match more and more modes to the point that, with the $N = 150$ run, we match the first ten modes to a very good precision. We also note that the $N = 150$ run gives evidence of some non-linear features in the spectrum, which shows considerable power also in its high-frequency part.

A more quantitative measure of the convergence of the power spectrum is shown in Fig. 6.13, where we report as a function of the resolution the relative difference between the measured frequencies for the overtones *H6*, *H7*, *H8* and *H9*, and the ones computed from perturbation theory. More specifically, the numerical frequencies were computed from the data during the first $5000 M_\odot$ by using the procedure proposed in [9]. Namely, we computed the discrete Fourier transform (DFT) of the data with an Hanning window and used a three point interpolation of the power spectrum to correct for the incoherency error and determine the correct eigenfrequencies (this procedure is conceptually equivalent to a Lorentzian fit of the peaks of the DFT). Note that the lower-order modes are not shown because their precision is such that the error on those frequencies is well below the nominal uncertainty of the DFT and is basically dominated by the error on our measure. (The values of the error is not shown for those resolutions/frequencies that could not be determined in a reliable way: namely, the *H8* and *H9* modes for the $N = 75$ run.). Overall, we find that the measured convergence rate depends somehow on the specific mode, but it is compatible with a fifth-order convergence rate (*cf.* the fifth and sixth-order convergence rates which are shown as dashed lines). As we will comment later on, this results indicates that our largest source of error on the frequencies is not dominated by the low-order FV approach used at the surface of the star. Rather, the behaviour in Fig. 6.13 shows that in the case of global quantities such as the oscillation frequencies of a TOV, the treatment of the surface is not the most critical element of the “error budget” in EDGES⁸.

As a concluding consideration we note that in a recent work Cerdá-Durán [80] has proposed to measure the numerical bulk viscosity of general relativistic hydrodynamical codes by looking at the damping time of the *F*-mode in the case of oscillating TOVs. In particular, the rate of change with resolution of the damping time (and which clearly increases with resolution) can be used as a measure of the convergence rate of the code. To explore this interesting suggestion we have computed the damping time by analysing more systematically the evolution of the central density. More specifically, using a sampling frequency of $\approx 1 M_\odot$, we have built a discrete signal which was then divided into chunks of 512 points with an overlap of 128 points. A DFT of each chunk was then computed using an Hanning window and the power of the signal at the frequency associated with each mode was computed with a linear interpolation of the absolute value of the DFT. In this way we obtained an estimate of the energy in the mode for each time window. Finally we performed a least square fit of an exponential function to determine the damping time. (We have estimated the accuracy of this procedure to be better than 10%, on the basis of tests performed with signals of known spectral properties.)

In Fig. 6.14 we show the results of this measure when made on different modes. The data for the $N = 50, 75$ and $N = 100$ runs have been obtained

⁸Note that the rapid decrease of the error in the estimate of *H6* and *H7* at high resolution is most probably due to the non-linear effects mentioned above which concentrate power in these higher-order modes, making them sharper and better resolved.




Figure 6.14: Estimated damping rates for various eigenmodes as a function of the resolution as computed by EDGES in 1D and by Whisky in 3D. Black triangles, blue squares and red circles refer respectively to the F , $H1$ and $H2$ modes, while filled/hollow points distinguish the estimates with EDGES from those with Whisky. Finally the dashed lines show the estimated convergence order.

by windowing the evolution in the interval $0 \leq t \leq 5000$ to avoid spurious values due to the fourth phase of the dynamics described above, while for the $N = 125, 150$ runs it was necessary to use the full data in the interval $0 \leq t \leq 20000 \approx 100$ ms to obtain a reliable estimate of the damping time for the F -mode. As a comparison, we also report the results obtained in 3D using the Whisky code [29, 145, 254, 289], when computed using a similar procedure. The simulations with Whisky were done using a PPM reconstruction and the HLLE Riemann solver for the hydrodynamics and using a fourth-order FD scheme for the evolution of the spacetime on a uniform grid also covering $0 \leq r \leq 15^9$.

We find that the convergence order of the damping time of the F -mode with Whisky is one, in agreement with the results reported by [80] for the COCONUT code [112, 113]. In [80] it was argued that the reason why the order reduces to first is that the damping is active mainly at the surface of the star, where the numerical methods are only first order. The results found with EDGES show however a different behaviour, with the order of the damping being 3, 2.5 and 2 for the F , $H1$ and $H2$ modes, respectively. Because the treatment of the surface in EDGES can be seen as a variant of the FV method used by both Whisky and COCONUT and is therefore only first-order accurate, our results suggest that the coefficient of the first-order (and surface-induced) error is much smaller than the coefficient of the error coming from the high-order filtering procedure, which

⁹Note that, since Whisky and EDGES use different gauges, the two resolutions are only roughly equivalent.




Figure 6.15: Evolution of the normalized central rest-mass density of an unstable TOV migrating towards the stable branch for a resolution of $N = 150$. The black solid line refers to a run employing an IPSV stabilization, while the red dashed line refers to a run employing a DLT filter; the blue horizontal line denotes the central rest-mass density of the stable model associated with the initial configuration. The inset shows a magnification of the dynamics around 1 ms, at the first peak in the central density.

then becomes the dominant source of damping. This is in agreement with the high convergence order found in the measurement of the eigenfrequencies (*cf.* Fig. 6.13).

As a side comment we want to point out that we are able to attain higher-than-second order convergence in our results, because the time-step that we use is small enough so that the errors in the spatial discretization are dominating over the errors due to the time evolution, which is only second order.

6.4.5 Nonlinear oscillations of spherical stars: the migration test

As a direct extension of the analysis carried in the previous Section on linear oscillations, we next study large non-linear oscillations which are produced as an equilibrium star model on the unstable branch of equilibria models migrates to the stable branch. This process has been used as a numerical test in 3D codes (see, *e.g.*, [131, 28, 29, 98]), has been analyzed extensively in the past [156, 239, 237] and has gained special interest recently when it was shown that it exhibits a critical behaviour [198, 265].

Here, in particular, we have considered a TOV constructed with a polytropic EOS having $K = 100$ and $\Gamma = 2$, central density $\rho_c = 7.0 \times 10^{-3}$, gravitational mass $M = 1.49 M_\odot$ and areal radius $R = 6 M_\odot \simeq 8.8$ km. The evolution, on the other hand, was made with an ideal-fluid EOS to properly take into account




Figure 6.16: Comparison in the evolution of the normalized central rest-mass density of the migrating star between a simulation in 1D with EDGES (dot-dashed red line), and a simulation in 3D with Whisky (solid black line). Also in this case, the blue dashed line denotes the central rest-mass density of the stable model associated with the initial configuration and the inset shows a magnification of the second peak (see main text for details).

shock-heating effects. The numerical grid covers the region $0 \leq r \leq 30$ and the simulations reported used a polynomial representation of the solution of degree five. Two different stabilization techniques were used. A first one employed an exponential filter of order six with $\mu = 40$, applied to the DLTs of the conserved variables and corrected with the “intrinsic” procedure. A second one used instead an IPSV stabilization with $\hat{Q}_k = 1 - \delta_{k0}$ and strength $\mu = 1$.

To trigger the migration on the stable branch, the star is perturbed with an *outgoing* velocity perturbation of the form

$$v(x) = \frac{A}{2}|x^3 - 3x|, \quad x = \frac{r}{R}, \quad (6.75)$$

where $A = 0.01$. Under the effect of this perturbation the star exhibits a violent expansion and migrates towards a new stable equilibrium configuration with a series of large-amplitude oscillations. During these violent oscillations the exterior layers of the star tend to infall with higher velocity than the interior layers and this leads to the formation of shock waves that heat the neutron star matter and result in the ejection of a small portion of the material of the star.

In Fig. 6.15 we show the evolution of the central density, normalized to its initial value, for two runs employing 150 elements and different stabilization techniques. Because this test does not have an analytic counterpart, we have compared it to the corresponding evolution performed with the Whisky code in a 3D simulation having $N = 100$ grid cells in each direction, PPM reconstruction and the HLLE Riemann solver.




Figure 6.17: Comparison of different stabilization techniques in the migration test. Shown in the two panels is the evolution of the central rest-mass density around the first peak. The left panel refers to a DLT filter on grids of $N = 100$ and $N = 200$ elements (blue solid and dashed red lines), while the right panel refers the IPSV stabilization technique with $N = 100$ and $N = 200$ elements (blue solid and cyan long-dashed lines). Note that the DLT-stabilized run reaches convergence with a smaller number of elements.

The comparison is offered in Fig. 6.16 and shows an extremely good agreement. We note that because the two codes have intrinsically different initial truncation errors, the expansion phase in *Whisky* is slightly delayed with respect to the dynamics produced by *EDGES*, and we account for this difference by shifting the data obtained by *Whisky* in order to obtain an approximate alignment at the first minimum of the density. Notwithstanding the very good agreement, the main difference between the two solutions is the presence of high-frequency oscillations in the central density computed by *EDGES*, and in particular for the DLT runs and near the maximum value of the density. These “spikes” can be tracked down to the propagation of shock waves which are formed in the outer layers of the star during the collapse phase. They have initially a small compression factor, but they also tend to sum up coherently as they travel towards the center, thanks to the assumption of spherical symmetry, so that they result in strong variation of the density near the center of the star. This phenomenon has also been observed in other works in spherical symmetry, employing standard finite-volume schemes, see *e.g.*, [98], but is not usually observed in 2D or 3D simulations and, indeed, it is not present in the results obtained with *Whisky*. In particular *Whisky* employs Cartesian coordinates, which inevitably introduce preferred directions in the grid, thus suppressing the constructive interference and focusing of the spherical waves produced at the surface. Another reason why these spikes are not observed in multi-dimensional solutions could be related to the larger numerical viscosity due to the necessity of using a much coarser resolution than in the 1D case. This is confirmed by the fact that the use of the IPSV stabilization techniques, which allows us to introduce a much larger numerical dissipation, is able to greatly suppress this phenomenon.

A more detailed comparison between the results obtained with the two stabilization techniques is shown in Fig. 6.17, where we offer a comparison between the evolution of the central density around the first bounce obtained with the two approaches and with different resolutions. Note that the DLT runs show signs of spikes in the density for all the resolutions (left panel), while the IPSV runs are much smoother (right panel). On the other hand, the results obtained with the DLT filters are already in the convergent regime and in fact the results obtained with $N = 100$ elements are very similar to those obtained by doubling the resolution. This is not the case for the solutions obtained with the IPSV filter, that seem to be only slowly approaching the DLT ones (*cf.* note in the right panel that the IPSV solutions approaches the DLT one as N goes from 100 to 200). The evidence that the simulations with the DLT filtering are already in the convergent regime and yet show spikes in the evolution can be interpreted, therefore, as a confirmation that the latter are mostly an artifact of the symmetry, which leads to a focusing of the waves travelling towards the center. This does not completely rule out the numerical origin of these spikes, but it suggests that they are originated from either numerical or unresolved, physical, perturbations at the location of the shock which are then amplified as they travel towards the center by the symmetry conditions.

Fortunately, because these perturbations actually carry only a very small energy and are amplified by the symmetry, the evolution of the spacetime variables is totally unaffected. This is shown in Fig. 6.18, where we report the evolution of a representative metric quantity, namely, the lapse function at the center. As can be seen from the figure, the lapse function shows no spikes or




Figure 6.18: Evolution of the lapse function at the center of an unstable TOV migrating towards the stable branch for two different runs using different stabilization techniques (black solid line for the IPSV stabilization, and red dashed line for the DLT filter) and $N = 150$. The inset shows a magnification of the dynamics around the first peak at 1 ms and highlights that the solution is always smooth.

spurious oscillations and appears to be smooth with both the DLT and the IPSV filters, the main difference being the value of the minimum attained during the first bounce with the former stabilization technique.

6.4.6 Gravitational collapse of unstable spherical stars

As a final test we consider another classical testbed in general relativistic hydrodynamics: namely, the gravitational collapse to black hole of an unstable TOV star. The problem of the gravitational collapse to a black hole has been already studied in great detail in a number of different conditions involving 1D, 2D and 3D simulations, as well as different physical conditions (see, e.g., [156, 278, 239, 237, 29, 118, 27, 30, 236, 99, 241, 319]), and has become a standard test of general relativistic codes. For this reason, we will not discuss the details of the dynamics of the collapse and concentrate instead on the quality of the results obtained with EDGES.

For the tests considered here we have evolved an unstable TOV built with a polytropic EOS with $K = 100$ and $\Gamma = 2$, having central density $\rho_c = 4.5 \times 10^{-3}$, gravitational mass $M = 1.6 M_\odot$ and areal radius $R = 6.9 M_\odot \simeq 10.2$ km. In contrast to the migration test, the collapse is triggered by introducing an *ingoing* velocity perturbation of the type (6.75) with $A = -0.01$. The evolution of this system is studied on a grid covering $0 \leq r \leq 15$, staggered about the origin, using polynomials of degree five and employing an exponential filter of order six, with intrinsic correction and strength $\mu = 40$, applied directly on the




Figure 6.19: Evolution of the lapse function at the center of an unstable TOV collapsing to a black-hole for different resolutions on. The solid black line, the dotted blue line and the dashed red line are associated with runs employing a grid of 50, 75 and 100 elements, respectively. The inset shows a magnification of the dynamics around 0.6 ms, highlighting some post-collapse dynamics in the lapse for the low-resolution run, which is however absent at higher resolutions.

conserved variables at every time step.

During the evolution and as expected, the central density of the star shows an exponential increase, which halts when the lapse function collapses to zero, signalling the formation of the black hole and “freezing” the evolution in the inner regions of the star. The maximum value attained by the central density does not have any physical meaning, partly because it depends entirely on the gauge conditions and partly because the gauge used will prevent its growth past a certain value. For this reason it is more interesting to look at the evolution of the lapse function at the center of the star, and which we show in Fig. 6.19. Note that for all the considered resolutions, the lapse function quickly collapses to zero and shows no appreciable subsequent evolution, with the exception of the lowest resolution case, for which a small growth appears again around 0.6 ms. This sudden evolution is the result of errors coming from the surface of the frozen star, where the spatial slice is stretched due to the fact that we fixed our shift gauge condition to zero and, thus, the metric functions present large spatial gradients. These errors, which cannot be compensated by simply increasing the resolution and are a shortcoming of the gauge used, propagate towards the inner regions of the frozen star and can induce a growth of the lapse, restarting the dynamics. This is clearly a resolution-dependent effect which disappears quickly by increasing the resolution (see inset in Fig. 6.19). Apart from this late-time, and low-resolution dynamics, the three curves are on top of each other, signaling a convergent regime, with a rate we measure to be slightly above fourth order before the collapse of the lapse. Afterwards

the convergence order starts to slowly decay, due to the fact that, as the lapse collapses to zero, the equations are only weakly hyperbolic. At the end of the simulation the convergence order reaches a value which is $\gtrsim 1.5$.

In summary, the solution of the gravitational collapse to a black hole with EDGES has been straightforward and indeed a test which is much less demanding than the migration one, at least in terms of the hydrodynamical variables. The only difficulty encountered is a well known one and has to do with the fact that with such a gauge the collapse of the lapse is not compensated by any change in the shift vector, which is identically zero. As a result, as the black hole is produced it quickly produces a stretching of the coordinates at the location immediately outside the “frozen star”. This, in turn, results in the development of strong gradients in the numerical solution which cannot be resolved indefinitely without suitable adaptivity.

6.5 Conclusions

Numerical relativistic hydrodynamics and MHD have seen a tremendous growth in the number and the quality of its results after the introduction of modern high-resolution shock capturing schemes [130]. These schemes have been proven to be of central importance in the modelling of complex systems involving strong gravity and/or high Lorentz factor. Yet, they suffer from important limitations that ultimately impact the accuracy of the obtained results [31]. For this reason the search for better numerical schemes is still on-going (and will always be).

Discontinuous Galerkin schemes were developed to overcome the previously mentioned limitations of finite-volume and finite-difference schemes, while maintaining important properties of these schemes, such as conservation and shock capturing, that made them so successful in a number of applications [93]. For this reason they are a natural candidate as an alternative to more traditional methods also in relativistic hydrodynamics.

In this chapter we have developed the necessary mathematical framework needed for the application of discontinuous Galerkin schemes to relativistic hydrodynamics in curved spacetimes. In particular, we have presented both a manifestly covariant weak formulation of relativistic hydrodynamics and a more traditional one obtained within a $3 + 1$ split. We have then specialized the latter formulation to the spherically symmetric case and implemented it in a new one-dimensional relativistic hydrodynamical code, EDGES, which uses a high-order spectral discontinuous Galerkin method.

The code was tested in a number of situations, including shock waves, spherical accretion, linear and non-linear oscillations of relativistic spherical stars and the gravitational collapse of unstable stars to black holes. Our results show that discontinuous Galerkin methods are able to sharply resolve shock waves and, at the same time, attain very high, spectral, accuracy in the case of smooth solutions. For this reason they constitute an excellent alternative to classical finite-volume and finite-difference schemes for relativistic hydrodynamics, especially in those situations in which shock waves as well as small-scale features of the flows have to be resolved. In light of the promising prospects shown with these tests and of their affinity with a pseudospectral solution of the Einstein equations, we anticipate that discontinuous Galerkin

methods could represent a new paradigm for the accurate modelling of, *e.g.*, oscillating neutron stars or the merger of binary neutron stars or mixed binaries in computational relativistic astrophysics.

As a final remark we note that the success of these methods in multi-dimensional implementations will ultimately depend on the development of techniques such as local timestepping, *hp*-adaptivity and load balancing [75], which are needed to take full advantage of the flexibility of these schemes [297]. This will represent the focus of our future work.

Part III

**Relativistic Radiation
Transport**

Chapter 7

The Filtered Spherical Harmonics Method

7.1 Introduction

One of the main difficulties in solving the radiation-transport equation arises from the multidimensionality of the problem. Radiation is described not only by the position of the radiation carriers, but also by their direction of propagation and energy (or, equivalently, by their momentum), making the problem $(6 + 1)$ -dimensional in the most general case (3 dimensions for the spatial coordinates, 2 for the angular direction, 1 for the energy, and 1 for the time). Another source of complication stems from the fact that many problems consist of regions of strongly-varying optical depth. For example, many astrophysical systems contain a central source of radiation, where the optical depth is high, surrounded by more transparent outer regions. Due to high opacity near the center, radiation migrates out of that region mostly via diffusion, while in the outer parts, it streams to infinity more freely without much interaction with matter. Correspondingly, in the limit of infinite optical depth, the transport equation acquires a parabolic character, while it maintains a hyperbolic one in all other regimes (see, *e.g.*, [227]). Radiation-transport approaches must therefore handle accurately not only these two distinct regimes, but also the transition between the two, which is generally the most difficult to treat. These features make the solution of the transport equation both complex and computationally expensive.

There are two commonly used ways of simplifying the solution of the transport equation. One approach consists of reducing the number of degrees of freedom by assuming spherical or axial symmetry. While this is a good simplification for some problems, there are many other situations in which the system does not possess any spatial symmetries and hence the transport equations need to be solved in three spatial dimensions. Another way of simplifying the problem is via approximating the form of the transport equation (this is equivalent to reducing dimensionality in the momentum space). One of the simplest examples is the diffusion approximation, where one approximates the transport equation with the diffusion equation (*e.g.*, [255]). This makes the equation far simpler and computationally less expensive to solve (*e.g.*, [227]).

However, there is a price to pay for this: Although the diffusion approximation is accurate at high optical depth, it leads to incorrect results in low optical depth, *i.e.*, “transparent”, regions. This can be improved by using flux limiters (*e.g.*, [78]), but the diffusion approximation still cannot correctly capture the anisotropy of radiation commonly encountered in those regions (see, *e.g.*, [245]). Moreover, the treatment of semi-transparent regimes is somewhat artificial because the radiation fluxes in those regions are usually calculated via interpolation between the two fluxes calculated assuming free-streaming and diffusive transport. There are several more accurate ways of approximating the transport (*e.g.*, two-moment schemes with analytic closures, etc.; *e.g.*, [68]), but the solution of the full (6+1)-dimensional transport equation is often not only desirable, but actually necessary for accurate modeling.

One of the most commonly used methods for solving the transport equation in the multidimensional (both in space and momentum) case is the discrete-ordinate (S_N) method, which solves the transport equation along several directions in each spatial zone [78, 245, 309, 148]. Unfortunately, this method has several drawbacks. Most importantly, it suffers from “ray effects” in multidimensional calculations [230]. Due to the discrete nature of the angular representation, in fact, radiation cannot reach regions between these discrete directions, leading to large spatial oscillations in the transport variables. In addition, time-implicit S_N methods require very complex solutions and parallelization procedures [7].

Monte-Carlo methods (*e.g.*, [129, 140, 5]) are often regarded as the most accurate method for radiation transport, but they are also not without drawbacks: Monte-Carlo solutions exhibit statistical noise due to the finite sampling of the phase space. Since this noise decreases only as $N^{-1/2}$, where N is the number of Monte Carlo particles, it can take many particles to produce a sufficiently smooth solution, making large simulations computationally very expensive.

Another approach to radiation transport is the spherical harmonics, or P_N , method. This method is based on an expansion of the radiation intensity (or distribution function) in angles using spherical harmonic basis functions. This results in a hyperbolic system of partial differential equations for the expansion coefficients, which represent angular moments in this basis. The spherical harmonics method has several interesting characteristics. Due to hyperbolicity, the P_N equations approximate radiation as a series of waves with velocity bounded by the speed of light [218]. This restriction is consistent with the transport equation, in contrast to diffusion methods where radiation propagates at infinite velocity. Moreover, the spherical harmonics expansion exhibits formal spectral convergence to the true solution. Such an expansion also preserves the rotational invariance of the solution, unlike S_N methods, where the absence of such invariance leads to the ray effects mentioned above.¹ Another advantage of the P_N method is that it generally uses less memory to model a given angular distribution at a given accuracy compared to, *e.g.*, the S_N method. A P_N truncation is roughly equivalent to a S_{N+1} solution, but, while the former has $(N+1)^2$ degrees of freedom, the latter has $2(N+1)^2$, thus it requires roughly twice as much memory. Given that memory requirements represent

¹Here, rotational invariance means that the operators of angular discretization and of rotation in space are commutative. In other words, the result of any rotation and then of a spherical harmonics angular discretization is the same as that of an angular discretization and then of a rotation [62]. In the S_N method, this is true only for those rotations which map the angular grid onto itself.

one of the main difficulties in 3D radiation-transport calculations, a factor two in memory saving represents a significant advantage.

However, spherical harmonics methods also have some negative aspects. Most importantly, in transparent regions, the solution to the P_N hyperbolic system exhibits non-physical oscillations. These oscillations are related to the so-called Gibbs phenomenon that occurs when non-smooth functions are approximated with smooth basis functions [62]. The worst consequence of such oscillations is that they can cause the radiation intensity to become negative, which may lead to negative matter temperatures when the radiation transport is coupled to the matter energy equation [218, 242]. There have been several attempts to address this problem [244, 242, 69, 215, 243, 164]. One of the most efficient, robust, and accurate approaches is the one by McClarren & Hauck [215], in which filters are proposed to remove oscillations of the radiation intensity. By filtering out the oscillations, McClarren & Hauck [215] were able to avoid negative solutions while maintaining high angular accuracy. Their filtering approach also has the advantage of being easily extendable to high-order P_N expansions, producing formal convergence to the transport solution and preserving the equilibrium diffusion limit.

