SISSA
Scuola
Internazionale
Superiore di
Studi Avanzati

Physics Area – PhD course in

Astrophysics and Cosmology

Relativistic Magnetohydrodynamic Modelling of Gamma-ray Burst Jets: From Early Propagation to Afterglow Emission

Candidate:

Emma Dreas

Advisors:

Prof. Annalisa Celotti

Dr. Riccardo Ciolfi

Dr. Om Sharan Salafia

Co-advisor:

Dr. Andrea Pavan

Academic Year 2025-26

For the ones who think for a living
For the minds that never quite switch off

To the people who look at the stars and wish
To the stars that listen—and the dreams that are answered

S. J. Maas

Then our hearts combined like
A neutron star collision

Neutron Star Collision, Muse

A million colours of hazel, golden and red
Oh, I won't be silent and I won't let go
I will hold on tighter 'til the afterglow

Afterglow, Ed Sheeran

Abstract

Gamma-ray bursts are among the most energetic transient phenomena in the Universe. Since their discovery, a dedicated branch of research has been developed to understand their origin and the physics involved. They are thought to arise from ultra-relativistic jets, launched by at least two distinct classes of progenitors: a binary neutron star merger, and the core collapse of a rotating, strongly magnetised massive star. In both cases, the jet is launched by a central engine, harbouring a rotating, newly formed compact object, into a dense and inhomogeneous progenitor environment: the binary neutron star merger ejecta or the collapsing stellar envelope. It is the interaction with this environment that shapes the jet as it reaches the larger scales where electromagnetic emission takes place.

The interpretation of observational data is most frequently performed through models that assume an axisymmetric jet configuration described by simple analytical functions. As a consequence, the jet-structure parameters, viewing angles and energetics inferred from afterglow fitting rest on assumptions whose validity in a genuinely three-dimensional setting remains untested. This thesis develops a framework that connects three-dimensional relativistic magnetohydrodynamic simulations of the early jet propagation to the afterglow signal produced at much larger scales, targeting both progenitor channels. The simulations resolve the jet–progenitor interaction during its early-time evolution and follow how the resulting features propagate outwards; the afterglow prediction is carried out within a flexible semi-analytic framework extended to include the full three-dimensional velocity and energy distribution of the jet material.

Focussing on the binary neutron star merger scenario, we characterise the dynamical evolution of a fiducial simulation, identifying the formation of a multi-shock structure whose outermost shell carries a non-negligible fraction of the total energy, and the most energetic region evolves into a spherical-cap geometry marked by pronounced angular asymmetries and radial stratification. For the resulting outflow structure to be meaningfully used as input for the afterglow emission predictions, the evolution must be computed up to a time when the jet structure has reached its asymptotic configuration. To this end, we introduce a remapping method that extends three-dimensional jet simulations up to a saturating regime approaching ballistic expansion. We then consider a set of analogous models in which the jet injection parameters are varied around the fiducial configuration in terms of injected energy, magnetisation and launching time. The main qualitative features persist across a range of injection conditions, and the parameter variations enable an assessment of their relative impact on the jet propagation and resulting structure. To bridge the gap between the scales accessible to the simulations and the region where the

afterglow is produced, the resulting jet structures are directly fed into the extended semi-analytic model. Accounting for the full three-dimensional structure reveals that the outer shell leaves an observational imprint at early times, only present when the radial velocity stratification is included. The deviations from axisymmetry produce order-of-magnitude differences even for small offsets in the viewing angle, a feature that axisymmetric models cannot reproduce. This azimuthal variation provides an additional degree of freedom when interpreting observational data. A proof-of-concept fit to the afterglow of GRB 170817A closes the chain, directly connecting jet structural features to the shape of the observed signal. The results show that the resulting jet structures can be brought into good agreement with the reference event within the available observational and theoretical constraints. This highlights the complexity of a direct jet-to-emission connection given the large degeneracy between the parameters involved.

The same strategy is then applied, at a preliminary stage, to the massive-star core-collapse channel, addressing the interaction between an incipient relativistic jet and a stellar envelope undergoing a magneto-rotational supernova. The goal is to investigate how this interaction shapes the jet's structure and stability, which conditions determine the successful propagation of the jet beyond the progenitor star, and how the concurrent supernova explosion affects this evolution. This preliminary step provides the basis for a systematic treatment of both classes of gamma-ray burst progenitors within a common framework.

Together, these results establish a route from three-dimensional progenitor-scale simulations to observable afterglow signatures, providing a framework for interpreting current and forthcoming gamma-ray burst observations in the multi-messenger era.

Contents

Abstract . . . . . **v**

Contents . . . . . **vii**

List of Figures . . . . . **xi**

List of Tables . . . . . **xv**

1 Introduction **1**
1.1 The Origin of Gamma-Ray Bursts 1
1.2 The Challenge: Connecting Progenitors to Observables 2
1.3 Thesis Objectives and Structure 3

2 Gamma-Ray Bursts **5**
2.1 Observational Phenomenology 5
2.1.1 Temporal variability and the compactness problem 5
2.1.2 Cosmological origin 6
2.1.3 Energetics and jet collimation 6
2.1.4 The prompt emission spectrum: the Band function 6
2.1.5 Duration classes and the T_{90} bimodality 8
2.1.6 Prompt emission morphology 9
2.1.7 Host galaxies 10
2.2 The Progenitors 11
2.2.1 GRBs and Collapsing Massive Stars 11
2.2.2 GRBs and Binary Mergers 14
2.3 The Central Engine 19
2.3.1 The Blandford–Znajek mechanism 19
2.3.2 Accretion regimes: SANE and MAD states 19
2.3.3 Neutrino annihilation 20
2.3.4 Nature of the remnant, implications for jet launching 20
2.3.5 From engine to simulation: the injection prescription 22
2.4 Landmark GRB events 23
2.4.1 GW170817: a multi-messenger event 23
2.4.2 Breaking the Long/Short Dichotomy 25
2.5 Jet Propagation and Structure 27
2.5.1 The jet–environment interaction 27
2.5.2 The jet head and its propagation velocity 27

2.5.3 Jet breakout and the choked jet scenario . . . 28
2.5.4 Jet structure: from top-hat models to realistic profiles . . . 30
2.6 Why Numerical Simulations? . . . 31
2.6.1 From the First 2D Simulations to the 3D Era . . . 31
2.6.2 Idealised vs Realistic Environments: Why It Matters . . . 32
2.7 The GRB Afterglow . . . 33
2.7.1 From Prompt Emission to External Shock . . . 33
2.7.2 Blast Wave Dynamics: the Blandford–McKee Solution . . . 34
2.7.3 Microphysics and the Synchrotron Spectrum . . . 35
2.7.4 Spectral Regimes and Light Curve Slopes . . . 36
2.7.5 Degeneracies and the Importance of Broadband Coverage . . . 38
2.8 Beyond the Standard Afterglow Model . . . 39
2.8.1 The Code . . . 40

3 Theoretical and Numerical Framework . . . 45
3.1 Relativistic Magneto–Hydrodynamics . . . 45
3.1.1 Conservation Laws in Covariant Form . . . 45
3.2 The PLUTO code . . . 47
3.2.1 Code Architecture and Available Modules . . . 48
3.2.2 The Godunov Framework . . . 48
3.2.3 The Divergence-Free Constraint and Hyperbolic Divergence Cleaning . . . 50
3.2.4 Equation of State . . . 51
3.2.5 The Simulation Setup . . . 52

4 Jet Evolution in Realistic Binary Merger Environments . . . 55
4.1 Simulation Setup . . . 55
4.1.1 The Post-Merger Environment . . . 55
4.1.2 Computational Grid . . . 59
4.1.3 Jet Injection Prescription . . . 59
4.1.4 Model Suite . . . 62
4.2 Early Dynamics: Jet–Ejecta Interaction (0–3 s) . . . 63
4.2.1 Jet Propagation and Cocoon Formation . . . 63
4.2.2 The Multi-Shock Structure at Breakout . . . 64
4.2.3 Energy Distribution . . . 66
4.3 Parameter Study: Impact of Injection Conditions . . . 67
4.3.1 Luminosity and Decay Timescale . . . 68
4.3.2 Magnetisation . . . 69
4.3.3 Launch Time . . . 70
4.4 Jet State at $t = 3$ s: A Comparative Overview . . . 70

5 Towards the Ballistic Regime . . . 73
5.1 Why the Ballistic Regime Matters . . . 73
5.2 The Grid Remapping Technique . . . 73
5.3 Assessment of the Magnetic Field Divergence . . . 76
5.4 Energy Conversion and Approach to Saturation . . . 77
5.4.1 Evolution of Energy Components . . . 77

5.4.2 The Saturation of the Jet Structure . . . 79
5.5 Jet Dynamics . . . 81

6 Connecting the Jet Structure to Electromagnetic Emission Across Models . . . 83
6.1 Final Angular and Radial Structure . . . 83
6.1.1 Front-View Angular Distributions . . . 83
6.1.2 Axisymmetrised profiles . . . 86
6.1.3 Radial Velocity Stratification . . . 86
6.2 From Simulation to Observable: the Modelling Pipeline . . . 90
6.3 Application to GW170817 . . . 90
6.3.1 Dataset and Fitting Procedure . . . 90
6.3.2 Best-Fit Results and Parameter Constraints . . . 92
6.4 Observational Signatures of the 3D Structure . . . 95
6.4.1 Impact of Velocity Stratification . . . 95
6.4.2 Impact of Residual Evolution . . . 96
6.4.3 Impact of Azimuthal Asymmetries . . . 97
6.4.4 Impact of Injection Parameters . . . 99

7 Gamma-ray Burst Jets in Core-collapse Supernova Environments 103
7.1 Introduction and Scientific Context . . . 103
7.2 Progenitor Model . . . 104
7.3 Core-collapse Simulation and Initial Data . . . 105
7.4 Analysis of the Initial Data . . . 106
7.4.1 Mass and Density Distribution . . . 107
7.4.2 Angular Momentum and Rotation . . . 108
7.4.3 Magnetic Field Configuration . . . 109
7.4.4 Shock Finding . . . 110
7.4.5 Excision Radius and Central-engine Parameters . . . 110
7.4.6 Equation-of-state Matching . . . 113
7.5 PLUTO simulations . . . 114
7.5.1 Grid Remapping and Setup . . . 114
7.5.2 Jet-off Evolution . . . 115
7.5.3 Jet Evolution . . . 115
7.5.4 Diagnostics and Results . . . 118

8 Conclusions and Future Perspectives . . . 123
8.1 Summary and Discussion . . . 123
8.2 Limitations of the Current Framework and Outlook . . . 126
8.2.1 Limitations due to resolution . . . 126

Appendices . . . 129

A Publications Related to This Thesis . . . 131

B Resolution Study . . . 133
B.1 Motivation and Setup . . . 133

B.2 Outcome Across Models 133
B.3 Case Study: Models *e1* and *e2* 134

C Detectability of the Outer-shell Feature 137

D On the Neglect of Lateral Expansion 141

Bibliography 143

List of Figures

2.1 Band function fit to the average spectrum of GRB911127, adapted from Band et al. [16]. 8
2.2 Distribution of durations for 95 GRBs from the Swift satellite. Adapted from Zhang and Choi [239]. 9
2.3 Light curves (15–350 KeV) of five typical sGRBs from the Swift/BAT catalogue, adapted from Li et al. [106]. 10
2.4 Schematic meridional cross-section (not to scale) of the post-merger environment of a BNS coalescence: shock-heated dynamical ejecta (outer shell) with neutron-rich tidal ejecta along the equatorial plane, oscillation-driven dynamical ejecta (inner shell), and secular outflows (innermost region) with a slow isotropic component and a faster collimated component along the spin axis. The central remnant (metastable NS or BH) is surrounded by an accretion disc. 17
2.5 Multi-wavelength afterglow light curves overlaid with different structured jet models, adapted from Ghirlanda et al. [68] 24
2.6 Schematic of the jet-environment interaction, inspired from Bromberg et al. [26]. 28
2.7 Schematic representation of top-hat (black line), power law (red line) and a Gaussian (blue line) jet profile, normalised to 1. 31
2.8 Representation of the synchrotron spectrum in the slow-cooling regime. 38

3.1 Schematic of PLUTO's algorithmic cycle. 50

4.1 3D rendering of the realistic post-merger environment imported in P23. The mass density distribution and magnetic field lines at 155 ms after merger are shown. Adapted from P23. 57
4.2 3D renderings of the system configuration at 560 ms, 735 ms and 2.385 s after merger, in terms of rest-mass density and Lorentz factor, spatially rescaled for comparison. Adapted from P23. 64
4.3 Meridional view of the energy density distribution excluding rest-mass component (left), and rest-mass density with superimposed Lorentz factor of the elements with $\Gamma > 2$ (right), at $t - t_j = 275$ ms, relative to model F (Table 4.1). 65
4.4 Meridional view of the energy density without rest mass component at $t - t_j = 3$ s for model F. White contours represent shocked surfaces. 65

4.5 Kinetic power computed at different radii at $t - t_j = 4.49\,\mathrm{s}$ for the jet model of P23 as described in D25. The blue, red, and green dot-dashed curves indicate 10%, 15%, and 20% of the peak dE/dt, respectively. The solid magenta curve represents the initial jet luminosity at injection (north side only). 66
4.6 Energy $E_{\rm sum}$(without rest-mass energy contribution) of the fluid elements of the jet cap with $\beta\Gamma$ larger than a threshold, as a function of the threshold. Analysis referred to the model of P23 as presented in D25. 68

5.1 Schematic representation of the grid remapping technique (not to scale). 75
5.2 Rest-mass density profile along the jet axis before and after the remap to the Cartesian domain at $t_{\rm beg}$, in different shades of blue. 76
5.3 Temporal evolution of the energy component of the simulation of D25. The kinetic-to-total ratio is extrapolated at later times with a polynomial fit. 78
5.4 Angular profile of the radially and azimuthally averaged Lorentz factor, together with its saturation value, at $t_{\rm beg}$ and $t_{\rm end}$. 80
5.5 Angular profiles of the azimuthally averaged jet cap Lorentz factor (left) and isotropic equivalent energy (right) at different evolutionary times for the simulation of D25. 81
5.6 Meridional view of Lorentz factor, rest-mass density and magnetic field intensity at three representative times for the simulation of D25. 82

6.1 Front-view angular distribution of isotropic equivalent energy and average Lorentz factor of the jet cap, at $t_{\rm end}$, for the simulation of D25. 84
6.2 3D rendering of the kinetic energy density at $t_{\rm end}$ for the simulation of D25. 85
6.3 Front-view distributions of the isotropic-equivalent energy of the jet cap for the jet models of the parameter study at $t_{\rm end}$. The red line shows the smallest contour containing 50% of the energy; the white cross points at the jet's energy barycentre. 85
6.4 Isotropic equivalent energy profiles of the azimuthally averaged jet cap for all models of the parameter study at $t_{\rm end}$. 87
6.5 Velocity profile of model *F*. The colorbar indicates different polar angles. 88
6.6 Jet energy distribution in the $u - \Theta$ plane. The dark blue shaded region identifies the interquartile range (25th–75th percentiles) while the range between the 5th and 95th percentiles is presented in light blue. 89
6.7 Multi-wavelength afterglow light curves for model *F*, at viewing angles $\theta_{\rm v} = 0°, 32°$. 91
6.8 Corner plots of MCMC fit results for models *F*, *a*, and *b* to the GW170817 data. Each panel shows marginalized one-dimensional histograms and two-dimensional contour posteriors of the afterglow jet-model parameters (*F*: top; *a*: bottom left; *b*: bottom right). Vertical black dashed lines mark the 5th and 95th percentiles, and red solid lines mark the medians. 93

6.9 Corner plots of MCMC fit results of models *c, d, e1* and *e2* to the GW170817 data. . . . 94
6.10 Afterglow light curves at 3GHz for model F (left) with or without including the velocity profile, and for model b (right) at $t_{\rm beg}$ and towards $t_{\rm end}$, for different observing angles. . . . 95
6.11 On-axis, slightly off-axis and off-axis (left to right) afterglow light curves at 3GHz for model F, for different polar and azimuthal observing angles. . . . 97
6.12 Same as figure 6.11, for all models of the parameter study. Horizontal lines on the off-axis light curve mark the 3σ sensitivities of the VLA (dashed) and ngVLA (dotted) at 3 GHz for a fiducial 1 hr on-source integration (see Appendix C). . . . 98
6.13 Corner plots of MCMC fit results of model F to the GW170817 data, considering the non-axisymmetrised jet structure and different azimuthal viewing angles, i.e. $\phi_{\rm v} = 0$ (upper left), $\phi_{\rm v} = \pi/2$ (upper right), $\phi_{\rm v} = \pi$ (bottom). . . . 100
6.14 On-axis and off-axis ($\theta_{\rm v} = 32°$) light curves for all models of the parameter study, for the fixed set of afterglow parameters corresponding to the best-fit of model F. . . . 101

7.1 Slices of the logarithmic mass density from the Aenus-ALCAR snapshot used as initial condition, on the meridional and equatorial planes, at different levels of zoom. . . . 107
7.2 Left: Radial mass density profile along the polar axis. The flat inner plateau at $r \lesssim 10^6$ cm corresponds to the proto-neutron star, which is not part of the physical problem addressed here. Centre: cumulative mass profile M_r as a function of radius. Right: Cumulative angular momentum J_r as a function of radius. . . . 108
7.3 Two-dimensional slice in the y–z plane of the magnetic field intensity (left), and 3D rendering of magnetic field streamlines around the PNS (right). . . . 109
7.4 Diagnostics used to bracket the excision radius on the meridional plane. Left: sonic condition $|v_r|/c_s > 1$, tracing material likely to be accreted onto the central compact object. Centre: radial velocity v_r/c. Right: shock indicator overplotted on the pressure distribution. . . . 111
7.5 Meridional slices of the logarithmic rest-mass density for the jet-off control (top) and the jet run (bottom), at $t = 0$ s (left), $t = 5$ s (middle) and $t = 10$ s (right). . . . 119
7.6 Meridional slices of the radial velocity for the jet-off control (top) and the jet run (bottom), at $t = 0$ s (left), $t = 5$ s (middle) and $t = 10$ s (right). . . . 119
7.7 Meridional slices of the magnetic field intensity for the jet-off control (top) and the jet run (bottom), at $t = 0$ s (left), $t = 5$ s (middle) and $t = 10$ s (right). . . . 120

7.8 Radius of the steepest pressure feature at 0°, 45° and 70° from the jet axis, for the jet run and the jet-off control, at $t = 10$ s. In the jet-off run this feature is the supernova shock at all angles; in the jet run it is the jet-driven forward shock at 0° and the supernova shock at 45° and 70°. Vertical dashed lines mark the identified shock radii for each angle. 121

B.1 Meridional view of the energy density without rest-mass component at 375 ms (top) and 3 s (bottom) after jet launching, for model *e1* (left) and *e2* (right) at the two resolutions. 135

B.2 Afterglow light curves at 3 GHz for model *e1* and *e2* (blue and magenta respectively) at the original resolution (full line) and increased resolution (dashed line), on-axis and with $\theta_{\rm v} = 32°$, computed for the best-fit values of model *F* for both resolutions, for comparison. 136

C.1 Off-axis afterglow light curves at 3 GHz for model *F* (axisymmetrised jet structure), for a set of viewing angles $\theta_v = 15°, 20°, 30°$ (colours) and ISM densities $n = 10^{-3}, 10^{-2}, 10^{-1}$ cm^{-3} (line styles), illustrating the dependence of the early shell-associated peak on these parameters. Horizontal dashed and dotted grey lines mark the 3σ VLA and ngVLA sensitivities, respectively, for a fiducial 1 h integration at 3 GHz, obtained from the corresponding official exposure calculators (see text). Grey stars with error bars and downward-pointing arrows show, respectively, GW170817 radio detections and 3σ upper limits at frequencies between 1 and 10 GHz, rescaled to 3 GHz, as compiled by Makhathini et al. [111] from Alexander et al. [6] and Hallinan et al. [81]. 140

D.1 Afterglow light curves at 3 GHz for model *F*, on-axis (solid) and off-axis, (dashed), isolating the effect of the velocity profile (VP) and of lateral expansion (LE) against the common baseline of neither ("no VP, no LE", blue). Grey: VP only. Magenta: LE only. 142

List of Tables

4.1 Injection parameters of the seven jet models. L_j is the initial injected luminosity, τ_d the exponential decay timescale of the luminosity, Σ_j the fractional magnetic contribution to the luminosity, and t_j the jet launching time after merger. The rightmost column summarises the main qualitative outcome of each model relative to the fiducial case *F*, discussed in detail in Sect. 4.3.1–4.3.3. 62

4.2 Jet cap properties at $t - t_j = 3\,\mathrm{s}$ for all models (north side). $E_{\mathrm{iso,max}}$ is the maximum isotropic-equivalent energy; $\overline{\Gamma}_{\mathrm{max}}$ is the maximum radially averaged Lorentz factor; $E_{\mathrm{sum}}/E_{\mathrm{inj}}$ is the fraction of injected energy retained in the outflow; $E_{\mathrm{kin}}/E_{\mathrm{sum}}$ is the kinetic energy fraction; Θ_{50} is the effective half-opening angle containing 50% of E_{sum}. Data from P25. 71

5.1 Jet cap properties at $t_{\mathrm{end}} = 9.5\,\mathrm{s}$ for all models. The jet cap properties at t_{end} are: kinetic energy fraction $E_{\mathrm{kin}}/E_{\mathrm{sum}}$, effective half-opening angle Θ_{50}, and offset of the jet barycentre from the injection axis Θ_{off}. Data from D26. 79

6.1 Results of the MCMC fitting of each model to GW170817 afterglow data. All fits use the axisymmetrised jet structure at t_{end}. Values are the median and 90% credible interval of the marginalised posterior. Fixed parameters: $\varepsilon_e = 0.1$, $p = 2.166$, $z = 0.0096$, $d_L = 40\,\mathrm{Mpc}$. . . . 92

B.1 Jet cap properties of models *e1* and *e2* at the fiducial resolution (from Table 2 of P25, north side) and at increased resolution ($e1_{\mathrm{HR}}$, $e2_{\mathrm{HR}}$) at $t - t_{\mathrm{j}} = 3\,\mathrm{s}$. 134

Chapter 1

Introduction

1.1 The Origin of Gamma-Ray Bursts

Gamma-ray bursts (GRBs) are the Universe's most violent outbursts, releasing within seconds an amount of energy similar to what our Sun will produce over the course of its entire lifetime. For decades, it was not clear whether these events were occurring within our own Galaxy or at the edge of the observable Universe. The answer, once it came, transformed GRBs into one of the most powerful probes of extreme astrophysics available to us.

Today we know that GRBs are of cosmological origin, characterised by an initial flash of gamma rays, the *prompt emission*, followed by a fading, broadband *afterglow* detectable from radio to γ-ray wavelengths for days to years [176, 235]. The extreme luminosities and rapid variability of the prompt emission, combined with the non-thermal spectral shape, point unambiguously to a relativistic origin: the radiation is produced due to the presence of a collimated outflow, a *jet*, moving towards the observer at typical bulk Lorentz factors $\Gamma \sim 10^2$–10^3 [97, 107]. Such an outflow requires a central engine – an accreting black hole or a rapidly spinning, strongly magnetised neutron star – capable of converting gravitational or rotational energy into a collimated relativistic flow on timescales of seconds [24, 213].

The search for GRB progenitors has a long history. Long-duration events ($\gtrsim 2\,\mathrm{s}$) were linked to the deaths of massive stars [109, 222], confirmed by the simultaneous detection of GRB 980425 and the supernova event SN 1998bw [63]. Short-duration events ($\lesssim 2\,\mathrm{s}$) were long suspected to arise from compact binary coalescence [50, 146, 162], a connection established definitively in August 2017 with the joint detection of gravitational waves from the binary neutron star merger GW170817 and the short gamma-ray burst GRB 170817A, detected just 1.74 s later [2, 3, 71, 199].

The electromagnetic afterglow associated with GW170817 proved to be as scientifically informative as the initial detection of gravitational waves and the prompt emission. The post-burst emission rose steadily for more than a hundred days before reaching its peak, a behaviour explained by a structured jet viewed at a significant angle from its axis, progressively revealing its more energetic core as the material decelerated [68, 101, 137]. This event made clear that understanding GRBs requires not just identifying the progenitor, but tracing the entire physical chain: from the merger to the jet structure that forms as the outflow drills through the

post-merger environment, to the afterglow signal that encodes that structure for a distant observer.

1.2 The Challenge: Connecting Progenitors to Observables

The central challenge this thesis addresses is the construction of a physically consistent connection between three things: the *progenitor*, the *three-dimensional jet structure* that emerges from the progenitor vestige, and the *multiwavelength afterglow emission* that observers actually measure.

Each element of this chain is difficult to model on its own; this thesis contributes to connecting all three, focusing mainly on binary neutron star (BNS) mergers.

Before the jet can produce observable emission, it must first navigate the post-merger ejecta: $\sim 10^{-3}$–$10^{-2}\,M_\odot$ of neutron-rich matter expelled during the merger at 0.1–0.3 c, distributed anisotropically and non-uniformly [180, 203]. As the jet propagates through this medium, it drives shocks, inflates a hot cocoon, and exchanges energy and momentum with its surroundings in a way that fundamentally shapes its angular and radial structure [26, 101]. The jet that eventually emerges is not a simple, idealised cone but a complex, stratified, three-dimensional structure whose properties depend on both the injection conditions at the central engine and the detailed properties of the environment it had to cross.

Most jet simulations in the literature adopt idealised, homogeneous post-merger density profiles, which miss the asymmetries and density gradients that a realistic BNS post-merger environment imposes on the propagating jet. A key innovation of the research group within which this thesis was developed is the use of post-merger environments *directly imported* from BNS merger simulations [165, 166], providing a physically self-consistent description of the ambient medium without invoking artificial density profiles.

Once the jet structure is established, the second challenge is to translate it into observable predictions. Afterglow emission is produced as the ejecta sweep up the circumburst medium, generating a forward and reverse shock that lead to broadband synchrotron radiation [122, 195]. Capturing the signatures of structured jets ideally requires modelling both the angular energy distribution and the radial velocity stratification of the outflow, since the former governs the viewing-angle dependence of the light curve and the latter controls the energy injection history of the blast wave. Standard afterglow models typically assume simple parametric profiles and a single bulk Lorentz factor, therefore missing the signature of these features.

A further, critical obstacle connects these two challenges. The simulation output can only serve as a reliable input to the afterglow code if the jet is *ballistic* (or coasting): the regime in which the conversion of thermal and magnetic energy into bulk kinetic energy is essentially complete and the angular structure is no longer evolving. In two-dimensional simulations this condition can be satisfied [101, 218]. In three dimensions, the computational cost has until now prevented simulations from reaching this phase: most existing 3D studies are terminated shortly after jet breakout, when the outflow is still actively accelerating and its structure is still

changing [91, 148, 149]. Feeding such premature jet profiles into afterglow codes can introduce systematic errors that are particularly severe in the early-time rising phase of the light curve, the part most sensitive to the jet structure and most useful for constraining the viewing geometry.

The simulations presented here are primarily motivated by binary neutron star (BNS) environments and jet parameters characteristic of the classical short GRB population. In addition to this main set of simulations, we consider a core-collapse environment to examine how the results extend to a qualitatively different progenitor configuration. Across these environments, we study jet propagation through realistic ejecta, the development of three-dimensional angular and radial structure, and the resulting afterglow emission. While the parameter space explored is motivated primarily by short GRBs, the underlying framework is not intrinsically restricted to BNS systems or to short GRB parameters. Its application to different progenitor environments and jet properties provides a means of investigating more generally how the medium through which a relativistic jet propagates influences its subsequent structure, dynamics, and observable radiation.

1.3 Thesis Objectives and Structure

This thesis addresses the open problems described above through five interconnected objectives.

1. **Reaching a nearly ballistic regime in 3D.** We develop and validate a grid-remapping technique that transitions the simulation from a spherical to a Cartesian computational domain, enabling the jet evolution to be tracked up to $\approx 10\,\mathrm{s}$ post-injection without loss of resolution, long enough for the kinetic energy fraction to reach $\gtrsim 95\%$ and the angular structure to be frozen [45].

2. **Exploring jet injection conditions.** We analyse a systematic parameter study of jet injection conditions [167], featuring variations in the jet luminosity, injection decay timescale, magnetisation, and launching time, quantifying how each parameter governs the jet's ability to penetrate the ejecta and shape the final outflow geometry [46].

3. **Characterising three-dimensional jet structure.** We map the deviations from axisymmetry and the radial velocity stratification that develop as the jet interacts with the asymmetric BNS environment, identifying features that challenge idealised jet models [45, 46].

4. **Predicting multi-wavelength afterglow light curves.** Using a semi-analytical afterglow framework extended to handle the full 3D angular and radial distributions from our simulations, we compute synthetic light curves and compare them against the observational data of GW170817/GRB 170817A [46].

5. **Extending the framework to different progenitor environments.** We extend the simulation methodology developed for the BNS case to a core-collapse environment, applying the grid-remapping of the realistic progenitor

and the prescribed jet injection techniques and analysing the resulting three-dimensional structure of the outflow. This provides a basis for extending the broader framework to different progenitor environments, with the longer-term goal of determining how the properties of the surrounding ejecta shape jet propagation, outflow structure, and observable emission.

The remainder of the thesis is organised as follows. Chapter 2 provides the astrophysical background on GRBs: observational phenomenology, progenitor channels, central engine physics, jet propagation, and afterglow theory. Chapter 3 describes the theoretical and numerical framework, covering the RMHD equations and the PLUTO code. Chapter 4 presents the jet dynamics and energetics during propagation through the post-merger environment and after the jet breakout. Chapter 5 presents the grid-remapping strategy and the extension of the simulation to the ballistic regime. Chapter 6 reports the systematic parameter study and the resulting afterglow light curves. Chapter 7 presents an extension of the developed pipeline to the core-collapse progenitor. Chapter 8 summarises the main conclusions and outlines directions for future work.

In the following, we refer to Pavan et al. [166, 167] and Dreas et al. [45, 46] as P23, P25 , D25 and D26 respectively.

Chapter 2

Gamma-Ray Bursts

2.1 Observational Phenomenology

Gamma-ray bursts were discovered in 1967 by the Vela military satellites, which were monitoring for nuclear tests. The data, declassified and published only in 1973 [93], revealed brief but intense flashes of gamma-ray radiation with no obvious terrestrial or solar origin. In the decades that followed, two fundamental questions dominated the field: what produces GRBs, and how far away are they? The answer turned out to require the resolution of a deep theoretical problem.

2.1.1 Temporal variability and the compactness problem

The prompt emission light curves of GRBs exhibit rapid variability, with flux doubling and halving on timescales as short as $\delta t \sim 10^{-3}\,\mathrm{s}$. By causality, the size of the emitting region cannot exceed $D \lesssim c\,\delta t$, which for millisecond variability gives $D \lesssim 300\,\mathrm{km}$. At the same time, the observed gamma-ray fluences imply enormous luminosities. For a source at a cosmological distance d_L, the isotropic-equivalent luminosity is

$$L_{\mathrm{iso}} = 4\pi d_L^2 F_\gamma \,, \tag{2.1}$$

where F_γ is the observed energy flux. For the bursts where the distance is known, L_{iso} is estimated to lie within the range $\sim 10^{50}$–$10^{53}\,\mathrm{erg\,s^{-1}}$. A source with such a high luminosity confined to a region of ~ 300 km would be optically thick to photon–photon pair production, because the observed spectrum extends well above the 511 keV threshold for $\gamma\gamma \to e^+e^-$. In a static, compact source, these high-energy photons would interact efficiently with the abundant lower-energy photons, producing e^+e^- pairs and driving the radiation toward a nearly thermal spectrum. This is the essence of the *compactness problem*: the observed non-thermal emission is difficult to reconcile with a static, compact, luminous source.

The optical depth to pair production scales as [107]

$$\tau_{\gamma\gamma} \propto \Gamma^{-(4+2\alpha)} \,, \tag{2.2}$$

where α is the photon spectral index and Γ is the bulk Lorentz factor of the emitting plasma. Requiring $\tau_{\gamma\gamma} \lesssim 1$ for the observed high-energy photons to escape places

a lower limit $\Gamma \gtrsim 10^2$–10^3 , providing the most direct model-independent argument for the ultra-relativistic nature of GRB outflows [107, 176]. Relativistic bulk motion alleviates the compactness problem because relativistic aberration and Doppler shifting make the radiation field in the comoving frame less energetic and less favourable for pair production than inferred from the observer-frame spectrum, substantially reducing the pair-production optical depth.

In a relativistically expanding flow, the causal size of the emitting region is also enhanced:

$$R_{\rm em} \lesssim 2\Gamma^2 c\,\delta t\,. \tag{2.3}$$

2.1.2 Cosmological origin

For two decades after their discovery, the distance scale of GRBs was unknown. The BATSE instrument on the *Compton Gamma-Ray Observatory* constrained the Galactic versus cosmological debate by demonstrating that GRBs are distributed isotropically on the sky, with no correlation with the Galactic plane [119]. The question was settled only when *BeppoSAX* began providing arcminute localisations rapid enough to be followed up at longer wavelengths. The first X-ray afterglow, detected for GRB 970228 [37], established that bursts leave a fading counterpart and led to the identification of a faint host galaxy. Two months later, optical spectroscopy of the afterglow of GRB 970508 revealed metal absorption lines setting a lower limit $z \geq 0.835$ on its redshift [126], placing the source several Gpc away and unambiguously establishing the cosmological nature of GRBs.

2.1.3 Energetics and jet collimation

A canonical event can produce isotropic-equivalent radiated energy $E_{\rm iso} = 4\pi d_L^2 \mathcal{F} \sim 10^{52} - 10^{53}$ erg, at least one order of magnitude higher than the total energy typically released in a supernova, but emitted over seconds rather than weeks. However, the emission is not isotropic. A characteristic achromatic steepening (*jet break*) appears in the afterglow light curve at all wavelengths; this feature arises when the edge of the collimated outflow enters the observer's line of sight as it decelerates and its relativistic beaming cone widens [186, 197]. The jet break constrains the half-opening angle θ_j, reducing the true radiated energy by a factor $(1 - \cos\theta_j) \approx \theta_j^2/2$, to values $E_\gamma \sim 10^{49}$–10^{51} erg.

2.1.4 The prompt emission spectrum: the Band function

The energy spectrum of GRB prompt emission is a direct observable encoding the physics of the emitting plasma. Its most striking feature is its *non-thermal* shape. Unlike a blackbody, which rises steeply towards the peak and then declines exponentially, GRB spectra have a much flatter low-energy slope and are therefore substantially broader around the spectral peak. At higher energies, the emission can extend to several MeV and is often described by a power-law tail, although an exponential-like cutoff can also be present. Empirically, the time-integrated spectra of GRBs are well described by the *Band function* [16], an analytic functional form

that smoothly connects two power laws:

$$N(E) = \begin{cases} A \left(\dfrac{E}{100\,\mathrm{keV}} \right)^{\alpha} \exp\!\left(-\dfrac{E}{E_0} \right), & E \leq (\alpha - \beta)\, E_0\,, \\ A \left[\dfrac{(\alpha - \beta)\, E_0}{100\,\mathrm{keV}} \right]^{\alpha - \beta} \exp(\beta - \alpha) \left(\dfrac{E}{100\,\mathrm{keV}} \right)^{\beta}, & E > (\alpha - \beta)\, E_0\,. \end{cases} \tag{2.4}$$

Here $N(E)\,dE$ is the number of photons per unit area per unit time in the energy range $[E,\, E + dE]$. The function has four free parameters:

- A: the *normalisation* at the pivot energy of 100 keV, in units of photons $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}\,\mathrm{keV}^{-1}$. It sets the overall flux level and does not affect the spectral shape.

- α: the *low-energy photon index*, describing the spectral slope below the break. Typical values are $\alpha \sim -1$, meaning the photon count rises slowly with decreasing energy.

- β: the *high-energy photon index*, describing the steep power-law decay above the break. Typical values are $\beta \sim -2.3$, meaning $N(E) \propto E^{-2.3}$ at high energies. The spectrum therefore falls sharply, and most of the radiated energy is concentrated near the break.

- E_0: the *break energy*, which sets the transition scale between the two power-law regimes.

However, these spectral parameters are found to vary from burst to burst, therefore not presenting universal values. For example, the energy at the spectral break, E_0, can vary from 100 keV up to 1 Mev. Figure 2.1 shows an example of fitting the Band function to a GRB spectrum, specifically GRB911127 [16].

The most physically meaningful quantity is the peak of the νF_ν (equivalently, $E^2 N(E)$) spectrum, where most of the radiated energy is emitted. This peak energy is related to the break energy by

$$E_{\mathrm{peak}} = (2 + \alpha)\, E_0\,, \tag{2.5}$$

and is directly measurable from the observed spectrum. Typical values are $E_{\mathrm{peak}} \sim$ 100–300 keV for collapsar-driven GRBs, while merger-driven events[1] tend to peak at higher energies, $E_{\mathrm{peak}} \sim$ 400–800 keV, consistent with their harder spectral shape [67, 150].

Both power-law indices are inconsistent with a thermal spectrum, confirming the non-thermal origin of the emission. The precise mechanism responsible – whether internal shocks dissipating kinetic energy, magnetic reconnection in a Poynting-dominated flow, or photospheric emission from below the jet photosphere – remains an open question [97].

[1] See Sections 2.1.5, 2.4.2 for a discussion on GRB classification.

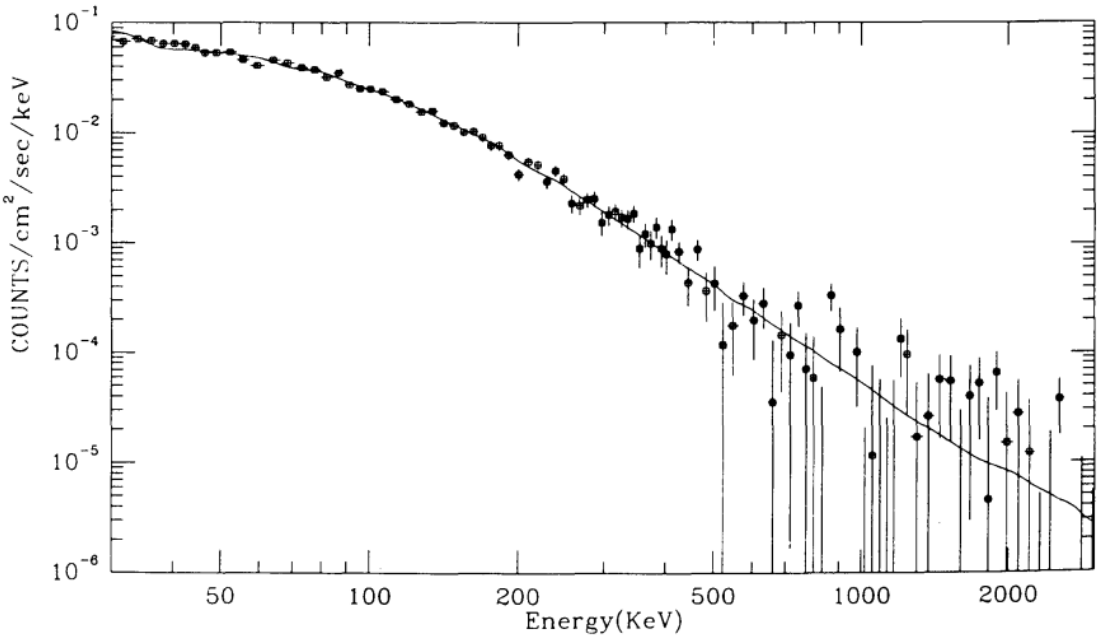


Figure 2.1: Band function fit to the average spectrum of GRB911127, adapted from Band et al. [16].

2.1.5 Duration classes and the T_{90} bimodality

The duration of a GRB is conventionally quantified by T_{90}: the time interval spanning the 5th to 95th percentile of the cumulative detector counts. In the BATSE catalogue, the distribution of T_{90} is broadly bimodal, with a conventional boundary at $T_{90} \approx 2\,\mathrm{s}$ [95]. Events shorter than this threshold are classified as *short GRBs*, while longer ones are *long GRBs*.

This bimodality was historically the first indication that two physically distinct populations might be part of the GRB class. Support for this interpretation came from three additional directions. First, the two populations differ in spectral hardness: short GRBs have on average higher E_{peak} and harder low-energy indices than long GRBs [95, 150]. Second, they occupy distinct regions in the *Amati relation* [11], an empirical correlation between peak photon energy and isotropic equivalent radiated energy,

$$E_{\mathrm{peak}} \propto E_{\mathrm{iso}}^{1/2}\,, \tag{2.6}$$

with short GRBs lying systematically above the long GRB track (higher E_{peak} for the same E_{iso}). Third, as discussed in Section 2.1.7, the two classes reside in markedly different host galaxy environments, pointing to young massive stars for long GRBs and old compact binary systems for short GRBs.

Taken together, these observations led the community to associate long GRBs with the core-collapse of a massive star and short GRBs with compact binary mergers. However, it has become increasingly clear that a clean one-to-one correspondence between duration and progenitor does not hold in general. The duration of the prompt emission depends not only on the progenitor type but also on detector sensitivity, redshift, and on the details of the central engine activity. Several observed

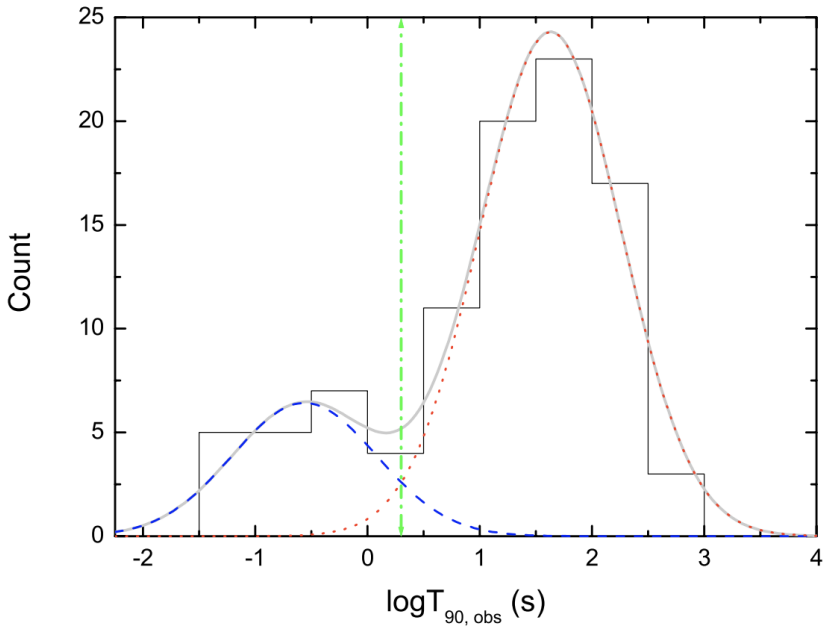


Figure 2.2: Distribution of durations for 95 GRBs from the Swift satellite. Adapted from Zhang and Choi [239].

events have provided compelling evidence for a more complex scenario, exhibiting long durations combined with kilonova emission signatures unambiguously associated with neutron-rich post-merger ejecta. These cases, and the physical mechanisms that can explain them, are discussed in detail in Section 2.4.2 after the central engine physics has been introduced. Figure 2.2 shows a visual representation of this bimodality.

Throughout this thesis, we therefore adopt the following terminology. We use *short* and *long* when referring to observational studies and population statistics based on the measured T_{90}, since in that context events are drawn from catalogues using duration as the selection criterion. When the distinction is based on the physical nature of the progenitor system, we refer instead to *merger-driven* and *collapsar-driven* GRBs. This work focuses mainly on merger-driven jets, which correspond to the short GRB population in the classical picture but whose observable properties may extend to longer durations in the cases discussed in Section 2.4.2.

2.1.6 Prompt emission morphology

Beyond the average spectral and temporal properties, the prompt light curves of individual GRBs are highly diverse: no two events are identical. The most commonly observed profile is the fast-rising, exponentially decaying (FRED) shape: a rapid rise followed by a slower exponential decline, sometimes repeated over multiple pulses. Individual pulse timescales are correlated with the overall burst duration [153] and range from milliseconds to tens of seconds.