Although the results of [215] strongly suggest the idea of applying a filter to spherical harmonics expansions is an efficient, robust, and accurate way of doing P_N radiation transport, there remain several open questions. For example, [215] considered only one type of filter, the so-called spherical-spline expansion [62]. There are other types of filters that have some interesting theoretical and numerical properties [62], which might be a more optimal choice for application to P_N transport. Moreover, McClarren & Hauck [215] presented results only for the 2D case, leaving open the question of how well their filtering approach performs in 3D. Also, this filtering scheme, as realized by [215], and as we will discuss in detail later in this chapter, does not have a clear continuum limit as the spatial resolution and/or timestep approaches zero, which implies that the solution cannot be studied for spatial convergence.

In this chapter, we reconsider the P_N scheme and extend the work by McClarren & Hauck [215] in a number of ways. Firstly, we reformulate the filtering procedure in such a way that it acquires a clear continuum limit. Secondly, we investigate a wider range of filter types and strength parameters. Finally, we perform calculations both in 2D and 3D. For this we have developed the new radiation transport code *Charon*, whose ultimate goal is that of performing 3D, general relativistic multi-energy, multi-angle and velocity-dependent radiation-transport calculations. In this chapter, we present the first step towards this goal.

Charon uses the semi-implicit scheme of McClarren et al. [217] for time integration. In this scheme, the streaming parts are treated explicitly with a second-order Runge-Kutta method, while the matter-coupling terms are treated implicitly because of the stiffness introduced by the coupling. Also, the implicit system of [217] is local to each spatial element, and for linearized matter-coupling terms this implicit integration scheme becomes trivial [217].

Although the timestep in the semi-implicit approach is limited by the Courant-Friedrichs-Lewy (CFL) condition based on the speed of light, this is not a serious drawback for the kind of applications that we have in mind. In radiation-hydrodynamics simulations, the timestep size would still be limited by the dependence of the matter properties (*e.g.*, opacities and emissivities) on

the matter temperature (and electron fraction, if we are dealing with neutrinos with conserved lepton number). Moreover, the semi-implicit approach is relatively easy to parallelize (*e.g.*, via domain decomposition) and has lower communication requirements compared to fully implicit schemes. This should translate to significant advantages in massively-parallel large-scale radiation-hydrodynamics simulations. Moreover, in many relativistic systems, the sound speed of the fluid is comparable to the speed of light, and thus the radiation and hydrodynamics timescales are comparable, reducing the extra cost in an explicit treatment of the radiation streaming. The spatial discretization in Charon is based on the asymptotic-preserving linear discontinuous Galerkin (DG) scheme [204, 217].

The chapter is organized as follows. In Section 7.2, we introduce the concept of the radiation distribution function and the relativistic Boltzmann equation. In Section 7.3, we describe the numerical schemes used in our code for frequency (Section 7.3.1), angular (Section 7.3.3), spatial (Section 7.3.5), and time discretization (Section 7.3.6). In Section 7.4, we present numerical tests of these schemes. Finally, we summarize our results and provide our conclusions in Section 7.5.

7.2 The relativistic Boltzmann equation

In this section we introduce the Boltzmann equation for radiation transport in special-relativity. This is a particular case of the general-relativistic transport equation discussed in Chapter 2, however, here, we are going to give an independent treatment, using a notation which is more common in the radiation transport literature.

7.2.1 The distribution function for radiation

Radiation is usually described in terms of the specific radiation intensity, I , defined such that

$$d\mathcal{E} = I \cos \theta dA dv d\Omega dt, \quad (7.1)$$

represents the energy of radiation in frequency range dv centered around ν , traveling in direction Ω confined to a solid angle element $d\Omega$, which crosses, within time interval dt , an area dA oriented such that θ is the angle between the normal to the surface dA and direction Ω (*e.g.*, [255]). In the case of neutrino transport and/or in the relativistic case, it is more convenient to work directly with the distribution function, F , which gives the density of radiation carriers at a given point in phase space. The reason for this is that (1) the distribution function is a Lorentz-invariant quantity [227], and (2), as we discuss in more detail later, the distribution function allows us to compute the number density and the energy density of the radiation in a more natural way.

In order to introduce the distribution function, we first discuss the notion of single-particle phase space in special relativity. In this picture, the particles are described in terms of their positions in spacetime, x^μ , and their momentum four-vector, p^μ , as measured in a fiducial inertial frame. Using the normalization condition for timelike vectors, the four-momentum has only three independent components, which can be expressed in terms of three spatial components, p^i ,

or in terms of radiation frequency, ν , and two angles (θ, ϕ) that describe the propagation direction:²

$$p^\mu = \frac{h\nu}{c} (1, \cos \phi \sin \theta, \sin \phi \sin \theta, \cos \theta). \quad (7.2)$$

Since we wish to define the distribution function in terms of p^i , or, equivalently, in terms of ν, ϕ and θ , we need to construct a Lorentz-invariant volume element, $d\Pi$, over the manifold of the allowed momentum four-vectors. This is accomplished with the choice [79]

$$d\Pi = \frac{dp^1 dp^2 dp^3}{-p_0} = \frac{h^2 \nu}{c^2} d\nu d\Omega. \quad (7.3)$$

The distribution function is then defined in such a way that the quantity

$$dN = F p^\mu t_\mu d^3x d\Pi = \frac{h^3 \nu^2}{c^2} F d^3x d\nu d\Omega, \quad (7.4)$$

is the total number density of radiation carriers in a spatial volume element d^3x and phase-space volume element $d\Pi$ with trajectories traversing a $t = \text{const}$ hypersurface with normal $\vec{t} = \partial_t$. Here, t_μ is the μ covariant component of the vector $\vec{t} = \partial_t$.³

Since $d^3x = dA \cos \theta dt$ and the energy per particle is given by $h\nu$, we have $d\mathcal{E} = h\nu dN$. Using this and comparing equation (7.1) with (7.4), we obtain:

$$I = \frac{h^4 \nu^3}{c^2} F. \quad (7.5)$$

Note that since dN is a scalar, F is also a scalar quantity.⁴

The relativistic Boltzmann equation

The special-relativistic Boltzmann equation can be written as [227]

$$p^\mu \frac{\partial F}{\partial x^\mu} = \mathbb{C}[F], \quad (7.6)$$

where $\mathbb{C}$ is the collisional term describing the interaction of radiation with matter, while the left-hand-side of the equation describes the propagation of radiation. In order to compute $\mathbb{C}$, we express it in terms of the absorption,

² Assuming the radiation carriers to be massless, their energy and frequency are simply related as $\varepsilon = h\nu$.

³ Note that d^3x is not a Lorentz-invariant quantity, while $p^\mu t_\mu d^3x$ is one.

⁴ Note that this relation slightly differs from the one frequently encountered in the neutrino-transport literature (e.g., [72]):

$$I = \frac{\varepsilon^3}{h^3 c^2} g F = \frac{\nu^3}{c^2} g F,$$

where g is the statistical weight of the particles ($g = 1$ for massless neutrinos, $g = 2$ for photons) and ε is the particle energy. This is due to three reasons: First, our specific intensity given by Eq. (7.1) is defined in terms of energy per frequency interval, in contrast to energy per energy interval, as used in the neutrino-transport literature. Moreover, our distribution function already contains the factor g , as can be seen from its relation to the total number density of radiation carriers given by Eq. (7.4). Finally we use the Lorentz-invariant volume element given by Eq. (7.3) instead of d^3p .

emission and scattering coefficients. To do that, we start by considering the evolution equation for the intensity of radiation [255],

$$\frac{1}{c} \frac{\partial I}{\partial t} + n^i \frac{\partial I}{\partial x^i} = \eta - \kappa I + \frac{\kappa_s}{4\pi} \int \frac{v}{v'} K(v', \vec{n}' \rightarrow v, \vec{n}) I(v', \vec{n}') d\Omega' dv', \quad (7.7)$$

where η represents the radiative emissivity of the matter, κ is the total extinction coefficient and combines the absorption and scattering coefficients⁵ κ_a and κ_s , *i.e.*, $\kappa = \kappa_a + \kappa_s$, and K is the scattering kernel, expressing the probability of scattering from a given angle and frequency over to another angle and frequency [255].

Using equations (7.2), (7.5), and (7.7), it is easy to obtain an equation for F in terms of the above extinction coefficients:

$$p^\mu \frac{\partial F}{\partial x^\mu} = \frac{c^2}{h^3} \frac{\eta}{v^2} - h\nu\kappa F + \frac{h\nu\kappa_s}{4\pi} \int \left(\frac{v'}{v}\right)^2 K(\vec{p}' \rightarrow \vec{p}) F(\vec{p}') h dv' d\Omega'. \quad (7.8)$$

Notice that, since F is a scalar, so is $\mathbb{C}[F]$, thus we find the classical result that η/v^2 , and $\nu\kappa$ are invariant quantities [227].

Our scheme is in principle able to handle any type of scattering kernel, but for simplicity, here we will only consider the case of elastic scattering, *e.g.*, scattering in which the radiation energy does not change. In this case, the scattering kernel can be expressed as [72],

$$K(v', \vec{n}' \rightarrow v, \vec{n}) = [1 + \sigma_a \vec{n} \cdot \vec{n}'] \delta(v - v'), \quad (7.9)$$

where the scattering anisotropy is modeled using only one coefficient σ_a .

7.3 The Charon Code

In general, the distribution function, F , is a function of 7 variables: the time and spatial coordinates, x^μ , the frequency ν and the angles of propagation φ and θ . These variables are usually defined either in the Eulerian (inertial) frame or in the comoving frame (*i.e.*, a set of frames, each of which has a velocity that instantaneously equals that of the matter element, *e.g.*, [227, 171]). In the case of static matter, which is the one considered in this chapter, these two frames are identical. In the scheme implemented by the Charon code, the distribution function is expanded in the spatial coordinates using the linear DG basis and in the angular variables using spherical harmonics. The frequency is treated using the multi-group approach. This yields a large system of ordinary differential equations that is then evolved in time using a semi-implicit time integrator. The details of the discretization are discussed in this Section.

7.3.1 Frequency discretization

We consider the case in which the distribution function has compact support in a frequency space given by the interval $[0, \nu_{\max}]$. Although this is not strictly valid in the general case, radiation usually has negligible contribution above

⁵Hereafter, the absorption or scattering extinction coefficients are defined as absorption or scattering opacities or inverse mean-free paths.

some cut-off frequency. Therefore, in many practical applications, one can choose $\nu_{\max}$ to be sufficiently large so that there is little radiation beyond this frequency. For simplicity of illustration, we introduce a uniform grid in this frequency space as $\nu_n = n \Delta\nu$, $n = 0, 1, \dots, N_\nu + 1$, where $\Delta\nu = \nu_{\max}/(N_\nu + 1)$ (an extension to a non-equidistant grid is conceptually trivial). The associated intervals $[\nu_n, \nu_{n+1}]$ are commonly called frequency or energy groups. Using these groups, we can construct an orthonormal basis $\{\chi_n\}_{n=0}^{N_\nu}$ as

$$\chi_n(\nu) = \begin{cases} 1/\sqrt{V_n}, & \text{if } \nu \in [\nu_n, \nu_{n+1}], \\ 0, & \text{otherwise,} \end{cases}, \quad V_n = \int_{\nu_n}^{\nu_{n+1}} h^3 \nu^2 d\nu = \frac{h^3}{3} (\nu_{n+1}^3 - \nu_n^3). \quad (7.10)$$

We then expand a function $f \in L^1(0, \nu_{\max})$ on this basis as (for clarity we report the summation symbols in the expressions below)

$$f_{N_\nu}(\nu) = \sum_{n=0}^{N_\nu} f^n \chi_n(\nu), \quad f^n = \frac{1}{\sqrt{V_n}} \int_{\nu_n}^{\nu_{n+1}} f(\nu) h^3 \nu^2 d\nu. \quad (7.11)$$

The truncated expansion, f_{N_ν} , is then a first-order accurate (in L^1 -norm) approximation of f . We point out that, thanks to our choice of basis (7.10), the final expansion of the distribution function will involve integrals performed with respect to the volume element in (7.3). This allows us to ensure exact conservation of the number of radiation particles in the numerical treatment of the transport equation. Also, this choice of the basis can be easily generalized to the case of curved spacetimes (which will be the subject of our future work), where ensuring conservation of radiation particles is particularly involved [77].

7.3.2 Real spherical harmonics

This section is dedicated to the derivation of the real spherical harmonics, whose implementation in Charon has been particularly advantageous. We start by recalling that the spherical harmonics, Y_ℓ^m are the eigenfunctions of the Laplace-Beltrami operator, Δ , on the unit 2-sphere:

$$\Delta Y_\ell^m = -\ell(\ell+1) Y_\ell^m.$$

Spherical harmonics are usually written, in complex form, as

$$Y_\ell^m(\varphi, \theta) = e^{im\varphi} P_\ell^m(\cos \theta),$$

where $-\ell \leq m \leq \ell$, P_ℓ^m are the associated Legendre functions, see e.g., [62],

$$P_\ell^m(x) = (1-x^2)^{m/2} C_{\ell-m}^{m+1/2}(x), \quad m \geq 0,$$

and C_n^α are the Gegenbauer polynomials of index α and degree n . We, also, use the standard convention that

$$P_\ell^{-m}(x) = (-1)^m \frac{(\ell-m)!}{(\ell+m)!} P_\ell^m(x). \quad (7.12)$$

The associated Legendre functions of index $m \geq 0$ are orthogonal in $[-1, 1]$ with unit weight [6],

$$\int_{-1}^1 P_\ell^m(x) P_{\ell'}^m(x) dx = \frac{2(\ell+m)!}{(2\ell+1)(\ell-m)!} \delta_{\ell\ell'},$$

while the associated Legendre functions of degree ℓ and index $m, m' \geq 0$ are orthogonal in $[-1, 1]$ with weight $(1-x^2)^{-1}$ [6],

$$\int_{-1}^1 P_\ell^m(x) P_{\ell'}^{m'}(x) \frac{dx}{1-x^2} = \frac{(\ell+m)!}{m(\ell-m)!} \delta_{mm'}.$$

The spherical harmonics with index $m, m' \geq 0$ are then orthogonal with unit weight on the sphere:

$$\int_{S_1} Y_\ell^m(\varphi, \theta) Y_{\ell'}^{m'}(\varphi, \theta) d\Omega = \int_0^{2\pi} e^{i(m-m')\varphi} d\varphi \int_0^\pi P_\ell^m(\cos \theta) P_{\ell'}^{m'}(\cos \theta) \sin \theta d\theta \quad (7.13)$$

$$= \int_0^{2\pi} e^{i(m-m')\varphi} d\varphi \int_{-1}^1 P_\ell^m(x) P_{\ell'}^{m'}(x) dx = \quad (7.14)$$

$$= 2\pi \frac{2(\ell+m)!}{(2\ell+1)(\ell-m)!} \delta_{mm'} \delta_{\ell\ell'}. \quad (7.15)$$

We can then redefine the spherical harmonics as

$$Y_\ell^m(\varphi, \theta) = \sqrt{\frac{(2\ell+1)}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} e^{im\varphi} P_\ell^m(\cos \theta) = N_\ell^m e^{im\varphi} P_\ell^m(\cos \theta).$$

It is easy to see, that for all m, m' , thanks to the normalization and the convention (7.12), we have

$$\int_{S_1} Y_\ell^m(\varphi, \theta) Y_{\ell'}^{m'}(\varphi, \theta) d\Omega = \delta_{mm'} \delta_{\ell\ell'}.$$

We can construct a real basis from the spherical harmonics by defining

$$Y_{\ell m} = \begin{cases} \frac{1}{\sqrt{2}}(Y_\ell^m + (-1)^m Y_\ell^{-m}), & \text{if } m > 0, \\ Y_\ell^0, & \text{if } m = 0, \\ \frac{1}{i\sqrt{2}}(Y_\ell^{-m} - (-1)^m Y_\ell^m) & \text{if } m < 0. \end{cases}$$

Using again the fact that, $N_\ell^{-|m|} P_\ell^{-|m|} = (-1)^{|m|} N_\ell^{|m|} P_\ell^{|m|}$, we obtain

$$Y_{\ell m}(\varphi, \theta) = \sqrt{2} N_\ell^m \cos(m\varphi) P_\ell^m(\cos \theta), \quad m > 0,$$

$$Y_{\ell m}(\varphi, \theta) = \sqrt{2} N_\ell^{|m|} \sin(|m|\varphi) P_\ell^{|m|}(\cos \theta), \quad m < 0,$$

which is the wanted expression for the real spherical harmonics.

We can construct a Gaussian quadrature associated with the spherical harmonic basis as a direct product of a uniform quadrature in the φ direction:

$$w = \frac{\pi}{M_\varphi + 1/2}, \quad \varphi_m = \frac{\pi}{M_\varphi + 1} m, \quad -M_\varphi \leq m \leq M_\varphi.$$

which is accurate to order $4M_\varphi + 2$ [62] and a Gauss-Legendre grid in the θ direction [74]:

$$\mu = \cos \theta, \quad w_\ell = \frac{2}{(1 - \mu_\ell^2) [P'_{M_\theta+1}(\mu_\ell)]^2}, \quad \mu_\ell = \{\mathcal{Z}[P_{M_\theta+1}(\mu)]\}_\ell, \quad 0 \leq \ell \leq M_\theta,$$

where P_M is the Legendre polynomial of degree M , i.e., $P_M = P_M^0$ and $\mathcal{Z}[p]$ denotes the set of the roots of p . This quadrature formula is then accurate up to order $2M_\theta + 1$. This means that if we want to obtain an exact representation of the scalar product of spherical harmonics up to order N_θ we have to choose $M_\theta = N_\theta$.

7.3.3 Angular discretization

As orthonormal basis on the unit 2-sphere, S_1 , we consider the real spherical harmonics, $Y_{\ell m}$ (see 7.3.2), whose orthonormality conditions are

$$\int_{S_1} Y_{\ell m}(\varphi, \theta) Y_{\ell' m'}(\varphi, \theta) d\Omega = \delta_{mm'} \delta_{\ell\ell'}.$$

As a result, any function $f \in L^2(S_1)$ can be expanded in spherical harmonics as

$$f_N(\varphi, \theta) = \sum_{\ell=0}^N \sum_{m=-\ell}^{\ell} f^{\ell m} Y_{\ell m}(\varphi, \theta), \quad f^{\ell m} = \int_{S_1} f(\varphi, \theta) Y^{\ell m}(\varphi, \theta) d\Omega, \quad (7.16)$$

where we have used the notation $Y^{\ell m}$ to denote the complex conjugate of $Y_{\ell m}$ (which is equal to $Y_{\ell m}$ since we are working with real spherical harmonics). If f is a smooth function, f_N will converge to f with spectral accuracy in the L^2 -norm [62].

7.3.4 The multi-group P_N scheme

We consider the following ansatz for the expansion of the distribution function:

$$F(x^\alpha, v, \varphi, \theta) = \sum_{n=0}^{N_v} \sum_{\ell=0}^N \sum_{m=-\ell}^{\ell} F^{n\ell m}(x^\alpha) \chi_n(v) Y_{\ell m}(\varphi, \theta), \quad (7.17)$$

where

$$F_{n\ell m}(x^\alpha) = \int_0^\infty h^3 v^2 dv \int_{S_1} d\Omega F(x^\alpha, v, \varphi, \theta) Y_{\ell m}(\varphi, \theta) \chi_n(v).$$

To simplify the notation, we introduce the multi-index $A = \{n, \ell, m\}$, and the basis functions,

$$\Psi_A(v, \varphi, \theta) := \chi_n(v) Y_{\ell m}(\varphi, \theta),$$

so that Eq. (7.17) becomes

$$F(x^\alpha, \epsilon, \varphi, \theta) = \sum_A F^A(x^\alpha) \Psi_A(\epsilon, \varphi, \theta) = F^A \Psi_A. \quad (7.18)$$

Note that the space spanned by our basis $\{\Psi_A\}$ is a vector space. We adopt the usual convention of denoting vector components with an upper index and co-vector components with a lower index. Linear operators acting on vectors and co-vectors will have upper and lower indices associated with their decomposition on an appropriate tensor product combination of the canonical basis $\{\Psi_A\}$ and its dual $\{\Psi^A\}$, defined by the requirement that $\int \Psi^A \Psi_B p^0 d\Pi = \delta^A_B$. Note that, thanks to the orthonormality of real spherical harmonics $\Psi^A = \Psi_A$.

Inserting Eq. (7.18) into Eq. (7.6), we obtain:

$$p^0 \frac{\partial F^B}{\partial t} \Psi_B + p^k \frac{\partial F^B}{\partial x^k} \Psi_B = \mathbb{C}[F]. \quad (7.19)$$

Multiplying Eq. (7.19) by Ψ^A and integrating with respect to $d\Pi$, we then obtain

$$\frac{\partial F^A}{\partial t} + \mathcal{P}^{kA}_B \frac{\partial F^B}{\partial x^k} = \mathbb{S}^A[F], \quad (7.20)$$

where we have again used Eq. (7.2) and exploited the orthonormality of the basis. We have also defined

$$\mathcal{P}^{kA}_B := \int p^k \Psi^A \Psi_B d\Pi, \quad (7.21)$$

and

$$\mathbb{S}^A[F] := \int \mathbb{C}[F] \Psi^A d\Pi. \quad (7.22)$$

The coefficients (7.21) can be computed exactly using a quadrature formula of high-enough order (see 7.3.2). Since they are independent of position and time, they can be pre-computed and stored for later usage.

The spectral decomposition of $\mathcal{P}^{kA}_B$, which determines the behavior of Eq. (7.20), is well known, see *e.g.*, [70]. In particular, it has been shown that the eigenvalues are strictly bounded by the speed of light c . While this implies that there is no superluminal propagation of radiation, it leads to slower-than-light motion of radiation waves for finite N (the radiation velocity converges to the correct value with increasing N) [69]. This is particularly evident in low-order P_N free-streaming solutions. For instance, the maximum propagation speed for P_1 is $c/\sqrt{3}$. Filtering can also affect the propagation velocity of radiation [215]. In the multidimensional case, $\mathcal{P}^{kA}_B$ has also zero-speed modes that have to be treated carefully in Godunov-based schemes to avoid numerical instabilities [70].

In the general case, the source term (7.22) has to be computed at run-time, but for the particular case of a source term in the form (7.9), and assuming that the opacity coefficients are constant in each of the energy groups (as commonly done in multi-group schemes [255]), the source terms can be pre-computed up to constant factors. Under these assumptions, the source term becomes

$$\mathbb{S}^A[F] = e^A - \kappa_n F^A - \kappa_{s,n} [\perp^A_B - \sigma_{a,n} \Delta^A_B] F^B = e^A + S^A_B F^B, \quad (7.23)$$

where⁶

$$e^A := \int \left[\frac{c^2 \eta}{h^4 v^3} \right] \Psi^A(v, \varphi, \theta) h^3 v^2 dv d\Omega, \quad (7.24)$$

⁶Notice that the term in square brackets in (7.24) is the number of emitted particles.




Figure 7.1: The grid structure for the spatial linear discontinuous Galerkin discretization. The white squares represent the cell centers, while the black dots show the cell interfaces. The dotted lines show the values of the Lagrangian basis $l_{i-1/2}, l_{i+3/2}$. Finally the + and - show the interfaces where the inter-element fluxes, $\mathcal{F}^\pm$, see (7.28), are computed.

$\perp^A_B$ is the projector perpendicular to Y_{00} ,

$$\perp^A_B := \delta^A_B - \delta^A_{n00} \delta^{n00}_B,$$

and Δ^A_B is the anisotropy matrix:

$$\Delta^A_B := \frac{1}{4\pi} \int h^3 v^2 dv d\Omega \Psi^A(v, \varphi, \theta) \int \Psi_B(v, \varphi', \theta') \vec{n}' \cdot \vec{n} d\Omega',$$

which can also be pre-computed. Note that we have denoted the opacity coefficients in the energy group $[v_n, v_{n+1}]$ with the subscript n .