Several additional features are observed in a fraction of events. *Precursors* are early, often softer emission episodes preceding the main burst by tens to hundreds of

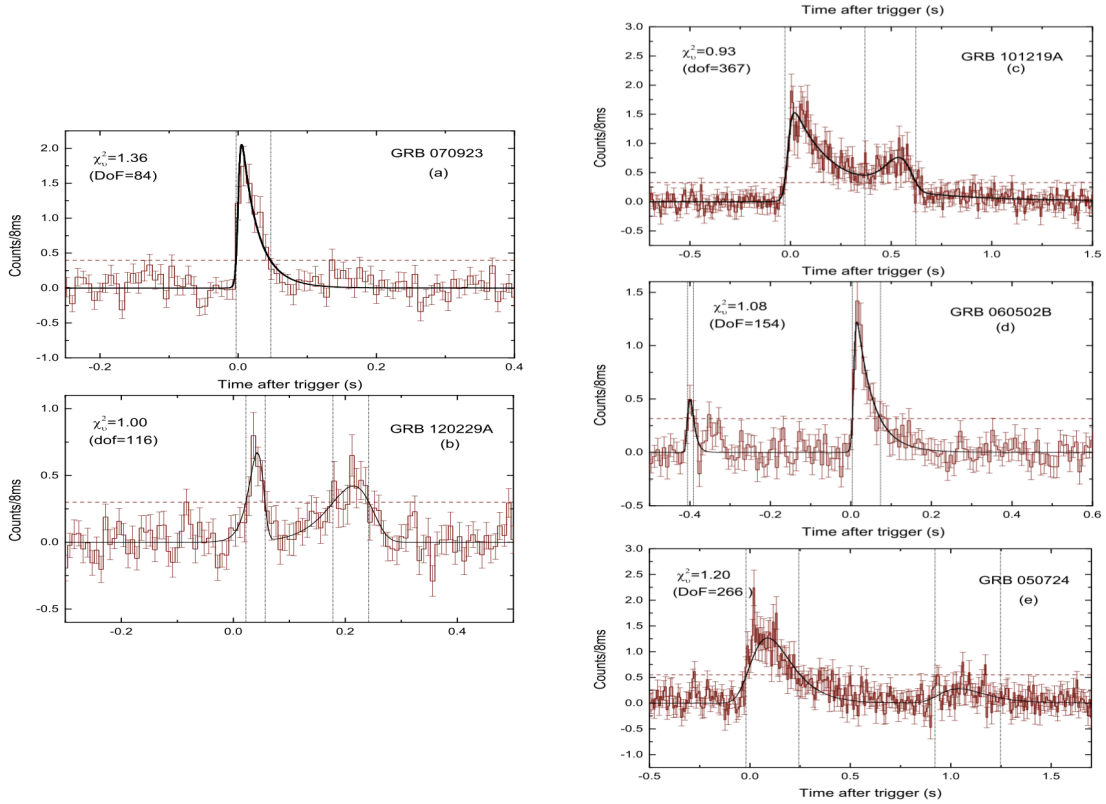


Figure 2.3: Light curves (15–350 KeV) of five typical sGRBs from the Swift/BAT catalogue, adapted from Li et al. [106].

seconds; their origin is debated, but they may trace the first interaction of the jet with the surrounding material. *Extended emission*, lasting up to $\sim 100\,\mathrm{s}$ beyond the main episode, is detected in some short GRBs and may point to prolonged central engine activity [152]. *X-ray flares* and a shallow *plateau* phase are commonly seen in *Swift*/XRT observations during the early afterglow; the latter is often interpreted as energy injection from a long-lived central engine, possibly a spinning-down magnetar [154, 237]. For reference, Figure 2.3 shows a set of GRB observations from the Swift/BAT instrument, as reported in Li et al. [106]. This diversity reflects the dependence of the emission on the central engine and on the environment crossed by the jet, which this thesis addresses, as well as on extrinsic factors such as the viewing angle.

2.1.7 Host galaxies

Long GRBs are preferentially found in actively star-forming, low-metallicity galaxies, specifically near the brightest star-forming regions [60, 104]. Short GRBs occur across all galaxy types, including passive ellipticals, and show significant spatial offsets from their host centres [21, 55]. These environmental differences are one of the primary observational bases for assigning distinct progenitors to the two duration classes.

2.2 The Progenitors

2.2.1 GRBs and Collapsing Massive Stars

The *collapsar* model, originally proposed by Woosley [222] and explored through hydrodynamic simulations by MacFadyen and Woosley [109], provides the leading explanation for the origin of long-duration GRBs. The essential picture is the following: a rapidly rotating massive star undergoes core collapse, and some of the infalling material forms an accretion disc around the central remnant. The gravitational energy released in this process can then power a relativistic jet along the rotation axis, which must successfully penetrate the stellar envelope before producing observable gamma-ray emission, a process described in the next Sections.

Wolf–Rayet progenitors

Wolf–Rayet (WR) stars are massive stars (typically $M \gtrsim 20$–$25\,M_\odot$) that have shed most of their hydrogen envelope through strong stellar winds or binary mass transfer, leaving behind a compact helium envelope. They are natural GRB jet progenitors for several reasons: they are likely to form a black hole at the end of their evolution, their compactness with respect to typical massive stars provides a relatively short path through the stellar material for a jet to successfully reach the surface, and their lack of extended hydrogen envelopes is consistent with the observational association of long GRBs with Type Ic supernovae [31, 84, 205, 223].

A persistent tension in this picture concerns angular momentum. The collapsar requires a rapidly rotating envelope to form a centrifugally supported disc. In a general relativistic framework, however, the boundary between disc formation and a direct radial plunge is set by the specific angular momentum at the innermost stable circular orbit (ISCO) of the newly born black hole [97]. This dynamical threshold is highly sensitive to the remnant's spin; prograde rotation contracts the ISCO, reducing the centrifugal barrier required for disc support. At solar metallicity, strong stellar winds that strip the progenitor's outer envelope carry away angular momentum through magnetic torques [224], therefore reducing the specific angular momentum of the helium core well below this spin-dependent dynamical threshold, preventing disc formation and forcing a direct, free-fall plunge into the event horizon. A resolution may lie in *low-metallicity* environments, where weaker winds preserve both the stripped envelope and sufficient rotation, consistent with the observed preference of long GRBs for low-metallicity host galaxies [224, 231]. Binary interactions offer an alternative channel: mass transfer or tidal spin-up from a companion star can endow the pre-collapse core with sufficient angular momentum even at higher metallicities [223].

Core collapse and remnant formation

When a massive star (typically $M \gtrsim 8\,M_\odot$) reaches the final stages of nuclear burning, it develops an iron-group core whose mass progressively increases as material processed by the surrounding burning shells is incorporated into the core. The core is initially supported against gravity primarily by electron degeneracy pressure and remains

approximately hydrostatic, while neutrino cooling and electron captures drive its gradual contraction. Electron captures on protons and heavy nuclei,

$$e^- + p \rightarrow n + \nu_e \tag{2.7}$$

progressively reduce the electron fraction $Y_e \equiv n_e/n_b$ where n_e is the electron number density and n_b is the baryon number density, and therefore progressively neutronise the matter. As Y_e decreases, the electron degeneracy pressure is reduced, while the photodisintegration of heavy nuclei,

$$(A, Z) + \gamma \rightarrow (A - 4, Z - 2) + {}^4\mathrm{He}\,, \tag{2.8}$$

absorbs thermal energy and further weakens the pressure support. The combined effects eventually render the core dynamically unstable, triggering gravitational collapse on a timescale of tens of milliseconds. As the central density approaches nuclear saturation density ($\rho \sim 10^{14}$–$10^{15}\,\mathrm{g\,cm^{-3}}$), the matter becomes increasingly neutron-rich and highly degenerate. At nuclear densities, the rapidly increasing pressure of neutron-rich matter, together with the repulsive component of the nuclear interaction, causes the equation of state to stiffen and decelerates the collapse of the inner core, leading to hydrodynamic bounce and the formation of an outward-propagating shock. The shock rapidly loses energy through the photodisintegration of heavy nuclei and neutrino emission, while the continued accretion of the outer core exerts a substantial ram pressure that opposes its outward propagation, eventually stalling the shock. Its eventual revival depends on the efficiency of neutrino heating by the intense neutrino flux emitted from the newly formed proto-neutron star. If neutrino heating becomes sufficiently efficient to overcome the energy losses and the accretion ram pressure, the shock can be revived and drive the ejection of the stellar envelope, producing a core-collapse supernova and leaving behind a compact remnant. If the mass of the central remnant exceeds the maximum stable neutron star mass – the Tolman–Oppenheimer–Volkoff limit, $M_{\mathrm{TOV}} \approx 2.2$–$2.4\,M_\odot$ for non-rotating stars [113] – a black hole forms. In the case that the progenitor star carries sufficient angular momentum, the infalling material does not fall directly into the black hole but instead forms a centrifugally supported accretion disc, providing the conditions required for jet launching.

Collapsar pathways

The transition from a collapsing core to an active central engine depends on how the stellar mantle falls onto the newly formed black hole. Three regimes are usually distinguished, depending on the progenitor and on the core-collapse shock.

In the prompt collapse scenario (historically referred to as a Type I collapsar), the iron core collapses directly into a black hole without launching a successful shock [109]. The highly rotating inner helium core accretes uninterrupted at strong rates of $\dot{M} \sim 0.01$–$1\,\mathrm{M_\odot\,s^{-1}}$ [109].

In the fallback-dominated scenario (historically referred to as a Type II collapsar), the core bounce launches a weak, sub-energetic shock that is insufficient to completely expel the outer stellar layers. A transient neutron star is briefly formed, and a small fraction of the envelope is ejected, but a large mass of stellar material (0.1–$5\,\mathrm{M_\odot}$)

fails to reach escape velocity and falls back onto the central remnant over minutes to hours [110]. Once this accumulated fallback mass exceeds the maximum stable limit, it triggers a delayed collapse into a black hole and establishes a fallback-driven accretion disc. Because this fallback material originates from larger radii and returns over a wider timescale, the accretion rates are moderately lower, typically spanning $\dot{M} \sim 10^{-3}$–$10^{-2}\,\mathrm{M}_\odot\,\mathrm{s}^{-1}$ [110].

A third, late-time accretion pathway occurs in massive, rotating supergiants where the inner core collapses directly to a black hole, but the outer hydrogen envelope contains substantial angular momentum [225]. As this envelope material slowly falls inward on free-fall timescales of days or weeks, it encounters a centrifugal barrier at very large radii ($r \gtrsim 10^{11}$ cm), forming a long-lived, low-density accretion disc with lower feeding rates ($\dot{M} \lesssim 10^{-4}\,\mathrm{M}_\odot\,\mathrm{s}^{-1}$) [225]. This mechanism can sustain a long-term relativistic jet, providing a potential explanation for ultra-long-duration high-energy transients.

Magneto-rotational supernovae

An alternative, physically distinct pathway for the terminal evolution of a rotating massive star is a *magneto-rotational supernova*. Originally proposed analytically by LeBlanc and Wilson [102] and Bisnovatyi-Kogan [22], this mechanism describes a core-collapse explosion where the decisive energy source is the rotational energy of the collapsing stellar core, extracted and redistributed by magnetic fields rather than neutrino heating alone [29].

The physical basis of this process lies in the reservoir of rotational kinetic energy stored in the rapidly spinning core after collapse. For a central core spun up to millisecond periods, this rotational reservoir can be estimated as:

$$E_{\mathrm{rot}} = \frac{1}{2}\, I\, \Omega^2 \approx 2 \times 10^{52} \left(\frac{P_0}{1\,\mathrm{ms}}\right)^{-2} \left(\frac{I}{10^{45}\,\mathrm{g\,cm^2}}\right) \mathrm{erg}\,, \tag{2.9}$$

where I is the core's moment of inertia and $\Omega = 2\pi/P_0$ is the angular velocity.

During collapse, differential rotation winds the seed poloidal magnetic field lines into a dominant toroidal configuration, while the magnetorotational instability (MRI) and convective dynamos exponentially amplify both field components on millisecond timescales, far exceeding the growth rate of simple winding [15, 141].

Because the magnetic pressure and hoop stresses are concentrated along the rotation axis, this process can produce highly aspherical, polar, jet-like outflows that can easily disrupt the star even if standard neutrino heating is inefficient [29, 200]. While the collapse can leave behind a stable, highly magnetised neutron star, continued accretion of infalling material can drive the central core past the maximum stable mass limit, collapsing it into a spinning black hole shortly after the explosion is launched. In such cases, the magnetorotational supernova represents the initial, explosive phase of a system that eventually transitions into an accreting black hole engine.

2.2.2 GRBs and Binary Mergers

Already in the 1980s–90s, the merger of compact binary systems containing at least one neutron star was identified as the most promising progenitor channel for GRBs, through fundamental theoretical contributions by Paczyński [162], Eichler et al. [50], and Narayan et al. [146].

Several observational clues support this picture. Most short duration GRBs (sGRBs) are found in host galaxies with little or no ongoing star formation, which is difficult to reconcile with a massive-star origin but is naturally explained by the long inspiral timescales ($\gtrsim$ Gyr) of compact binary systems. These systems of massive stars evolve through the formation of two neutron stars, and only merge after losing orbital energy to gravitational-wave emission. Since the coalescence time depends strongly on the orbital separation, even a small spread in separation can result in a broad range of merger timescales, extending up to ~ 1 Gyr [21]. Moreover, sGRBs typically show a significant spatial offset from the centre of their host galaxy, suggesting that the progenitor system migrated from its birth site before the merger occurred, a feature that can be naturally attributed to the natal kicks that a binary system receives during each supernova event [55].

From two massive stars to a bound compact binary

A neutron star is the typical outcome of the core collapse of a star with initial mass above $\approx 8\,M_\odot$, while black holes become increasingly common at higher masses (Section 2.2.1). The correspondence between initial mass and remnant type is, however, not a simple threshold. Whether a star explodes and which remnant it leaves depend on the structure of its core at collapse, so that neutron-star- and black-hole-forming progenitors can alternate over a wide mass range [159, 207]. Metallicity, through wind mass loss, and binary interaction, through envelope stripping and mass transfer, further alter the core structure and therefore the outcome [98, 201]. For a pair of neutron stars to end up gravitationally bound to each other, both stars must therefore follow evolutionary paths that end in neutron-star formation, and the binary must survive not one but two supernova explosions.

The dominant formation channel for field BNS systems is the evolution of an isolated massive binary [177, 211].

The more massive of the two stars evolves first, exhausts its fuel, and explodes as a supernova, leaving a neutron star. As the companion subsequently evolves and expands, it can overflow its Roche lobe – the equipotential surface within which material remains gravitationally bound to the star – transferring mass onto the neutron star, or it may engulf the compact remnant entirely in a shared envelope of gas. This *common envelope* phase is critical: as the two objects orbit within the shared envelope, dynamical friction rapidly shrinks their separation, ejecting the envelope and leaving a tight binary consisting of the neutron star and the stripped helium core of the companion. If this second star also avoids direct collapse to a black hole and instead explodes as a (typically hydrogen-poor) supernova, the result is a binary neutron star system.

This evolutionary path imposes stringent requirements. The system must survive each supernova explosion without being disrupted, which is only possible if the

explosion ejects no more than approximately half of the system's total mass and does not impart a kick velocity large enough to unbind the orbit [177]. Because of these constraints, only a fraction of massive binaries survive to become bound double neutron star systems: a fact directly relevant for predicting merger rates and, by extension, the expected rate of merger-driven GRBs. An alternative formation channel operates in dense stellar environments such as globular clusters, where neutron stars formed independently can later become gravitationally bound through close dynamical encounters [211].

Inspiral: the long road to coalescence

Once a bound BNS system has formed, general relativity predicts that its time-varying mass quadrupole moment due to the orbital motion of the two neutron stars causes it to emit gravitational waves. These waves carry away energy and angular momentum, driving the binary to lose orbital energy and start the *inspiral* phase.

For a circular orbit, the time remaining until coalescence can be estimated analytically as [175]

$$\tau_{\rm insp} = \frac{5\,c^5\,a_0^4}{256\,G^3\,M^2\mu}\,, \tag{2.10}$$

where a_0 is the initial orbital separation, $M = m_1 + m_2$ is the total mass of the system, and $\mu = m_1 m_2/M$ is the reduced mass. This timescale is extraordinarily sensitive to separation: a system with an orbital period of a few hours and negligible eccentricity may take of order 10^9 years to merge, while the same system with high eccentricity can coalesce in 10^6 years or less. Because gravitational wave emission also efficiently circularises eccentric orbits well before coalescence, by the time a BNS system reaches the final stages of inspiral its orbit is essentially circular regardless of its initial eccentricity [189].

This inspiral phase dominates the lifetime of the binary system: from formation to merger, BNS systems typically spend the majority of their existence in this slow, gravitational-wave driven spiral, with the dynamically violent merger itself lasting tens to hundreds of milliseconds, depending on the definition of the merger phase and the properties of the remnant. This vast separation of timescales is the reason that the gravitational-wave "chirp" signal, the characteristic rise in frequency and amplitude as the stars approach coalescence, can be used as a clock and distance indicator, and is the channel through which BNS mergers were first directly detected [2].

Merger and the immediate remnant

The coalescence itself is conventionally divided into three phases: *inspiral*, *merger*, and *ringdown* [17]. While the inspiral can be modelled accurately with post-Newtonian or effective-one-body approximations, the merger phase is intrinsically non-linear and strong-field, and can only be captured through numerical relativity simulations solving the full Einstein field equations coupled to relativistic hydrodynamics.

Most observed and theoretically expected BNS systems have nearly equal component masses, with total mass typically in the range 2.5–2.9 $M_\odot$ [161]. For such near-equal-mass systems, the two stars do not undergo significant tidal disruption

until immediately before contact, and the majority of the combined mass ends up bound in the post-merger remnant rather than being ejected. The nature of this remnant depends sensitively on the total mass of the system relative to a critical threshold $M_{\rm crit} \approx 2.6$–$3.4\,M_\odot$, set by the neutron star equation of state (EoS) [18].

The outcome of a BNS merger depends on the mass of the remnant relative to different EoS-dependent characteristic masses. The maximum mass of a non-rotating neutron star is given by the Tolman–Oppenheimer–Volkoff limit, $M_{\rm TOV}$, while uniform rotation increases the maximum supported mass to $M_{\rm rot,max}$, typically by $\sim$ 20%–30%. For a remnant that does not undergo prompt collapse, differential rotation immediately after merger provides additional support against gravity, allowing the remnant to temporarily exceed the uniformly rotating maximum mass. Angular-momentum transport progressively damps the differential rotation on a timescale of order ~ 100 ms, driving the remnant towards uniform rotation. If its mass exceeds $M_{\rm rot,max}$, the remnant cannot be supported by uniform rotation and undergoes gravitational collapse to a black hole, corresponding to the *hypermassive* neutron-star scenario. If instead its mass satisfies $M_{\rm TOV} < M < M_{\rm rot,max}$, the remnant can reach a uniformly rotating configuration without immediate collapse, forming a *supramassive* neutron star. The subsequent loss of angular momentum through electromagnetic spin-down eventually removes the rotational support, causing collapse on the corresponding spin-down timescale. Finally, if the remnant mass is below $M_{\rm TOV}$, it can evolve towards a stable neutron-star configuration. At sufficiently high total binary masses, however, the merger can undergo *prompt collapse* before a long-lived neutron-star remnant is established. These thresholds depend on the neutron-star EoS, and the remnant mass is determined by the detailed dynamics of the merger and is therefore not simply equal to the sum of the gravitational masses of the two initial stars. Consequently, the boundaries between these different outcomes are not described by single values of the component masses, but instead define transition curves in the progenitor mass parameter space.

The coalescence is also a violently turbulent event. At the moment of contact, the shearing interface between the two stellar surfaces is unstable to the Kelvin–Helmholtz instability, which efficiently amplifies the magnetic field on millisecond timescales [92, 178]. This rapid amplification is a key ingredient for the eventual production of a large-scale, organised magnetic field capable of launching a relativistic jet, as discussed in Section 2.3.4.

The post-merger environment

Not all of the binary's mass is retained by the central object resulting from the merging process. A fraction of the stellar material is expelled during and after the coalescence, establishing a highly structured and anisotropic circumburst environment. The latter constitutes the physical medium that the newly launched relativistic jet must subsequently propagate through and interact with. The mass-outflow processes are classified into several physically distinct channels operating on different timescales and characteised by different geometric, thermodynamic, and nucleosynthetic properties. A schematic representation of the various components described below is shown in Figure 2.4.

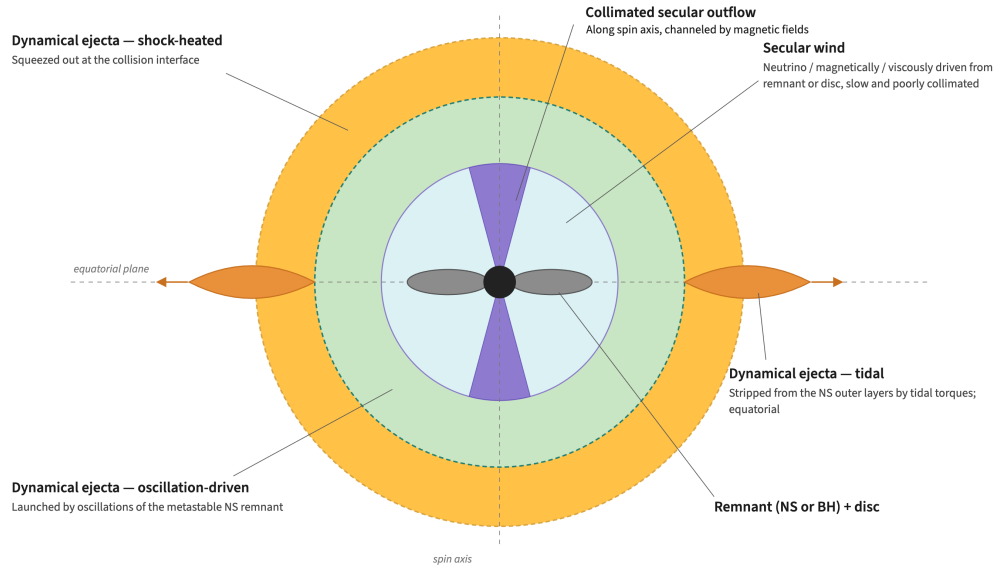


Figure 2.4: Schematic meridional cross-section (not to scale) of the post-merger environment of a BNS coalescence: shock-heated dynamical ejecta (outer shell) with neutron-rich tidal ejecta along the equatorial plane, oscillation-driven dynamical ejecta (inner shell), and secular outflows (innermost region) with a slow isotropic component and a faster collimated component along the spin axis. The central remnant (metastable NS or BH) is surrounded by an accretion disc.

Dynamical ejecta The dynamical ejecta are unbound material launched during the final orbits of the inspiral and the merger, on timescales of order $\sim 1\,\mathrm{ms}$ [180, 203]. This component carries a total mass of 10^{-4}–$10^{-2}\,\mathrm{M}_\odot$ at velocities spanning 0.1–0.3 c and is driven by two distinct mechanisms [198]: The tidal ejecta, driven by gravitational tidal torques that strip material from the outer layers of the neutron stars. This process is highly efficient in asymmetric mass-ratio mergers, where the tidally disrupted lower-mass companion forms long tidal tails. Tidal ejecta is highly concentrated in the equatorial orbital plane and remains extremely neutron-rich (electron fraction $Y_e \lesssim 0.15$–0.2), making it the primary site for robust rapid neutron-capture (r-process) nucleosynthesis of heavy lanthanides and actinides [180, 203]. The shock-heated ejecta instead is squeezed out from the contact interface as the two stellar cores collide. The intense shock heating increases the local temperature and triggers the emission of a high neutrino flux, which irradiates the outgoing material. This neutrino irradiation converts neutrons to protons via electron neutrino absorption, raising the electron fraction to higher values ($Y_e \gtrsim 0.25$–0.4). Shock-heated ejecta is distributed over a broader angular range, showing a characteristic polar profile.

Oscillation-driven outflows If the merger avoids prompt collapse and leaves behind a meta-stable (hypermassive or supramassive) neutron star remnant, the system undergoes strong, high-amplitude radial and rotational oscillations that can alter the post-merger mass-loss dynamics. The rotational oscillations are dominated

by non-axisymmetric, quadrupole-like deformations (such as the f-mode or bar-mode instabilities) that rotate rapidly within the orbital plane [203]. This time-varying, non-axisymmetric gravitational potential exerts torques on the disc, exciting strong spiral density waves. These waves efficiently transport angular momentum outward, expanding the disc and driving a massive, relatively fast equatorial outflow that can unbind $\sim 10^{-2}\,\mathrm{M}_\odot$ of matter on timescales of tens to hundreds of milliseconds. Simultaneously, the violent, quasi-radial oscillations of the remnant periodically compress and expand the stellar surface. This high-frequency radial pulsation acts as an acoustic piston, driving strong hydrodynamic shock waves that deposit thermal energy into the dense torus, accelerating the material outward as a hot, oscillation-driven wind. Because these oscillation-driven components are ejected while the remnant is actively cooling via neutrino emission, they are subjected to intense, continuous neutrino irradiation [198].

Secular post-merger outflows On longer timescales of $\sim 0.1 - 1\,\mathrm{s}$ and beyond, the system evolves via secular mass-loss channels driven by the thermal, magnetic, and viscous evolution of the central massive neutron star or, in case a black hole is formed, of the surrounding accretion disc (e.g. Fernández and Metzger [53]). The neutrino-driven (e.g. Dessart et al. [44]) and magnetically-driven outflows (e.g. Ciolfi [35] and refs. therein) characterising this phase of the evolution can expel much more material than what is dynamically ejected during and shortly after the merger. Nonetheless, only part of this material is unbound and will effectively contribute to mass ejection. Moreover, a significant (if not dominant) portion of these outflows is slowly expanding ($\sim 0.1\,c$) and poorly collimated, but a faster ($\sim 0.2 - 0.3\,c$) and rather collimated component can emerge along the spin axis of the differentially rotating massive neutron star remnant (before its eventual collapse, if any), channeled by the growing magnetic field. In merging BNS systems where black hole formation occurs no earlier than a few tens of ms after merger, secular outflows represent the main obstacle that a later incipient jet has to drill through and, as such, play a dominant role in shaping the angular structure and energetics of the final escaping jet.

Multi-messenger signatures: gravitational waves and the kilonova

The multi-messenger nature of BNS mergers makes them rich astrophysical laboratories: the same event can be observed simultaneously through gravitational waves and across the electromagnetic spectrum, with each channel probing a different underlying physics.

Gravitational waves are emitted throughout the inspiral, merger, and ringdown, with the inspiral phase producing the well-known *chirp* signal: a wave of increasing frequency and amplitude as the stars approach coalescence [2]. The post-merger gravitational wave signal, dominated by the oscillation modes of the (hyper- or supra-massive) remnant if one forms, carries direct information on the neutron star EoS, although it has not yet been constrained due to its high characteristic frequency and intrinsically weaker amplitude [203].

The neutron-rich material ejected during and after the merger provides the ideal

site for rapid neutron capture nucleosynthesis (the *r-process*), synthesising heavy elements up to and beyond the lanthanides [99, 124]. The radioactive decay of these freshly synthesised, unstable isotopes powers a thermal transient known as the *kilonova*: a quasi-isotropic optical and infrared glow that peaks days after the merger and fades over roughly a week to ten days [105, 124].

Because the kilonova luminosity and colour evolution depend on the ejecta mass, velocity, and lanthanide content, kilonova observations provide a direct, independent probe of the ejecta properties.

2.3 The Central Engine

Regardless of whether the progenitor is a collapsar or a compact binary merger, a necessary condition for the production of a GRB is the formation of a *central engine* capable of sustaining an outflow that accelerates to the high bulk Lorentz factors required to power the observed prompt emission.

2.3.1 The Blandford–Znajek mechanism

The Blandford–Znajek (BZ) mechanism [24] is the most widely invoked process for relativistic jet launching, and is considered responsible for a broad variety of astrophysical jets, including those from active galactic nuclei and microquasars. In this process, the magneto-rotational energy of a spinning (Kerr) black hole is extracted in the presence of a large-scale magnetic field threading the event horizon. Due to frame-dragging effects produced by the rotating spacetime, the magnetic field lines are twisted, generating a toroidal component from an originally predominantly poloidal configuration. The resulting magnetic pressure gradient, combined with the presence of a low-density funnel along the black hole rotation axis, drives material outward along the polar direction, powering a Poynting-flux-dominated jet confined by the hoop stress of the toroidal field.

The power extracted by the BZ mechanism can be written as [24, 212]

$$P_{\rm BZ} \simeq \frac{\kappa}{4\pi c}\,\Phi_{\rm BH}^2\,\Omega_H^2\,f(\Omega_H)\,, \tag{2.11}$$

where $\kappa \approx 0.044$ is a numerical factor that depends on the magnetic field geometry at the horizon, $\Phi_{\rm BH}$ is the total magnetic flux threading one hemisphere of the black hole horizon, $\Omega_H = ac/(2r_H)$ is the angular velocity of the horizon (with a the dimensionless spin parameter and r_H the horizon radius), and $f(\Omega_H)$ is a function that encapsulates higher-order spin corrections and approaches unity in the slow-rotation limit. The power scales strongly with both the black hole spin and the accumulated magnetic flux: a faster-spinning black hole and a stronger magnetic field both contribute to a more powerful jet.

2.3.2 Accretion regimes: SANE and MAD states

The efficiency of jet launching depends on the magnetic field configuration near the event horizon, and three-dimensional general-relativistic magnetohydrodynamic

(GRMHD) simulations distinguish two limiting accretion states: standard and normal evolution (SANE) and magnetically arrested disc (MAD) [172, 213]. In weakly magnetised SANE discs, the field is mostly disordered, angular momentum transport is driven by MRI turbulence, and jet efficiencies, defined as $\eta_{\rm jet} = P_{\rm jet}/(\dot{M}c^2)$, are comparatively low ($\eta_{\rm jet} \sim 0.1$–0.2) [213]. When large-scale poloidal flux is advected onto the black hole, the accumulated magnetic pressure becomes comparable to the ram pressure of the infalling gas and the flow enters the MAD state [147, 213]. Accretion does not stop in this regime, but proceeds in an intermittent and strongly magnetised form, with gas reaching the horizon through magnetic interchange instabilities and episodic flux eruptions [118]; in the GRB context, this intermittency has been linked to the variability of the resulting jets [85, 86]. In MAD states the jet efficiency can exceed unity [213]. For GRB engines, efficiencies of this order combined with the high accretion rates expected would imply jet energies larger than those typically inferred, and a rapid spin-down of the black hole [108] that may be difficult to reconcile with the high spins and small accreted masses required. To address this, Wu et al. [226] proposed a unified framework modelling the transition between the standard thin disc and MAD regimes as a function of the magnetic flux ratio, demonstrating that the maximum jet energy peaks at an intermediate, sub-saturated magnetisation state, where a moderate black hole spin and magnetic flux combine to power observed GRB jets while avoiding the rapid spin-down and severe accretion constraints of a fully arrested disc.

2.3.3 Neutrino annihilation

Another mechanism capable of converting accretion energy into an outflow is neutrino–antineutrino annihilation ($\nu\bar{\nu} \rightarrow e^+e^-$) above the poles of the remnant [50]. In this process, the intense neutrino flux emitted by the hot accretion disc partially annihilates along the rotation axis, depositing thermal energy and driving a thermally-powered outflow. The annihilation luminosity scales approximately as $L_{\nu\bar{\nu}} \propto \dot{M}^{9/4}$ [233], but for the same black hole properties the resulting luminosity is typically orders of magnitude lower than the BZ power [89]. The two mechanisms can coexist: neutrino annihilation is expected to contribute predominantly in the earliest accretion phase, when $\dot{M}$ is largest, but it is insufficient to sustain GRB jets on its own, and the BZ mechanism remains the dominant engine throughout most of the activity timescale.

2.3.4 Nature of the remnant, implications for jet launching

The nature of the compact remnant formed after the collapse depends on the total mass of the system and on the neutron star equation of state, a relation between pressure, density, and temperature inside dense nuclear matter that remains highly uncertain.[2] In both the collapsar and binary merger scenarios, the outcome may

[2]The EoS determines the maximum mass a neutron star can support against gravitational collapse. A "stiffer" EoS (higher pressure at a given density) allows heavier and larger neutron stars, increasing the threshold mass for black hole formation. Current constraints from gravitational wave observations and neutron star mass measurements suggest $M_{\rm TOV} \approx 2.2$–$2.4\,M_\odot$ for non-rotating stars [113], with the exact value still debated.

be either a rapidly rotating, highly magnetised neutron star – a *magnetar* – that eventually collapses into a black hole as it accretes further material, or a direct and prompt collapse to a black hole.

The accreting black hole engine. Black holes formed in collapsars require a sufficiently strong, coherent poloidal magnetic flux threading the horizon to enter the magnetically arrested disc (MAD) regime and power a Blandford–Znajek jet. The origin of this large-scale flux remains uncertain: it may be inherited from the magnetised proto-neutron star formed immediately before black-hole formation, provided the accretion flow can confine the nascent magnetosphere, or it may be generated in situ by a dynamo operating in the collapsar disc. Recent 3D GRMHD simulations show that even predominantly toroidal fields inherited from stellar evolution can generate coherent poloidal magnetic loops at large radii, which are subsequently advected onto the black hole [33, 76].

Black holes produced in BNS mergers, by contrast, typically form with high spin parameters and in a predominantly toroidal magnetic field configuration. In this case, a large-scale poloidal component must grow dynamically through turbulent amplification and the *magnetorotational instability* (MRI) [15]. The MRI acts on a differentially rotating, weakly magnetised plasma in which the angular velocity decreases outwards. Consider two fluid elements on neighbouring orbits, connected by a weak magnetic field line. The inner element orbits faster and runs ahead of the outer one, stretching the field line in the azimuthal direction. The resulting magnetic tension acts like a spring joining the two elements: it brakes the inner one and drags the outer one forward. Having lost angular momentum, the inner element sinks to a smaller radius, where it still orbits faster; having gained angular momentum, the outer element drifts further out, where it orbits more slowly. Their separation therefore grows, the field line is stretched further, and the tension that produced the displacement increases, so the perturbation runs away and the field is amplified exponentially on a timescale of the order of the orbital period. The consequent turbulence is thought to be the primary agent of angular momentum transport in accretion discs, and, in combination with the underlying shear, can build and maintain a large-scale field from an initially disordered configuration [15, 92]. Once the poloidal flux is sufficiently amplified, the conditions for the BZ mechanism and for the eventual transition to a fully or partially arrested state are met.

The magnetar engine. A magnetar remnant is also considered a possible central engine, although it is generally regarded as less likely than the accreting black hole scenario. For a magnetar to power a GRB jet, it must survive long enough to accelerate the outflow to the required Lorentz factor before collapsing. If the total mass of the remnant exceeds $M_{\rm TOV}$, collapse occurs on a short timescale, due to angular momentum losses. However, if the total mass of the system is low or the EoS is particularly stiff, the massive neutron star can survive the collapse for a significantly longer time, supported by rotation [113]. During this phase, the rapid differential rotation drives a dynamo that amplifies both the poloidal and toroidal magnetic field components, and the MRI can further enhance turbulent field amplification [92, 141].

A critical obstacle to the magnetar engine is *baryon loading*. In the immediate post-merger or post-collapse phase, the proto-neutron star loses mass through neutrino-driven winds: the intense neutrino luminosity ($L_\nu \sim 10^{53}\,\mathrm{erg\,s^{-1}}$) deposits energy in the surface layers of the star, driving a baryonic outflow that fills the polar region [44]. The maximum achievable Lorentz factor for an outflow loaded with baryons is $\Gamma_\infty \lesssim \eta = L/\dot{M}c^2$, where L is the available power and $\dot{M}$ is the baryon mass-loading rate; if the polar region is contaminated, Γ_∞ remains too low to power a canonical GRB. The polar funnel can be emptied of baryons only after the star has cooled and the neutrino-driven winds have died down, potentially allowing a relativistic outflow. By this time, however, the remnant may have already collapsed to a black hole.

Despite these difficulties, there is observational evidence compatible with a magnetar engine, most notably the shallow X-ray plateau phase observed by *Swift*/XRT in a large fraction of both long and short GRBs [154, 237]. This feature requires a sustained energy injection on timescales of 10^3–10^4 s, which is difficult to explain by standard accretion disc timescales around a black hole, but is naturally produced by the gradual spin-down of a long-lived magnetar radiating via magnetic dipole emission [125]. Furthermore, three-dimensional numerical simulations of long-lived BNS merger remnants have demonstrated that magnetised, differentially rotating neutron stars can launch collimated outflows, although with Lorentz factors typically insufficient for canonical sGRBs [34, 35].

In general, although not excluded, this scenario is disfavoured as the primary engine for the majority of GRBs due to the baryon loading problem and the low Lorentz factors produced in simulations.

2.3.5 From engine to simulation: the injection prescription

When modelling jet launching through numerical simulations, two complementary approaches exist. The first starts from a full central engine setup, a black hole accreting from a disc, and models jet formation self-consistently from first principles, but typically employs idealised, homogeneous environments. The second prescribes the jet analytically as a boundary condition, mimicking the effect of the engine without resolving it directly in the simulation, but can be coupled with more realistic environments. The two approaches are largely complementary but difficult to combine. The choice is primarily determined by the scientific question at hand. In this work, we are mainly interested in the jet structure and its impact on the observable afterglow emission, and we therefore adopt the second approach as developed in Pavan et al. [165], injecting the jet as an inner boundary condition in a realistic post-merger environment (see Chapter 4).

The most direct injection parameter is the jet *luminosity*, which is set by the BZ power of the central engine. GRBs have typical injection luminosities in the range 10^{49}–$10^{51}\,\mathrm{erg\,s^{-1}}$. The jet injection is sustained for a duration determined by the disc accretion timescale, which sets a *decay timescale* for the injected power. Another key timescale for the merger scenario is the *jet launching delay*: the jet is not launched immediately after the merger, but only after the polar region has been sufficiently cleared of baryonic contamination and the magnetic field has grown and organised on a large scale. Finally, the *magnetisation* of the injected jet – the ratio

of electromagnetic to total energy flux – is a fundamental parameter that determines the jet's stability properties, its interaction with the cocoon, and ultimately the Lorentz factor it can achieve. The specific injection profiles and the parameter space explored in this work are described in detail in Chapter 6.

Independently of the precise nature of the central engine and the jet launching mechanism, the relativistic outflow must immediately navigate a dense, anisotropic post-merger (or stellar envelope) environment. The physics of this propagation is the subject of the following chapters.

2.4 Landmark GRB events

2.4.1 GW170817: a multi-messenger event

A fundamental event in the field of merger driven GRBs occurred in August 2017. The simultaneous detection of GW170817 and GRB 170817A represented the long-awaited smoking gun linking sGRB jets with binary neutron star mergers.

The gravitational-wave signal detected by the LIGO–Virgo collaboration was consistent with the inspiral and merger of two neutron stars with component masses $M = 1.1\text{–}1.6\,M_\odot$ [1]. Unlike previous gravitational-wave detections, GW170817 was accompanied by a rich set of electromagnetic counterparts spanning the entire spectrum [2]. Approximately 1.74 s after the merger signal, the short gamma-ray burst GRB 170817A was independently detected by both Fermi/GBM and INTEGRAL [71, 199]. In the following days, observations revealed a rapidly evolving optical and infrared transient consistent with *kilonova* emission powered by the radioactive decay of heavy elements synthesised through r-process nucleosynthesis [39, 90, 209, 220]. In addition, a long-lasting non-thermal afterglow was monitored from radio to X-ray frequencies for more than three years, providing an unprecedented view of the dynamics of the relativistic outflow [14, 80, 217].

GW170817 was hosted in NGC 4993, a galaxy at ≈ 40 Mpc [32, 163], making it substantially closer than typical cosmological sGRBs. The prompt gamma-ray counterpart, GRB 170817A, exhibited several unusual properties compared to the cosmological sGRB population. Despite the remarkably small distance of the source, the observed gamma-ray emission was faint and intrinsically sub-luminous, while its temporal and spectral characteristics remained broadly consistent with those of short GRBs [2, 71]. These observations suggested that the geometry of the relativistic outflow and the observer's viewing angle played a major role in shaping the observed signal, highlighting the importance of angular structure in relativistic jets. Early analyses combining gravitational-wave measurements with the host-galaxy localisation indicated a viewing angle of approximately $\theta_v \approx 32^{\circ}{}^{+10}_{-13}$ relative to the jet axis, implying a significantly off-axis line of sight [54]. Subsequent analyses taking into account the radio centroid motion prefer instead smaller viewing angles, in the range 15° [68] – 28° [221].

The subsequent evolution of the afterglow provided even stronger evidence for this picture. Rather than fading immediately after the burst, the broadband afterglow gradually brightened for approximately 120–160 days before reaching its peak luminosity [40, 217]. Such behaviour is naturally explained if the relativistic

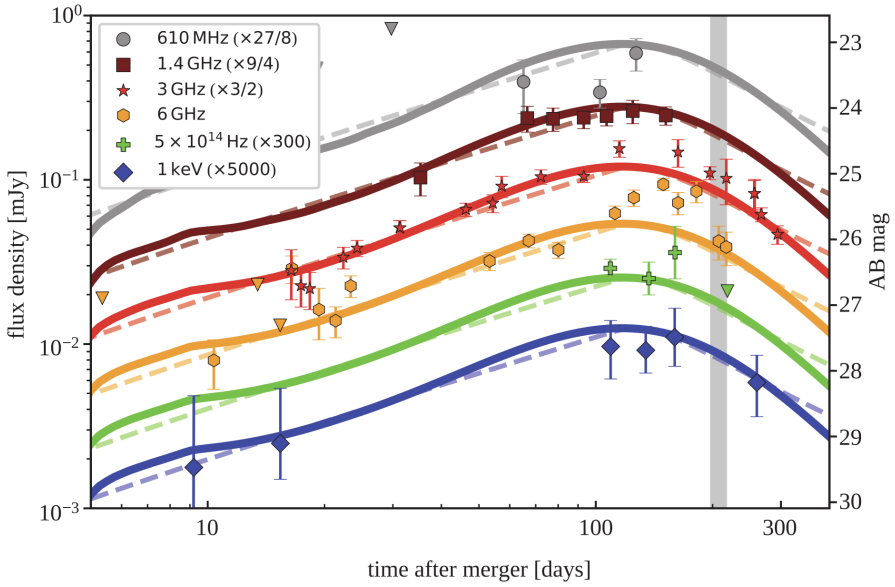


Figure 2.5: Multi-wavelength afterglow light curves overlaid with different structured jet models, adapted from Ghirlanda et al. [68]

ejecta possess an angular structure, with a narrow, energetic ultra-relativistic core surrounded by slower and less energetic material at larger angles, and the system is viewed from outside the jet axis [68, 101]. As the shock decelerates and its beaming cone widens, progressively more energetic regions enter the observer's line of sight, producing the characteristic delayed rise of the light curve.

For reference, Figure 2.5 shows multi-wavelength afterglow detections compared to different jet models, as reported in Ghirlanda et al. [68].

Independent confirmation of a relativistic jet came from very long baseline interferometry (VLBI) observations, which revealed both a compact radio source and apparent superluminal motion: both signatures of a narrowly collimated relativistic outflow observed at a significant angle from its axis [68, 137, 139]. Furthermore, modelling of the afterglow shows that, while the rising phase is primarily sensitive to the angular structure of the ejecta, the peak and subsequent decline are governed by the deceleration of an ultra-relativistic jet core with a kinetic energy comparable to that inferred for typical sGRBs [68, 190].

The success of structured-jet models in explaining GW170817 has had implications extending well beyond this single event. Numerical simulations indicate that angular structure naturally arises during the interaction between a relativistic jet and its surrounding environment, whether the stellar envelope or the merger ejecta [26, 47]. Consequently, structured outflows are now considered a potentially generic feature of GRB jets, capable of influencing the observed properties of bursts detected at a wide range of viewing angles. Growing evidence suggests that such effects may contribute to the diverse phenomenology observed in the afterglows of cosmological GRBs, including unusual early-time behaviour and departures from standard temporal evolution [13, 20, 49, 160].

Non-trivial jet structures are likely the outcome of two factors: the details of jet

launching, which determine the jet power and therefore its ability to successfully emerge from the surrounding material, as well as its magnetisation, which in turn influences stability and can leave imprints on the outflow; and the interaction with the surrounding environment.

The post-merger ejecta is a physically rich and intrinsically anisotropic environment through which the relativistic jet must travel before producing observable emission. This interaction strongly shapes both the angular structure of the jet and, ultimately, its electromagnetic signatures.

2.4.2 Breaking the Long/Short Dichotomy

The commonly used boundary at $T_{90} = 2\,\mathrm{s}$ should not be interpreted as a fundamental physical division between two distinct classes of GRBs. It is instead a conventional threshold chosen to separate two populations that nevertheless exhibit a clear bimodality in their observed duration distribution. As established in Section 2.1.5, the bimodality itself is a genuine feature of the GRB population, while the precise location of the dividing boundary is somewhat arbitrary. The duration-based classification therefore remains a useful operational tool for organising and comparing GRB populations, particularly in statistical studies of observed catalogues.

The observed prompt duration depends not only on the progenitor but also on detector sensitivity, redshift, and on the details of the central engine activity and accretion history. A collapsar-driven event with rapid jet quenching could in principle fall below 2 s; likewise, extended central engine activity in a merger could, in some circumstances, lead to prompt emission exceeding the conventional 2 s threshold. Over the past two decades, observations have increasingly highlighted the limitations of interpreting the long/short classification as a one-to-one mapping onto progenitor type.

Anomalous events: long duration without a supernova

The first major challenge to the canonical picture came with **GRB 060614**, detected by the *Swift* satellite on 14 June 2006 [43, 61, 62, 65].

With $T_{90} \approx 102\,\mathrm{s}$, the event was unambiguously long in duration and therefore expected to be collapsar-driven. Yet deep photometric follow-up found no supernova emission down to a limit more than one hundred times fainter than the faintest Type Ic supernova ever observed at comparable redshifts: a striking absence. The host galaxy, a faint dwarf with little ongoing star formation, was further incompatible with the massive-star hypothesis. The spectral properties of the prompt emission, moreover, more closely resembled those of the short, hard class than of canonical long GRBs.

Years later, near-infrared excess emission above the standard afterglow decay was identified in archival data, consistent with a kilonova component and therefore with r-process nucleosynthesis from post-merger ejecta [87, 228]. GRB 060614 thus became the first well-documented *hybrid* event: a long-duration GRB with the environmental and spectral signatures of a merger-driven progenitor, requiring a fundamental revision of the simple duration–progenitor correspondence.

GRB 211211A and GRB 230307A: the kilonova signature confirmed

More recently, two further events have made the case for long-duration, merger-driven GRBs essentially conclusive.

GRB 211211A, detected in December 2021, had a duration of $T_{90} \approx 51\,\mathrm{s}$ and was located in a galaxy at $\approx 350\,\mathrm{Mpc}$ [72, 182, 216, 230].

Despite its long duration, the prompt emission spectrum was harder than typical collapsar-driven events, and no supernova was detected. In contrast, the late-time infrared emission was clearly in excess of standard afterglow models and well described by a kilonova component powered by r-process nucleosynthesis: a signature of neutron-rich post-merger ejecta. The event also displayed a late-time GeV excess, potentially explained by inverse Compton scattering of kilonova photons by the electrons of a very long-lasting jet injection activity, possibly due to fallback accretion [120].

GRB 230307A, detected in March 2023, is among the brightest GRBs ever observed and had a duration of $T_{90} \approx 35\,\mathrm{s}$ [70, 103, 229].

The proximity of the event and the availability of the *James Webb Space Telescope* (JWST) allowed for an exceptionally detailed spectroscopic study of the afterglow: a broad emission feature in the near-infrared spectrum at around one month post-burst revealed the presence of tellurium and other lanthanide-rich material from r-process nucleosynthesis, providing direct spectroscopic evidence of a kilonova. The event also showed a spatial offset of $\approx 40\,\mathrm{kpc}$ from the centre of the identified host galaxy, a signature consistent with the natal-kick migration of a compact binary system over its long inspiral lifetime.