7.3.5 Spatial discretization

We discretize the system of Eqs. (7.20) in space using the asymptotic-preserving (AP) linear DG scheme, see *e.g.*, [204, 217]. The scheme we present here is restricted to the case of orthogonal grids, but we point out that discontinuous Galerkin methods can be (and have been) extended to general, unstructured grids [93]. We recall that a scheme is AP if it reproduces a discretization of the diffusion limit of the continuum transport equation in the limit of small mean-free-path. This is an important property since it guarantees that the diffusion of radiation has a correct rate even if the mean-free-path is small compared to the spatial grid size. If a scheme is not asymptotically preserving, then the diffusion rate becomes unphysical when the mean-free-path is unresolved.

For simplicity of notation, we consider a simplified 1D version of Eq. (7.20):

$$\frac{\partial F^A}{\partial t} + \mathcal{P}^{1A}_B \frac{\partial F^B}{\partial x} = 0,$$

where the multidimensional case will be discussed at the end of this Section. Furthermore, we employ a uniform numerical grid $x_i = i\delta x$, while extension to a non-uniform grid is conceptually trivial. We then choose the points of the grid that are used to construct elements of width $\Delta x = 2\delta x$, as shown in Fig. 7.1. This rather uncommon grid structure has been adopted to ease the integration of Charon with existing general relativistic hydrodynamic codes that use traditional finite-volume schemes.

In the classical linear DG scheme, the degrees of freedom are usually identified with the one-sided limits of the solution at the points $x_{i-1/2}$ and $x_{i+3/2}$ from the interior of the element, while in our case we evolve the cell-centered values at x_i and x_{i+1} . The semi-discrete equations for the evolution of these cell-centered values can be easily obtained, since any function u in the finite-element space can be written as

$$u(x) = u_{i-1/2}l_{i-1/2}(x) + u_{i+3/2}l_{i+3/2}(x),$$

where

$$l_{i-1/2}(x) = 1 - \frac{x - x_{i-1/2}}{x_{i+3/2} - x_{i-1/2}}, \quad l_{i+3/2}(x) = \frac{x - x_{i-1/2}}{x_{i+3/2} - x_{i-1/2}}, \quad (7.25)$$

so that

$$u_{i-1/2} = \frac{3}{2}u_i - \frac{1}{2}u_{i+1}, \quad u_{i+3/2} = -\frac{1}{2}u_i + \frac{3}{2}u_{i+1}, \quad (7.26)$$

and

$$u_i = \frac{3}{4}u_{i-1/2} + \frac{1}{4}u_{i+3/2}, \quad u_{i+1} = \frac{1}{4}u_{i-1/2} + \frac{3}{4}u_{i+3/2}. \quad (7.27)$$

The resulting numerical scheme is

$$\Delta x \frac{dF_i^A}{dt} = \mathbb{F}_i^A,$$

where the flux terms in the element $[x_{i-1/2}, x_{i+3/2}]$ are given by

$$\mathbb{F}_i^A := \frac{3}{2}\mathcal{F}^- - \overline{\mathcal{F}} - \frac{1}{2}\mathcal{F}^+, \quad \mathbb{F}_{i+1}^A := \frac{1}{2}\mathcal{F}^- + \overline{\mathcal{F}} - \frac{3}{2}\mathcal{F}^+. \quad (7.28)$$

Here,

$$\overline{\mathcal{F}} := \frac{1}{2}[(\mathcal{P}^{1A}_B)_i F_{i-1/2}^B + (\mathcal{P}^{1A}_B)_{i+1} F_{i+3/2}^B],$$

is the average flux, while

$$\mathcal{F}^- := \frac{1}{2}[\mathcal{P}^{1A}_B(F_L^B + F_R^B) - \mathcal{R}^{1A}_C \max(v, |\Lambda^{1C}_D|) \mathcal{L}^{1D}_B(F_R^B - F_L^B)],$$

is the flux computed from the exact solution of the Riemann problem at $x_{i-1/2}$ with the left “ L ” and right “ R ” states F_L^A and F_R^A . The term $\mathcal{F}^+$ is the flux across $x_{i+3/2}$ and is calculated analogously to $\mathcal{F}^-$. In the previous equations, we have decomposed $\mathcal{P}^{1A}_B$ as

$$\mathcal{P}^{1A}_B = \mathcal{R}^{1A}_C \Lambda^{1C}_D \mathcal{L}^{1D}_B,$$

where $\mathcal{R}^{1A}_C$ are the matrices of the right eigenvectors, while Λ^{1C}_D and $\mathcal{L}^{1D}_B$ are the eigenvalues and the left eigenvectors of $\mathcal{P}^{1A}_B$, respectively. We have also introduced $v > 0$ which is taken to be the first positive abscissa of the adopted Legendre quadrature and it is used to introduce extra numerical dissipation on the zero speed modes similarly to what is done in [70]. Notice again that the spectral decomposition of $\mathcal{P}^{1A}_B$, discussed in the previous Section, can also be pre-computed at the beginning of the calculations.

Having described the numerical scheme for the 1D problem, we can construct the multidimensional numerical scheme for Eq. (7.20) by repeating the same construction in every direction:

$$\frac{dF^A_{i,j,k}}{dt} = e^A + S^A_B F^B + \frac{1}{\Delta x} \mathbb{F}^A_{i,j,k} + \frac{1}{\Delta y} \mathbb{G}^A_{i,j,k} + \frac{1}{\Delta z} \mathbb{H}^A_{i,j,k}, \quad (7.29)$$

where the fluxes in the y and z direction, $\mathbb{G}$ and $\mathbb{H}$, are computed analogously to the ones in the x direction.

To avoid creation of false extrema in the numerical solution we use the slope limiting technique [93]. Among the different limiters that we have implemented are (1) the so-called “step limiter”, which simply reduces the scheme to a first order discontinuous Galerkin scheme, (2) the “minmod” and (3) the asymptotic-preserving “minmod2” limiters [216]. The reason for using these particular limiters is that they have been well tested in the context of the transport equation (e.g., [216]). Furthermore, the minmod2 limiter has been studied in detail in the context of linear DG methods, where it has been shown that it does not affect smooth solutions away from local extrema [92], thus yielding a scheme with a very small dissipation.

7.3.6 Time discretization

For the time integration, we use the predictor-corrector method proposed by McClarren et al. [217]. In this approach, the streaming terms that model the transport of radiation are treated explicitly, while the source terms responsible for interaction with matter are treated implicitly. The use of this particular time integrator is motivated by the fact that this yields a relatively inexpensive, stable and asymptotic-preserving scheme. As discussed in Section 7.1, the fact that the streaming terms are treated explicitly makes this scheme particularly easy to parallelize, while the associated CFL constraint is not particularly demanding for applications involving fluid moving at relativistic velocities and general relativistic gravity.

In order to simplify the notation, we rewrite Eq. (7.29) as

$$\frac{dF^A}{dt} = e^A + S^A_B F^B + \mathcal{A}^A[F], \quad (7.30)$$

where $\mathcal{A}^A[F]$ is a shorthand for the treatment of the spatial flux terms. For the time integration of Eq. (7.30), we use the following two-step semi-implicit asymptotic-preserving scheme. Given the solution F^A_k at time $k\Delta t$, we first perform a predictor step

$$\frac{F^A_{k+1/2} - F^A_k}{\Delta t/2} = \mathcal{A}^A[F_k] + e^A_k + S^A_B F^B_{k+1/2},$$

to obtain the solution at time $(k + 1/2)\Delta t$ and then a corrector step:

$$\frac{F_{k+1}^A - F_k^A}{\Delta t} = \mathcal{A}[F_{k+1/2}] + e_{k+1/2}^A + S_B^A F_{k+1}^B,$$

to obtain the solution at time $(k + 1)\Delta t$. At both stages, the absorption, emission and scattering terms are treated implicitly, while the streaming terms are treated explicitly. The explicit part of this scheme is second-order accurate in time, while its implicit part is first-order accurate [217].

7.3.7 Filtering

Filtering is a common procedure to reduce the effects of the Gibbs phenomenon in spectral methods for numerical solution of partial differential equations [76, 263]. Filtered spherical harmonics expansions have been successfully used in, *e.g.*, meteorology (see *e.g.*, [62] and references therein) and the effects of filtering on the truncation error of a spectral expansion are now reasonably well understood [327]. However, the use of filters to mitigate (and, in most situations, remove) the occurrence of negative solutions in P_N schemes has only been proposed recently by McClarren and Hauck [215].

In their work, the authors propose to filter the spherical harmonic expansion of the solution after each timestep using a spherical-spline filter. Applying this suggestion to the spherical harmonic expansion of F we obtain (for clarity we report here the summation symbols)

$$[\mathcal{F}(F)](\varphi, \theta) = \sum_{\ell=0}^N \sum_{m=-\ell}^{\ell} \left[\frac{1}{1 + \alpha \ell^2 (1 + \ell^2)} \right] F^{\ell m} Y_{\ell m}(\varphi, \theta), \quad (7.31)$$

where

$$\alpha := \frac{c\Delta t}{\Delta x} \frac{1}{N^2} \frac{1}{(\sigma_t L + N)^2}, \quad (7.32)$$

and L is a characteristic length scale used to make α dimensionless, while σ_t is chosen to be of the same order of magnitude as κ .

The filtered spherical harmonics, FP_N , scheme has several interesting properties. Filtering has been found to be very effective and robust in removing numerical oscillations in P_N solutions, while preserving the rotational invariance of the scheme. Furthermore, for this particular choice of α , filtering turns off automatically in the limit $N \rightarrow \infty$, thus it does not spoil the convergence of the spherical harmonics expansion.

However, one important drawback, also remarked by McClarren and Hauck [216], is that the filtered P_N scheme does not have a clear continuum limit as $\Delta x, \Delta t \rightarrow 0$. This is unfortunate because it implies that the filtered P_N scheme, FP_N , cannot be interpreted as a system of partial differential equations. This in turn means that the quality of the FP_N solution will depend on the spatial grid resolution in a way that is hard to predict. The ultimate and most important implication is that an FP_N solution cannot be studied for spatial convergence.

To solve this problem, we propose a modification/generalization of the FP_N scheme as follows. We introduce a strength parameter, $s \geq 0$, to be specified

later, and construct the filtered expansion as

$$\left[\mathcal{F}(F)\right](\varphi, \theta) = \sum_{\ell=0}^N \sum_{m=-\ell}^{\ell} \left[\sigma\left(\frac{\ell}{N+1}\right) \right]^s F^{\ell m} Y_{\ell m}(\varphi, \theta), \quad (7.33)$$

where $\sigma(\eta)$ is a *filter function* of order p , that is, a function $\sigma \in C^p(\mathbb{R}^+; [0, 1])$ such that⁷

$$\sigma(0) = 1, \quad \sigma^{(k)}(0) = 0, \text{ for } k = 1, 2, \dots, p-1. \quad (7.34)$$

Notice that, since the filter strength depends only on ℓ , this does not destruct the rotational invariance of the scheme.⁸ Furthermore, as the order of the spherical harmonics, N , increases, the effect of filtering automatically decreases, so that the convergence of the scheme for $N \rightarrow \infty$ is retained. More specifically, for a filter of order p , we expect a convergence order of $\sim p-1$, as suggested by Vandeven's theorem for Fourier expansion [327], see also [167].

In our analysis we have considered two second-order and two fourth-order filters. The first one is the classical second-order Lanczos filter⁹:

$$\sigma_{\text{Lanczos}}(\eta) = \frac{\sin \eta}{\eta}, \quad (7.35)$$

while the second and third choices are given by the ErfcLog filter [61],

$$\sigma_{\text{ErfcLog}}(\eta) = \frac{1}{2} \text{Erfc} \left\{ 2p^{1/2} \left(|\eta| - \frac{1}{2} \right) \sqrt{\frac{-\log[1 - 4(\eta - 1/2)^2]}{4(\eta - 1/2)^2}} \right\}, \quad (7.36)$$

of order $p = 2, 4$. Finally, the fourth filter is the fourth-order spherical-spline filter

$$\sigma_{\text{SSpline}}(\eta) = \frac{1}{1 + \eta^4}. \quad (7.37)$$

We point out that, with our definition, the spherical-spline filter (7.37) is very similar to the one used in [215], but is *not* exactly equivalent. The reason for using a slightly different filter form is that the filter of [215] is not compatible with the form of Eq. (7.33) as it cannot be written in terms of a function $\sigma(\cdot)$ of $\ell/(N+1)$.

In addition, we have considered only even-order filters since the truncation error associated with these filters can be interpreted as a numerical viscosity of order higher than two [221], while for odd-order filters the leading truncation error is of the dispersion type [221]. Moreover, we also do not consider filters of orders higher than 4. This is because, as we will see later, the fourth-order

⁷ Here we ignore the technical requirement for Vandeven's theorem that $\sigma^{(k)}(1) = 0$ for $k = 0, 1, 2, \dots, p-1$, which is not satisfied by our filters (nor by the one proposed by [215]). This is a condition that does not influence the formal accuracy of the filtered expansion with respect to the unfiltered truncated expansion, but it is mainly needed to prove the convergence of the filtered expansion in the case in which the unfiltered expansion is not converging point-wise (for instance due to the presence of discontinuities) [152, 314].

⁸ This is a consequence of the classical addition theorem for spherical harmonics (e.g., [61]).

⁹ Note that the Lanczos filter is usually defined as $\sigma(\eta) = \sin \pi \eta / \pi \eta$ to have a first-order zero at $\eta = 1$, as discussed in footnote 7. Our modified Lanczos filter yields a more uniform damping of high-order modes and works very well in our experiments.

filters are already too weak to completely remove oscillations, suggesting that even higher order will be even less efficient.

In our scheme, we filter the solution after each sub-step of the time integrator. This yields the following scheme

$$\frac{F^A_* - F^A_k}{\Delta t/2} = \mathcal{A}^A[F_k] + e^A_k + S^A_B F^B_{k+1/2}, \quad (7.38)$$

$$F^A_{k+1/2} = \mathcal{F}^A_B F^B_*, \quad (7.39)$$

$$\frac{F^A_{**} - F^A_k}{\Delta t} = \mathcal{A}^A[F_{k+1/2}] + e^A_{k+1/2} + S^A_B F^B_{k+1}, \quad (7.40)$$

$$F^A_{k+1} = \mathcal{F}^A_B F^B_{**}, \quad (7.41)$$

where $\mathcal{F}^A_B$ is a diagonal matrix representing the filtering operation.

We should remark that both our scheme and the one by [215], cannot be interpreted as a continuum problem, in the sense that the equations (7.38)–(7.41) do not, in general, represent a discretization of any system of partial differential equations. The main reason is that $\mathcal{F}^A_B$ is not idempotent, *i.e.*, $\mathcal{F}^A_C \mathcal{F}^C_B \neq \delta^A_B$, so that the filtering operations in Eqs. (7.39) and (7.41) do not have a well-defined behavior in the limit $\Delta t \rightarrow 0$. In the case in which $\mathcal{F}^A_B$ is idempotent, the scheme has indeed a continuum limit, but it can be easily demonstrated that the FP_N method is just the P_M method for some $M \leq N$.¹⁰

This problem can be solved by making the strength parameter s (and $\mathcal{F}^A_B$ with it) depend on the timestep. In order to see that, we consider the behavior of our scheme for a given mode $u = F^{\ell m}$ with $\ell \neq 0$. Let $q = \sigma(\ell/(N+1))$, where σ is a filter function. Then the effect of filtering on u in each of the two filtering steps in, *e.g.*, Eq. (7.41) is simply:

$$u_{k+1} = q^s u_{**}.$$

This can be rewritten as

$$\frac{u_{k+1} - u_{**}}{\Delta t/2} = \frac{1}{\Delta t/2} [q^s - 1] u_{**}.$$

If we let $s = \beta \Delta t$, then, in the limit of $\Delta t \rightarrow 0$, we obtain

$$\frac{du}{dt} = \beta \log q u. \quad (7.42)$$

In other words, we can interpret filtering as a first-order, operator split, discretization of the system of equations

$$\frac{\partial F^A}{\partial t} + \mathcal{P}^{kA}_B \frac{\partial F^B}{\partial x^k} = e^A + S^A_B F^B + \beta L^A_B F^B, \quad (7.43)$$

where L^A_B is a diagonal matrix with coefficients $\log \sigma(\ell/(N+1))$. This is the desired continuum limit. The physical interpretation is that filtering is equivalent to a forward-peaked scattering operator (notice that $\sigma(0) = 1$). Finally, we note that we can estimate the filter effective opacity by looking at the dissipation rate for the highest-order multipole of the expansion as

$$\sigma_{\text{eff}} = -\beta \log \sigma(N/(N+1)).$$

¹⁰The reason is that the only idempotent filter is the cut-off filter, that is, the filter that simply sets to zero all the modes with $\ell > M$ for some M , while leaving unaffected all the modes with $\ell \leq M$.

7.4 Tests

In this Section, we present some tests of the numerical schemes described above as implemented in our Charon code. Charon uses 3D Cartesian coordinates in space and is currently parallelized employing hybrid OpenMP/MPI parallelization using the domain decomposition method. It uses the open-source Cactus Computational Toolkit [150, 2], which provides MPI parallelization, input/output, and restart capabilities.

7.4.1 1D diffusion of a step function

In this first test, we primarily focus on verifying the ability of our scheme to handle diffusion of radiation when the opacity is high and the mean-free-path is small compared to the grid spacing. In this limit, the continuous hyperbolic transport equation displays parabolic character to leading order [216]. Despite this, there is no guarantee that a numerical scheme for solving the hyperbolic system will be AP, that is, will reproduce a valid discretization of the asymptotic limit of the continuous equations (cf. the discussion in Section 7.3.5).

Consider therefore the following initial conditions for the radiation energy density $E := \int I d\Omega dv$, in a non-moving infinite medium with a constant (isotropic and elastic) scattering opacity:

$$E(z, t = 0) = H(z + 1/2)H(1/2 - z), \quad (7.44)$$

where $H(\cdot)$ is the Heaviside step function. If the scattering opacity is high, the solution of the transport equation is well approximated by the solution of the diffusion equation. The corresponding diffusion equation has the following analytic solution to problem (7.44)

$$E(z, t) = \frac{1}{2} \left[\text{Erf} \left(\frac{z + 1/2}{2\sqrt{t/\tau}} \right) - \text{Erf} \left(\frac{z - 1/2}{2\sqrt{t/\tau}} \right) \right], \quad (7.45)$$

where $\text{Erf}(\cdot)$ is the error function and $\tau := 3\kappa_s/c$ is the diffusion timescale, where κ_s is the total scattering opacity, which we set to $\kappa_s = 10^5$ [e.g., 216].

We employ five different schemes for this test: the step scheme (*i.e.*, a DG scheme with step-limiter, or, equivalently a first-order FV scheme), two linear DG methods employing minmod and minmod2 limiters and two finite-volume (FV) methods also employing minmod and minmod2 limiters. In our implementation, the finite-volume scheme is obtained from the linear DG scheme by simply replacing the linear DG slope with the one obtained from the reconstruction procedure.

In all of our runs, we use the P_1 scheme because in 1D there are no negative solutions and thus filtering is not necessary, and because the radiation is nearly isotropic in such a diffusive regime so that the P_1 scheme should be sufficiently accurate. We perform calculations using three different resolutions $\Delta z = 0.16, 0.08$ and $\Delta z = 0.04$, with the grid ranging from -2 to 2 , and imposing periodic boundary conditions at the outer boundaries. We choose the CFL factor to be 0.25 and we recall that the maximum CFL factor that guarantees the L^2 -stability of our scheme is $1/3$ in 1D [93]. In all of our tests, the CFL factor is mainly chosen for convenience in order to have a sufficient number of




Figure 7.2: Radiation energy density as a function of the z coordinate at time $100/c$ for the 1D diffusion of the step function (7.44). The thick black line represents the analytical solution of the corresponding diffusion equation, the line with red circles corresponds to the linear DG with minmod2 limiter solution, the line with orange squares represents the solution from the finite-volume scheme with minmod2 reconstruction, while the lines with blue diamonds and with green triangles show the results obtained with linear DG with minmod and step limiters, respectively. The symbols also mark the values of the numerical solution at each grid point (*i.e.*, we show two points for each element).

timesteps within a given time interval. Moreover, in many radiation-transport calculations, in the absence of hydrodynamical equations, the truncation error due to the time discretization is expected to be small compared to other sources of error (*e.g.*, the angular and spatial discretization).

Figure 7.2 shows the radiation energy density as a function of z coordinate for the run with $\Delta z = 0.04$ at time $100/c$. The thick black line corresponds to the analytic solution, while the other lines show numerical results obtained with the above schemes. The line with red circles corresponds to the linear DG with minmod2 limiter solution. The line with orange squares represents the solution from the finite-volume scheme with minmod2 reconstruction. Finally, the lines with blue diamonds and green triangles show the results obtained with linear DG with minmod and step limiters, respectively. Note that the different symbols also mark the value of the numerical solution at each grid point (*i.e.*, we show two points for each element).

The linear DG method with minmod2 agrees well with the analytical result. This is expected since this scheme has the correct asymptotic limit [216]. In contrast, all other schemes overestimate the diffusion rate. In particular, the step scheme produces the worst results. It reaches stationarity already at time $t/\tau \sim 10^{-4}$, which is much smaller than the diffusion timescale for this problem. The linear DG and finite-volume schemes with minmod yield identical results and for this reason we show only the results from the linear DG scheme. Both are only marginally better than the step algorithm. These results are in overall agreement with the ones reported by [216] for a very similar test.

The FV scheme with the minmod2 reconstruction produces results that are relatively accurate compared to the linear DG and finite-volume schemes




Figure 7.3: L^∞ -norm of the deviation of the numerical result from the asymptotic solution as a function of numerical resolution for the diffusion problem of a sine function. The lines with red circles and blue squares correspond to the linear DG schemes with minmod2 and minmod limiters, respectively, while the line with orange diamonds represents the FV scheme with minmod2 reconstruction. Finally, the line with green triangles corresponds to the step scheme.

with the minmod limiter, even though the observed diffusion timescale (at the current resolution) is still orders of magnitude larger than the physical one. Finally, we point out that the results obtained from the other lower-resolution runs (not shown here) are in overall agreement with the ones presented here.

7.4.2 1D diffusion of a sine wave

Next, we explore convergence of the numerical solution to the asymptotic one. To this scope we consider the 1D diffusion of a sine wave and thus adopt as initial conditions the energy density given by

$$E(z, t = 0) = 3\sqrt{4\pi} \left[\sin\left(\frac{\pi z}{3}\right) + 1 \right], \quad (7.46)$$

which has the following analytic solution in the diffusion limit

$$E(z, t) = 3\sqrt{4\pi} \left[1 + \exp\left(-\frac{\pi^2}{9} \frac{t}{\tau}\right) \sin\left(\frac{\pi z}{3}\right) \right]. \quad (7.47)$$

For this test, our computational domain ranges from -3 to 3 and we use eight different resolutions ranging from $\Delta z = 0.3$ to $\Delta z = 0.009375$. The CFL factor is chosen to be 0.3 .