Physical explanations for long-duration merger-driven GRBs

The growing body of evidence for long-duration, merger-driven GRBs demands a physical explanation for how a compact binary merger can sustain central engine activity on timescales of tens to hundreds of seconds, well beyond the few seconds typically associated with BNS merger remnants. Several mechanisms have been proposed, and they are not mutually exclusive.

Long-lived magnetar. A millisecond magnetar formed in the merger can inject energy following its spin-down timescale, $t_{\mathrm{sd}} \propto B^{-2} P_0^2$, where P_0 is the initial spin period and B is the surface dipole magnetic field strength of the neutron star. For typical magnetar parameters ($P_0 \sim 1\,\mathrm{ms}$, $B \sim 10^{15}\,\mathrm{G}$), t_{sd} can reach hundreds of seconds [125, 232]. This scenario naturally produces extended central engine activity and is independently motivated by X-ray plateau observations in a fraction of sGRBs (see Section 2.3).

Fallback accretion. A fraction of the dynamical ejecta that is gravitationally bound to the remnant can fall back onto the central object on timescales that scale as $t_{\mathrm{fb}} \propto r^{3/2}$ [123, 188].

For the outer, slower shells, this timescale extends to minutes or longer, providing a fuel supply that could sustain accretion and jet activity well beyond the initial dynamical phase.

Magnetic flux and accretion disc lifetime. The magnetisation of the accretion disc plays a key role in setting the engine duration. In the MAD state, the strong magnetic flux threading the horizon can mediate episodic accretion outbursts over extended timescales [75]. The spin-down of a strongly magnetised disc through magnetic winds may also contribute to the late-time energy budget.

Neutron star – white dwarf mergers. An alternative channel involves the merger of a neutron star with a white dwarf rather than another neutron star. In this case, the tidal disruption radius is significantly larger, the mass ratio is more extreme, and the circularisation of the debris disc can lead to a more extended and more gradual accretion history, naturally producing longer-duration events [234].

2.5 Jet Propagation and Structure

2.5.1 The jet–environment interaction

From the moment the central engine switches on, the jet does not propagate into a vacuum. Whether it is drilling through the stellar envelope of a collapsar or the post-merger ejecta of a BNS coalescence, the outflow immediately encounters a dense, structured medium that resists its passage and shapes its final properties.

As the relativistic jet digs into the ambient material, a system of shocks develops at the interface. A forward shock is driven outward into the unperturbed medium, while a reverse shock propagates back into the incoming jet material. The two shocks are separated by a contact discontinuity. The region between the forward and reverse shock is commonly referred to as the *jet head*: the outermost region of the jet that is in direct contact with the environment, and the boundary across which kinetic energy is transferred from the jet to the surroundings. Behind the jet head, the unshocked jet material, has not yet interacted with the environment and continues to flow outward at essentially its injection velocity.

The material processed by both shocks, a combination, often separated, of jet and ambient fluid that has been heated and compressed, flows sideways from the jet head under its own pressure. This expelled material accumulates around the jet, inflating a hot, over-pressured *cocoon* [19, 26], (see Figure 2.6 for a reference).

The cocoon plays a dual role. On one hand, it acts as a pressure wall that collimates the jet: as long as the cocoon pressure exceeds the transverse jet ram pressure at its boundary, an oblique *recollimation shock* forms at the base of the jet, narrowing its opening angle below the injection value [26, 101]. On the other hand, it acts as a reservoir of energy that can contribute to the eventual outflow structure and afterglow emission [47, 101].

2.5.2 The jet head and its propagation velocity

The velocity at which the jet head advances through the medium is a fundamental parameter: it sets the timescale over which the jet can cross the ambient medium and determines whether the engine can sustain the jet long enough for it to break out.

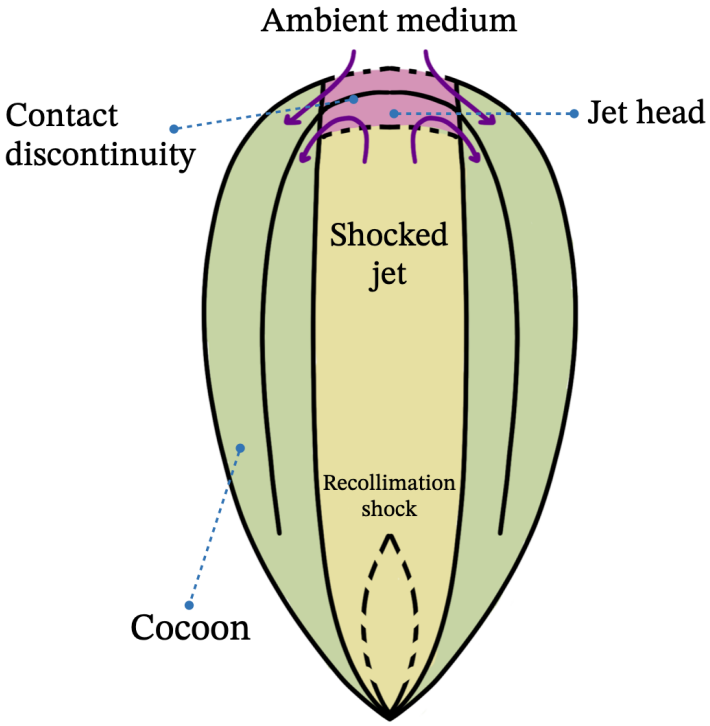


Figure 2.6: Schematic of the jet-environment interaction, inspired from Bromberg et al. [26].

In the frame of the jet head, pressure balance must hold across the contact discontinuity between the incoming jet material and the shocked ambient medium. For a relativistic jet moving at Lorentz factor Γ_j and expanding into an ambient medium of rest-mass density ρ_a, this balance gives [26, 116]

$$\beta_h \approx \frac{\tilde{L}^{1/2}}{1+\tilde{L}^{1/2}}\,, \qquad \tilde{L} \equiv \frac{L_j}{\rho_a\, c^3\, \theta_j^2\, r^2}\,, \tag{2.12}$$

where $\beta_h = v_h/c$ is the jet head velocity in units of c, L_j is the one-sided jet luminosity, θ_j is the injection half-opening angle, and r is the distance from the central engine. The dimensionless parameter $\tilde{L}$ measures the ratio of jet ram pressure to ambient ram pressure. When $\tilde{L} \gg 1$ (powerful jet or rarefied medium), the head advances relativistically, $\beta_h \to 1$. When $\tilde{L} \ll 1$ (weak jet or dense medium), the head is sub-relativistic and the jet struggles to advance.

This analytic estimate assumes a steady injection, but it captures the essential competition between jet power and environmental resistance that governs the dynamics at all times.

2.5.3 Jet breakout and the choked jet scenario

The propagation described above has a finite destination: the outer edge of the ambient medium, whether the stellar surface or the outer boundary of the merger ejecta. *The jet breakout* is defined as the moment the jet head exits this boundary,

after which it is free to accelerate and expand into the near-vacuum of the circumburst medium.

Defining the breakout time as $t_{\rm bo}$, taking place at a distance $r_{\rm bo}$, the fundamental condition for a successful GRB jet expressed in the laboratory frame is [26]

$$t_{\rm bo} - \frac{r_{\rm bo}}{\beta_{\rm j} c} \leq t_{\rm eng}\,, \tag{2.13}$$

where $t_{\rm eng}$ is the duration of central engine activity, $r_{\rm bo}/(\beta_{\rm j} c)$ is the laboratory-frame propagation time of the jet from the central engine to the breakout radius and $\beta_{\rm j}$ the jet velocity. In the collapsar case, the breakout time depends on the stellar radius and density profile, and scales approximately as $t_{\rm bo} \propto R_*/\beta_h$ [27], with $R_* \sim 10^5 - 10^6$ km for a typical Wolf-Rayet star. In our reference simulation of P23, performed in a BNS merger environment, the relevant scale is much smaller, the ejecta extend to only $\sim 10^4$ km at the time of jet launching, and the breakout times is of order hundreds of milliseconds, e.g., $t_{\rm bo} \approx 350$ ms . In general, for BNS environments, this process strongly depends on the jet launching time and velocity propagation, since the merger ejecta expand with time.

Aloy et al. [8] demonstrated through simulations that the propagation of a relativistic jet through a massive-star progenitor depends sensitively on the jet luminosity and the structure of the stellar envelope. Subsequent analytical work, in particular Aloy et al. [10], has formalised this dependence in terms of a *minimum jet luminosity* for successful jet injection and propagation. For a given progenitor structure, sufficiently low jet luminosities produce a slow and poorly collimated jet that can become trapped within the stellar material, without successfully emerging from the stellar surface. The corresponding threshold luminosity depends on the properties of the ambient medium, including its density and pressure structure, and on the jet injection parameters. In addition, successful breakout requires the central engine to remain active for at least the breakout time, so that a jet with a luminosity above the threshold is not necessarily successful if the engine shuts off too early.

If the jet fails to break out before the engine shuts off, the reverse shock propagates back through the jet and the jet is said to be *choked*. The energy injected by the jet is then deposited predominantly into the surrounding cocoon, driving a hot, wide-angle outflow with a range of velocities that can extend into the mildly relativistic regime. In the absence of an ultra-relativistic jet emerging from the ejecta, a classical GRB is not expected; however, the subsequent breakout of the cocoon shock can produce weaker, softer high-energy emission, distinct from a classical afterglow [73, 144].

Some statistical analyses of the long GRB T_{90} distribution have been interpreted as evidence that a fraction of collapsars do indeed lead to choked jets, with the majority of events having central engine duration insufficient for a successful breakout [27]. In the BNS case, the injection time relative to the merger is an additional key parameter, since the ejecta density decreases rapidly with time as the medium expands; a jet injected later encounters a more extended environment, which can significantly alter the breakout dynamics. This dependence is explored systematically in the parameter study of Chapter 6.

2.5.4 Jet structure: from top-hat models to realistic profiles

Initially, the most widely used model for GRB jet structure was the *top-hat* profile: a conical outflow with uniform Lorentz factor Γ and isotropic-equivalent energy per solid angle $dE/d\Omega$ within a sharp cone of half-opening angle θ_j:

$$\frac{dE}{d\Omega}(\theta) = \begin{cases} E_0/2\pi(1-\cos\theta_j) & \theta \leq \theta_j\,, \\ 0 & \theta > \theta_j\,, \end{cases} \tag{2.14}$$

with E_0 the total one-sided jet energy. This idealisation was adequate for explaining the achromatic jet break in afterglow light curves and gave reasonable estimates of the jet opening angle and total energy. However, it predicts an abrupt transition from full emission to no emission at $\theta = \theta_j$, which is physically unrealistic: a sharp discontinuity cannot persist in a flow that has interacted with its environment along the entire length of its journey.

The off-axis observation of GW170817 made this limitation impossible to ignore. The observed data required a *structured* jet: a smooth distribution of energy and velocity with angle, with a bright, narrow core and decreasing energy at larger angles, so that progressively more energetic emission entered the line of sight as the jet decelerated and its beaming cone widened.

The idea that jets should be structured had already been explored in the early 2000s [96, 187, 236], but GW170817 provided the first direct measurement of the angular profile. Afterglow modelling of GRB 170817A constrained the jet to have a narrow core with half-opening angle $\theta_c \approx 2°$–$6°$[68, 190].

Analytically, structured jets are commonly described by two parametric families. A *power-law* profile takes

$$\frac{dE}{d\Omega}(\theta) = \frac{E_c}{1+(\theta/\theta_c)^s}\,, \tag{2.15}$$

where E_c is the core energy per solid angle, θ_c is the core half-opening angle, and $s > 0$ sets the steepness of the wing decay. A *Gaussian* profile takes

$$\frac{dE}{d\Omega}(\theta) = E_c \exp\left(-\frac{\theta^2}{2\theta_c^2}\right)\,, \tag{2.16}$$

which concentrates most of the energy within $\theta \lesssim 2\theta_c$ and decays smoothly to zero at large angles. Both families allow a continuous transition across the jet boundary and produce afterglow light curves broadly consistent with GW170817 and other off-axis events [20, 190]. For reference, Figure 2.7 represents graphically such analytical jet profiles.

While analytically convenient, these parametric profiles are fundamentally limited, and simulations indicate that the structure can be more complex. The origin and significance of the jet's angular structure therefore gained a growing interest from the GRB community, as extensively discussed in Salafia and Ghirlanda [192], and is a pivotal element guiding this thesis.

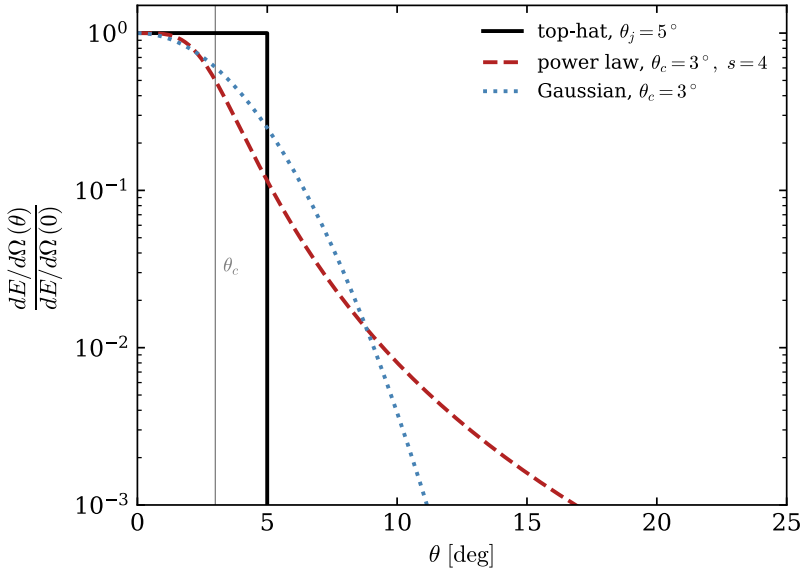


Figure 2.7: Schematic representation of top-hat (black line), power law (red line) and a Gaussian (blue line) jet profile, normalised to 1.

2.6 Why Numerical Simulations?

The physical processes described in the preceding sections are fundamentally nonlinear and multi-scale. Analytical treatments such as the Bromberg et al. [26] framework provide powerful order-of-magnitude estimates, but they necessarily rely on simplifying assumptions (uniform density profiles, fixed pressure ratios, instantaneous mixing) that break down in realistic environments. Numerical simulations are therefore an essential tool: they allow one to follow the full nonlinear dynamics of the jet from injection to breakout, and to extract the angular and velocity structure of the outflow in a self-consistent way. Over the past two decades, progress in the field has followed the growth of computational resources and of observational data.

2.6.1 From the First 2D Simulations to the 3D Era

The computational investigation of relativistic jet propagation precedes gamma-ray-burst-specific numerical simulations by several years. In the 1990s, a series of 2D axisymmetric simulations of relativistic jets in external media (e.g., including Martí et al. [116], Duncan and Hughes [48], Martí et al. [117], and Komissarov and Falle [94]) established the basic morphological picture that would later become the foundation of GRB jet theory: an advancing jet head bounded by a forward and reverse shock, surrounded by an over-pressured cocoon of shocked material. These studies demonstrated that material entering the jet head is deflected transversely and fills a hot, overpressurised, cocoon around the jet, that collimates the jet itself. However, they were conducted in the context of extragalactic jets at kiloparsec scales, not in GRB progenitors.

One of the first multidimensional special-relativistic hydrodynamic simulations

of GRB jet propagation through a stellar progenitor were performed by Aloy et al. [8], who used the GENESIS special-relativistic hydrodynamics code to simulate an axisymmetric jet injected into a collapsing rotating massive star. They demonstrated that a relativistic jet could remain collimated while drilling through the stellar envelope and reach breakout with a maximum Lorentz factor of $\Gamma \approx 33$. It established the main morphological elements of the jet-in-medium system: a terminal bow shock, a narrow cocoon, a contact discontinuity, and a hot spot, confirming that the phenomenology developed in the AGN context carries over to the GRB regime, and providing the first quantitative picture of how a relativistic outflow drills through a stellar envelope.

Over the following decade, a generation of 2D axisymmetric simulations refined this picture substantially. Zhang et al. [238], Mizuta et al. [135], and Morsony et al. [140] explored jet propagation in collapsars over a wide range of parameters – luminosity, initial Lorentz factor, injection geometry – and began connecting the jet properties at injection to the post-breakout angular structure. A key finding of Morsony et al. [140] was that the jet opening angle is not fixed at injection but evolves continuously during propagation, reflecting the cumulative effect of cocoon confinement: the angular structure of the jet in the afterglow phase thus encodes the full history of the jet–cocoon interaction together with the beam pattern set by the central engine. Subsequent studies extended this programme to GRBs from BNS mergers and to different progenitor profiles and injection geometries [e.g. 100, 134, 142], providing the numerical foundation for the analytical understanding of jet confinement and breakout.

These 2D studies shared, however, a fundamental limitation beyond the imposed axisymmetry. It was later recognised that axisymmetric simulations produce a numerical artefact: a polar plug of artificially high pressure and density along the jet axis [74, 238], which can bias the jet propagation speed and the cocoon structure in ways that are absent in 3D. Moving to three dimensions was therefore not only a matter of capturing azimuthal asymmetries, but of removing a systematic bias that affects many 2D GRB jet simulations. The growth in computational power and improvements in algorithms eventually made 3D modelling of relativistic jets tractable, and these studies confirmed that 3D jets can develop qualitatively different behaviour from their 2D counterparts, including non-axisymmetric instabilities, turbulent mixing at the jet–cocoon interface, and a richer post-breakout angular structure.

For GRBs from BNS mergers, the modelling challenge is a dynamical, anisotropic ejecta, whose properties depend on the outcome of the merger itself. The transition from idealised to realistic environments represents one of the most important methodological advances in the field, and is discussed in the following section.

2.6.2 Idealised vs Realistic Environments: Why It Matters

Early simulations of GRB jets used simplified, parametric environments: typically power-law or exponential density profiles designed to approximate the expected ejecta distribution [e.g. 9, 25, 143]. These setups offered full control over the parameters and were computationally tractable, but they cannot capture two key properties of

the actual medium: its anisotropy and its magnetic field structure.

A jet propagating through this complex medium does not encounter the same resistance in all directions, and the resulting asymmetry is directly imprinted on the jet structure. No parametric environment, however carefully designed, can reproduce this effect without input from a progenitor simulation. Furthermore, the post-merger, or stellar, environment is magnetised. The large-scale magnetic field amplified during the merger or the collapse permeates the ejecta and can influence the jet–medium interface, possibly changing the physical character of the interaction at the jet boundary.

Pavan et al. [165] performed the first 3D RMHD simulation of a sGRB jet in a realistic, non-magnetised post-merger environment extracted from a GRMHD BNS merger simulation, demonstrating that the realistic medium produces a substantially more asymmetric jet structure than idealised setups. P23 extended this work by reporting the first 3D RMHD simulation of a magnetised sGRB jet propagating through a realistic *magnetised* post-merger environment, where the magnetic field components were directly imported from the reference merger simulation. A quantitative comparison between realistic and simplified environments in that work revealed major differences in jet structure and breakout conditions, emphasising that incorporating the former is essential for accurate modelling. These two simulations are the direct precursors of the work presented in this thesis, which builds on the P23 setup.

The inclusion of magnetic fields in jet simulations is not only a matter of completeness. Magnetic fields play a key role in jet breakout and collimation. The toroidal magnetic field of the jet suppresses Kelvin-Helmholtz instabilities at the jet–cocoon interface, maintaining a smoother boundary and a better-collimated structure compared to purely hydrodynamical jets: an effect that is not captured in HD simulations and that becomes quantitatively significant even at the moderate magnetisations explored in this thesis.

Three-dimensional simulations of GRB jets in realistic BNS environments have begun to appear in recent years [91, 148, 149, 165, 166], confirming that the assumption of axisymmetry misses qualitatively important features of the flow. However, as discussed in Section 1.2, most existing 3D studies have been terminated before reaching the ballistic regime, limiting their ability to provide reliable input for afterglow models. The development of a strategy to overcome this limitation is the primary technical contribution of this thesis, presented in Chapter 5.

2.7 The GRB Afterglow

2.7.1 From Prompt Emission to External Shock

At smaller radii, the main energy dissipation mechanisms proposed for the prompt γ-ray emission include internal shocks between shells of jet material and magnetic reconnection, in which magnetic energy is converted into particle kinetic energy and radiation through the rearrangement and dissipation of magnetic field lines. The relative importance of these mechanisms remains an open question and is the subject of ongoing studies.

At larger radii, the interaction between the jet and the ambient medium becomes increasingly important. As the relativistic jet propagates into the circumburst medium, it drives a forward shock into the external material, sweeping up and accelerating an increasing amount of ambient matter. The shocked material behind this forward shock forms a relativistically expanding *blast wave*, whose dynamics evolve as an increasing fraction of the jet energy is transferred to the swept-up medium. The transition from approximately coasting to significant deceleration is not instantaneous, but occurs around a characteristic scale, the *deceleration radius*, which marks the distance at which the swept-up mass becomes dynamically important. Consider a shell of isotropic-equivalent energy $E_{\rm iso}$ expanding into a cold, uniform medium predominantly composed of hydrogen, with particle number density n. The blast wave begins to decelerate when the rest-mass energy of the swept-up material becomes comparable to $E_{\rm iso}/\Gamma^2$, giving

$$R_{\rm dec} \simeq \left(\frac{3E_{\rm iso}}{4\pi\, n\, m_p\, c^2\, \Gamma_0^2} \right)^{1/3}, \tag{2.17}$$

where Γ_0 is the initial Lorentz factor and m_p is the proton mass [122]. For representative short-GRB parameters, with $E_{\rm kin,iso} \sim 10^{51}$–$10^{52}$ erg and circumburst densities of $n \sim 10^{-3}$–10^{-2} cm^{-3} [56], the deceleration radius is typically of order 10^{17}–10^{18} cm, depending also on the initial Lorentz factor. For example, adopting $E_{\rm kin,iso} = 10^{51}$ erg, $n = 10^{-2}$ cm^{-3}, and $\Gamma_0 = 100$ gives $R_{\rm dec} \simeq 1.2 \times 10^{17}$ cm. The corresponding observer-frame deceleration time, $t_{\rm dec} \simeq \frac{R_{\rm dec}}{2c\Gamma_0^2}$, is of order a few minutes for these parameters. This timescale marks the onset of significant deceleration of the blast wave, while the observed onset of afterglow emission can occur over a broader range of times depending on the observing band and the microphysical parameters.

Beyond $R_{\rm dec}$, the dynamics is primarily governed by the expansion of the blast wave into the ambient medium, and the resulting radiation is produced by the shock-accelerated particles therein. In the standard picture, the afterglow at this stage retains little memory of the detailed workings of the central engine: the jet angular structure, the total energy budget, and the circumburst density are the dominant inputs [121, 183]. In practice, however, the situation is more intricate: the afterglow light curve can still carry imprints of the central engine activity through mechanisms such as energy injection from late-time engine restarts or the velocity stratification of the ejecta (see Section 2.8), and disentangling these from the external-shock contribution remains an active area of research.

2.7.2 Blast Wave Dynamics: the Blandford–McKee Solution

The self-similar solution for an ultra-relativistic spherical blast wave propagating into a uniform medium was derived by Blandford and McKee [23]. In the adiabatic limit, valid as long as the radiative losses are a small fraction of the total energy, the Lorentz factor of the forward shock evolves as

$$\Gamma(R) \propto R^{-3/2}, \qquad \text{or equivalently} \qquad \Gamma(t) \propto t^{-3/8}, \tag{2.18}$$

where t denotes the observer-frame photon arrival time measured from the burst. Although the Lorentz factor continues to decrease with time, its rate of decrease

becomes progressively smaller as the blast wave expands. Because the swept-up mass grows as R^3 in a uniform medium, the Lorentz factor follows the self-similar scaling above.

A second relevant density profile is the stellar wind case, $n(r) \propto r^{-2}$, which is commonly associated with GRBs from massive-star progenitors. In this case the temporal evolution of Γ is shallower: $\Gamma(t) \propto t^{-1/4}$. For GRBs from BNS mergers, which are the main focus of this thesis, the wind profile is not physically motivated: compact binary systems migrate over timescales of Gyr from their star-forming birth sites, typically ending up in low-density environments well described by a uniform ISM [56]. All afterglow calculations in this work therefore assume a constant circumburst density n.

In a spherically symmetric description, the deceleration of the blast wave is entirely determined by three macroscopic quantities: the total energy $E_{\rm tot}$, the ambient density n, and the initial Lorentz factor Γ_0. These alone, however, are not sufficient to predict the electromagnetic emission, which additionally requires knowledge of how energy is distributed between particles and magnetic fields at the shock front: a problem that cannot currently be solved from first principles and must instead be parametrised.

2.7.3 Microphysics and the Synchrotron Spectrum

At the shock front, two microphysical processes determine the electromagnetic output: the acceleration of charged particles and the amplification of the magnetic field.

Particle acceleration in relativistic shocks is understood to proceed via the *Fermi mechanism*: electrons scattered back and forth across the shock front by magnetic field irregularities gain energy at each crossing, producing a non-thermal energy distribution. In the test-particle limit, this process naturally generates a power-law spectrum of electron Lorentz factors, $N(\gamma_e) \propto \gamma_e^{-p}$ for $\gamma_e \geq \gamma_m$, where $p \approx 2$–2.5 is the spectral index and γ_m is the minimum Lorentz factor of the accelerated population [23, 64]. The fraction of the post-shock internal energy deposited into relativistic electrons is parametrised as ε_e.

Simultaneously, the shock compresses and amplifies the ambient magnetic field via plasma instabilities, which generates small-scale magnetic field fluctuations in the downstream plasma. The resulting field strength is parametrised by ε_B, the fraction of post-shock internal energy in the amplified field. The average magnetic field strength behind the shock is parametised as $B = (8\pi\,\varepsilon_B\,e)^{1/2}$, where e is the post-shock internal energy density and ε_B is the fraction of the internal energy density carried by the magnetic field. For a strong ultra-relativistic shock propagating into a cold medium of proper number density n, the internal energy density immediately behind the shock is given by the relativistic shock jump conditions [23] as $e = 4\Gamma^2 n m_p c^2$, where Γ is the Lorentz factor of the shocked fluid measured in the rest frame of the unshocked medium. ε_e and ε_B are phenomenological parameters whose values has to be inferred from observations or constrained by simulations. Assuming that a fraction ε_e of the post-shock internal energy is transferred to the electrons, and that the accelerated electrons are described by a power-law distribution $N(\gamma_e) \propto \gamma_e^{-p}$

above a minimum Lorentz factor γ_m, the latter is given by

$$\gamma_m = \frac{p-2}{p-1}\,\frac{m_p}{m_e}\,\varepsilon_e\,\Gamma, \tag{2.19}$$

where $p > 2$ is the power-law index of the electron distribution and Γ is the bulk Lorentz factor of the shocked fluid. The four microphysical parameters of the standard afterglow model are therefore $\{\varepsilon_e, \varepsilon_B, p, n\}$. Together with $E_{\rm tot}$ and the viewing angle θ_v, these constitute the six free parameters that are typically fitted to multiwavelength afterglow observations. A systematic characterisation of these parameters across a sample of ten well-monitored bursts was carried out by Panaitescu and Kumar [164], who found $\varepsilon_e \sim 0.01$–0.1, $\varepsilon_B \sim 10^{-5}$–$10^{-2}$, and $p \sim 1.4$–2.8 as representative values.

The relativistic electrons gyrating in the amplified magnetic field emit *synchrotron radiation*: a relativistic electron with Lorentz factor γ_e moving in a field B radiates predominantly at a characteristic frequency $\nu \propto \gamma_e^2 B$. Because the electron population spans a power-law range in γ_e, the integrated synchrotron emission is itself a power law in frequency, broken at the characteristic frequencies of specific electron subpopulations. The resulting spectrum is a series of power-law segments that join at three break frequencies [196]:

- ν_m, the synchrotron frequency of an electron with $\gamma_e = \gamma_m$ (the peak of the injected electron distribution);
- ν_c, the *cooling frequency*, defined as the synchrotron frequency of an electron whose radiative cooling time equals the dynamical time of the system;
- ν_a, the *self-absorption frequency*, below which the plasma is optically thick to synchrotron self-absorption and the spectrum turns over.

Physically, ν_m tells us where the injected population peaks, ν_c separates electrons that have had time to cool efficiently from those that have not, and ν_a marks the transition from optically thin to thick emission, the latter being relevant mainly at radio frequencies.

2.7.4 Spectral Regimes and Light Curve Slopes

The relative ordering of the break frequencies determines the spectral shape, and because all three frequencies evolve in time, the spectrum also evolves. For an adiabatic, ultra-relativistic blast wave in a uniform medium, the temporal scalings are [196]:

$$\nu_m \propto t^{-3/2}, \qquad \nu_c \propto t^{-1/2}, \qquad F_{\nu,\max} \propto t^{0}. \tag{2.20}$$

Since ν_m decreases faster than ν_c, the two frequencies cross at a time $t_{\rm cross}$: before this crossing the system is in *fast cooling* ($\nu_c < \nu_m$), and after it is in *slow cooling* ($\nu_m < \nu_c$).

For sGRBs observed at days to weeks after the trigger, the timescales relevant for GW counterpart follow-up, the system is generally in the slow-cooling regime,

which is also the regime most relevant for this thesis. The slow-cooling spectrum takes the form

$$F_\nu \propto \begin{cases} \nu^2 & \nu < \nu_a, \\ \nu^{1/3} & \nu_a < \nu < \nu_m, \\ \nu^{-(p-1)/2} & \nu_m < \nu < \nu_c, \\ \nu^{-p/2} & \nu > \nu_c. \end{cases} \tag{2.21}$$

(See Figure 2.8 for a reference). At a given observing frequency $\nu_{\rm obs}$, the resulting light curve becomes a piecewise (multiply) broken power law as a function of time, whose slope depends on which spectral segment $\nu_{\rm obs}$ falls in and on which way the break frequencies are moving through it. In the standard on-axis case, this leads to a temporal decay $F_\nu \propto t^{-\alpha}$, where the temporal decay index α is not a free parameter but is uniquely determined by the spectral index β (defined by $F_\nu \propto \nu^{-\beta}$) and by which spectral regime the observing frequency falls in. These *closure relations* (for example, $\alpha = 3\beta/2$ for $\nu_m < \nu_{\rm obs} < \nu_c$ in the slow-cooling ISM case) follow directly from the blast-wave dynamics and the synchrotron spectrum, and provide a powerful self-consistency check when multi-band data are available [79]: measuring α and β independently and verifying that they satisfy the expected relation confirms that the standard model applies, while deviations signal additional physics such as energy injection or a non-standard density profile. This statement holds, however under the implicit assumption that the microphysical afterglow parameters do not vary with time, which is not necessarily true.

A useful quantitative anchor is provided by the explicit expressions for ν_m and ν_c in the slow-cooling regime [196]:

$$\nu_m \simeq 5.7 \times 10^{11} \left(\frac{\varepsilon_B}{10^{-2}}\right)^{1/2} \left(\frac{\varepsilon_e}{0.1}\right)^{2} E_{52}^{1/2}\, t_{\rm d}^{-3/2}\, (1+z)^{1/2} \quad \text{Hz}, \tag{2.22}$$

$$\nu_c \simeq 2.7 \times 10^{15} \left(\frac{\varepsilon_B}{10^{-2}}\right)^{-3/2} E_{52}^{-1/2}\, n_1^{-1}\, t_{\rm d}^{-1/2}\, (1+z)^{-1/2} \quad \text{Hz}. \tag{2.23}$$

where $E_{52} = E/10^{52}$ erg, $n_1 = n/1$ cm^{-3}, and $t_{\rm d}$ is the observer-frame time in days and z is the source redshift. These expressions show, at a glance, how the observing frequency and the burst parameters conspire to determine the observed light curve shape: a burst with high ε_B or low energy has ν_c at lower frequencies, potentially placing the X-ray band in the fast-cooling regime even at late times.

One particularly clean observational consequence of this framework is the *jet break*. As long as $\Gamma \gtrsim \theta_j^{-1}$, the radiation arriving at the observer comes from a patch of angular size $\sim 1/\Gamma$ around the line of sight, and the blast wave appears approximately identical to an isotropic explosion. When Γ drops below θ_j^{-1}, the observer becomes aware of the finite angular extent of the jet: the relativistic beaming cone widens beyond the jet opening angle, causing the afterglow light curve to steepen. The magnitude of this steepening, however, is not universal: it depends on the spectral regime, the density profile of the circumburst medium, and the degree of lateral expansion of the jet [186, 197]. For example, in a uniform external medium and in the absence of significant lateral spreading, the edge effect alone produces a steepening $\Delta\alpha \simeq 3/4$, whereas lateral spreading can lead to a substantially larger steepening.

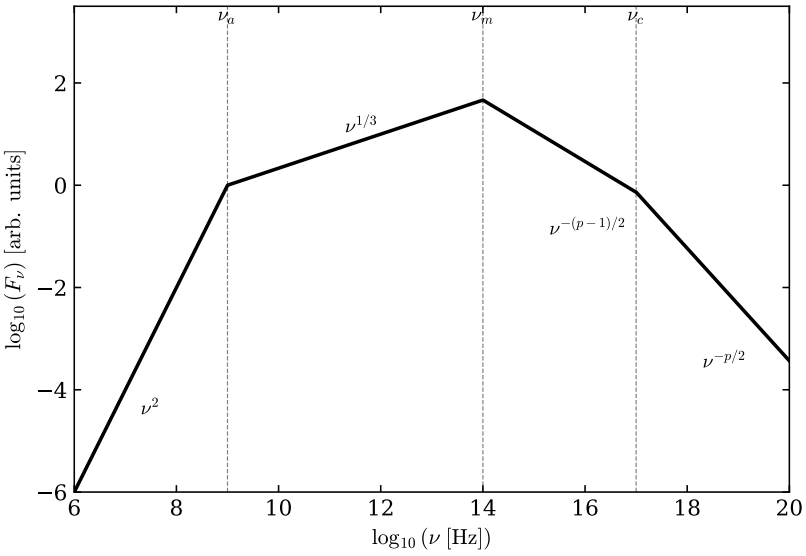


Figure 2.8: Representation of the synchrotron spectrum in the slow-cooling regime.

This achromatic break, appearing simultaneously at all wavelengths, is the cleanest observational signature of jet collimation, and it allows one to estimate θ_j directly from the break time:

$$\theta_j \simeq 0.057\, t_{j,\mathrm{d}}^{3/8} \left(\frac{1+z}{2}\right)^{-3/8} \left(\frac{E}{10^{52}\,\mathrm{erg}}\right)^{-1/8} n_1^{1/8} \quad \mathrm{rad}, \tag{2.24}$$

where $t_{j,\mathrm{d}}$ is the jet break time in days. Applying this relation to a sample of long GRBs, Frail et al. [59] found that the collimation-corrected gamma-ray energy, $E_\gamma = E_{\gamma,\mathrm{iso}}(1-\cos\theta_j)$, was narrowly clustered around $\sim 10^{51}$ erg. This result was interpreted as evidence for a quasi-universal energy budget and became an influential early argument for strongly collimated GRB outflows. However, this conclusion has not remained robust with the advent of *Swift*. The identification of achromatic jet breaks has proven considerably more difficult than suggested by pre-*Swift* samples, with many afterglow light curves exhibiting chromatic breaks, plateaus, or other features that cannot be straightforwardly attributed to jet geometry. Consequently, the evidence for a narrowly clustered, quasi-universal beaming-corrected energy is now considerably weaker [30, 179].

2.7.5 Degeneracies and the Importance of Broadband Coverage

While the standard afterglow model outlined above provides a simple and well-defined framework for interpreting the broadband emission, it also comes with a well-known limitation: parameter degeneracies. Different combinations of $\{E_{\mathrm{tot}}, n, \varepsilon_e, \varepsilon_B\}$ can produce virtually identical light curves at a single wavelength, making it difficult to constrain the physical parameters from single-band observations alone.

The key to reducing these degeneracies is broad spectral and temporal coverage. Panaitescu and Kumar [164] demonstrated the constraining power of simultaneous

radio, optical, and X-ray observations by modelling a sample of ten well-monitored long GRBs. For example, radio observations can constrain the self-absorption frequency and the overall normalization of the synchrotron spectrum, the optical band often probes the vicinity of ν_m, and the X-ray band can provide constraints on ν_c and the electron spectral index p. However, broadband coverage does not by itself make the inference problem deterministic. The standard afterglow model relies on a chain of simplifying assumptions concerning the shock dynamics, the magnetic field, the electron acceleration process, and the shape of the electron distribution. These assumptions are not all on equally firm physical ground, and a fully self-consistent description of the microphysics is still lacking. In particular, the electron distribution at low Lorentz factors, which strongly affects the radio and self-absorption regime, is not robustly established from first principles. Consequently, different modelling assumptions or implementations can lead to different inferred values of the physical parameters even for well-sampled broadband afterglows. Moreover, convergence of a fit does not guarantee that the solution is unique or that the assumed model provides an unbiased description of the data. Broadband observations should therefore be regarded as essential for constraining afterglow physics, but the resulting parameter estimates remain intrinsically model-dependent.

For the case of GW170817, the afterglow was monitored in the X-ray (*Chandra*), optical (*HST*) and radio (VLA) bands for ~ 940 days. Over this whole time and frequency range, however, the emission is consistent with a single power-law spectral segment, $\nu_m < \nu < \nu_c$, with $p \approx 2.17$ [57, 111]: the slope fixes p, but, since no spectral break is observed, only a combination of the remaining parameters is constrained. As a result, the parameter constraints remain broad: the strong degeneracies between n, ε_B, and the jet structure mean that different combinations of these parameters can fit the data comparably well, and the inferred viewing angle carries uncertainties of several degrees depending on the assumed jet model [e.g. 68, 136].

2.8 Beyond the Standard Afterglow Model

The framework described so far, a single shell with a given kinetic energy at each angle, decelerating in a uniform medium, is adequate as a zeroth-order description of on-axis afterglows at late times. It fails, however, to capture several physical effects that become especially important for structured jets observed off-axis.

The most significant of these effects is *velocity stratification*: a realistic outflow does not carry all its energy in a single, coherent shell moving at a single Lorentz factor. Instead, energy is distributed over a broad range of four-velocities $u = \Gamma\beta$, reflecting the layered structure that emerges from the jet–cocoon interaction and the conversion of internal and magnetic energy into kinetic energy over the duration of the propagation. Slower shells, which lag behind the fast-moving core, eventually catch up with the decelerating forward shock and inject additional energy into the blast wave: a process known as *refreshed shocks* [184, 194].

The observational consequence of refreshed shocks is a modification of the afterglow light curve, particularly at early times and for off-axis observers. Instead of the sharp rise expected from a simple shell model, the light curve evolves more gradually, as successive shells contribute at progressively later times.

In the context of this thesis, velocity stratification is not a free parameter but a direct output of the RMHD simulations: the distribution of energy over four-velocities at each angle is extracted from the simulation data and fed into the afterglow model. The ability to propagate this information to observable quantities is one of the key advances of the model described in this Section, and its consequences for the afterglow light curves are discussed in detail in Chapters 5 and 6.

2.8.1 The Code

The standard afterglow model outlined in Sections 2.7.2–2.7.4 is built on two implicit simplifying assumptions: the outflow is a single, uniform shell at each angle, and it is axially symmetric. Both assumptions become problematic when the goal is to model structured jets observed at significant off-axis angles.

The semi-analytical framework developed by Salafia et al. [193] addresses this by computing the afterglow emission from a jet with an *arbitrary* angular structure, defined as the distribution of isotropic-equivalent energy $E_{\rm iso}(\Theta)$ and Lorentz factor $\Gamma(\Theta)$ as functions of the polar angle from the jet axis. The blast wave at each angular coordinate (Θ, Φ) is evolved independently, however neglecting lateral energy transport between neighbouring angles, an approximation well-justified at early times [78].

The dynamics at each angle are governed by energy conservation. For a single ejecta shell with isotropic-equivalent kinetic energy $E_{\rm iso}$ expanding into a uniform ISM of number density n, the swept-up mass within radius r is $m_{\rm sw}(r) = 4\pi r^3 n m_p/3$. In the ultra-relativistic, adiabatic limit, the total energy of the system is [23]

$$E_{\rm iso} = \Gamma^2\, m_{\rm sw}(r)\, c^2 = \frac{4\pi}{3}\,\Gamma^2\, r^3\, n\, m_p\, c^2, \tag{2.25}$$

where the factor Γ^2 arises from the Lorentz transformation of the rest-mass energy and the kinetic energy of the shocked material into the lab frame. Solving for Γ immediately gives

$$\Gamma(r) = \left(\frac{3\,E_{\rm iso}}{4\pi\, n\, m_p\, c^2}\right)^{1/2} r^{-3/2} \;\propto\; r^{-3/2}. \tag{2.26}$$

To convert to observer time, we note that a photon emitted at radius r on-axis ($\theta_v = 0$, $\beta \approx 1$) arrives at the observer at

$$t_{\rm obs} \approx t_{\rm lab}\,(1-\beta) \approx \frac{r}{2c\Gamma^2}, \tag{2.27}$$

so that $r \approx 2c\Gamma^2 t_{\rm obs}$. Substituting back into Eq. (2.26) and solving for Γ gives the standard Blandford–McKee temporal scaling [23]:

$$\Gamma(t_{\rm obs}) \propto t_{\rm obs}^{-3/8}. \tag{2.28}$$

This power-law deceleration governs the afterglow dynamics at each independent angle Θ, with the normalisation set by $E_{\rm iso}(\Theta)$.

The resulting synchrotron emissivity at each point on the blast wave is computed following the spectral model of Section 2.7.4. The total flux at the observer is

obtained by integrating this emissivity over the *equal-arrival-time surface* (EATS) – the locus of emission points that contribute to the same observer-frame time $t_{\rm obs}$ – accounting for Doppler boosting and light-travel-time effects [192, 193]. For an off-axis observer ($\theta_v > 0$), the EATS geometry couples contributions from different angles and radii, producing the characteristic rising light curve that precedes the afterglow peak.

Extension I: Radial Velocity Stratification

In the version of the model employed in this thesis [developed by O. S. Salafia and described in 46], the baseline framework of Salafia et al. [193] is extended in two important aspects that are directly motivated by the structure of the outflows emerging from the 3D RMHD simulations.

The first extension concerns the radial distribution of energy within the ejecta. The extended model incorporates velocity stratification by replacing the single-Lorentz-factor prescription with the full distribution $dE_{\rm iso}/du\,(\Theta, \Phi)$ extracted directly from the simulation output at $t_{\rm end}$. The blast-wave dynamics are then governed by an energy-conservation equation that accounts for the cumulative contribution of all ejecta faster than a given four-velocity $u = \Gamma\beta$.

Defining the isotropic-equivalent energy in ejecta with four-velocity greater than u,

$$E_{\rm iso}(>u) = \int_u^\infty \frac{dE_{\rm iso}}{du'}(u')\,du', \tag{2.29}$$

and the corresponding isotropic-equivalent rest mass,

$$m_{\rm ej,iso}(>u) = \int_u^\infty \frac{1}{c^2\sqrt{1+u'^2}}\frac{dE_{\rm iso}}{du'}(u')\,du', \tag{2.30}$$

conservation of energy per unit solid angle when the blast-wave downstream four-velocity is u gives [Appendix B of D26]

$$E_{\rm iso}(>u) + m_{\rm iso}(u)\,c^2 = \Gamma[m_{\rm ej,iso}(>u) + m_{\rm iso}(u)]\,c^2 + \Gamma_{\rm eff}[(\Gamma-1)\,m_{\rm iso}(u)]\,c^2, \tag{2.31}$$

where $m_{\rm iso}(u) = 4\pi r^3(u)\,n\,m_p/3$ is the swept-up ISM mass, $\Gamma = \sqrt{1+u^2}$, and the effective Lorentz factor $\Gamma_{\rm eff}$ accounts for the post-shock internal energy in the lab frame [151, 171]:

$$\Gamma_{\rm eff} = \frac{\hat{\gamma}_{\rm ad}\,\Gamma^2 - \hat{\gamma}_{\rm ad} + 1}{\Gamma}, \tag{2.32}$$

with $\hat{\gamma}_{\rm ad}$ the adiabatic index of the blast-wave downstream. Equation (2.31) has the analytic solution

$$m_{\rm iso}(u) = \frac{E_{\rm iso}(>u)/c^2 - \Gamma(u)\,m_{\rm ej,iso}(>u)}{[\Gamma(u)-1][\Gamma_{\rm eff}(u)+1]}, \tag{2.33}$$

which, combined with $m_{\rm iso} = 4\pi r^3 n m_p/3$, gives a one-to-one mapping between shock radius r and downstream four-velocity u without requiring numerical time-integration. In the single-shell limit, $dE_{\rm iso}/du \to E_{\rm iso}\,\delta(u-u_0)$, this reduces to the standard Blandford–McKee dynamics; the full distribution generalises this to a continuous family of shells, each contributing to the blast-wave energy at the epoch when the forward shock decelerates to their velocity.

Extension II: Full 3D Geometry

The second extension relaxes the assumption of axial symmetry that is otherwise universal in semi-analytical afterglow models.

For standard afterglow calculations, the code can take as input the azimuthally averaged profiles $E_{\rm iso}(\Theta)$ and $\Gamma(\Theta)$, treating the outflow as rotationally symmetric around the jet axis. This is a natural starting point for comparing with 2D simulations, where axisymmetry is imposed by construction, and with 3D simulations with idealised, uniform environments. However, for jets that have propagated through the anisotropic, non-uniform ejecta of a BNS merger, where the three-dimensional structure of the environment leaves a clear imprint on the angular and azimuthal distribution of energy in the outflow (see Section 2.5.4), this is a limitation.

In the extended version of the code, the EATS integration is performed directly over the full three-dimensional distributions $E_{\rm iso}(\Theta, \Phi)$ and $dE_{\rm iso}/du\,(\Theta, \Phi)$, preserving all azimuthal information. Light curves are then computed as a function of both the polar viewing angle θ_v and the azimuthal viewing angle ϕ_v.