Figure 7.3 shows the L^∞ -norm of the deviation of the numerical result from the asymptotic solution as a function of numerical resolution at time $t = 1000/c$. As expected, the linear DG scheme with minmod2 (line with red circles) exhibits approximately second-order convergence for the entire range of resolutions shown in the plot, while the linear DG with minmod (line with blue squares) starts converging only when $\Delta z \sim 10^{-1}$, afterwards it converges with order ≈ 1.28 . The step scheme (line with green triangles) does not show

any sign of convergence¹¹. These results are again consistent with what was observed in [216]. Finally, the finite-volume with minmod2 reconstruction (line with orange diamonds) exhibits second-order convergence even though this scheme is not asymptotic preserving.

7.4.3 The 2D line-source problem

As a first multidimensional problem used to benchmark different implementations of the filtered spherical harmonics discretization schemes we consider the so-called “line-source” problem, where we have initial conditions given by¹²

$$I(x, y, z, \Omega, t = 0) = \frac{E^0}{4\pi} \delta(x) \delta(y), \quad (7.48)$$

which represents an isotropic pulse of radiation with a total energy E^0 concentrated along the z -axis in vacuum [68]. This radiation field propagates in vacuum as it evolves in time according to the following analytical solution

$$E(x, y, t) = \frac{E^0}{2\pi} \frac{H(ct - r)}{ct \sqrt{c^2 t^2 - r^2}}, \quad (7.49)$$

where $r := \sqrt{x^2 + y^2}$ is the distance from the z -axis and $H(\cdot)$ is the unit step function and $\delta(\cdot)$ is the Dirac delta function. According to this solution, the radiation field consists of a front that forms a cylindrical shell that travels outwards at the speed of light, while in the interior, the radiation energy density smoothly decreases along the radial direction.

We point out that, while this test is actually one-dimensional in cylindrical coordinates, it becomes particularly challenging for radiation-transport codes, except for Monte-Carlo codes, when solved on a two-dimensional Cartesian grid (as we do). In these coordinates, the radiation beam, which originates from a single spatial grid zone, has a very forward-peaked distribution in angle. This is a huge challenge for both spatial and angular discretization schemes. Moreover, such a configuration favors negative solutions in the P_N scheme. Indeed, the analytical P_N solution to this problem was shown to have regions with negative values of the energy [68, 218, 215], while the P_1 solution even exhibits a negative singularity [214]. For all the results presented here, we use a grid with resolution $\Delta x = \Delta y = 0.02$ and a CFL factor of 0.0625. Furthermore, we choose $E_0 = \sqrt{4\pi}$.

Figure 7.4 displays the colormap of the radiation energy density in the $x - y$ plane at $t = 1/c$. The upper left panel shows the analytic solution, while the upper right panel shows the pure P_7 solution (note the considerable difference in scale). As expected, the P_7 solution exhibits unphysical oscillations in the radial direction that are absent in the analytical solution to the full transport problem (7.49). The lower left panel of Fig. 7.4 shows instead the FP_7 solution

¹¹Notice that for the step scheme is expected to reach the convergence regime only when the numerical diffusion becomes smaller than the physical one. This means that the step scheme can be expected to be accurate only when the mean-free-path of radiation is resolved, *i.e.*, at a resolution which is unfeasible in many practical applications. For instance, for the problem we considered here, we expect the step scheme to reach the convergence regime for $\Delta z \sim 10^{-6}$.

¹²The initial conditions are 3D but we exploit the cylindrical symmetry to solve the problem on the (x, y) plane only.




Figure 7.4: Colormaps of the radiation energy density on the $x - y$ plane at $t = 1/c$ for the line problem with different methods. The upper left panel shows the analytic solution, while the upper right panel shows the pure P_7 solution (note the considerable difference in scale). The lower left panel shows the FP_7 solution with spherical-spline filter with effective opacity $\sigma_{\text{eff}} = 20$, while the lower right panel shows the FP_7 solution with the Lanczos filter with the same effective opacity.




Figure 7.5: The radiation energy density as a function of the x coordinate at $t = 1/c$ for the line problem test. The thick black line corresponds to the analytic solution, while the rest of the lines represent the results from P_7 calculations without filter (blue line), with fourth-order ErfcLog filter (green line), with the second-order ErfcLog filter (red line), with the Lanczos filter (cyan line), and with the fourth-order spherical-spline filter (magenta line).

with spherical-spline filter with effective opacity $\sigma_{\text{eff}} = 20$ (the dependence of the solution on the order N and on the value of the filter strength σ_{eff} will be discussed below). In this case, the radial oscillations are significantly reduced compared to the unfiltered P_7 solution, similar to what was found in [215]. Finally, the lower right panel shows the FP_7 solution with the Lanczos filter with the same effective opacity. In this case, the amount of oscillations is even smaller and we get a result that is closer to the analytical one. The reason seems to be that the Lanczos filter, being a second-order filter, is more effective in reducing the appearance of oscillations. The solution obtained with the spherical-spline filter is still characterized by the presence of a ring structure that resembles the more oscillatory (unfiltered) P_N solution. This structure does not disappear even for values of the filter strength as high as $\sigma_{\text{eff}} = 10^4$, suggesting that this is a result of a shortcoming of this particular filter. We point out that we have repeated these runs with the second-order (fourth-order) ErfcLog filter and we obtain a result similar to the one with the second-order Lanczos (fourth-order ErfcLog) filter, suggesting that the order of the filter plays the most important role, at least for this test.

A more quantitative measure of this test is shown in Fig. 7.5, where we plot a 1D cut of the radiation energy density as a function of the x -coordinate at $t = 1/c$. The thick black line again corresponds to the analytic solution, while the rest of the lines represent the results from P_7 calculations without filter (blue line), with the fourth-order ErfcLog filter (green line), with the second-order ErfcLog filter (red line), with the Lanczos filter (cyan line), and with the spherical-spline filter (violet line). All of these runs with filters are performed using a filter strength of $\sigma_{\text{eff}} = 20$. We can easily notice again the presence of large oscillations in the unfiltered P_7 solution (which are larger than the scale of the plot). The fourth-order spherical-spline and ErfcLog filters are able to suppress most of the oscillations and remove the negative values. Nevertheless,




Figure 7.6: Radiation energy density as a function of the x coordinate at $t = 1/c$ obtained from the FP_7 solution with the second-order Lanczos filter of varying strength σ_{eff} . The thick black line corresponds to the analytic solution, while the rest of the lines represent the numerical solutions corresponding to different σ_{eff} .

both solutions are still affected by the oscillations (although to a much smaller extent compared to the P_7 solution). The second-order Lanczos and ErfcLog filters, on the other hand, are able to remove most of the oscillations and give the best numerical solutions, overall.

Figure 7.6, which reports the radiation energy density as a function of the x coordinate at $t = 1/c$, highlights how the quality of the solution varies with the (Lanczos) filter effective opacity. The black line again corresponds to the analytic solution, while the rest of the lines represent the FP_7 solution with the Lanczos filter of varying strength σ_{eff} . In the case of weak filters (e.g., $\sigma_{\text{eff}} = 1$ or $\sigma_{\text{eff}} = 5$), there are significant oscillations, whose amplitude is significantly reduced as we increase the filter strength. For example, for $\sigma_{\text{eff}} = 20$, there are tiny oscillations, while for $\sigma_{\text{eff}} = 50$ there are no noticeable oscillations. However, the quality of the solution with $\sigma_{\text{eff}} = 50$ is actually worse than the one with $\sigma_{\text{eff}} = 20$ (e.g., the radiation wavefront lags significantly behind the real solution) as a result of the large smearing of the radiation beam produced by the excessive filtering (this is even more evident for $\sigma_{\text{eff}} = 100$ or $\sigma_{\text{eff}} = 1000$). Therefore, the filter strength needs to be chosen large enough to damp oscillations and small enough to avoid excessive smearing of the solution. We have repeated these runs with the second-order ErfcLog filter and we again get similar results.

Figure 7.6 should be contrasted with Fig. 7.7, where we show a study of the effect of different filter opacities for the same problem, but using the spherical-spline filter. The first thing to notice is that the spherical-spline filter is never able to completely remove the “ring” structure in the solution, even when the filter strength is so strong that the result resembles the solution of the diffusion equation for this problem. Secondly, the dependence of the filter behavior on the filter strength does not seem to be easily predictable: at first, as we increase filter strength, negative solutions disappear (for $\sigma_{\text{eff}} \lesssim 100$), then they reappear for higher values of σ_{eff} around 1000. We have repeated these runs with the




Figure 7.7: Same as in Figure 7.6 but for the FP_7 solution with the fourth-order spherical-spline filter of varying strength σ_{eff} .

fourth-order ErfcLog filter and we obtained similar results.

Finally, Fig. 7.8 shows the FP_N solutions for $N = 3, 5, 7, 9$ and 11 with the Lanczos filter with $\sigma_{\text{eff}} = 20$, and can be used to study the convergence of the FP_N approximation to the analytical solution. In this plot, we can distinguish three different types of errors: (1) An error in the position of the radiation front, which is particularly evident for small N , is mainly related to the fact that the propagation speed of radiation is smaller than c for low N ; (2) An error in the profile of the radiation energy density behind the front, which is again particularly pronounced for small N , and is due to the fact that high angular resolution is needed to properly describe the very forward-peaked angular distribution of radiation; (3) A relatively large spreading of the radiation beam in space compared to the $x^{-1/2}$, $x = ct - r$, singularity in the analytical solution (7.49). This is an artifact of the spatial discretization and mainly stems from the fact that the radiation beam originates from one spatial element and results in the presence of a “precursor” in the radiation front for high-order FP_N (e.g., FP_{11}) solutions, where the spatial discretization scheme propagates the radiation front superluminally, despite the fact that the characteristic speeds of the FP_N system are always smaller than c . Superluminal propagation of sharp features in numerical solutions of hyperbolic PDEs is a well known artifact of the spatial and temporal discretization of the equations. An in-depth explanation of this phenomenon for the case of the Maxwell equations can be found in [313]. As we can see in the figure, the FP_N solution nevertheless approaches the analytical one as we increase the order N and, in particular, the errors associated with the angular discretization decrease to the point that the FP_{11} solution yields only a small improvement with respect to the FP_9 one. In particular, a large contribution to the error in the FP_{11} solution comes from the presence of the superluminal precursor discussed above. Since this precursor can only be attributed to the spatial discretization error, we can conclude that, at this particular resolution and order, the spatial discretization error is already comparable with the angular discretization ones. We have repeated these runs with the second-order ErfcLog filter and at half of the resolution and we again get




Figure 7.8: Radiation energy density as a function of the x coordinate at $t = 1/c$ obtained from the FP_N solution with the second-order Lanczos filter with $\sigma_{\text{eff}} = 20$ for different values of order N . The thick black line represents that analytical solution, while the rest of the lines represent numerical solutions for different N . Clearly, the FP_N solution approaches the analytical one as we increase the order N .

similar results. These results show that our filtering strategy is able to recover the convergence of the P_N scheme for this particular case, and that the second-order filters, unlike the fourth-order ones, do not require a delicate fine-tuning of the effective scattering opacity. In particular, σ_{eff} can be chosen on the basis of the physics and geometry of the problem, in a way that is independent of the order of the employed P_N scheme.

Overall, the results of this test confirm that the filtering approach is effective and robust in suppressing unphysical oscillations in the P_N solution, even with moderately low-order N . Moreover, we find that the second-order Lanczos and ErfcLog filters produce significantly better numerical results compared to the fourth-order spherical-spline and ErfcLog filters.

7.4.4 A lattice problem

Next, we consider another 2D problem consisting of a chessboard of highly scattering and highly absorbing square regions located around a central emitting square region. Although this geometry is not expected to be present in the astrophysical scenarios we are most interested in, it nevertheless represents a very demanding test of the capabilities of the different numerical schemes in complicated geometries.

In our calculation we use a setup illustrated in the upper left panel of Fig. 7.9, which consists of a central emitting square (shown in white) and 11 absorbing squares (shown in black) with a constant absorption opacity $\kappa_a = 10$ surrounding the central emitting square. The space between the squares (shown in gray) and the central emitting region has a small uniform scattering opacity of $\kappa_s = 1$. Each square has a linear size of 1. The size of the computational domain is 7 along both axes. We choose a spatial resolution of $\Delta x = \Delta z = 0.035$ and the CFL factor was set to ≈ 0.14 , with outgoing boundary conditions imposed at




Figure 7.9: Upper-left panel: illustration of the setup of the 2D lattice problem. The rest of the panels: colormap of the $\log_{10}$ of the radiation energy density at time $t = 3.2/c$ as obtained with the P_7 scheme (upper-right panel) and FP_N schemes with the Lanczos filter with opacity 5 for different order N (middle and bottom panels). The time $t = 3.2/c$ corresponds to the moment when the radiation front first reaches the outer boundary of the computational domain.

the outer boundary.

The remaining panels in Fig. 7.9 show the $\log_{10}$ of the radiation energy density produced by different schemes at a time $t = 3.2/c$, which roughly corresponds to the moment when the radiation reaches the outer boundary of the computational domain. The upper right panel corresponds to the P_7 solution. Not surprisingly, this solution has regions of negative energy density (shown in white), although the negative values reach at most a relatively small magnitude of $\sim 10^{-5}$. All other solutions are computed with the Lanczos filter with effective opacity $\sigma_{\text{eff}} = 5$. We find that $\sigma_{\text{eff}} \gtrsim 5$ is necessary to avoid the appearance of negative solutions. The middle-left panel represents the FP_1 solution, which does not have negative regions, but where the radiation wavefront has reached only half of the computational domain. This is again due to the fact that the $N = 1$ wave travels at a slower velocity of $\simeq c/\sqrt{3}$ [244] (see the discussion in Section 7.3.3). The middle-right panel shows the FP_3 solution, which also does not have regions with negative energy density and where the front has travelled enough to cover $\sim 95\%$ of the computational domain, as a result of larger propagation speeds with higher N (of course, the velocity is always bounded by the speed of light). The FP_5 solution shown in the lower-left panel is very similar to the FP_3 case, with the only noticeable difference being the slightly faster propagation velocity in the FP_5 case. Finally, the lower-right panel shows the FP_7 solution, which looks almost indistinguishable from the FP_5 one.

We complete the analysis of this test by showing in Fig. 7.10 equivalent snapshots of the energy density at a later time of $t = 16/c$, when the radiation field has reached a stationary state. In these conditions, the P_N solution is not expected to have any negative values [68]. This is indeed confirmed by the upper left panel of Fig. 7.10, which shows the P_7 solution without negative regions. The remaining panels in Fig. 7.10 report the FP_1 , FP_3 , and FP_7 solutions, respectively. Note that the FP_3 , and FP_7 solutions are very similar to the P_7 one, underlining that the filter we use does not compromise the accuracy of the solution. However, the FP_1 solution appears to be significantly different from the FP_3 and FP_7 solutions, implying that $N = 1$ is not a sufficiently accurate approximation for this problem.

7.4.5 3D Homogeneous sphere

Finally, we consider the 3D homogeneous sphere problem, which is frequently employed to test radiation transport codes [303, 256, 266, 5]. This problem consists of a static homogeneous and isothermal sphere of radius R that radiates in vacuum. Inside the sphere, the radiation interacts with the background matter only via isotropic absorption and thermal emission. Despite the rather simple setup, the sharp discontinuity at the surface of the sphere is a model for astrophysical phenomena with rapidly varying opacity. This represents a major challenge for finite-difference methods (although, it is less challenging for Monte Carlo methods; see, e.g., [5]).

We assume that the sphere of radius R has a constant absorption opacity κ_a and emissivity B in the interior, while in the ambient vacuum at $r > R$, we have $\kappa_a = B = 0$. For this problem, the transport equation can be solved analytically




Figure 7.10: Colormaps of the $\log_{10}$ of the radiation energy density at time $t = 16/c$ obtained with the P_7 scheme (upper-left panel) and FP_1 (lower-left panel), FP_7 (upper-right panel), and FP_3 (lower-right panel) schemes with the Lanczos filter with opacity 5. The time $t = 16/c$ corresponds to the time by which the radiation field reaches the stationary state. Since the P_N scheme is less likely to exhibit negative solutions in the stationary state, we do not observe such solutions in our numerical result (upper-left panel), similarly to the filtered P_N solutions.




Figure 7.11: Radiation energy density as a function of the radial coordinate for the homogeneous sphere problem. The thick black line shows the analytical solution, while the rest of the lines show the FP_N solution with the Lanczos filter with $\sigma_{\text{eff}} = 1$ but with different values of order N . Clearly, the numerical result approaches the analytical solution as we increase N .

and has solution [303]

$$I(r, \mu) = \frac{B}{\kappa_a} \Big|_{r=0} [1 - \exp(-\kappa_a s(r, \mu))] , \quad (7.50)$$

where $r := \sqrt{x^2 + y^2 + z^2}$, $\mu := \cos \theta$ and

$$s(r, \mu) = \begin{cases} r\mu + Rg(r, \mu) & \text{if } r < R, \quad -1 \leq \mu \leq 1, \\ 2Rg(r, \mu) & \text{if } r \geq R, \quad \sqrt{1 - (R/r)^2} \leq \mu \leq 1, \\ 0 & \text{otherwise,} \end{cases} \quad (7.51)$$

and

$$g(r, \mu) = \sqrt{1 - \left(\frac{r}{R}\right)^2 (1 - \mu^2)}. \quad (7.52)$$

Note that this solution depends only on three parameters: κ_a , R , and B , where the latter acts as a scale factor for the solution.

We perform simulations in full 3D with Cartesian coordinates and use the following computational setup. We set $R = 1$ and cover the interior of the sphere with 40 elements in diameter along each coordinate direction, with the outer boundary being located at $5R$. The absorption opacity is chosen to be $\kappa_a = 10$ and the CFL factor is set to be ~ 0.12 .

It is useful to remark that, although the matter distribution is spherically symmetric, this is a genuinely 3D test due to the Cartesian geometry of our spatial grid. Indeed, it leads to the propagation of radiation from one spatial zone to another not only in the radial direction, but also in the angular directions, and the degree of sphericity of the numerical solution can be taken as a measure of the accuracy.




Figure 7.12: L^1 -norm of the deviation of the FP_N solution from the analytic result as a function of the order N for the homogeneous sphere problem. The black dots show the error as computed in a sphere of radius $R = 4.5$, while the black line shows the fitted convergence rate.




Figure 7.13: Radiation energy density as a function of the r coordinate for the homogeneous sphere problem. The thick black line shows the analytical solution, the blue line corresponds to the unfiltered P_7 solution, while the rest of the lines represent the FP_7 solutions obtained with different filters and different values of σ_{eff} .

Figure 7.11 shows the radiation energy density along the diagonal direction for the analytical solution and the FP_N solutions of different orders ranging from 1 to 11. The results are shown for the time when the radiation field has reached a stationary state.¹³ These runs are performed with the Lanczos filter of $\sigma_{\text{eff}} = 1$. Interestingly, all of the FP_N solutions produce the correct result in the interior of the sphere. This is not surprising since the radiation inside the sphere is nearly isotropic and the low-order FP_N solutions are already accurate enough. Outside the sphere, radiation streams freely outwards with a highly forward-peaked distribution in angle, which is a challenge for low-order FP_N schemes. Indeed, the FP_1 result deviates significantly from the analytical solution in that region. However, the solution clearly becomes more accurate everywhere in the computational domain as we increase the order of the scheme. Figure 7.12 shows the L^1 -norm of the deviation of the FP_N solution from the analytic result in a sphere of radius $R = 4.5$.¹⁴ As we can see from the plot, the FP_N scheme starts to converge already for $N \approx 3$, an order with only $4^2 = 16$ angular degrees of freedom. The convergence order is ≈ 1.16 , which is consistent with what expected from the theory of spectral filtering.

Figure 7.13 shows the radiation energy density along the diagonal direction for the analytical solution and different numerical solutions at the stationary state. These are the unfiltered P_7 solution, the three FP_7 solutions computed using respectively the second-order Lanczos, the fourth-order spherical-spline and ErfcLog filters with $\sigma_{\text{eff}} = 1$. Also shown are the two FP_7 solutions computed using the Lanczos filter with $\sigma_{\text{eff}} = 0.1$ and $\sigma_{\text{eff}} = 10$. Although the P_7 solution does not exhibit any negative solutions, it shows large oscillations in the free streaming region. The FP_7 solution with the Lanczos filter with $\sigma_{\text{eff}} = 0.1$ also yields a somewhat oscillatory solution, suggesting that the filter effective opacity is too low for this problem. As σ_{eff} is increased, the spurious oscillations disappear and all the filters that we have tried yield solutions of very similar quality for $\sigma_{\text{eff}} = 1$. Finally, the FP_7 solution with the Lanczos filter with $\sigma_{\text{eff}} = 10$ is similar in quality to the FP_3 solution with the same filter but with $\sigma_{\text{eff}} = 1$. This is due to the excessive damping of the high-order multipoles of the solution by the filtering procedure.

Finally, Fig. 7.14 shows the colormaps of the $\log_{10}$ of the radiation energy density from the FP_1 (left panel) and FP_{11} (right panel) solutions with the Lanczos filter with $\sigma_{\text{eff}} = 1$ at $t = 3.75/c$, which corresponds to the time when the radiation front almost reaches the outer boundary. Clearly, and as observed also in the previous tests, the FP_1 radiation front lags significantly behind the FP_{11} one because of its slower propagation speed (*cf.* the discussion in Section 7.3.7). Moreover, we can also see that both solutions maintain a high level of spherical symmetry despite the Cartesian geometry of the spatial grid. We obtain similar results in tests where the sphere was covered by 20, 10 and 5 elements.

¹³Note that stationarity is reached at different times depending on the different schemes used. For this reason and given the high computational costs, we did not evolve all the models up to the same time. Instead we report the solution as obtained as soon as stationarity is reached. In all cases the computations are performed up to at least $t = 20/c$.

¹⁴We compute the error inside $R = 4.5$ instead of $R = 5$ in order to exclude effects due to boundary conditions. We also normalize the L^1 -norm by dividing it by $\frac{4}{3}\pi R^3$.




Figure 7.14: Colormaps of the radiation energy density at time $t = 3.75/c$ obtained with the FP_1 and FP_7 schemes using the Lanczos filter with opacity $\sigma_{\text{eff}} = 1$. The time $t = 3.75/c$ corresponds to the moment when the radiation front almost reaches the outer boundary of the computational domain.

7.5 Conclusions

We have presented an extension of the filtered spherical harmonics method by McClarren and Hauck [215], the FP_N scheme, to three dimensions. We have developed the new 3D/multigroup radiation transport code Charon, built within the Cactus Computational Toolkit [150, 2]. Charon uses an asymptotic-preserving linear discontinuous Galerkin discretization scheme in space [216] and a semi-implicit time integration scheme [217] (*cf.* Section 7.3).

Our filtering scheme differs from the one presented by [215] in one important aspect: we reformulate the filtering procedure so that it acquires a well-defined continuum limit. In particular, we have shown that in the limit where the spatial and time steps are reduced to zero, our filtering scheme can be interpreted as the addition of a forward-peaked artificial scattering term to the P_N equations. The filtering procedure is also constructed in such a way as to retain the convergence of the FP_N solution to the solution of the transport equation as $N \rightarrow \infty$.

We have tested our scheme against a few challenging benchmark problems for radiation transport using four different filtering kernels: the fourth order spherical-spline filter, which is similar in spirit to the filter used by [215], the fourth-order and second-order ErfcLog filters [61], and the classical second-order Lanczos filter [61]. Our findings indicate that the FP_N scheme behaves well also in the three-dimensional case. In addition, we have shown that the second-order filters are more robust and accurate and require somewhat less tuning of the filter strength when compared to the fourth-order spherical-spline and the ErfcLog filters. Since the order of a filter is one of its most important properties, this result is likely to apply also to several other second- and fourth-order filters.

In future work, we plan to extend our numerical algorithms to include velocity dependence, the coupling to hydrodynamics, and, eventually, general relativity.

Part IV

Conclusions

Chapter 8

Conclusions

General-relativistic hydrodynamics and radiation-hydrodynamics are indispensable tools in the modeling of many high-energy astrophysical phenomena and sources of gravitational waves, such as binary neutron star mergers, active galactic nuclei or core-collapse supernovae, involving general-relativistic gravity and/or high-Lorentz factor flows. Given the highly non-linear nature of the equations involved, their analytic treatment is available only in special cases, or if certain simplifying assumptions are made, while a numerical treatment is needed to solve them in general. For this reason, while analytic and semi-analytical models have greatly advanced the comprehension of phenomena such as gamma-ray bursts or core-collapse supernovae, it is clear that the actual mechanism originating them can only be demonstrated by means of fully non-linear simulations.

Two key issues need to be addressed, however, in order for this goal to be achieved. First of all, numerical relativistic hydrodynamics needs to transition from the current state where it can provide qualitative or semi-quantitative results to a state where fully-quantitative results can be achieved reliably. Secondly, it is necessary to improve the quality of the models in use in current codes: for instance with the inclusion of realistic, nuclear theory based, equation of state and neutrino physics. The goal of this thesis is to contribute on both of these aspects.

Concerning the issue of the numerical accuracy of numerical relativity codes, we argue that switching away from the commonly adopted second-order schemes towards more advanced methods is necessary to make quantitative predictions, at least in scenarios such as the inspiral and merger of two neutron stars in quasi-circular orbits or turbulence, where it is necessary to accurately transport features of the solution over long distances/time, without having them overwhelmed by the numerical dissipation of the scheme. For this reason a large part of my thesis work has been devoted to the development of higher order numerical schemes.

The most direct and efficient way to extend the current generation of numerical-relativity codes is by adopting higher-order, high-resolution shock-capturing, finite-differencing schemes. For this reason we started developing a modular infrastructure for the solution of balance laws in multiple spatial dimensions inside the Cactus computational toolkit. The first code built with this infrastructure is the Templated Hydrodynamics Code, THC. This code solves

the equations of classical and special-relativistic hydrodynamics using higher-order methods. The high accuracy of this code has been demonstrated in a series of stringent tests involving the propagation of shock waves, as well as a study of the linear and non-linear development of the relativistic Kelvin-Helmholtz instability in two and three spatial dimensions. As a first application of THC we studied the statistical properties of turbulence in an ultra-relativistic gas and, thanks to the superior accuracy of the code, we were able to show, for the first time, similarities and differences between turbulence in Newtonian physics and turbulence in relativity.

The general-relativistic extension of THC, *WhiskyTHC*, also part of this thesis work, is the first genuinely higher than second order general-relativistic hydrodynamics code. We benchmarked this code using a series of classical tests involving the evolution of stable and unstable isolated neutron stars. We showed that *WhiskyTHC* can produce high order of convergence even in highly dynamical situations, such as the gravitational collapse of a non-rotating neutron star to black-hole. *WhiskyTHC* is also the first code able to obtain high order ($\geq$ third) in the phase of the gravitational wave generated in a binary neutron star merger. As a first application of *WhiskyTHC* we showed a comparison between the numerical-relativity waveform computed for the merger of binary neutron stars in quasi-circular orbit and the prediction of the post-Newtonian theory. In particular, we find that, in the case of binaries with high compactness stars, tidal contributions lead only to a modest de-phasing with respect to the point-particle post-Newtonian prediction. This de-phasing can be further reduced to the point of being below the systematic uncertainties of our evolution with the inclusion of tidal contributions up to the next-leading order terms.

In the quest for even more accurate numerical schemes we also studied the application of discontinuous Galerkin methods to general-relativistic hydrodynamics. Discontinuous Galerkin methods offer spectral-like accuracy for smooth flows and nearly-optimal scalability in parallel applications, due to their very compact stencils. For these reasons they are a natural candidate for the next generation of numerical relativity codes. In order to assess their viability we developed the first one-dimensional general-relativistic hydrodynamics code based on these methods. We showed that, indeed, they offer superior performance with respect to other methods, even though some aspects, in particular the implementation of robust flattening procedures and of local adaptivity, need to be addressed before their potential can be fully tapped for numerical relativity.

In addition we also worked on the inclusion of neutrino radiation transport in general-relativistic hydrodynamics codes. In particular we started the development of a new code, *Charon*, based on filtered spherical-harmonics and discontinuous Galerkin methods for radiation transport. *Charon* is a multi-angle, multi-energy, full-Boltzmann solver. As such, it is not based on any approximate model for the radiation transport, *e.g.*, the flux-limited diffusion or the Eddington approximations, but its solution can be made arbitrarily accurate by increasing the order of the spherical harmonics expansion used in the code. The work on *Charon* resulted in the development of a new, improved, version of the spherical harmonics scheme, which employs more general filters and has more clear convergence properties with respect to the original method. The *Charon* code has been benchmarked in a number of tests and it is, to the best of my knowledge, the only non Monte Carlo code for which convergence

to the analytic transport solution has been shown for some non-trivial problems of astrophysical interest, such as the homogeneous sphere test. In the future we plan to extend Charon to include velocity dependence and general-relativistic corrections so that it can be finally coupled with `WhiskyTHC` in order to construct a highly-accurate general-relativistic radiation-hydrodynamics code.

As a final remark we would like to point out that the work presented in this thesis constitutes the foundation of a new generation of accurate codes for computational relativistic astrophysics. These codes can be and will be applied to the simulation of compact binary mergers, to study core-collapse supernovae and to further explore the properties of turbulence in relativity.

Bibliography

- [1] <http://www.oonumerics.org/blitz/>.
- [2] <http://www.cactuscode.org>.
- [3] URL <http://www.lorene.obspm.fr>.
- [4] URL <http://www.cct.lsu.edu/~eschnett/McLachlan/index.html>.
McLachlan, a Public BSSN Code.
- [5] E. Abdikamalov, A. Burrows, C. D. Ott, F. Löffler, E. O'Connor, J. C. Dolence, and E. Schnetter. A New Monte Carlo Method for Time-Dependent Neutrino Radiation Transport. *Astrophys. J.*, 755:111, August 2012.
- [6] M. Abramowitz and I. A. Stegun. *Handbook of mathematical functions with formulas, graphs and mathematical tables*. 1968.
- [7] Marvin L. Adams and Edward W. Larsen. Fast iterative methods for discrete-ordinates particle transport calculations. *Progress in Nuclear Energy*, 40(1):3, 2002.
- [8] Oscar Agertz, Ben Moore, Joachim Stadel, Doug Potter, Francesco Miniati, Justin Read, Lucio Mayer, Artur Gawryszczak, Andrey Kravtsov, Åke Nordlund, Frazer Pearce, Vicent Quilis, Douglas Rudd, Volker Springel, James Stone, Elizabeth Tasker, Romain Teyssier, James Wadsley, and Rolf Walder. Fundamental differences between SPH and grid methods. *Mon. Not. R. Astron. Soc.*, 380(3):963–978, August 2007. ISSN 00358711. doi: 10.1111/j.1365-2966.2007.12183.x. URL <http://doi.wiley.com/10.1111/j.1365-2966.2007.12183.x>.
- [9] Dusan Agrez. Dynamics of frequency estimation in the frequency domain. *IEEE Transactions on Instrumentation and Measurement*, 56(6):2111, 2007. doi: 10.1109/TIM.2007.908240.
- [10] Miguel Alcubierre, Gabrielle Allen, Bernd Brügmann, Edward Seidel, and Wai-Mo Suen. Towards an understanding of the stability properties of the 3+1 evolution equations in general relativity. *Phys. Rev. D*, 62: 124011, 2000.
- [11] Miguel Alcubierre, Bernd Brügmann, Peter Diener, Michael Koppitz, Denis Pollney, Edward Seidel, and Ryoji Takahashi. Gauge conditions for long-term numerical black hole evolutions without excision. *Phys. Rev. D*, 67(8):084023, Apr 2003. doi: 10.1103/PhysRevD.67.084023.

- [12] D. Alic, C. Bona-Casas, C. Bona, L. Rezzolla, and C. Palenzuela. Conformal and covariant formulation of the Z4 system with constraint-violation damping. *Phys. Rev. D*, 85:064040, 2012.
- [13] D. Alic, W. Kastaun, and L. Rezzolla. Constraint-damping of the CCZ4 formulation in simulations of binary neutron stars. *ArXiv e-prints*, July 2013.
- [14] M. A. Aloy, J. A. Pons, and J. M. Ibáñez. *Comput. Phys. Commun.*, 120:115, 1999.
- [15] Miguel A. Aloy and Luciano Rezzolla. A powerful hydrodynamic booster for relativistic jets. *Astrophys. J.*, 640:L115–L118, 2006.
- [16] Luigi Ambrosio, Nicola Fusco, and Diego Pallara. *Functions of Bounded Variation and Free Discontinuity Problems*. Clarendon Press, 2000.
- [17] Nils Andersson and Gregory L. Comer. Relativistic Fluid Dynamics: Physics for Many Different Scales. *Living Rev. Relativ.*, 10(1):1–83, 2007. URL <http://www.livingreviews.org/lrr-2007-1>.
- [18] A. M. Anile. *Relativistic Fluids and Magneto-fluids*. Cambridge University Press, February 1990.
- [19] M. Ansorg, A. Kleinwächter, and R. Meinel. Highly accurate calculation of rotating neutron stars. *Astron. Astrophys.*, 381:L49, 2002. [astro-ph/0111080](https://arxiv.org/abs/astro-ph/0111080).
- [20] Luis Antón, Olindo Zanotti, Joan. A. Miralles, José M. Martí, José M. Ibáñez, José A. Font, and José A. Pons. Numerical 3+1 general relativistic magnetohydrodynamics: a local characteristic approach. *Astrophys. J.*, 637:296, 2006.
- [21] D. N. Arnold. An Interior Penalty Finite Element Method with Discontinuous Elements. *SIAM J. Numer. Anal.*, 19:742, 1982.
- [22] D. N. Arnold, F. Brezzi, B. Cockburn, and L. D. Marini. Unified analysis of discontinuous Galerkin methods for elliptic problems. *SIAM J. Numer. Anal.*, 39:1749, 2002.
- [23] Richard Arnowitt, Stanley Deser, and Charles W. Misner. The dynamics of general relativity. In L. Witten, editor, *Gravitation: An introduction to current research*, pages 227–265. John Wiley, New York, 1962.
- [24] Benjamin Aylott et al. Testing gravitational-wave searches with numerical relativity waveforms: Results from the first Numerical INjection Analysis (NINJA) project. *Class. Quantum Grav.*, 26:165008, 2009. doi: 10.1088/0264-9381/26/16/165008.
- [25] L. Baiotti, B. Giacomazzo, and L. Rezzolla. Accurate evolutions of inspiralling neutron-star binaries: Prompt and delayed collapse to a black hole. *Phys. Rev. D*, 78(8):084033, October 2008. doi: 10.1103/PhysRevD.78.084033.

-
- [26] L. Baiotti, T. Damour, B. Giacomazzo, A. Nagar, and L. Rezzolla. Accurate numerical simulations of inspiralling binary neutron stars and their comparison with effective-one-body analytical models. *Phys. Rev. D*, 84(2):024017, July 2011. doi: 10.1103/PhysRevD.84.024017.
 - [27] Luca Baiotti and Luciano Rezzolla. Challenging the paradigm of singularity excision in gravitational collapse. *Phys. Rev. Lett.*, 97:141101, 2006.
 - [28] Luca Baiotti, Ian Hawke, Pedro Montero, and Luciano Rezzolla. A new three-dimensional general-relativistic hydrodynamics code. In R. Capuzzo-Dolcetta, editor, *Computational Astrophysics in Italy: Methods and Tools*, volume 1, page 210, Trieste, 2003. MSAIt.
 - [29] Luca Baiotti, Ian Hawke, Pedro J. Montero, Frank Löffler, Luciano Rezzolla, Nikolaos Stergioulas, José A. Font, and Ed Seidel. Three-dimensional relativistic simulations of rotating neutron star collapse to a Kerr black hole. *Phys. Rev. D*, 71:024035, 2005.
 - [30] Luca Baiotti, Ian Hawke, and Luciano Rezzolla. On the gravitational radiation from the collapse of neutron stars to rotating black holes. *Class. Quantum Grav.*, 24:S187–S206, 2007.
 - [31] Luca Baiotti, Bruno Giacomazzo, and Luciano Rezzolla. Accurate evolutions of inspiralling neutron-star binaries: assessment of the truncation error. *Class. Quantum Grav.*, 26:114005, 2009.
 - [32] Luca Baiotti, Thibault Damour, Bruno Giacomazzo, Alessandro Nagar, and Luciano Rezzolla. Analytic modelling of tidal effects in the relativistic inspiral of binary neutron stars. *Phys. Rev. Lett.*, 105:261101, 2010. doi: 10.1103/PhysRevLett.105.261101.
 - [33] Luca Baiotti et al. Three-dimensional relativistic simulations of rotating neutron star collapse to a Kerr black hole. *Phys. Rev. D*, 71:024035, 2005.
 - [34] D. S. Balsara. Self-adjusting, positivity preserving high order schemes for hydrodynamics and magnetohydrodynamics. *J. of Comput. Phys.*, 231:7504–7517, September 2012. doi: 10.1016/j.jcp.2012.01.032.
 - [35] D. S. Balsara and C.-W. Shu. Monotonicity Preserving Weighted Essentially Non-oscillatory Schemes with Increasingly High Order of Accuracy. *J. of Comput. Phys.*, 160:405–452, May 2000. doi: 10.1006/jcph.2000.6443.
 - [36] F. Banyuls, J. A. Font, J. M. Ibáñez, J. M. Martí, and J. A. Miralles. Numerical 3+1 general-relativistic hydrodynamics: A local characteristic approach. *Astrophys. J.*, 476:221, 1997.
 - [37] Thomas W. Baumgarte and Stuart L. Shapiro. On the numerical integration of Einstein’s field equations. *Phys. Rev. D*, 59:024007, 1998.
 - [38] Thomas W. Baumgarte and Stuart L. Shapiro. Numerical relativity and compact binaries. *Physics Reports*, 376(2):41–131, March 2003.

- [39] Kris Beckwith and James M. Stone. A Second-Order Godunov Method for Multi-Dimensional Relativistic Magnetohydrodynamics. *Astrophys. J. Suppl.*, 193(1):6, March 2011. ISSN 0067-0049. doi: 10.1088/0067-0049/193/1/6. URL <http://stacks.iop.org/0067-0049/193/i=1/a=6?key=crossref.dbd6ffd5d95382531554173d56e257c8>.
- [40] Matania Ben-Artzi and Philippe G. LeFloch. Well-posedness theory for geometry-compatible hyperbolic conservation laws on manifolds. *Annales de l'Institut Henri Poincaré (C) Non Linear Analysis*, 24(6):989 – 1008, 2007. ISSN 0294-1449. doi: <http://dx.doi.org/10.1016/j.anihpc.2006.10.004>. URL <http://www.sciencedirect.com/science/article/pii/S0294144906001132>.
- [41] R. Benzi, S Ciliberto, R. Tripiccone, C. Baudet, F. Massaioli, and S. Succi. Extended self-similarity in turbulent flows. *Phys. Rev. E*, 48(1):29–32, 1993. URL <http://link.aps.org/doi/10.1103/PhysRevE.48.R29>.
- [42] Roberto Benzi, Luca Biferale, Robert Fisher, Leo Kadanoff, Donald Lamb, and Federico Toschi. Intermittency and Universality in Fully Developed Inviscid and Weakly Compressible Turbulent Flows. *Phys. Rev. Lett.*, 100(23):234503, June 2008. ISSN 0031-9007. doi: 10.1103/PhysRevLett.100.234503. URL <http://link.aps.org/doi/10.1103/PhysRevLett.100.234503>.
- [43] Marsha Berger, Micheal J. Aftosmis, and Scott M. Murman. Analysis of slope limiters on irregular grids. NAS Technical Report 2005-0490, American Institute of Aeronautics and Astronautics, May 2005.
- [44] Marsha J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.*, 82:64–84, 1989.
- [45] Marsha J. Berger and Joseph Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.*, 53:484–512, 1984.
- [46] S. Bernuzzi, A. Nagar, M. Thierfelder, and B. Brügmann. Tidal effects in binary neutron star coalescence. *Phys. Rev. D*, 86(4):044030, August 2012. doi: 10.1103/PhysRevD.86.044030.
- [47] Sebastiano Bernuzzi and David Hilditch. Constraint violation in free evolution schemes: comparing BSSNOK with a conformal decomposition of Z4. *Phys. Rev. D*, 81:084003, 2010. doi: 10.1103/PhysRevD.81.084003.
- [48] Sebastiano Bernuzzi, Marcus Thierfelder, and Bernd Brügmann. Accuracy of numerical relativity waveforms from binary neutron star mergers and their comparison with post-Newtonian waveforms. *Phys. Rev. D*, 85(10), May 2012. ISSN 1550-7998.
- [49] H. A. Bethe. Supernova mechanisms. *Reviews of Modern Physics*, 62: 801–866, October 1990. doi: 10.1103/RevModPhys.62.801.
- [50] Rupak Biswas, Karen D. Devine, and Joseph E. Flaherty. Parallel, adaptive finite element methods for conservation laws. *Applied Numerical Mathematics*, 14:255, 1994.

-
- [51] G. Bodo, A. Mignone, and R. Rosner. Kelvin-Helmholtz instability for relativistic fluids. *Phys. Rev. E*, 70(3):1–4, September 2004. ISSN 1539-3755. doi: 10.1103/PhysRevE.70.036304. URL <http://link.aps.org/doi/10.1103/PhysRevE.70.036304>.
 - [52] Gianluigi Bodo, Fausto Cattaneo, Attilio Ferrari, Andrea Mignone, and Paola Rossi. Symmetries, scaling laws, and convergence in shearing-box simulations of magneto-rotational instability driven turbulence. *Astrophys. J.*, 739(2):82, October 2011. ISSN 0004-637X. doi: 10.1088/0004-637X/739/2/82. URL <http://arxiv.org/abs/1106.5727><http://stacks.iop.org/0004-637X/739/i=2/a=82?key=crossref.6bfb57f977d332ca1edb1b03cc2549cc>.
 - [53] S. Boldyrev. Kolmogorov-Burgers model for star-forming turbulence. *Astrophys. J.*, 569:841, 2002. URL <http://iopscience.iop.org/0004-637X/569/2/841>.
 - [54] C. Bona, J. Massó, E. Seidel, and J. Stela. New Formalism for Numerical Relativity. *Phys. Rev. Lett.*, 75:600–603, July 1995.
 - [55] C. Bona, T. Ledvinka, C. Palenzuela, and M. Zacek. General-covariant evolution formalism for numerical relativity. *Phys. Rev. D*, 67:104005, 2003.
 - [56] C. Bona, C. Bona-Casas, and C. Palenzuela. Action principle for Numerical Relativity evolution systems. *Phys. Rev. D*, 82:124010, 2010. doi: 10.1103/PhysRevD.82.124010.
 - [57] Carles Bona and Carlos Palenzuela-Luque. *Elements of Numerical Relativity*. Springer-Verlag, Berlin, 2005. ISBN 3-540-25779-9. doi: 10.1007/b135928.
 - [58] S. Bonazzola, E. Gourgoulhon, and J. A. Marck. Relativistic formalism to compute quasidequilibrium configurations of nonsynchronized neutron star binaries. *Phys. Rev. D*, 56:7740–7749, 1997.
 - [59] R Borges, M Carmona, B Costa, and W Don. An improved weighted essentially non-oscillatory scheme for hyperbolic conservation laws. *J. Comp. Phys.*, 227(6):3191–3211, March 2008. ISSN 00219991. doi: 10.1016/j.jcp.2007.11.038. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999107005232>.
 - [60] J.P. Boris and Naval Research Lab Washington D C. *A Fluid Transport Algorithm that Works*. Defense Technical Information Center, 1971. URL <http://books.google.de/books?id=UQ2JNwAACAAJ>.
 - [61] J. P. Boyd. The Erfc-Log filter and the asymptotics of the Euler and Vandeven sequence accelerations. In L. Ridgway Scott A. V. Illin, editor, *Proceedings of the Third International Conference on Spectral and High Order Methods*, pages 267–276. Houston Journal of Mathematics, 1996.
 - [62] John P. Boyd. *Chebyshev and Fourier Spectral Methods (Second Edition, Revised)*. Dover Publications, New York, 2001. ISBN 0-486-41183-4.

- [63] Michael Boyle and Abdul H. Mroue. Extrapolating gravitational-wave data from numerical simulations. 2009.
- [64] Michael Boyle, Duncan A. Barrow, Lawrence E. Kidder, Abdul H. Mroué, Harald P. Pfeiffer, Mark A. Scheel, Gregory B. Cook, and Saul A. Teukolsky. High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions. *Phys. Rev. D*, 76:124038, 2007. doi: 10.1103/PhysRevD.76.124038.
- [65] Axel Brandenburg and Ake Nordlund. Astrophysical turbulence modeling. *Reports on Progress in Physics*, 74(4):046901, April 2011. ISSN 0034-4885. doi: 10.1088/0034-4885/74/4/046901. URL <http://stacks.iop.org/0034-4885/74/i=4/a=046901?key=crossref.00acfdb5bc1021fe1d27a3c6476a8179>.
- [66] David Brown, Peter Diener, Olivier Sarbach, Erik Schnetter, and Manuel Tiglio. Turduckening black holes: An analytical and computational study. *Phys. Rev. D*, 79(4), Feb 2009. URL <http://dx.doi.org/10.1103/PhysRevD.79.044023>.
- [67] S. W. Bruenn. Stellar core collapse: Numerical model and infall epoch. *Astrophys. J. Suppl.*, 58:771–841, 1985.
- [68] T. A. Brunner. Forms of approximate radiation transport. *Sandia Report*, 2002.
- [69] Thomas A. Brunner and James Paul Holloway. One-dimensional Riemann solvers and the maximum entropy closure. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 69(5):543 – 566, 2001. ISSN 0022-4073. doi: 10.1016/S0022-4073(00)00099-6. URL <http://www.sciencedirect.com/science/article/pii/S0022407300000996>.
- [70] Thomas a. Brunner and James Paul Holloway. Two-dimensional time dependent Riemann solvers for neutron transport. *J. Comp. Phys.*, 210(1):386–399, November 2005. ISSN 00219991. doi: 10.1016/j.jcp.2005.04.011. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999105002275>.
- [71] N. Bucciantini and L. Del Zanna. General relativistic magnetohydrodynamics in axisymmetric dynamical spacetimes: the X-ECHO code. *Astron. Astrophys.*, 528:A101, April 2011. doi: 10.1051/0004-6361/201015945.
- [72] Adam Burrows, Timothy Young, Philip Pinto, Ron Eastman, and Todd A Thompson. A New Algorithm for Supernova Neutrino Transport and Some Applications. *Astrophys. J.*, 539(2):865–887, August 2000. ISSN 0004-637X. doi: 10.1086/309244. URL <http://stacks.iop.org/0004-637X/539/i=2/a=865>.
- [73] Manuela Campanelli, Carlos O. Lousto, Pedro Marronetti, and Yosef Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Phys. Rev. Lett.*, 96:111101, 2006.
- [74] C. Canuto, M.Y. Hussaini, A. Quarteroni, and T.A. Zang. *Spectral Methods: Fundamentals in Single Domains*. Springer, Berlin, 2006.