This has an important observational consequence: two observers at the same polar angle θ_v but different ϕ_v will, in general, see different afterglow light curves. The differences are most pronounced at early times and for $\theta_v \gtrsim$ a few degrees, when the observer's line of sight does not intersect the jet core directly and the received flux is dominated by the asymmetric wings. In D26, this is quantified explicitly: at $\theta_v = 24°$, close to the inferred viewing angle of GW170817, the flux density at early times varies by up to an order of magnitude depending on ϕ_v (see Fig. 11 therein).

This result has a broader implication: any afterglow fit performed under the assumption of axisymmetry implicitly fixes the azimuthal structure of the jet. Allowing for azimuthal asymmetry introduces an additional degree of freedom that can affect the interpretation of the data and the inferred jet structure. Quantifying the impact of this additional degree of freedom is one of the central results of D26 and is discussed further in Chapter 6.

Current Limitations

For completeness, we note the main approximation that remains in the current implementation: the model does not include lateral expansion of the blast wave. For a smooth jet, lateral spreading becomes important once $\Gamma \lesssim \theta_j^{-1}$, i.e. around the jet break, after which the independent-angle approximation breaks down. The same criterion, however, applies locally: a region of angular size $\delta\theta$ is smoothed out once its shocked material has had time to expand sideways over that scale at the local speed of sound, when $\Gamma \lesssim (\delta\theta)^{-1}$. The jet structures extracted from our simulations show fluctuations in energy and Lorentz factor on angular scales much smaller than θ_j, in both the polar and azimuthal directions (Section 6.1.1). Fluctuations could therefore be smeared out once before the jet break, driving the blast wave towards a smoother configuration. Local smoothing is, however, likely less efficient in the presence of a broad velocity profile. The downstream sound speed decreases as the shock decelerates, so the anisotropy introduced by late energy injection from slower shells can be preserved to a larger extent [191]. This can affect the shape of the light curve, as discussed in Appendix D, where we address the separate impact of the

velocity profile and of a standard treatment of lateral expansion on the afterglow emission of the fiducial model. A joint implementation is left to future work.

Chapter 3

Theoretical and Numerical Framework

Modelling the propagation of a relativistic magnetised jet through a post-merger environment requires a theoretical framework that can simultaneously accommodate three physical ingredients: bulk velocities approaching the speed of light, dynamically significant magnetic fields, and strong discontinuities such as shocks and contact surfaces. The most suitable framework under these conditions is special-relativistic ideal magnetohydrodynamics (RMHD), and it is the one adopted in this thesis. Throughout this Chapter we adopt Heaviside–Lorentz units with $c = 1$; factors of c and 4π are absorbed into the definitions of the electromagnetic fields.

This chapter is organised in two parts. The first presents the RMHD equations in covariant form and discusses the divergence-free constraint on the magnetic field. The second part describes the PLUTO code [131], its modular architecture, and the Godunov-type numerical framework through which the RMHD equations are solved, the strategy used to enforce the divergence free constraint numerically, and introduces the equation of state adopted for the simulations. The chapter closes with the general simulation setup that is common to all the models presented in Chapters 4–6.

3.1 Relativistic Magneto–Hydrodynamics

The dynamics of a relativistic magnetised plasma, in the regime relevant to GRB jets, is governed by three sets of equations: the conservation of baryon number, the conservation of energy and momentum, and Maxwell's equations for the electromagnetic field. In this section, we derive the conservative form of these equations as solved by PLUTO, following the treatment of Anile [12], Bucciantini [28] and the formulation of Mignone and Bodo [127].

3.1.1 Conservation Laws in Covariant Form

In a flat Minkowski spacetime with signature $(-,+,+,+)$, the fundamental equations governing a magnetised ideal fluid are baryon number conservation,

$$\partial_\mu N^\mu = 0\,, \tag{3.1}$$

energy-momentum conservation,

$$\partial_\mu T^{\mu\nu} = 0\,, \tag{3.2}$$

and the homogeneous Maxwell equations,

$$\partial_\mu F^{*\mu\nu} = 0\,, \tag{3.3}$$

where $N^\mu = \rho u^\mu$ is the baryon number 4-current, $T^{\mu\nu}$ is the total stress-energy tensor, $F^{*\mu\nu}$ is the dual of the electromagnetic field tensor, ρ is the rest-mass density, and u^μ is the fluid 4-velocity normalised as $u_\mu u^\mu = -1$.

The stress-energy tensor receives contributions from both the fluid and the electromagnetic field:

$$T^{\mu\nu} = T^{\mu\nu}_{\rm fluid} + T^{\mu\nu}_{\rm em}\,. \tag{3.4}$$

For a perfect fluid with isotropic pressure P, the matter contribution takes the standard form

$$T^{\mu\nu}_{\rm fluid} = \rho h\, u^\mu u^\nu + P\, g^{\mu\nu}\,, \tag{3.5}$$

where $h = 1 + \epsilon + P/\rho$ is the specific enthalpy, ρ is the rest-mass density, ϵ is the specific internal energy, i.e., the internal energy per unit rest mass, and $g^{\mu\nu}$ is the metric tensor.

To express the electromagnetic contribution in a form useful for MHD, we introduce the magnetic field 4-vector as measured in the comoving frame of the fluid, $b^\mu = u_\nu F^{*\nu\mu}$. Its spatial components give the proper magnetic field, while its time component satisfies $b^0 = \Gamma(\mathbf{v}\cdot\mathbf{B})$, where $\Gamma = (1 - |\mathbf{v}|^2)^{-1/2}$ is the Lorentz factor and $\mathbf{B}$ is the magnetic field measured in the laboratory frame. The comoving magnetic energy density is $b^2/2 = \frac{1}{2} b_\mu b^\mu$.

A key simplification comes from the ideal MHD condition: in a highly conducting astrophysical plasma, resistive effects operate on scales far below the grid resolution, and the electric field vanishes in the fluid rest frame. Formally, this condition reads

$$u_\mu F^{\nu\mu} = 0\,. \tag{3.6}$$

Under this assumption, the electromagnetic stress-energy tensor simplifies to

$$T^{\mu\nu}_{\rm em} = b^2 u^\mu u^\nu - b^\mu b^\nu + \frac{b^2}{2}\, g^{\mu\nu}\,. \tag{3.7}$$

Combining Eqs. (3.5) and (3.7), the total stress-energy tensor becomes

$$T^{\mu\nu} = \left(\rho h + b^2\right) u^\mu u^\nu - b^\mu b^\nu + \left(P + \frac{b^2}{2}\right) g^{\mu\nu}\,. \tag{3.8}$$

The magnetic field enters the dynamics through two channels: it contributes an isotropic magnetic pressure $b^2/2$ (analogous to thermal pressure) and an anisotropic tension $-b^\mu b^\nu$ that acts to straighten field lines. It is convenient to define a total enthalpy $w_t = \rho h + b^2$ and a total pressure $P_t = P + b^2/2$, in terms of which the stress-energy tensor takes the compact form

$$T^{\mu\nu} = w_t u^\mu u^\nu - b^\mu b^\nu + P_t g^{\mu\nu}\,. \tag{3.9}$$

To arrive at a system suitable for numerical integration, the covariant equations are projected onto the laboratory frame through a standard 3+1 decomposition. The result is a set of conservation laws in the form

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{T}(\mathbf{U}) = \mathbf{S}(\mathbf{U})\,, \tag{3.10}$$

where $\mathbf{U}$ is the vector of conserved variables, $\mathbf{T}(\mathbf{U})$ is the flux tensor, and $\mathbf{S}(\mathbf{U})$ contains source terms arising from curvilinear coordinates or from external forces. In Cartesian coordinates and in the absence of external forces, $\mathbf{S} = 0$ and the equations take a conservative form.

The conserved variables and their relation to the primitive variables $\mathbf{V} = (\rho, \mathbf{v}, P, \mathbf{B})^T$ are

$$\mathbf{U} = \begin{pmatrix} D \\ \mathbf{m} \\ \mathbf{B} \\ E \end{pmatrix}\,, \tag{3.11}$$

with

$$D = \Gamma \rho\,, \tag{3.12}$$

$$\mathbf{m} = \left(\rho h \Gamma^2 + |\mathbf{B}|^2\right)\mathbf{v} - (\mathbf{v}\cdot\mathbf{B})\,\mathbf{B}\,, \tag{3.13}$$

$$E = \rho h \Gamma^2 - P + \frac{|\mathbf{B}|^2}{2} + \frac{|\mathbf{v}|^2|\mathbf{B}|^2 - (\mathbf{v}\cdot\mathbf{B})^2}{2}\,, \tag{3.14}$$

where D is the laboratory-frame mass density, $\mathbf{m}$ is the total (fluid plus electromagnetic) momentum density, E is the total energy density. The magnetic field $\mathbf{B}$ is evolved through the induction equation, which is the spatial projection of Eq. (3.3):

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{v} \times \mathbf{B}) = 0\,. \tag{3.15}$$

The induction equation preserves the solenoidal constraint $\nabla \cdot \mathbf{B} = 0$ analytically. Numerical discretisation, however, introduces truncation errors that violate this constraint and can corrupt the field topology; the strategy adopted in this thesis to control such errors is described in Section 3.2.3, in the context of the PLUTO framework. An important computational aspect of the RMHD system, absent in Newtonian hydrodynamics, is that the mapping from conserved to primitive variables, $\mathbf{U} \to \mathbf{V}$, cannot be inverted analytically. Recovery of the primitive variables at each cell and each time-step requires a numerical root-finding procedure, which adds a non-trivial computational cost to every update cycle. The algorithm implemented in PLUTO for this inversion is described in Mignone et al. [131].

3.2 The PLUTO code

The numerical code used for all simulations presented in this thesis, PLUTO [131], is a freely available, modular code designed for the numerical integration of conservation laws in astrophysical fluid dynamics. PLUTO is written in C and supports parallel

execution through MPI, making it well suited for large-scale three-dimensional computations on distributed-memory clusters. In the following subsections we describe the code architecture, the Godunov-type numerical framework at the core of the solver, and the specific algorithmic choices adopted for the simulations in this thesis.

3.2.1 Code Architecture and Available Modules

PLUTO is built around a modular design in which the choice of physics, numerical algorithms, and coordinate geometry can be combined independently. The physics modules currently available cover four regimes of increasing complexity: classical hydrodynamics (HD), classical magnetohydrodynamics (MHD, ideal or resistive), special-relativistic hydrodynamics (RHD), and special-relativistic magnetohydrodynamics (RMHD). Each module defines the corresponding set of conserved variables, fluxes, and source terms, while the numerical infrastructure (reconstruction, Riemann solver, time stepping) is shared across modules and can be selected by the user at configuration time.

All simulations in this thesis use the RMHD module, which solves the system of equations presented in Section 3.1. The code supports Cartesian, cylindrical, and spherical coordinate systems, handling the curvilinear geometric source terms automatically. This flexibility is exploited in the present work, where the jet propagation through the post-merger environment is followed on a spherical grid for the first three seconds of evolution and subsequently remapped to a Cartesian grid for the extension to the ballistic regime (see Chapter 5).

3.2.2 The Godunov Framework

At the core of PLUTO lies a Godunov-type finite-volume method (see Figure 3.1 for a schematic representation). The computational domain is partitioned into cells, and the solution is represented by cell-averaged values of the conserved variables. The fundamental update equation advances the solution from time level t^n to t^{n+1} by accounting for the net flux of conserved quantities through each cell interface. In one spatial dimension, this reads

$$\mathbf{U}_i^{n+1} = \mathbf{U}_i^n + \frac{\Delta t}{\Delta x}\left[\mathbf{F}^*_{i-1/2} - \mathbf{F}^*_{i+1/2}\right], \tag{3.16}$$

where $\mathbf{F}^*_{i\pm1/2}$ are the numerical fluxes at the cell interfaces and Δx is the cell width. This is a discrete treatment of exact conservation: the change in the total amount of any conserved quantity within a cell equals the net flux entering or leaving through its boundaries. The extension to multiple dimensions in the unsplit formulation adopted by PLUTO applies the flux contributions from all directions simultaneously, avoiding the directional bias introduced by dimensional splitting.

The computation of the numerical fluxes at each interface proceeds through a three-step cycle that repeats at every time-step.

Reconstruction. The first step converts cell-averaged conserved variables into primitive variables and reconstructs left and right states at each cell interface. The

reconstruction is performed on the primitive variables $\mathbf{V}$ instead of the conserved variables $\mathbf{U}$, which produces more robust results. The accuracy of the reconstruction depends on the order of the interpolating polynomial: a piecewise constant reconstruction yields first-order accuracy but is excessively diffusive, a piecewise linear reconstruction (PLM) achieves second-order accuracy, and the piecewise parabolic method [PPM; 36] reaches third-order accuracy in smooth regions. To prevent spurious oscillations near discontinuities, the reconstruction is combined with slope limiters that enforce monotonicity. All simulations in this thesis adopt the PPM reconstruction, which provides a good balance between accuracy in smooth flows and sharpness at shocks.

Riemann solver. Given the left and right states at each interface, the numerical flux is obtained by solving (approximately) the Riemann problem that arises from the discontinuous initial data. The exact Riemann problem in RMHD involves seven characteristic waves and is computationally expensive to solve. Approximate solvers replace the exact solution with a simplified wave structure. The RMHD module of PLUTO implements several such solvers, spanning a range of numerical dissipation and computational complexity. The HLL solver [82] retains only the fastest leftward and rightward waves and collapses the intermediate wave structure into a single averaged state, providing a robust and inexpensive but relatively dissipative flux. HLLC [127] restores the contact discontinuity as an additional wave, capturing shear structures more sharply; HLLD [132] further restores the Alfvén discontinuities, resolving five of the seven RMHD waves explicitly, at the price of a more elaborate primitive-variable inversion that can become fragile in strongly magnetised intermediate states. A different strategy is followed by the GFORCE centred scheme [214, 215], which builds the intercell flux as a convex combination of fluxes with a weight tied to the local Courant number; it is more diffusive than the upwind HLL-family solvers of comparable order but does not require an eigendecomposition of the RMHD Jacobian and remains robust across degenerate wave configurations in which characteristic speeds coincide. The signal speed estimates required by these solvers are computed from the eigenvalues of the RMHD Jacobian matrix as implemented in PLUTO.

Time integration. The spatial discretisation described above reduces the system of partial differential equations to a system of ordinary differential equations (ODEs) in time, following the method of lines. The time integration is then performed using a strong-stability-preserving (SSP) Runge-Kutta scheme. The third-order SSP-RK3 scheme, adopted here, advances the solution through three stages:

$$\mathbf{U}^{*} = \mathbf{U}^{n} + \Delta t\, \mathcal{L}^{n}\,, \tag{3.17}$$

$$\mathbf{U}^{**} = \frac{1}{4}\left[3\mathbf{U}^{n} + \mathbf{U}^{*} + \Delta t\, \mathcal{L}^{*}\right]\,, \tag{3.18}$$

$$\mathbf{U}^{n+1} = \frac{1}{3}\left[\mathbf{U}^{n} + 2\mathbf{U}^{**} + 2\Delta t\, \mathcal{L}^{**}\right]\,, \tag{3.19}$$

where $\mathcal{L}$ denotes the spatial discretisation operator (obtained from the right-hand side of Eq. 3.10 after isolating the time derivative), and $\mathcal{L}^{n}, \mathcal{L}^{*}, \mathcal{L}^{**}$ are evaluated on

$\mathbf{U}^n, \mathbf{U}^*, \mathbf{U}^{**}$ respectively. Each stage requires the evaluation of the spatial operator, meaning that three sets of reconstruction and Riemann solves are performed per time-step. This additional cost is offset by the higher accuracy and the SSP property, which guarantees that the total variation of the solution does not increase across the time-step, a desirable property in the presence of shocks.

The time-step Δt is constrained by the Courant-Friedrichs-Lewy (CFL) condition,

$$\Delta t = C_{\rm CFL} \min_d \left(\frac{\Delta l^d_{\rm min}}{|\lambda^d_{\rm max}|} \right) , \tag{3.20}$$

where $\Delta l^d_{\rm min}$ is the smallest cell dimension along direction d, $|\lambda^d_{\rm max}|$ is the maximum signal speed in that direction, and $C_{\rm CFL}$ is the Courant number (< 1 for stability). A typical value of $C_{\rm CFL} = 0.3$ is used in the simulations presented in this thesis.

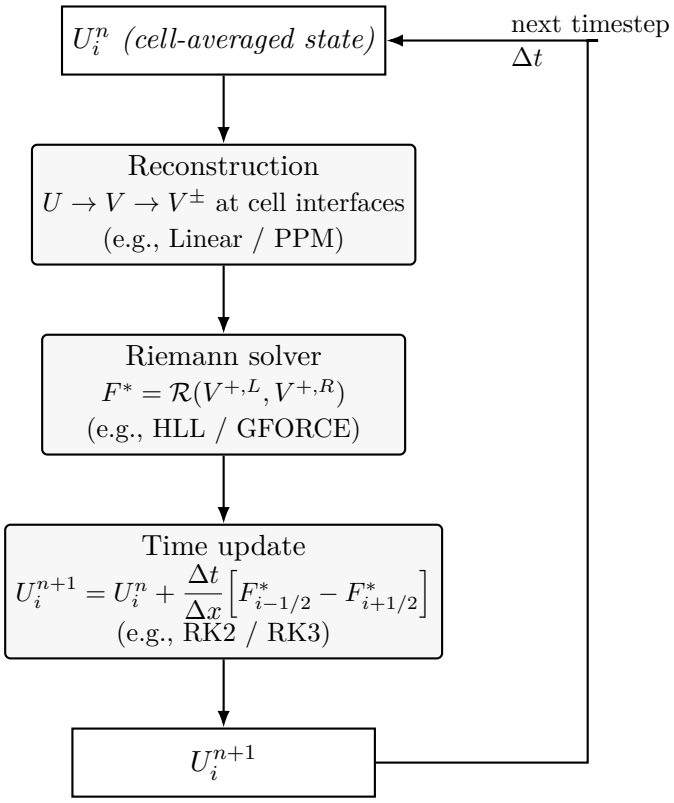


Figure 3.1: Schematic of PLUTO's algorithmic cycle.

3.2.3 The Divergence-Free Constraint and Hyperbolic Divergence Cleaning

The magnetic field must satisfy the solenoidal constraint

$$\nabla \cdot \mathbf{B} = 0 \tag{3.21}$$

at all times. This condition, which expresses the absence of magnetic monopoles, acts as an initial constraint: if it is satisfied at $t = 0$ and the induction equation (Eq. 3.15) is integrated exactly, it remains satisfied for all subsequent times. In practice, however,

finite-difference and finite-volume discretisations introduce truncation errors that can generate non-zero $\nabla \cdot \mathbf{B}$. Even small violations of the constraint produce unphysical monopole source terms in the momentum equation, leading to spurious forces parallel to $\mathbf{B}$ that can corrupt the field topology and, in magnetised relativistic flows, trigger numerical instabilities.

Several strategies have been developed to control this problem. Constrained transport methods [52] employ a staggered grid arrangement that preserves $\nabla \cdot \mathbf{B} = 0$ to machine precision, but require a fundamentally different data layout from a cell-centred scheme. Projection methods solve a Poisson equation at each time-step to subtract the divergent component, but this global elliptic solve can be expensive in three dimensions. The approach adopted in this thesis is hyperbolic divergence cleaning [42], as implemented in PLUTO by Mignone et al. [133].

The idea is to couple the induction equation to a scalar field ψ that absorbs and transports divergence errors. The modified system replaces Eq. (3.15) with

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{v} \times \mathbf{B}) + \nabla \psi = 0 \,, \tag{3.22}$$

and introduces the evolution equation

$$\frac{\partial \psi}{\partial t} + c_h^2 \, \nabla \cdot \mathbf{B} = -\frac{c_h^2}{c_p^2} \, \psi \,, \tag{3.23}$$

where c_h is the propagation speed of divergence errors and c_p controls the damping rate. In the implementation of Mignone et al. [133], c_h is set to the maximum signal speed allowed by the CFL condition, $c_h = C_{\rm CFL} \times \Delta l_{\rm min} / \Delta t^n$, while $c_p = \sqrt{\Delta l_{\rm min} \, c_h / \alpha}$ with α a free parameter for which optimal values lie in the range $0.05 \leq \alpha \leq 0.3$ and Δt the time step at the n^{th} snapshot.

Any $\nabla \cdot \mathbf{B} \neq 0$ generated during a time-step propagates away at speed c_h and is simultaneously damped exponentially. The scheme preserves the cell-centred variable arrangement, keeps the system in conservative form, and integrates seamlessly into the Godunov framework described in Section 3.2.2.

3.2.4 Equation of State

The system of conservation laws presented above must be closed by an equation of state (EoS) relating pressure, density, and internal energy.

The closure is selected at configuration time in PLUTO, independently of the physics module and of the numerical algorithms. For the classical HD and MHD modules several options are available: an isothermal law, an ideal law with constant adiabatic index, and a temperature-dependent law that accounts for the chemical composition. In the relativistic modules only two of these are implemented, namely the ideal constant-$\hat{\gamma}$ law and the Taub–Matthews law described below. The ideal law is the direct relativistic generalisation of the familiar polytropic gas, and prescribes the specific enthalpy as

$$h = 1 + \frac{\hat{\gamma}}{\hat{\gamma} - 1} \, \mathcal{T} \,, \qquad \mathcal{T} \equiv \frac{P}{\rho} \,, \tag{3.24}$$

where $\mathcal{T}$ is the dimensionless temperature and the unit term accounts for the rest-mass contribution.

The choice of EoS is not trivial in the context of relativistic jet simulations, because the plasma spans a wide range of thermodynamic regimes within the same computational domain. In the ultra-relativistic jet core, where the thermal energy per particle greatly exceeds the rest-mass energy, the effective adiabatic index approaches $\hat{\gamma} \to 4/3$. In the cold, non-relativistic post-merger ejecta and cocoon wings, the appropriate value is $\hat{\gamma} \to 5/3$. Using a constant-$\hat{\gamma}$ EoS forces the entire domain into a single thermodynamic regime and systematically misrepresents the other: $\hat{\gamma} = 5/3$ overestimates the pressure support in the hot jet spine, while $\hat{\gamma} = 4/3$ underestimates it in the surrounding cold material. The exact closure for a single-species relativistic perfect gas in thermal equilibrium is the Synge EoS [208], in which the specific enthalpy is expressed through ratios of modified Bessel functions of the temperature. This form transitions smoothly and correctly between the two limits, but its evaluation at every cell and every stage of the time integration makes it impractical in production runs.