-
- [75] C. Canuto, M.Y. Hussaini, A. Quarteroni, and T.A. Zang. *Spectral Methods: Evolution to Complex Geometries and Applications to Fluid Dynamics*. Springer, 2008.
 - [76] Claudio Canuto, M. Yousuff Hussani, Alfio Quarteroni, and Thomas A. Zang. *Spectral Methods in Fluid Dynamics*. Springer-Verlag, New York and Berlin, 2nd printing edition, 1988. ISBN 3-540-52205-0 (Berlin), 0-387-52205-0 (New York).
 - [77] Christian Cardall and Anthony Mezzacappa. Conservative formulations of general relativistic kinetic theory. *Phys. Rev. D*, 68(2), 2003. doi: 10.1103/PhysRevD.68.023006.
 - [78] J. I. Castor. *Radiation Hydrodynamics*. Cambridge University Press, Cambridge, UK, November 2004.
 - [79] C. Cercignani and G. M. Kremer. *The relativistic Boltzmann equation: theory and applications*. Boston: Birkhäuser, 2002. Progress in mathematical physics; v. 22, 2002.
 - [80] P Cerdá-Durán. Numerical viscosity in hydrodynamics simulations in general relativity, 2010. URL <http://stacks.iop.org/0264-9381/27/i=20/a=205012>.
 - [81] GQ Chen. Euler equations and related hyperbolic conservation laws. *Handbook of Differential Equations: Evolutionary Equations*, 2, 2006. doi: 10.1016/S1874-5717(06)80004-6. URL <http://www.sciencedirect.com/science/article/pii/S1874571706800046>.
 - [82] Gui-Qiang Chen and Hermano Frid. Extended Divergence-Measure Fields and the Euler Equations for Gas Dynamics. *Comm. Math. Phys.*, 236(2):251, 2003. doi: 10.1007/s00220-003-0823-7.
 - [83] Gui-Qiang Chen, William P. Ziemer, and Monica Torres. Gauss-green theorem for weakly differentiable vector fields, sets of finite perimeter, and balance laws. *Comm. Pure Appl. Math.*, 62(2):242, 2009. doi: 10.1002/cpa.20262.
 - [84] Chernikov. *Dokl. Akad. Nauk SSSR*, (144), 1962.
 - [85] H. H. Chiu. Relativistic gasdynamics. *Physics of Fluids*, 16(6):825, 1973. ISSN 00319171. doi: 10.1063/1.1694434. URL <http://link.aip.org/link/PFLDAS/v16/i6/p825/s1&Agg=doi>.
 - [86] B. Cockburn. The Runge Kutta Discontinuous Galerkin Method for Conservation Laws V Multidimensional Systems. *J. Comp. Phys.*, 141:199–224, April 1998. doi: 10.1006/jcph.1998.5892.
 - [87] B. Cockburn. Discontinuous Galerkin methods. *Zamm*, 83(11):731–754, November 2003. ISSN 0044-2267. doi: 10.1002/zamm.200310088. URL <http://doi.wiley.com/10.1002/zamm.200310088>.
 - [88] B. Cockburn and C.-W. Shu. TVB Runge Kutta Local Projection Discontinuous Galerkin Finite Element Method for Scalar Conservation Laws II: General Framework. *Math. Comp.*, 52:411, 1989.

- [89] B. Cockburn, S.-Y. Lin, and C.-W. Shu. TVB Runge Kutta Local Projection Discontinuous Galerkin Finite Element Method for Conservation Laws III: One-Dimensional Systems. *J. Comp. Phys.*, 84:90, September 1989. doi: 10.1016/0021-9991(89)90183-6.
- [90] B. Cockburn, S. How, and C.-W. Shu. TVB Runge Kutta Local Projection Discontinuous Galerkin Finite Element Method for Conservation Laws IV: The Multidimensional Case. *Math. Comp.*, 54:545, 1990.
- [91] B. Cockburn, G. E. Karniadakis, and C.-W. Shu. *Discontinuous Galerkin Methods: Theory, Computation and Applications*. Lecture Notes on Computational Science and Engineering. Springer, 2000.
- [92] Bernardo Cockburn. The Runge-Kutta local projection P1-discontinuous Galerkin finite element method for scalar conservation laws. *Math. Model. Numer. Anal. (MMAN)*, 25:337–361, 1991. URL <http://128.101.10.22/preprints/Jan87Dec87/388.pdf>.
- [93] Bernardo Cockburn and Chi-Wang Shu. Runge-Kutta discontinuous Galerkin methods for convection-dominated problems. *J. Scientific Comp.*, 16(3):173, 2001.
- [94] P. Colella and M. D. Sekora. A limiter for PPM that preserves accuracy at smooth extrema. *J. Comp. Phys.*, 227:7069–7076, July 2008. doi: 10.1016/j.jcp.2008.03.034.
- [95] P. Colella and P. R. Woodward. *J. Comput. Phys.*, 54:174, 1984.
- [96] Phillip Colella and Paul R. Woodward. The Piecewise Parabolic Method (PPM) for gas-dynamical simulations. *J. Comp. Phys.*, 54(1): 174–201, 1984. ISSN 0021-9991. doi: DOI:10.1016/0021-9991(84)90143-8. URL <http://www.sciencedirect.com/science/article/B6WHY-4DD1PHM-SJ/2/13d69a59afba3d6a5d6bbf1144d860aa>.
- [97] Edward Conway and Joel Smoller. Global solutions of the cauchy problem for quasi-linear first-order equations in several space variables. *Comm. Pure Appl. Math.*, 19(1):95–105, September 1966. ISSN 00103640. doi: 10.1002/cpa.3160190107. URL <http://doi.wiley.com/10.1002/cpa.3160190107>.
- [98] I. Cordero-Carrión, P. Cerdá-Durán, H. Dimmelmeier, J. L. Jaramillo, J. Novak, and E. Gourgoulhon. Improved constrained scheme for the Einstein equations: An approach to the uniqueness issue. *Phys. Rev. D*, 79(2):024017, January 2009. doi: 10.1103/PhysRevD.79.024017.
- [99] I. Cordero-Carrión, P. Cerdá-Durán, and J. María Ibáñez. Dynamical spacetimes and gravitational radiation in a Fully Constrained Formulation. *Journal of Physics Conference Series*, 228(1):012055, May 2010. doi: 10.1088/1742-6596/228/1/012055.
- [100] Michael G. Crandall and Andrew Majda. Monotone Difference Approximations for Scalar Conservation Laws. *Math. Comp.*, 34(149):1, January 1980. ISSN 00255718. doi: 10.2307/2006218. URL <http://www.jstor.org/stable/2006218?origin=crossref>.

-
- [101] Michael G. Crandall and Luc Tartar. Some relations between non-expansive and order preserving mappings. *Proceedings of the American Mathematical Society*, 78(3):385–385, March 1980. ISSN 0002-9939. doi: 10.1090/S0002-9939-1980-0553381-X. URL <http://www.ams.org/jourcgi/jour-getitem?pii=S0002-9939-1980-0553381-X>.
 - [102] Thibault Damour, Alessandro Nagar, and Loïc Villain. Measurability of the tidal polarizability of neutron stars in late-inspiral gravitational-wave signals. *Phys. Rev. D*, 85:123007, Jun 2012. doi: 10.1103/PhysRevD.85.123007. URL <http://link.aps.org/doi/10.1103/PhysRevD.85.123007>.
 - [103] Timothy A. Davis. UMFPACK: A set of routines for solving sparse linear systems via LU factorization. URL <http://www.cise.ufl.edu/research/sparse/umfpack/>. UMFPACK.
 - [104] Timothy A. Davis. A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. Technical Report TR-02-001, Univ. of Florida, CISE Dept., Gainesville, FL, January 2002.
 - [105] Timothy A. Davis. Algorithm 8xx: UMFPACK V3.2, an unsymmetric-pattern multifrontal method with a column pre-ordering strategy. Technical Report TR-02-002, Univ. of Florida, CISE Dept., Gainesville, FL, January 2002. (<http://www.cise.ufl.edu/tech-reports>. Submitted to *ACM Trans. Math. Softw.*).
 - [106] Timothy A. Davis and Iain S. Duff. An unsymmetric-pattern multifrontal method for sparse LU factorization. *SIAM J. Matrix Anal. Applic.*, 18(1): 140–158, 1997.
 - [107] Timothy A. Davis and Iain S. Duff. A combined unifrontal/multifrontal method for unsymmetric sparse matrices. *ACM Trans. Math. Softw.*, 25(1): 1–19, 1999.
 - [108] F Debbasch and W van Leeuwen. General relativistic Boltzmann equation, I: Covariant treatment. *Physica A: Statistical Mechanics and its Applications*, 388(7):1079, 2009. doi: 10.1016/j.physa.2008.12.023.
 - [109] F. Debbasch and W.A. van Leeuwen. General relativistic Boltzmann equation, II: Manifestly covariant treatment. *Physica A: Statistical Mechanics and its Applications*, 388(9):1818, 2009. doi: 10.1016/j.physa.2009.01.009.
 - [110] L. Del Zanna, N. Bucciantini, and P. Londrillo. An efficient shock-capturing central-type scheme for multidimensional relativistic flows. II. Magnetohydrodynamics. *Astron. Astrophys.*, 400:397–413, March 2003. doi: 10.1051/0004-6361:20021641.
 - [111] L. Del Zanna, O. Zanotti, N. Bucciantini, and P. Londrillo. ECHO: a Eulerian conservative high-order scheme for general relativistic magnetohydrodynamics and magnetodynamics. *Astron. Astrophys.*, 473:11–30, October 2007. doi: 10.1051/0004-6361:20077093.

- [112] H. Dimmelman, J. A. Font, and E. Müller. Relativistic simulations of rotational core collapse. I. Methods, initial models, and code tests. *Astron. Astrophys.*, 388:917–935, 2002.
- [113] H. Dimmelman, J. Novak, J. A. Font, J. M. Ibáñez, and E. Müller. Combining spectral and shock-capturing methods: A new numerical approach for 3D relativistic core collapse simulations. *Phys. Rev. D*, 71(6): 064023, March 2005. doi: 10.1103/PhysRevD.71.064023.
- [114] Ronald J. DiPerna. Measure-valued solutions to conservation laws. *Archive for Rational Mechanics and Analysis*, 88(3):223–270, 1985. ISSN 0003-9527. doi: 10.1007/BF00752112. URL <http://www.springerlink.com/index/10.1007/BF00752112>.
- [115] R. Donat and A. Marquina. Capturing shock reflections: an improved flux formula. *J. Comput. Phys.*, 125:42, 1996.
- [116] 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.
- [117] Matthew D. Duez, Yuk Tung Liu, Stuart L. Shapiro, and Branson C. Stephens. Relativistic magnetohydrodynamics in dynamical spacetimes: Numerical methods and tests. *Phys. Rev. D*, 72:024028, 2005. astro-ph/0503420.
- [118] Matthew D. Duez, Yuk Tung Liu, Stuart L. Shapiro, Masaru Shibata, and Branson C. Stephens. Collapse of magnetized hypermassive neutron stars in general relativity. *Phys. Rev. Lett.*, 96:031101, 2006. doi: 10.1103/PhysRevLett.96.031101.
- [119] Matthew D. Duez et al. Evolving black hole-neutron star binaries in general relativity using pseudospectral and finite difference methods. *Phys. Rev. D*, 78:104015, 2008. doi: 10.1103/PhysRevD.78.104015.
- [120] Paul C Duffell and Andrew I. MacFadyen. TESS: A RELATIVISTIC HYDRODYNAMICS CODE ON A MOVING VORONOI MESH. *Astrophys. J. Suppl.*, 197(2):15, December 2011. ISSN 0067-0049. doi: 10.1088/0067-0049/197/2/15. URL <http://arxiv.org/abs/1104.3562><http://stacks.iop.org/0067-0049/197/i=2/a=15?key=crossref.db5541359416a159060286f1329e82a6>.
- [121] M. Dumbser and O. Zanotti. Very high order PNPM schemes on unstructured meshes for the resistive relativistic MHD equations. *J. Comp. Phys.*, 228:6991–7006, October 2009. doi: 10.1016/j.jcp.2009.06.009.
- [122] M. Dumbser, M. Kaeser, V. A. Titarev, and E. F. Toro. Quadrature-free non-oscillatory finite volume schemes on unstructured meshes for non-linear hyperbolic systems. *J. Comp. Phys.*, 226:204–243, September 2007. doi: 10.1016/j.jcp.2007.04.004.
- [123] M. Dumbser, C. Enaux, and E. F. Toro. Finite volume schemes of very high order of accuracy for stiff hyperbolic balance laws. *J. Comp. Phys.*, 227:3971–4001, April 2008. doi: 10.1016/j.jcp.2007.12.005.

-
- [124] J. Ehlers. General relativity and kinetic theory. In *General Relativity and Cosmology*, pages 1–70. Proceedinss Varenna Summer School on Relativistic Astrophysics 1969 (Academic Press), 1971.
 - [125] B. Einfeldt. On Godunov-type methods for gas dynamics. *SIAM J. Numer. Anal.*, 25:294–318, 1988.
 - [126] V. Eswaran and SB Pope. An examination of forcing in direct numerical simulations of turbulence. *Computers & Fluids*, 16(3):257–278, 1988. URL <http://www.sciencedirect.com/science/article/pii/0045793088900138>.
 - [127] B. D. Farris, T. K. Li, Y. T. Liu, and S. L. Shapiro. Relativistic radiation magnetohydrodynamics in dynamical spacetimes: Numerical methods and tests. *Phys. Rev. D*, 78(2):024023, July 2008. doi: 10.1103/PhysRevD.78.024023.
 - [128] S. E. Field, J. S. Hesthaven, S. R. Lau, and A. H. Mroue. Discontinuous Galerkin method for the spherically reduced Baumgarte-Shapiro-Shibata-Nakamura system with second-order operators. *Phys. Rev. D*, 82(10):104051, November 2010. doi: 10.1103/PhysRevD.82.104051.
 - [129] J. A. Fleck, Jr. and J. D. Cummings, Jr. An implicit monte carlo scheme for calculating time and frequency dependent nonlinear radiation transport. *J. Comput. Phys.*, 8:313–342, December 1971.
 - [130] J. A. Font. Numerical hydrodynamics and magnetohydrodynamics in general relativity. *Living Rev. Relativ.*, 6:4, 2008. URL <http://www.livingreviews.org/lrr-2008-7>.
 - [131] J. A. Font, T. Goodale, S. Iyer, M. Miller, L. Rezzolla, E. Seidel, N. Stergioulas, W. M. Suen, and M. Tobias. Three-dimensional general relativistic hydrodynamics. II. Long-term dynamics of single relativistic stars. *Phys. Rev. D*, 65:084024, 2002.
 - [132] T. Frankel. *The geometry of physics : an introduction*. Cambridge University Press, 2004.
 - [133] H. Friedrich. On the hyperbolicity of Einstein’s and other gauge field equations. *Comm. Math. Phys.*, 100:525–543, 1985.
 - [134] K O Friedrichs. Symmetric hyperbolic linear differential equations. *Comm. Pure Appl. Math.*, 7(2):345–392, 1954. ISSN 1097-0312. doi: 10.1002/cpa.3160070206. URL <http://dx.doi.org/10.1002/cpa.3160070206>.
 - [135] Michael Gabler, Ulrich Sperhake, and Nils Andersson. Non-linear radial oscillations of neutron stars. *Phys. Rev. D*, 80:064012, 2009. doi: 10.1103/PhysRevD.80.064012.
 - [136] F. Galeazzi, W. Kastaun, L. Rezzolla, and J. A. Font. Implementation of a simplified approach to radiative transfer in general relativity. *arXiv:1306.4953*, January 2013.
 - [137] Filippo Galeazzi. Modelling fluid interfaces in numerical relativistic hydrodynamics. Master’s thesis, Università degli studi di Padova, 2008.

- [138] Charles F. Gammie, Jonathan C. McKinney, and G. Tóth. Harm: A numerical scheme for general relativistic magnetohydrodynamics. *Astrophys. J.*, 589:458, 2003.
- [139] Gregor Gassner, Michael Dumbser, Florian Hindenlang, and Claus-Dieter Munz. Explicit one-step time discretizations for discontinuous Galerkin and finite volume schemes based on local predictors. *J. Comp. Phys.*, 230(11):4232–4247, May 2011. ISSN 00219991. doi: 10.1016/j.jcp.2010.10.024. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999110005802>.
- [140] N. A. Gentile. Implicit Monte Carlo Radiation Transport in Multi-Physics Simulations. In *In Proc. International Conference on Mathematics, Computational Methods & Reactor Physics (M & C 2009)*, volume May 3-7, 2009, page 1170. American Nuclear Society, LaGrange Park, IL, 2009.
- [141] G.a. Gerolymos, D. Sénéchal, and I. Vallet. Very-high-order weno schemes. *J. Comp. Phys.*, 228(23):8481–8524, December 2009. ISSN 00219991. doi: 10.1016/j.jcp.2009.07.039. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999109003908>.
- [142] Axel Geyer and Heinz Herold. Slicing the Schwarzschild spacetime: Harmonic versus maximal slicing. *Phys. Rev. D*, 52(10):6182–6185, Nov 1995. doi: 10.1103/PhysRevD.52.6182.
- [143] B. Giacomazzo, L. Rezzolla, and L. Baiotti. Can magnetic fields be detected during the inspiral of binary neutron stars? *Mon. Not. R. Astron. Soc.*, 399:L164–L168, October 2009. doi: 10.1111/j.1745-3933.2009.00745.x.
- [144] Bruno Giacomazzo and Luciano Rezzolla. The Exact Solution of the Riemann Problem in Relativistic MHD. *Journal of Fluid Mechanics*, 562: 223–259, 2006.
- [145] Bruno Giacomazzo and Luciano Rezzolla. WhiskyMHD: a new numerical code for general relativistic magnetohydrodynamics. *Class. Quantum Grav.*, 24:S235, 2007.
- [146] Bruno Giacomazzo, Luciano Rezzolla, and Luca Baiotti. Accurate evolutions of inspiralling and magnetized neutron stars: Equal-mass binaries. *Phys. Rev. D*, 83(4):044014, Feb 2011. doi: 10.1103/PhysRevD.83.044014.
- [147] James Glimm. Solutions in the large for nonlinear hyperbolic systems of equations. *Comm. Pure Appl. Math.*, 18(4):697–715, November 1965. ISSN 00103640. doi: 10.1002/cpa.3160180408. URL <http://doi.wiley.com/10.1002/cpa.3160180408>.
- [148] William F. Godoy and Xu Liu. Parallel Jacobian-free Newton Krylov solution of the discrete ordinates method with flux limiters for 3d radiative transfer. *J. Comp. Phys.*, 231(11):4257, 2012.
- [149] S. K. Godunov. A difference method for numerical calculations of discontinuous solutions of the equations of hydrodynamics. *Mat. Sb.*, 47:271, 1959. in Russian.

-
- [150] T. Goodale, G. Allen, G. Lanfermann, J. Massó, T. Radke, E. Seidel, and J. Shalf. The Cactus framework and toolkit: Design and applications. In *Vector and Parallel Processing – VECPAR'2002, 5th International Conference, Lecture Notes in Computer Science*, Berlin, 2003. Springer.
 - [151] J.B. Goodman and R.J. LeVeque. On the accuracy of stable schemes for 2D scalar conservation laws. *Math. Comp.*, 45(171):15–21, 1985. URL <http://www.ams.org/mcom/1985-45-171/S0025-5718-1985-0790641-4/S0025-5718-1985-0790641-4.pdf>.
 - [152] D. Gottlieb. On the Gibbs phenomenon and its resolution. *SIAM review*, 39(4):644–668, 1997. URL <http://www.jstor.org/stable/10.2307/2132695>.
 - [153] David Gottlieb and Eitan Tadmor. The CFL condition for spectral approximations to hyperbolic initial-boundary value problems. *Math. Comp.*, 56(194):565–588, 1991. ISSN 00255718. URL <http://www.jstor.org/stable/2008395>.
 - [154] Sigal Gottlieb, David Ketcheson, and Chi-Wang Shu. High order strong stability preserving time discretizations. *Journal of Scientific Computing*, 38:251–289, 2009. ISSN 0885-7474. URL <http://dx.doi.org/10.1007/s10915-008-9239-z>.
 - [155] E. Gourgoulhon. Simple equations for general relativistic hydrodynamics in spherical symmetry applied to neutron star collapse. *Astron. Astrophys.*, 252:651–663, dec 1991.
 - [156] E Gourgoulhon. 1d numerical relativity applied to neutron star collapse. *Class. Quantum Grav.*, 9(S):S117, 1992. doi: 10.1088/0264-9381/9/S/005.
 - [157] E. Gourgoulhon, P. Grandclément, K. Taniguchi, J. A. Marck, and S. Bonazzola. Quasiequilibrium sequences of synchronized and irrotational binary neutron stars in general relativity: Method and tests. *Phys. Rev. D*, 63:064029, 2001.
 - [158] Eric Gourgoulhon. An introduction to relativistic hydrodynamics. *EAS Publications Series*, 21:43–79, 2006. doi: 10.1051/eas:2006106. URL <http://dx.doi.org/10.1051/eas:2006106>.
 - [159] P. Grandclément and J. Novak. Spectral methods for numerical relativity. *Living Rev. Relativ.*, 12(1), 2009. URL <http://www.livingreviews.org/lrr-2009-1>.
 - [160] Carsten Gundlach, Jose M. Martin-Garcia, G. Calabrese, and I. Hinder. Constraint damping in the Z4 formulation and harmonic gauge. *Class. Quantum Grav.*, 22:3767–3774, 2005.
 - [161] A. Harten, J. M. Hyman, P. D. Lax, and B. Keyfitz. On finite-difference approximations and entropy conditions for shocks. *Comm. Pure Appl. Math.*, 29(3):297–322, May 1976. ISSN 00103640. doi: 10.1002/cpa.3160290305. URL <http://doi.wiley.com/10.1002/cpa.3160290305>.

- [162] A. Harten, P. D. Lax, and B. van Leer. On upstream differencing and godunov-type schemes for hyperbolic conservation laws. *SIAM Rev.*, 25: 35, 1983.
- [163] A. Harten, B. Engquist, S. Osher, and S. R. Chakravarthy. Uniformly High Order Accurate Essentially Non-oscillatory Schemes III. *J. Comp. Phys.*, 71:231–303, August 1987. doi: 10.1016/0021-9991(87)90031-3.
- [164] Cory Hauck and Ryan McClarren. Positive P_N Closures. *SIAM Journal on Scientific Computing*, 32(5):2603–2626, January 2010. ISSN 1064-8275. doi: 10.1137/090764918. URL <http://www.cscamm.umd.edu/frgUploads/FRG-2009-Hauck-Cory.ppn2.pdf><http://epubs.siam.org/doi/abs/10.1137/090764918>.
- [165] Ian Hawke. *Computational Ultrarelativistic Hydrodynamics*. PhD thesis, University of Cambridge, 2001.
- [166] Andrew K. Henrick, Tariq D. Aslam, and Joseph M. Powers. Mapped weighted essentially non-oscillatory schemes: Achieving optimal order near critical points. *J. Comp. Phys.*, 207(2):542–567, August 2005. ISSN 00219991. doi: 10.1016/j.jcp.2005.01.023. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999105000409>.
- [167] Jan S. Hesthaven and Robert M. Kirby. Filtering in Legendre spectral methods. *Math. Comput.*, 77:1425, 2008. doi: 10.1090/S0025-5718-08-02110-8.
- [168] J.S. Hesthaven and T. Warburton. *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications*. Texts in Applied Mathematics. Springer, 2007. ISBN 9780387720654. URL <http://books.google.it/books?id=APQkD0mwykC>.
- [169] K. Hotokezaka, K. Kyutoku, and M. Shibata. Exploring tidal effects of coalescing binary neutron stars in numerical relativity. *Phys. Rev. D*, 87(4):044001, February 2013. doi: 10.1103/PhysRevD.87.044001.
- [170] Xiangyu Y. Hu, Nikolaus a. Adams, and Chi-Wang Shu. Positivity-preserving method for high-order conservative schemes solving compressible Euler equations. *J. Comp. Phys.*, 242:169–180, June 2013. ISSN 00219991. doi: 10.1016/j.jcp.2013.01.024. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999113000557>.
- [171] I. Hubeny and A. Burrows. A New Algorithm for Two-Dimensional Transport for Astrophysical Simulations. I. General Formulation and Tests for the One-Dimensional Spherical Case. *Astrophys. J.*, 659:1458–1487, April 2007. doi: 10.1086/512179.
- [172] Tsuyoshi Inoue, Katsuaki Asano, and Kunihiro Ioka. Three-dimensional simulations of magnetohydrodynamic turbulence behind relativistic shock waves and their implications for gamma-ray bursts. *Astrophys. J.*, 734(2):77, 2011. URL <http://stacks.iop.org/0004-637X/734/i=2/a=77>.

-
- [173] W. Israel and J. M. Stewart. Transient relativistic thermodynamics and kinetic theory. *Ann. Phys.*, 118:341, 1979.
 - [174] Werner Israel. Relativistic kinetic theory of a simple gas. *Journal of Mathematical Physics*, 4(9):1163, 1963. doi: 10.1063/1.1704047.
 - [175] H.-T. Janka, K. Langanke, A. Marek, G. Martínez-Pinedo, and B. Müller. Theory of core-collapse supernovae. *Physics Reports*, 442:38–74, April 2007. doi: 10.1016/j.physrep.2007.02.002.
 - [176] Guang-Shan Jiang and Chi-Wang Shu. Efficient Implementation of Weighted ENO Schemes. *J. Comput. Phys.*, 126:202–228, 1996.
 - [177] Eric Johnsen, Johan Larsson, Ankit V. Bhagatwala, William H. Cabot, Parviz Moin, Britton J. Olson, Pradeep S. Rawat, Santhosh K. Shankar, Björn Sjögreen, and H.C. Yee. Assessment of high-resolution methods for numerical simulations of compressible turbulence with shock waves. *J. Comp. Phys.*, 229(4):1213, 2010. doi: 10.1016/j.jcp.2009.10.028.
 - [178] W. Kastaun, F. Galeazzi, D. Alic, L. Rezzolla, and J. A. Font. On the black hole from merging binary neutron stars: how fast can it spin? *arXiv:1301.7348*, January 2013.
 - [179] Wolfgang Kastaun. High-resolution shock capturing scheme for ideal hydrodynamics in general relativity optimized for quasistationary solutions. *Phys. Rev. D*, 74(12):124024, 2006. doi: 10.1103/PhysRevD.74.124024.
 - [180] Wolfgang Kastaun. *Developing a code for general relativistic hydrodynamics with application to neutron star oscillations*. PhD thesis, University of Tübingen, 2007.
 - [181] S. Kitsionas, C. Federrath, R. S. Klessen, W. Schmidt, D. J. Price, L. J. Dursi, M. Gritschneider, S. Walch, R. Piontek, J. Kim, a. K. Jappsen, P. Ciecielag, and M.-M. Mac Low. Algorithmic comparisons of decaying, isothermal, supersonic turbulence. *Astron. Astrophys.*, 508(1):541–560, December 2009. ISSN 0004-6361. doi: 10.1051/0004-6361/200811170. URL <http://www.aanda.org/10.1051/0004-6361/200811170>.
 - [182] S. Koide, K. Shibata, and T. Kudoh. Relativistic Jet Formation from Black Hole Magnetized Accretion Disks: Method, Tests, and Applications of a General Relativistic Magnetohydrodynamic Numerical Code. *Astrophys. J.*, 522:727–752, September 1999. doi: 10.1086/307667.
 - [183] A. N. Kolmogorov. The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds Numbers. *Proc. R. Soc. A: Math. Phys. Eng. Sciences*, 434(1890):9–13, July 1991. ISSN 1364-5021. doi: 10.1098/rspa.1991.0075. URL <http://rspa.royalsocietypublishing.org/cgi/doi/10.1098/rspa.1991.0075>.
 - [184] S. S. Komissarov. General relativistic MHD simulations of monopole magnetospheres of black holes. *Mon. Not. R. Astron. Soc.*, 350:1431, 2004.
 - [185] A. Konigl. Relativistic gasdynamics in two dimensions. *Physics of Fluids*, 23:1083–1090, June 1980. doi: 10.1063/1.863110.

- [186] Heinz Otto Kreiss and Joseph Oliger. *Methods for the approximate solution of time dependent problems*. GARP publication series No. 10, Geneva, 1973.
- [187] Lilia Krivodonova. Limiters for high-order discontinuous Galerkin methods. *J. Comp. Phys.*, 226(1):879, 2007. doi: 10.1016/j.jcp.2007.05.011.
- [188] S N Kružkov. FIRST ORDER QUASILINEAR EQUATIONS IN SEVERAL INDEPENDENT VARIABLES. *Mathematics of the USSR-Sbornik*, 10(2):217–243, February 1970. ISSN 0025-5734. doi: 10.1070/SM1970v010n02ABEH002156. URL <http://stacks.iop.org/0025-5734/10/i=2/a=A06?key=crossref.4eecd9c5a212905d35fc0614fd12cc90>.
- [189] A.G. Kulikovskii, N.V. Pogorelov, and A. Yu. Semenov. *Mathematical aspects of numerical solution of hyperbolic systems*. Chapman & Hall/CRC, 2001.
- [190] A. Kurganov and E. Tadmor. New high-resolution central schemes for nonlinear conservation laws and convection-diffusion equations. *J. Comput. Phys.*, 160:241, 2000.
- [191] L. D. Landau and E. M. Lifshitz. *The Classical Theory of Fields, Course of Theoretical Physics, Volume 2*. Elsevier Butterworth-Heinemann, Oxford, 2004.
- [192] P. D. Lax. Hyperbolic systems of conservation laws II. *Comm. Pure Appl. Math.*, 10(4):537–566, 1957. ISSN 00103640. doi: 10.1002/cpa.3160100406. URL <http://doi.wiley.com/10.1002/cpa.3160100406>.
- [193] P. D. Lax and R. D. Richtmyer. Survey of the stability of linear finite difference equations. *Comm. Pure Appl. Math.*, 9(2):267–293, May 1956. ISSN 00103640. doi: 10.1002/cpa.3160090206. URL <http://doi.wiley.com/10.1002/cpa.3160090206>.
- [194] P. D. Lax and B. Wendroff. Systems of conservation laws. *Commun. Pure Appl. Math.*, 13:217–237, 1960.
- [195] Peter D Lax. Weak solutions of nonlinear hyperbolic equations and their numerical computation. *Comm. Pure Appl. Math.*, 7(1):159–193, 1954. ISSN 1097-0312. doi: 10.1002/cpa.3160070112. URL <http://dx.doi.org/10.1002/cpa.3160070112>.
- [196] R. J. Leveque. *Numerical Methods for Conservation Laws*. Birkhauser Verlag, Basel, 1992.
- [197] R. J. Leveque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, New York, 2002.
- [198] S. L. Liebling, L. Lehner, D. Neilsen, and C. Palenzuela. Evolutions magnetized and rotating neutron stars. *Phys. Rev. D*, 81(12):124023, June 2010. doi: 10.1103/PhysRevD.81.124023.
- [199] L. Lindblom and B. Szilágyi. Improved gauge driver for the generalized harmonic Einstein system. *Phys. Rev. D*, 80(8):084019, October 2009. doi: 10.1103/PhysRevD.80.084019.

-
- [200] Lee Lindblom, Mark A. Scheel, Lawrence E. Kidder, Robert Owen, and Oliver Rinne. A new generalized harmonic evolution system. *Class. Quantum Grav.*, 23:S447–S462, 2006.
 - [201] Richard W. Lindquist. Relativistic transport theory. *Annals of Physics*, 37(3):487–518, 1966. ISSN 0003-4916. doi: DOI:10.1016/0003-4916(66)90207-7. URL <http://www.sciencedirect.com/science/article/B6WB1-4DD36ST-TB/2/38359f63281be542b707fb56ec03a469>.
 - [202] X.-D. Liu, S. Osher, and T. Chan. Weighted Essentially Non-oscillatory Schemes. *J. Comp. Phys.*, 115:200–212, November 1994. doi: 10.1006/jcph.1994.1187.
 - [203] Frank Löffler, Joshua Faber, Eloisa Bentivegna, Tanja Bode, Peter Diener, Roland Haas, Ian Hinder, Bruno C. Mundim, Christian D. Ott, Erik Schnetter, Gabrielle Allen, Manuela Campanelli, and Pablo Laguna. The Einstein Toolkit: A Community Computational Infrastructure for Relativistic Astrophysics. *Class. Quantum Grav.*, 29(11):115001, 2012. doi: doi:10.1088/0264-9381/29/11/115001.
 - [204] R B Lowrie and J E Morel. Methods for hyperbolic systems with stiff relaxation. *International Journal for Numerical Methods in Fluids*, 40(3-4): 413–423, September 2002. ISSN 0271-2091. doi: 10.1002/fld.321. URL <http://doi.wiley.com/10.1002/fld.321>.
 - [205] Y. Maday, S. M. O. Kaber, and E. Tadmor. Legendre pseudospectral viscosity method for nonlinear conservation laws. *SIAM J. Numer. Anal.*, 30(2):321–342, 1993.
 - [206] P.J. Mann. Some mixed finite element-finite difference methods for spherically symmetric relativistic collapse. *J. Comp. Phys.*, 58(3):377–394, 1985. URL <http://www.sciencedirect.com/science/article/pii/002199918590169X>.
 - [207] Antonio Marquina. Local piecewise hyperbolic reconstruction of numerical fluxes for nonlinear scalar conservation laws. *SIAM Journal on Scientific Computing*, 15(4):892–915, Jul 1994. doi: 10.1137/0915054. URL <http://dx.doi.org/10.1137/0915054>.
 - [208] J. M. Martí and E. Müller. The Analytical Solution of the Riemann Problem in Relativistic Hydrodynamics. *J. Fluid Mech.*, 258:317–333, 1994.
 - [209] J. M. Martí and E. Müller. Numerical Hydrodynamics in Special Relativity. *Living Rev. Relativ.*, 6:7, 2003.
 - [210] J. M. Martí, J. M. Ibáñez, and J. A. Miralles. Numerical relativistic hydrodynamics: Local characteristic approach. *Phys. Rev. D*, 43:3794, 1991.
 - [211] M. P. Martín, E. M. Taylor, M. Wu, and V. G. Weirs. A bandwidth-optimized WENO scheme for the effective direct numerical simulation of compressible turbulence. *J. Comp. Phys.*, 220(1):270–289, 2006. ISSN 0021-9991. doi: DOI:10.1016/j.jcp.2006.05.009. URL <http://www.sciencedirect.com/science/article/B6WHY-4KCXJRV-1/2/320ffa1e875f5c80c5934e442b412abd>.

- [212] Andrea Maselli, Leonardo Gualtieri, Francesco Pannarale, and Valeria Ferrari. On the validity of the adiabatic approximation in compact binary inspirals. *Phys. Rev. D*, 86:044032, Aug 2012. doi: 10.1103/PhysRevD.86.044032. URL <http://link.aps.org/doi/10.1103/PhysRevD.86.044032>.
- [213] M. M. May and R. H. White. Hydrodynamic calculations of general relativistic collapse. *Phys. Rev.*, 141:1232, 1966.
- [214] R. G. McClarren, J. P. Holloway, and T. A. Brunner. Analytic P_1 solutions for time-dependent, thermal radiative transfer in several geometries. *J. Quantitative Spectroscopy and Radiative Transfer*, 109:389–403, February 2008. doi: 10.1016/j.jqsrt.2007.08.006.
- [215] Ryan G. McClarren and Cory D. Hauck. Robust and accurate filtered spherical harmonics expansions for radiative transfer. *J. Comput. Phys.*, 229(16):5597–5614, August 2010. ISSN 0021-9991. doi: 10.1016/j.jcp.2010.03.043. URL <http://dx.doi.org/10.1016/j.jcp.2010.03.043>.
- [216] Ryan G. McClarren and Robert B. Lowrie. The effects of slope limiting on asymptotic-preserving numerical methods for hyperbolic conservation laws. *J. Comp. Phys.*, 227(23):9711–9726, December 2008. ISSN 00219991. doi: 10.1016/j.jcp.2008.07.012. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999108004002>.
- [217] Ryan G. McClarren, Thomas M. Evans, Robert B. Lowrie, and Jeffery D. Densmore. Semi-implicit time integration for PNP thermal radiative transfer. *J. Comp. Phys.*, 227(16):7561–7586, August 2008. ISSN 00219991. doi: 10.1016/j.jcp.2008.04.029. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999108002489>.
- [218] Ryan G. McClarren, James Paul Holloway, and Thomas A. Brunner. On solutions to the P_n equations for thermal radiative transfer. *J. Comp. Phys.*, 227(5):2864–2885, February 2008. ISSN 00219991. doi: 10.1016/j.jcp.2007.11.027. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999107005153>.
- [219] Colin P. McNally, Wladimir Lyra, and Jean-claude Passy. A well-posed kelvin-helmholtz instability test and comparison. *The Astrophysical Journal Supplement Series*, 201(2):18, August 2012. ISSN 0067-0049. doi: 10.1088/0067-0049/201/2/18.
- [220] D. L. Meier. Multidimensional astrophysical structural and dynamical analysis. i. development of a nonlinear finite element approach. *Astrophys. J.*, 518(2):788, 1999. doi: 10.1086/307292.
- [221] A. Meister, S. Ortleb, and Th. Sonar. On spectral filtering for discontinuous Galerkin methods on unstructured triangular grids. http://cms.uni-kassel.de/unicms/fileadmin/groups/w_180000/prep/prep0904.pdf, 2009.
- [222] R. Menikoff and B. J. Plohr. The Riemann problem for fluid flow of real materials. *Reviews of Modern Physics*, 61:75–130, January 1989. doi: 10.1103/RevModPhys.61.75.

-
- [223] F. C. Michel. Accretion of matter by condensed objects. *Astrophys. Spa. Sci.*, 15:153, 1972.
 - [224] A. Mignone, T. Plewa, and G. Bodo. The piecewise parabolic method for multidimensional relativistic fluid dynamics. *Astrophys. J. Supp.*, 160(1): 199, 2005. doi: 10.1086/430905.
 - [225] A. Mignone, M. Ugliano, and G. Bodo. A five-wave Harten-Lax-van Leer Riemann solver for relativistic magnetohydrodynamics. *Mon. Not. R. Astron. Soc.*, 393(4):1141–1156, March 2009. ISSN 00358711. doi: 10.1111/j.1365-2966.2008.14221.x. URL <http://doi.wiley.com/10.1111/j.1365-2966.2008.14221.x>.
 - [226] A. Mignone, P. Tzeferacos, and G. Bodo. High-order conservative finite difference GLM-MHD schemes for cell-centered MHD. *J. Comp. Phys.*, 229:5896–5920, aug 2010. doi: 10.1016/j.jcp.2010.04.013.
 - [227] D. Mihalas and B. Mihalas. *Foundations of radiation hydrodynamics*. 1984.
 - [228] S. T. Millmore and I. Hawke. Numerical simulations of interfaces in relativistic hydrodynamics. *Class. Quantum Grav.*, 27(1):015007, January 2010. doi: 10.1088/0264-9381/27/1/015007.
 - [229] Charles W. Misner, Kip S. Thorne, and John A. Wheeler. *Gravitation*. W. H. Freeman, San Francisco, 1973.
 - [230] J. E. Morel, T. A. Wareing, R. B. Lowrie, and D. K. Parsons. Analysis of ray-effect mitigating techniques. *Nucl. Sci. Eng.*, 144:1, 2003.
 - [231] A. H. Mroue and *et al.* A catalog of 171 high-quality binary black-hole simulations for gravitational-wave astronomy. *arxiv:1304.6077*, April 2013.
 - [232] Ilia Musco, John C. Miller, and Luciano Rezzolla. Computations of primordial black-hole formation. *Class. Quantum Grav.*, 22:1405–1424, 2005.
 - [233] Takashi Nakamura, Ken-ichi Oohara, and Yasufumi Kojima. General relativistic collapse to black holes and gravitational waves from black holes. *Prog. Theor. Phys. Suppl.*, 90:1–218, 1987.
 - [234] D Neilsen, E W Hirschmann, and R S Millward. Relativistic MHD and black hole excision: Formulation and initial tests. *Class. Quantum Grav.*, 23:S505–S527, 2006.
 - [235] H. Nessyahu and E. Tadmor. Non-oscillatory Central Differencing for Hyperbolic Conservation Laws. *J. Comp. Phys.*, 87:408, April 1990. doi: 10.1016/0021-9991(90)90260-8.
 - [236] S. C. Noble and M. W. Choptuik. Type II critical phenomena of neutron star collapse. *Phys. Rev. D*, 78(6):064059, September 2008. doi: 10.1103/PhysRevD.78.064059.
 - [237] Scott C. Noble. *A Numerical Study of Relativistic Fluid Collapse*. PhD thesis, University of Texas at Austin, 2003. URL <http://www.citebase.org/abstract?id=oai:arXiv.org:gr-qc/0310116>.

- [238] W Noh. Errors for calculations of strong shocks using an artificial viscosity and an artificial heat flux. *J. Comp. Phys.*, 72(1):78, 1987. doi: 10.1016/0021-9991(87)90074-X.
- [239] Jerome Novak. Velocity-induced collapses of stable neutron stars. *Astron. Astrophys.*, 376:606, 2001. URL <http://www.citebase.org/abstract?id=oai:arXiv.org:gr-qc/0107045>.
- [240] M. Obergaulinger, M. A. Aloy, and E. Müller. Local simulations of the magnetized Kelvin-Helmholtz instability in neutron-star mergers. *Astron. Astrophys.*, 515:A30, June 2010. doi: 10.1051/0004-6361/200913386.
- [241] E. O'Connor and C. D. Ott. A new open-source code for spherically symmetric stellar collapse to neutron stars and black holes. *Class. Quantum Grav.*, 27(11):114103, June 2010. doi: 10.1088/0264-9381/27/11/114103.
- [242] Gordon L. Olson. Second-order time evolution of PN equations for radiation transport. *J. Comput. Phys.*, 228(8):3072–3083, May 2009. ISSN 0021-9991. doi: 10.1016/j.jcp.2009.01.012. URL <http://dx.doi.org/10.1016/j.jcp.2009.01.012>.
- [243] Gordon L. Olson. Alternate closures for radiation transport using Legendre polynomials in 1d and spherical harmonics in 2d. *J. Comput. Phys.*, 231(7):2786–2793, April 2012. ISSN 0021-9991. doi: 10.1016/j.jcp.2011.12.013. URL <http://dx.doi.org/10.1016/j.jcp.2011.12.013>.
- [244] Gordon L. Olson, Lawrence H. Auer, and Michael L. Hall. Diffusion, P1, and other approximate forms of radiation transport. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 64(6):619–634, 2000. ISSN 0022-4073. doi: 10.1016/S0022-4073(99)00150-8. URL <http://www.sciencedirect.com/science/article/pii/S0022407399001508>.
- [245] C. D. Ott, A. Burrows, L. Dessart, and E. Livne. Two-Dimensional Multigroup, Multigroup Neutrino Radiation-Hydrodynamic Simulations of Postbounce Supernova Cores. *Astrophys. J.*, 685:1069–1088, October 2008. doi: 10.1086/591440.
- [246] Jayandran Palaniappan, Robert B. Haber, and Robert L. Jerrard. A space-time discontinuous Galerkin method for scalar conservation laws. *Comput. Methods Appl. Mech. Engrg.*, 193:3607, 2004. doi: 10.1016/j.cma.2004.01.028.
- [247] F. Pannarale, L. Rezzolla, F. Ohme, and J. S. Read. Will black hole-neutron star binary inspirals tell us about the neutron star equation of state? *Phys. Rev. D*, 84(10):104017, November 2011. doi: 10.1103/PhysRevD.84.104017.
- [248] Philippos Papadopoulos and José A. Font. Relativistic hydrodynamics on spacelike and null surfaces: Formalism and computations of spherically symmetric spacetimes. *Phys. Rev. D*, 61:024015, December 15 1999. doi: 10.1103/PhysRevD.61.024015.
- [249] A Patera. A spectral element method for fluid dynamics: Laminar flow in a channel expansion. *J. Comp. Phys.*, 54(3):468, 1984. doi: 10.1016/0021-9991(84)90128-1.

-
- [250] M. Perucho, J. M. Martí, J. M. Cela, M. Hanaasz, R. de La Cruz, and F. Rubio. Stability of three-dimensional relativistic jets: implications for jet collimation. *Astron. Astrophys.*, 519:A41, September 2010. doi: 10.1051/0004-6361/200913012.
 - [251] Manuel Perucho, Michal Hanaasz, José-María Martí, and Juan-Antonio Miralles. Resonant Kelvin-Helmholtz modes in sheared relativistic flows. *Phys. Rev. E*, 75(5):1–10, May 2007. ISSN 1539-3755. doi: 10.1103/PhysRevE.75.056312. URL <http://link.aps.org/doi/10.1103/PhysRevE.75.056312>.
 - [252] T Plewa and E. Mueller. The consistent multi-fluid advection method. *Astron. Astrophys.*, 342:179–191, 1999. URL <http://arxiv.org/abs/astro-ph/9807241>.
 - [253] Eric Poisson. *Advanced General Relativity*. <http://www.physics.uoguelph.ca/poisson/research/agr.pdf>, 2002. ISBN 0521830915.
 - [254] Denis Pollney, Christian Reisswig, Luciano Rezzolla, Béla Szilágyi, Marcus Ansorg, Barret Deris, Peter Diener, Ernst Nils Dorband, Michael Koppitz, Alessandro Nagar, and Erik Schnetter. Recoil velocities from equal-mass binary black-hole mergers: a systematic investigation of spin-orbit aligned configurations. *Phys. Rev. D*, 76:124002, 2007.
 - [255] Gerard C. Pomraning. *The Equations of Radiation Hydrodynamics*. Pergamon Press, Oxford and New York, 1973.
 - [256] J. A. Pons, J. M. Ibanez, and J. A. Miralles. Hyperbolic character of the angular moment equations of radiative transfer and numerical methods. *Monthly Notices of the Royal Astronomical Society*, 317(3): 550–562, Sep 2000. doi: 10.1046/j.1365-8711.2000.03679.x. URL <http://dx.doi.org/10.1046/j.1365-8711.2000.03679.x>.
 - [257] 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.
 - [258] D. Porter, A. Pouquet, and P. Woodward. Measures of intermittency in driven supersonic flows. *Phys. Rev. E*, 66(2):1–12, August 2002. ISSN 1063-651X. doi: 10.1103/PhysRevE.66.026301. URL <http://link.aps.org/doi/10.1103/PhysRevE.66.026301>.
 - [259] J. Qiu and C. Shu. Runge Kutta Discontinuous Galerkin Methods using WENO limiters. *SIAM J. Sci. Comp.*, 26:907–929, 2005.
 - [260] Jianxian Qiu and Chi-Wang Shu. Hermite WENO schemes and their application as limiters for Runge-Kutta discontinuous Galerkin method: one-dimensional case. *J. Comp. Phys.*, 193(1):115–135, January 2004. ISSN 00219991. doi: 10.1016/j.jcp.2003.07.026. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999103004212>.
 - [261] A. Quarteroni and A. Valli. *Numerical Approximation of Partial Differential Equations*. Springer, 1997.

- [262] J.J. Quirk. A contribution to the great Riemann solver debate. *International Journal for Numerical Methods in Fluids*, 18(6):555–574, 1994. URL <http://onlinelibrary.wiley.com/doi/10.1002/flid.1650180603/abstract>.
- [263] D. Radice and L. Rezzolla. Discontinuous Galerkin methods for general-relativistic hydrodynamics: Formulation and application to spherically symmetric spacetimes. *Phys. Rev. D*, 84(2):024010, July 2011. doi: 10.1103/PhysRevD.84.024010.
- [264] D. Radice and L. Rezzolla. THC: a new high-order finite-difference high-resolution shock-capturing code for special-relativistic hydrodynamics. *Astron. Astrophys.*, 547:A26, June 2012. doi: 10.1051/0004-6361/201219735.
- [265] D. Radice, L. Rezzolla, and T. Kellerman. Critical phenomena in neutron stars: I. Linearly unstable nonrotating models. *Class. Quantum Grav.*, 27(23):235015, December 2010. doi: 10.1088/0264-9381/27/23/235015.
- [266] Markus Rampp and H.-Thomas Janka. Radiation hydrodynamics with neutrinos: Variable Eddington factor method for core-collapse supernova simulations. *Astron. Astrophys.*, 396:361, 2002.
- [267] W. H Reed and T. R. Hill. Triangular mesh methods for the neutron transport equation. Technical report, Los Alamos Scientific Laboratory, 1973.
- [268] C. Reisswig and Pollney D. Notes on the integration of numerical relativity waveforms. *Class. Quantum Grav.*, 28:195015, 2011.
- [269] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilagyi. Unambiguous determination of gravitational waveforms from binary black hole mergers. *Phys. Rev. Lett.*, 103:221101, 2009.
- [270] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilagyi. Characteristic extraction in numerical relativity: binary black hole merger waveforms at null infinity. *Class. Quantum Grav.*, 27:075014, 2010.
- [271] L. Rezzolla and J. C. Miller. Relativistic radiative transfer for spherical flows. *Class. Quantum Grav.*, 11:1815–1832, July 1994. doi: 10.1088/0264-9381/11/7/018.
- [272] L. Rezzolla and O. Zanotti. New Relativistic Effects in the Dynamics of Nonlinear Hydrodynamical Waves. *Phys. Rev. Lett.*, 89(11):114501, August 2002. doi: 10.1103/PhysRevLett.89.114501.
- [273] L. Rezzolla, O. Zanotti, and J. A. Pons. An Improved Exact Riemann Solver for Relativistic Hydrodynamics with non-zero tangential velocities. *Journ. of Fluid Mech.*, 479:199, 2003.
- [274] L. Rezzolla, B. Giacomazzo, L. Baiotti, J. Granot, C. Kouveliotou, and M. A. Aloy. The missing link: Merging neutron stars naturally produce jet-like structures and can power short Gamma-Ray Bursts. *Astrophys. J.*, 732(11):L6, May 2011.
- [275] Luciano Rezzolla and Olindo Zanotti. *Relativistic Hydrodynamics*. Oxford University Press. Oxford, UK, 2013.

-
- [276] W Rider. Revisiting wall heating. *J. Comp. Phys.*, 162(2):395, 2000. doi: 10.1006/jcph.2000.6544.
 - [277] P. L. Roe. Approximate Riemann solvers, parameter vectors and difference schemes. *J. Comput. Phys.*, 43:357, 1981.
 - [278] J. V. Romero, J. M. Ibáñez, J. M. Martí, and J. A. Miralles. A new spherically symmetric general relativistic hydrodynamical code. *Astrophys. J.*, 462: 839–854, 1996.
 - [279] Stephan Rosswog. Conservative, special-relativistic smoothed particle hydrodynamics. *J. Comp. Phys.*, 229(22):8591, 2010. doi: 10.1016/j.jcp.2010.08.002.
 - [280] M. Ruffert and H.-Th. Janka. Colliding neutron stars, gravitational waves, neutrino emission, and gamma-ray bursts. *Astron. Astrophys.*, 338:535–555, 1998.
 - [281] Milton Ruiz, David Hilditch, and Sebastiano Bernuzzi. Constraint preserving boundary conditions for the Z4c formulation of general relativity. *Phys. Rev. D*, 83:024025, 2011. doi: 10.1103/PhysRevD.83.024025.
 - [282] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, 1996.
 - [283] L. Santamaría et al. Matching post-Newtonian and numerical relativity waveforms: Systematic errors and a new phenomenological model for nonprecessing black hole binaries. *Phys. Rev. D*, 82(6):064016, September 2010. doi: 10.1103/PhysRevD.82.064016.
 - [284] Shigeo Sasaki. On the differential geometry of tangent bundles of Riemannian manifolds. *Tohoku Mathematical Journal*, 10(3):338, 1958. doi: 10.2748/tmj/1178244668.
 - [285] Shigeo Sasaki. On the differential geometry of tangent bundles of Riemannian manifolds, II. *Tohoku Mathematical Journal*, 14(2):146, 1962. doi: 10.2748/tmj/1178244169.
 - [286] B. S. Sathyaprakash and B. F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Rev. Relativ.*, 12:2, 2009.
 - [287] W Schmidt, W Hillebrandt, and J Niemeyer. Numerical dissipation and the bottleneck effect in simulations of compressible isotropic turbulence. *Computers & Fluids*, 35(4):353–371, May 2006. ISSN 00457930. doi: 10.1016/j.compfluid.2005.03.002. URL <http://linkinghub.elsevier.com/retrieve/pii/S0045793005000563>.
 - [288] V. Schneider, U. Katscher, D. H. Rischke, B. Waldhauser, J. A. Maruhn, and C. D. Munz. *J. Comput. Phys.*, 105:92, 1993.
 - [289] Erik Schnetter, Scott H. Hawley, and Ian Hawke. Evolutions in 3D numerical relativity using fixed mesh refinement. *Class. Quantum Grav.*, 21(6):1465–1488, 21 March 2004.
 - [290] Bernard F. Schutz. *A first course in general relativity*. Cambridge University Press, 1985.

- [291] Y. Sekiguchi. Stellar Core Collapse in Full General Relativity with Microphysics – Formulation and Spherical Collapse Test. *Progress of Theoretical Physics*, 124:331–379, August 2010.
- [292] Y. Sekiguchi, K. Kiuchi, K. Kyutoku, and M. Shibata. Gravitational waves and neutrino emission from the merger of binary neutron stars. *Phys. Rev. Lett.*, 107:051102, May 2011.
- [293] Z.S. She and E. Leveque. Universal scaling laws in fully developed turbulence. *Phys. Rev. Lett.*, 72(3):336–339, 1994. URL <http://link.aps.org/doi/10.1103/PhysRevLett.72.336>.
- [294] Masaru Shibata and Takashi Nakamura. Evolution of three-dimensional gravitational waves: Harmonic slicing case. *Phys. Rev. D*, 52:5428, 1995.
- [295] C. W. Shu. Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws. Lecture notes ICASE Report 97-65; NASA CR-97-206253, NASA Langley Research Center, 1997. URL http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19980007543_1998045663.pdf.
- [296] C. W. Shu. High Order ENO and WENO Schemes for Computational Fluid Dynamics. In T. J. Barth and H. Deconinck, editors, *High-Order Methods for Computational Physics*, pages 439–582. Springer, 1999.
- [297] C. W. Shu. High order finite difference and finite volume WENO schemes and discontinuous Galerkin methods for CFD. Technical Report ICASE Report 2001-11; NASA CR-2001-210865, NASA Langley Research Center, 2001.
- [298] C.-W. Shu and S. Osher. Efficient Implementation of Essentially Non-oscillatory Shock-Capturing Schemes, II. *J. Comp. Phys.*, 83:32, July 1989. doi: 10.1016/0021-9991(89)90222-2.
- [299] C. W. Shu and S. J. Osher. Efficient Implementation of Essentially Non-oscillatory Shock-Capturing Schemes. *J. Comput. Phys.*, 77:439, 1988.
- [300] C.W. Shu. High order finite difference and finite volume WENO schemes and Discontinuous Galerkin methods for CFD. *International Journal of Computational fluid dynamics*, 17:107–118, 2003.
- [301] C.W. Shu. High order numerical methods for time dependent Hamilton-Jacobi equations. *Lecture Notes Series, Institute for Mathematical Sciences, National University of Singapore*, 11:47–91, 2007.
- [302] S. Siegler and H. Riffert. Smoothed particle hydrodynamics simulations of ultrarelativistic shocks with artificial viscosity. *Astrophys. J.*, 531(2): 1053, 2000. doi: 10.1086/308482.
- [303] J. M. Smit, J. Cernohorsky, and C. P. Dullemond. Hyperbolicity and critical points in two-moment approximate radiative transfer. *Astron. Astrophys.*, 325:203, September 1997.

-
- [304] J. Smoller and B. Temple. Global solutions of the relativistic Euler equations. *Comm. Math. Phys.*, 156:67–99, September 1993. doi: 10.1007/BF02096733.
 - [305] G. A. Sod. A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. *J. Comp. Phys.*, 27:1–31, April 1978. doi: 10.1016/0021-9991(78)90023-2.
 - [306] Carlos F. Sopuerta and Pablo Laguna. A finite element computation of the gravitational radiation emitted by a point-like object orbiting a non-rotating black hole. *Phys. Rev. D*, 73:044028, 2006. doi: 10.1103/PhysRevD.73.044028.
 - [307] Carlos F. Sopuerta, Pengtao Sun, Pablo Laguna, and Jinchao Xu. A toy model for testing finite element methods to simulate extreme-mass-ratio binary systems. *Class. Quantum Grav.*, 23:251–286, 2006. doi: 10.1088/0264-9381/23/1/013.
 - [308] Volker Springel. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Mon. Not. R. Astron. Soc.*, 401(2):791–851, January 2010. ISSN 00358711. doi: 10.1111/j.1365-2966.2009.15715.x. URL <http://doi.wiley.com/10.1111/j.1365-2966.2009.15715.x>.
 - [309] K. Sumiyoshi and S. Yamada. Neutrino Transfer in Three Dimensions for Core-collapse Supernovae. I. Static Configurations. *Astrophys. J. Suppl.*, 199:17, March 2012.
 - [310] A. Suresh and H. T. Huynh. Accurate Monotonicity-Preserving Schemes with Runge-Kutta Time Stepping. *J. Comp. Phys.*, 136(1):83–99, 1997. ISSN 0021-9991. doi: DOI:10.1006/jcph.1997.5745. URL <http://www.sciencedirect.com/science/article/B6WHY-45V7FSX-6/2/d88d3b5c02364ae5aa69b21d0e0787e7>.
 - [311] J. L. Synge. *The relativistic gas*. North-Holland Publishing, Amsterdam, 1957.
 - [312] Eitan Tadmor. Approximate solutions of nonlinear conservation laws. *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations*, (1991), 1998. doi: 10.1007/BFb0096352. URL <http://www.springerlink.com/index/G3M50642X4512V53.pdf>.
 - [313] A. Taflove and S. C. Hagness. *Computational Electrodynamics: The Finite-Difference Time-Domain Method, 3rd edition*. Norwood, MA: Artech House, 2005.
 - [314] Jared Tanner. Optimal filter and mollifier for piecewise smooth spectral data. *Math. Comp.*, 75(254):767–791, January 2006. ISSN 0025-5718. doi: 10.1090/S0025-5718-06-01822-9. URL <http://www.ams.org/journal-getitem?pii=S0025-5718-06-01822-9>.
 - [315] G. E. Tauber and J. W. Weinberg. Internal state of a gravitating gas. *Phys. Rev.*, 122(4):1342, 1961. doi: 10.1103/PhysRev.122.1342.

- [316] Ellen M. Taylor, Minwei Wu, and M. Pino Martín. Optimization of nonlinear error for weighted essentially non-oscillatory methods in direct numerical simulations of compressible turbulence. *J. Comp. Phys.*, 223(1):384–397, 2007. ISSN 0021-9991. doi: DOI:10.1016/j.jcp.2006.09.010. URL <http://www.sciencedirect.com/science/article/B6WHY-4M877S3-1/2/984e4f76b52d2a76ac9e13b5201fdef8>.
- [317] Alexander Tchekhovskoy, Jonathan C. McKinney, and Ramesh Narayan. wham: a WENO-based general relativistic numerical scheme ? I. Hydrodynamics. *Mon. Not. R. Astron. Soc.*, 379(2):469–497, August 2007. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2007.11876.x. URL <http://doi.wiley.com/10.1111/j.1365-2966.2007.11876.x>.
- [318] Alexander Tchekhovskoy, Jonathan C. McKinney, and Ramesh Narayan. WHAM: a WENO-based general relativistic numerical scheme. I. Hydrodynamics. *Mon. Not. R. Astron. Soc.*, 379(2):469, 2007. doi: 10.1111/j.1365-2966.2007.11876.x.
- [319] M. Thierfelder, S. Bernuzzi, D. Hilditch, B. Bruegmann, and L. Rezzolla. The trumpet solution from spherical gravitational collapse with puncture gauges. *Phys. Rev. D*, 83:064022, December 2010.
- [320] Marcus Thierfelder, Sebastiano Bernuzzi, and Bernd Brügmann. Numerical relativity simulations of binary neutron stars. *Phys. Rev. D*, 84(4): 1–30, August 2011. ISSN 1550-7998. URL <http://link.aps.org/doi/10.1103/PhysRevD.84.044012>.
- [321] K. S. Thorne. Relativistic radiative transfer - Moment formalisms. *Mon. Not. R. Astron. Soc.*, 194:439–473, February 1981.
- [322] Robert F. Tooper. Adiabatic fluid spheres in general relativity. *Astrophys. J.*, 142:1541, 1965.
- [323] E. F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer-Verlag, 1999.
- [324] H. Touil, M.Y. Hussaini, and M. Sussman. Tracking discontinuities in hyperbolic conservation laws with spectral accuracy. *J. Comput. Phys.*, 225:1810, 2007. doi: 10.1016/j.jcp.2007.02.016.
- [325] B. J. van Leer. Towards the ultimative conservative difference scheme I. the quest for monotonicity. *Lecture Notes in Physics*, 18:163–168, 1973.
- [326] James R. van Meter, John G. Baker, Michael Koppitz, and Dae-Il Choi. How to move a black hole without excision: gauge conditions for the numerical evolution of a moving puncture. *Phys. Rev. D*, 73:124011, 2006.
- [327] Hervé Vandeven. Family of spectral filters for discontinuous problems. *Journal of Scientific Computing*, 6(2):159–192, June 1991. ISSN 0885-7474. doi: 10.1007/BF01062118. URL <http://www.springerlink.com/index/10.1007/BF01062118>.

-
- [328] Todd L. Veldhuizen. Arrays in blitz++. In *Proceedings of the Second International Symposium on Computing in Object-Oriented Parallel Environments*, ISCOPE '98, pages 223–230, London, UK, 1998. Springer-Verlag. ISBN 3-540-65387-2. URL <http://portal.acm.org/citation.cfm?id=646894.709708>.
 - [329] Justin Vines, Éanna É. Flanagan, and Tanja Hinderer. Post-1-newtonian tidal effects in the gravitational waveform from binary inspirals. *Phys. Rev. D*, 83:084051, Apr 2011. doi: 10.1103/PhysRevD.83.084051. URL <http://link.aps.org/doi/10.1103/PhysRevD.83.084051>.
 - [330] Robert M. Wald. *General relativity*. The University of Chicago Press, Chicago, 1984. ISBN 0-226-87032-4 (hardcover), 0-226-87033-2 (paperback).
 - [331] Z Wang. Spectral (finite) volume method for conservation laws on unstructured grids. basic formulation basic formulation. *J. Comp. Phys.*, 178(1):210, 2002. doi: 10.1006/jcph.2002.7041.
 - [332] Andreas Weyhausen, Sebastiano Bernuzzi, and David Hilditch. Constraint damping for the Z4c formulation of general relativity. *Phys. Rev.*, D85:024038, 2012. doi: 10.1103/PhysRevD.85.024038.
 - [333] J. R. Wilson. Numerical study of fluid flow in a kerr space. *Astrophys. J.*, 173:431, apr 1972. doi: 10.1086/151434.
 - [334] P. Woodward and P. Collela. *J. Comput. Phys.*, 54:115–173, 1984.
 - [335] Daoqi Yang. *C++ and Object-Oriented Numeric Computing for Scientists and Engineers*. Springer-Verlag, 2000.
 - [336] S. Yoshida and Y. Eriguchi. Quasi-radial modes of rotating stars in general relativity. *Mon. Not. R. Astron. Soc.*, 322:389, 2001.
 - [337] O. Zanotti, C. Roedig, L. Rezzolla, and L. Del Zanna. General relativistic radiation hydrodynamics of accretion flows - I. Bondi-Hoyle accretion. *Mon. Not. R. Astron. Soc.*, 417:2899–2915, November 2011. doi: 10.1111/j.1365-2966.2011.19451.x.
 - [338] W. Zhang, A. MacFadyen, and P. Wang. Three-Dimensional Relativistic Magnetohydrodynamic Simulations of the Kelvin-Helmholtz Instability: Magnetic Field Amplification by a Turbulent Dynamo. *Astrophys. J.*, 692:L40–L44, February 2009. doi: 10.1088/0004-637X/692/1/L40.
 - [339] Weiqun Zhang and A.I. MacFadyen. RAM: A relativistic adaptive mesh refinement hydrodynamics code. *Astrophys. J. Suppl.*, 164:255, 2006. URL <http://iopscience.iop.org/0067-0049/164/1/255>.
 - [340] X. Zhang and C.-W. Shu. On positivity-preserving high order discontinuous Galerkin schemes for compressible Euler equations on rectangular meshes. *J. Comp. Phys.*, 229:8918–8934, November 2010. doi: 10.1016/j.jcp.2010.08.016.

- [341] X. Zhang and C.-W. Shu. Maximum-principle-satisfying and positivity-preserving high-order schemes for conservation laws: survey and new developments. *Proc. R. Soc. A: Math., Phys. Eng. Sciences*, 467(2134):2752–2776, May 2011. ISSN 1364-5021. doi: 10.1098/rspa.2011.0153. URL <http://rspa.royalsocietypublishing.org/cgi/doi/10.1098/rspa.2011.0153>.
- [342] Xiangxiong Zhang and Chi-Wang Shu. Positivity-preserving high order discontinuous Galerkin schemes for compressible Euler equations with source terms. *J. Comp. Phys.*, 230(4):1238–1248, February 2011. ISSN 00219991. doi: 10.1016/j.jcp.2010.10.036. URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999110006017>.
- [343] Tie Zhou, Yinfan Li, and Chi-Wang Shu. Numerical comparison of weno finite volume and runge–kutta discontinuous galerkin methods. *Journal of Scientific Computing*, 16(2):145, 2001. doi: 10.1023/A:1012282706985.
- [344] Yosef Zlochower, Marcelo Ponce, and Carlos O. Lousto. Accuracy issues for numerical waveforms. *Phys. Rev. D*, 86(10):104056, November 2012. ISSN 1550-7998. doi: 10.1103/PhysRevD.86.104056. URL <http://link.aps.org/doi/10.1103/PhysRevD.86.104056>.
- [345] Jonathan Zrake and Andrew MacFadyen. Numerical Simulations of Driven Supersonic Relativistic MHD Turbulence. pages 102–105, September 2011. doi: 10.1063/1.3621748. URL <http://arxiv.org/abs/1109.6294><http://link.aip.org/link/APCPCS/v1358/i1/p102/s1&Agg=doi>.
- [346] Jonathan Zrake and Andrew I. MacFadyen. Numerical simulations of driven relativistic magnetohydrodynamic turbulence. *Astrophys. J.*, 744(1):32, January 2012. ISSN 0004-637X. doi: 10.1088/0004-637X/744/1/32. URL <http://arxiv.org/abs/1108.1991><http://stacks.iop.org/0004-637X/744/i=1/a=32?key=crossref.8d123aad5745527a52069623fb7f1309>.
- [347] G Zumbusch. Finite element, discontinuous Galerkin, and finite difference evolution schemes in spacetime. *Class. Quantum Grav.*, 26(17):175011, 2009. doi: 10.1088/0264-9381/26/17/175011.

Curriculum Vitæ

Personal

Name	David Radice
Citizenship	Italy
Date of Birth	03.06.1984

Education

2010 – Today	PhD Student in Gravitational Wave Astronomy Numerical Relativity Group Max Planck Institute for Gravitational Physics Supervisors Prof. Bernard F. Schutz, bernard.schutz@aei.mpg.de Prof. Luciano Rezzolla, rezzolla@aei.mpg.de
2006 – 2009	Master's degree in Mathematical Engineering Department of Mathematics, Politecnico di Milano Supervisors Prof. Luciano Rezzolla, rezzolla@aei.mpg.de Prof. Giulio Magli, giulio.magli@polimi.it Thesis title <i>Numerical Simulations of Critical Phenomena in Neutron Star Collapse</i> Grade 110/110 cum laude
2003 – 2006	Bachelor's degree in Mathematical Engineering Department of Mathematics, Politecnico di Milano Supervisor Prof. Alessandro Veneziani, ale@mathcs.emory.edu Thesis title <i>Analisi numerica di un nuovo modello di traffico automobilistico</i> (Numerical Analysis of a New Model of Traffic Flow) Grade 108/110

Invited Talks

February 2013	<i>A New Spherical-Harmonics Scheme for Multi-Dimensional Radiation Transport</i> Talk at SIAM Computational Science and Engineering 2013, Boston, USA
June 2012	<i>Turbulence and Relativity</i> Seminar at TAPIR, California Institute of Technology, USA
January 2011	<i>Discontinuous Galerkin Methods for General Relativistic Hydrodynamics</i> Seminar: Albert Einstein Institute and Lusiana State University shared video-seminar, Potsdam, Germany

Publications

D. Radice, L. Rezzolla and F. Galeazzi. *Beyond second-order convergence in simulations of binary neutron stars in full general-relativity*, arXiv:1306.6052

D. Radice, E. Abdikamalov, L. Rezzolla and C.D. Ott, *A New Spherical-Harmonics Scheme for Multi-Dimensional Radiation Transport I: Static Matter Configurations*, J. of Comp. Physics **242**, 648 (2013), arXiv:1209.1634

D. Radice and L. Rezzolla, *Universality and Intermittency in Relativistic Turbulent Flows of Hot Plasma*, Astrophys. J. Letters **766**, L10 (2013), arXiv:1209.2936

D. Radice and L. Rezzolla, *THC: a New High-Order Finite-Difference High-Resolution Shock-Capturing Code for Special-Relativistic Hydrodynamics*, A&A **547**, A26 (2012), arXiv:1206.6502

D. Radice and L. Rezzolla, *Discontinuous Galerkin Methods for General-Relativistic Hydrodynamics: Formulation and Application to Spherical Symmetric Spacetimes*, Phys. Rev. D **84**, 024010 (2011), arXiv:1103.2426

T. Kellerman, L. Rezzolla and **D. Radice**, *Critical Phenomena in Neutron Stars II: Head-on Collisions*, Class. and Quantum Grav. **27**, 235016 (2010), arXiv:1007.2797

D. Radice, L. Rezzolla and T. Kellerman, *Critical Phenomena in Neutron Stars I: Linearly Unstable Nonrotating Models*, Class. and Quantum Grav. **27**, 235015 (2010), arXiv:1007.2809