An accurate and computationally inexpensive alternative is the Taub–Matthews (TM) equation of state [128, 130], which provides an analytic approximation to the Synge gas and departs from it by less than a few parts in a thousand across the whole temperature range. The TM EoS follows from the relativistic Taub inequality [210],

$$(h - \mathcal{T})\,(h - 4\mathcal{T}) \geq 1\,, \tag{3.25}$$

which any physically consistent relativistic perfect gas must satisfy, by adopting the equality case. Solving Eq. (3.25) for h then gives

$$h = \frac{5}{2}\,\mathcal{T} + \sqrt{\frac{9}{4}\,\mathcal{T}^2 + 1}\,, \tag{3.26}$$

The correct asymptotic behaviour is recovered in both limits: $h \to 1 + \frac{5}{2}\mathcal{T}$, that is $\hat{\gamma} = 5/3$, for $\mathcal{T} \ll 1$; and $h \to 4\mathcal{T}$, that is $\hat{\gamma} = 4/3$, for $\mathcal{T} \gg 1$. The transition between the two regimes is smooth, which is particularly important for the cocoon region where relativistic and non-relativistic material are in direct contact across the jet boundary, and where a fixed index would introduce a systematic error on one side of that boundary regardless of the value chosen.

For these reasons the TM equation of state is adopted for all the simulations in binary neutron star merger environments presented in Chapters 4–6. The preliminary simulations in the magnetorotational supernova environment of Chapter 7 adopt instead an ideal law with $\hat{\gamma} = 4/3$, for reasons of consistency with the thermodynamic closure of the core-collapse simulation that provides the initial data; the matching procedure and the tests supporting this choice are described in Section 7.4.6.

3.2.5 The Simulation Setup

The simulations in the context of BNS mergers presented in this thesis (Chapters 4-6) share a common set of algorithmic choices, summarised here for reference. The RMHD module of PLUTO is employed in all cases, with piecewise parabolic

reconstruction (PPM), third-order Runge-Kutta time stepping (SSP-RK3), and a Riemann solver that differs between the two evolution phases: the GFORCE centred scheme [214] for the spherical phase, and the HLL solver [82] for the Cartesian phase. HLL is a more diffusive choice, adopted to accommodate the steeper primitive-variable gradients produced by the increased effective resolution and by the new outer boundary condition of the Cartesian grid; we consider this acceptable given that the ballistic evolution amounts essentially to free expansion of the outflow, with no remaining jet-environment interaction to resolve accurately. The divergence-free condition on the magnetic field is maintained through the hyperbolic divergence cleaning technique described in Section 3.2.3. Thermodynamic closure is provided by the Taub-Matthews equation of state (Section 3.2.4). Newtonian gravity from the central remnant (mass $M = 2.596\,M_\odot$) is included as a source term in the momentum and energy equations.

Two distinct computational grids are used across the simulations, corresponding to two phases of the jet evolution. In the first phase, covering the jet propagation through the post-merger environment and breakout ($t \lesssim 3$ s), the evolution is carried out on a three-dimensional spherical grid (r, θ, ϕ). The geometry, resolution, and boundary conditions of this grid are described in Section 4.1.2. In the second phase, a grid-remapping technique transfers the solution to a three-dimensional Cartesian grid, enabling the evolution to continue to $t \approx 9.5$ s without loss of resolution in the expanding outflow (Section 5.2).

The specific choices of grid dimensions, resolution, boundary conditions, and the jet injection prescription differ across the models investigated in this thesis and are detailed in the chapters where each model is introduced. Reconstruction and time integration are kept identical across all simulations, so that differences in the results can be attributed to the physical parameters being varied rather than to changes in the numerical treatment.

Chapter 4

Jet Evolution in Realistic Binary Merger Environments

As outlined in Section 2.5, immediately after its launch the relativistic jet must propagate through the dense environment produced by the BNS merger. This chapter describes the setup and main results of the reference simulations following the early-time evolution of the jet through this environment, together with a systematic parameter study investigating how variations in the jet injection conditions influence the resulting dynamics. These simulations were performed prior to and independently of the present work and provide the starting point for the investigations presented in this thesis. They adopt a realistic post-merger environment obtained from a GRMHD BNS merger simulation [34, 35], following the approach originally developed by Pavan et al. [165] and extended to include magnetic fields in [P23]. The results presented in Sec. 4.2 constitute new analysis carried out as part of the present work, whereas the remainder of the chapter presents the reference simulations and their original analysis.

4.1 Simulation Setup

What follows is a brief description of the numerical and physical setup developed by Pavan et al. [165] and [P23], which is relevant for interpreting the results presented in this thesis. The reader is referred to the cited works for the full description.

4.1.1 The Post-Merger Environment

Reference merger simulation

The post-merger environment is drawn from the GRMHD simulation of Ciolfi [34], which models a merging BNS system with chirp mass matching the inferred value for GW170817 and a mass ratio $q \simeq 0.9$. The neutron star matter is described using a piecewise-polytropic approximation of the APR4 equation of state [5]. The two stars are initialised with a purely poloidal magnetic field reaching up to 5×10^{15} G, a value higher than typical isolated neutron stars, but necessary in global simulations to compensate for unresolved small-scale amplification mechanisms [e.g., 92]. Overall,

in BNS mergers, different mechanisms amplify the magnetic fields, including the Kelvin–Helmholtz instability, magnetorotational instability, and magnetic winding. Together, these mechanisms can increase the total magnetic energy up to $E_{\rm mag} \sim 10^{51}$ erg [35].

The outcome of the merger is a highly magnetised and differentially rotating neutron star of mass $M \simeq 2.596\, M_\odot$, surrounded by a complex, asymmetric distribution of dynamical ejecta and post-merger winds. A notable result of the simulation is the spontaneous formation of a magnetically driven, mildly collimated outflow along the remnant spin axis [34], powered by the differential rotation of the remnant itself. However, this outflow is still heavily baryonic, with Lorentz factors $\Gamma \sim 1.05$, too low for a canonical sGRB jet. This is consistent with the baryon loading problem discussed in Section 2.3.4: a canonical relativistic jet requires a cleaner polar region which can likely emerge only after a black hole and accretion disk have formed. This is however still debated: whether a long-lived, differentially rotating magnetar remnant could instead sustain the required Lorentz factor before eventual collapse depends on the competition between its rotational support and the neutrino-driven baryon pollution of the polar funnel [113, 125].

Since the collapse to a black hole does not occur within the ~ 255 ms temporal window covered by the reference merger simulation, in the setup of P23 the collapse time is treated as a free parameter. Specifically, using an appropriate extrapolation procedure (see below), a collapse time of $t_c = 355$ ms after merger is adopted, with a further delay of 30 ms between the collapse and the onset of jet injection, giving $t_{\rm jet} = 385$ ms. The physical reasoning is that the formation of the accretion disk, the re-organisation of the magnetic field, and the clearing of the polar region all require a finite time after the collapse itself.

Importing the environment into PLUTO

The merger simulation data in terms of the rest-mass density ρ, pressure P, velocity $\mathbf{v}$, and magnetic field $\mathbf{B}$ are imported into the PLUTO code [131], constituting what we refer to as the *initial conditions* for our jet simulations. Specifically, the import involves four key steps:

i) The merger data, lying on a hierarchical 3D Cartesian grid, are first remapped onto a uniform Cartesian grid with resolution the same as the original grid at 380 km radial distance, and then interpolated onto a suitable 3D spherical grid (Section 4.1.2);

ii) A 90° rotation of the coordinate system is applied to the imported data. In the original merger simulation, the BNS orbital axis, which coincides with the remnant's spin axis, is aligned with the vertical z-axis of the Cartesian grid. In PLUTO, however, the z-axis corresponds to the polar axis of the spherical coordinate system (r, θ, ϕ), which has coordinate singularities at $\theta = 0$ and $\theta = \pi$. Injecting a relativistic jet along this axis would therefore introduce numerical artefacts precisely in the region of highest physical interest. The 90° rotation instead realigns the remnant spin axis with the Cartesian y-axis, which is free of coordinate singularities, allowing the jet to be injected away from the singular regions of the spherical grid;

iii) An excision radius of $r_{\rm exc} = 380$ km is applied, removing the central region where general-relativistic corrections not implemented in PLUTO would be required. At this distance from the remnant, the Newtonian approximation is well justified

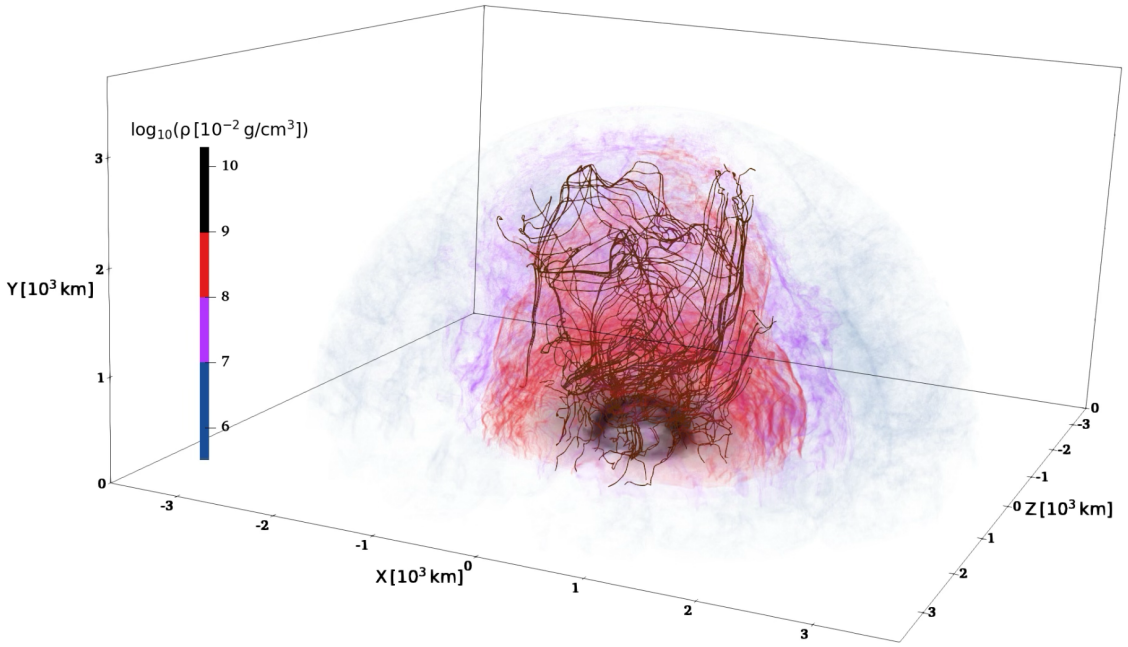


Figure 4.1: 3D rendering of the realistic post-merger environment imported in P23. The mass density distribution and magnetic field lines at 155 ms after merger are shown. Adapted from P23.

for the jet evolution we study. The excision surface then serves as the inner radial boundary through which the jet is manually injected;

iv) The constant density floor $\rho_{\rm floor} = 6.3 \times 10^4$ g cm^{-3} imposed in the merger simulation for numerical stability, corresponding to a total atmospheric mass of $3.5 \times 10^{-3}\, M_\odot$ within 3×10^3 km, is replaced by a more tailored prescription. Within $r < 2000$ km the original floor value is retained; beyond this radius it transitions to a profile decreasing as $r^{-6.5}$, representing a static, non-magnetised atmosphere. This choice prevents the jet from having to sweep up an artificially large baryonic reservoir at large radii, while remaining consistent with the inner density structure of the merger simulation.

Figure 4.1 shows a 3D rendering of the imported environment at 155 ms after merger, in terms of its mass density distribution and magnetic field lines.

Time evolution of the environment

After the import, the post-merger environment is not held fixed: it continues to evolve in PLUTO with radial boundary conditions at the excision surface set to reproduce the angle-averaged time evolution of all physical quantities as extrapolated from the merger simulation data [166]. This allows the full ~ 255 ms of available data to be used, after which the environment is evolved for a further 100 ms until the assumed collapse time (t_c). The result is a dynamically evolving medium that retains the complex angular structure and density stratification of the actual BNS merger aftermath. This environment, and the physical setup of the jet injection, are the starting point for all the simulations described in the chapters that follow.

Gravity and the collapse prescription

Two effects are included in the PLUTO simulations that are frequently omitted in the literature, despite their non-negligible impact on the jet dynamics. The first is the Newtonian gravitational pull of the central remnant, added to the equations of motion as

$$\mathbf{g} = -\frac{GM_0}{r^2}\,\hat{\mathbf{r}}\,, \tag{4.1}$$

with $M_0 = 2.596\,M_\odot$ the gravitational mass of the BNS merger remnant. This is a Newtonian approximation to the full general-relativistic problem, but it captures the dominant dynamical effects at the scales of interest here: e.g., the fallback of slow ejecta towards the excision boundary, which affects the pressure balance at the base of the jet. Neglecting this term leaves the surrounding material in artificial outward expansion, which suppresses Kelvin–Helmholtz mixing and overestimates both jet collimation and energy retention.

The second effect concerns the transition to a black hole. Since the collapse itself does not occur within the temporal window of the reference merger simulation, in PLUTO it is treated artificially, through an additional acceleration term that progressively removes the radial pressure support at the excision boundary. This reproduces the loss of pressure gradients associated with the collapse and the onset of disc accretion, and is implemented via an effective mass parameter (coupled to M_0) that decays exponentially after the collapse time t_c:

$$M_{\rm eff}(t) = M_{\rm eff}(t_c)\,\exp\!\left(-\frac{t-t_c}{\tau}\right)\,, \tag{4.2}$$

where τ combines the black hole–disc accretion timescale $\tau_d \simeq 0.3\,\mathrm{s}$ and the jet launching delay $\tau_j = 30\,\mathrm{ms}$ (relative to the collapse). This provides a smooth and physically motivated transition between the neutron-star-supported ejecta configuration and the post-collapse conditions, after which the jet is injected.

What changes with a realistic environment

Pavan et al. [165] showed that the choice of environment affects the jet dynamics not only quantitatively, but at the level of the qualitative outcome. Idealised setups typically assume an isotropic post-merger medium, with power-law density and pressure profiles and homologous radial expansion, which is analytically convenient but, as described in Section 2.6.2, does not capture relevant structural features of the actual BNS merger aftermath (see Figure 4.1). When a realistic environment is adopted, a low density funnel offers a preferential path for jet propagation, and the outflow breaks out[1] earlier and with a higher terminal Lorentz factor than in the isotropic environment case. The resulting jet structure, both its angular energy distribution and its degree of asymmetry, also differs qualitatively, with the realistic case yielding a more deflected and turbulent outflow, consistent with the genuinely three-dimensional nature of the environment.

[1]The *break-out* is the moment in which the jet front successfully pierces the boundary of the post-merger environment.

The role of magnetic fields

Magnetic fields influence the jet dynamics in two complementary ways. First, magnetic pressure gradients provide an additional outward force that accelerates the jet. Second, the toroidal magnetic field exerts a hoop stress which, together with the thermal pressure of the cocoon, helps confine the jet laterally. This magnetic tension also suppresses Kelvin–Helmholtz instabilities at the jet–cocoon interface, which would otherwise promote episodic baryon entrainment and intermittently decelerate the jet core. The combined effect is a faster, more stable, and better collimated relativistic jet, as will be shown in Section 4.3.2.

4.1.2 Computational Grid

Up to $t - t_j \simeq 3$ s after jet launching, the evolution is carried out over a three-dimensional spherical grid (r, θ, ϕ). The radial direction extends from an inner excision radius $r_{\rm exc} = 380$ km to an outer boundary $r_{\rm max} = 2.5 \times 10^6$ km. The polar angle is limited to $\theta \in [0.1, \pi - 0.1]$, while the azimuthal angle spans the full range $\phi \in [0, 2\pi]$.

The grid resolution is $768 \times 256 \times 512$ cells along r, θ, and ϕ, respectively, with uniform angular spacing and logarithmic radial spacing. The latter choice reflects where the relevant small-scale physics actually occurs, close to the inner boundary at early times. The logarithmic grid concentrates resolution there, yielding a radial spacing of $\simeq 4.4$ km at $r_{\rm exc}$, matching the resolution of the original Cartesian GRMHD merger data, while coarsening smoothly at larger radii to keep the total cell count computationally tractable.

As mentioned in Section 4.1.1, user-defined radial boundary conditions are imposed at $r_{\rm exc}$. Specifically, from 155 ms after merger, the three-dimensional angular distribution of each variable at $r_{\rm exc}$ is preserved and rescaled over time according to its angle-averaged evolution in the reference GRMHD simulation, thereby allowing the anisotropic BNS merger environment to feed self-consistently into the computational domain up to the adopted collapse time of 355 ms after merger. After collapse, the same excision surface serves as the jet injection boundary, with the hydrodynamic and magnetic field variables overwritten within the injection cone according to the jet prescription described in Section 4.1.3.

At $r = r_{\rm max}$, "outflow" boundary conditions are applied: all primary variables are extrapolated with a zero radial gradient ($\partial_r X = 0$), so that the expanding environment and the eventual blast wave can exit the domain without generating spurious reflections.

Along ϕ, periodic boundary conditions are imposed. At the polar boundaries $\theta = 0.1$ and $\theta = \pi - 0.1$, outflow boundary conditions are adopted, handling the excluded axis consistently with the surrounding flow and avoiding any unphysical accumulation or loss of mass and energy at the grid edges.

4.1.3 Jet Injection Prescription

The jet is injected at the inner radial boundary $r = r_{\rm exc}$ using a prescription similar to that of Geng et al. [66] (see also Martí [115]), which describes an axisymmetric, mag-

netised, and uniformly rotating jet in transverse equilibrium. Notably, axisymmetry at the injection boundary allows the jet profiles to be specified without azimuthal dependence, but is not enforced during the subsequent evolution: propagation through the non-uniform, three-dimensional environment naturally breaks this symmetry.

We remark that the jet is injected along the y-axis, which is orthogonal to the spherical polar axis and aligned with the remnant spin axis following the 90° rotation described in Section 4.1.1. To avoid confusion, in the following we denote by Θ and Φ the polar and azimuthal angles, respectively, measured with respect to the jet injection axis. The jet profiles are then defined in terms of these angles, which differ from the standard θ and ϕ angles of the spherical coordinate system.

Eight angular functions specify the jet at injection: the comoving-frame density ρ_j and pressure P_j, the three-velocity components v^r, v^Θ, and v^Φ, and the laboratory-frame magnetic field components B^r, B^Θ, and B^Φ. These functions are defined only within the jet injection cone, characterised by a half-opening angle $\Theta_j = 10°$. In particular, the polar components of the velocity and magnetic field are taken to be negligible compared to the radial and azimuthal components, such that $v^\Theta = B^\Theta = 0$. For the remaining components, uniform radial velocity and rigid rotation are adopted, namely

$$v^r = c\sqrt{1 - \frac{1}{\Gamma_j^2}}\,, \tag{4.3}$$

$$v^\Phi(\Theta) = \bar{\Omega}\, r_{\rm exc} \sin\Theta\,, \tag{4.4}$$

where $\Gamma_j = 3$ is the initial bulk Lorentz factor, and $\bar{\Omega}$ is the mean angular velocity of the surrounding environment evaluated at the jet edges (i.e. the jet is assumed to co-rotate with its surroundings at injection).[2] The magnetic field structure then follows the prescription of Geng et al. [66] (see also Martí [115]), that is

$$B^\Phi = \frac{2B^\Phi_{j,m}(\Theta/\Theta_{j,m})}{1 + (\Theta/\Theta_{j,m})^2}\,, \tag{4.5}$$

$$B^r = B_{\rm ratio}\, B^\Phi_{j,m}\,, \tag{4.6}$$

with $B^\Phi_{j,m}$ being the maximum strength of the toroidal magnetic field (which varies across the models depending on the magnetisation level; see below), $\Theta_{j,m} = 0.4\,\Theta_j$, and $B_{\rm ratio} = 0.5$. The functional form rises linearly from the injection axis, peaks at $\Theta_{j,m}$, and falls off as $1/\Theta$ beyond it. The adopted prescription therefore provides a smooth toroidal field profile that vanishes on the jet axis, as required for regularity, reaches a maximum at a finite angular distance, and subsequently decreases towards the jet boundary, where the surrounding environment is expected to be substantially

[2] In a self-consistent model, conservation of angular momentum would imply that the jet co-rotates with the central engine, e.g. the black hole–accretion disc system. Such a system is expected to rotate substantially faster than the surrounding environment at $r_{\rm exc} = 380\,$km. In the present setup, however, the black hole–accretion disc system is not directly formed in the reference merger simulation, and its rotation rate is therefore not available. Assuming co-rotation between the jet and its surroundings at injection provides a physical prescription consistent with the information imported into PLUTO.

less magnetised than the jet. Moreover, from Ampère's law, the adopted prescription implies uniform axial current density for $\Theta < \Theta_{j,m}$, while the outer region is approximately current-free (see Martí [115]).

Finally, the comoving-frame density and pressure profiles are determined by the transverse equilibrium condition along Θ, together with the specification of the jet's two-sided luminosity and magnetisation level at $r_{\rm exc}$ (see [P23] for the full derivation). Specifically, the two-sided luminosity is defined as

$$L_j = 2\int_0^{\Theta_j}\int_0^{2\pi} T^{0r}(r_{\rm exc},\Theta,\Phi)\, c\, r_{\rm exc}^2 \sin\Theta {\rm d}\Theta{\rm d}\Phi, \tag{4.7}$$

where

$$T^{0r} = \rho h^* \Gamma_j^2 c^2 \beta^r - \Gamma(\boldsymbol{\beta}\cdot\boldsymbol{B})\left(\frac{B^r}{\Gamma} + \beta^r(\boldsymbol{\beta}\cdot\boldsymbol{B})\right) \tag{4.8}$$

is the time-radius component of the energy-momentum tensor, corresponding to the radial energy flux, with $\boldsymbol{\beta} = \boldsymbol{v}/c$ and h^* being the total relativistic enthalpy (including magnetic field contribution). Here, $h^* \equiv \Gamma_\infty/\Gamma_j$ where $\Gamma_\infty = 300$ is the assumed terminal Lorentz factor of the jet. The magnetisation level is then given by[3]

$$\Sigma_j = \frac{L_j - L_{j,\rm HD}}{L_j}, \tag{4.9}$$

where $L_{j,\rm HD}$ is the purely hydrodynamic luminosity obtained without the magnetic field contribution, and corresponds to 8% for the simulation presented in P23.

As in the previous setup, the jet luminosity is imposed to decay exponentially over time, $L(t) = L_j\, e^{-t/\tau_d}$, where τ_d is the decay timescale (the same as the accretion mechanism described in Section 4.1.1). Keeping ρ and v^Φ fixed in time, the remaining quantities evolve according to

$$v^r(t) = \frac{\Gamma_j v^r e^{-t/2\tau_d}}{\Gamma(t)}, \tag{4.10}$$

$$B^\Phi(t) = B^\Phi\sqrt{\frac{v^r}{v^r(t)}}\, e^{-t/2\tau_d}, \tag{4.11}$$

$$B^r(t) = B^r\sqrt{\frac{v^r}{v^r(t)}}\, e^{-t/2\tau_d}, \tag{4.12}$$

$$P(t) = \frac{1}{4}\left[\rho\,(h^*(t)-1) - b^2(t)\right], \tag{4.13}$$

where

$$\Gamma(t) = \sqrt{1+\left(\Gamma_j v^r e^{-t/2\tau_d}\right)^2}, \tag{4.14}$$

$$h^*(t) = \frac{h^*\Gamma_j e^{-t/2\tau_d}}{\Gamma(t)}, \tag{4.15}$$

and

$$b^2 = \frac{B^2}{\Gamma^2} + (\boldsymbol{\beta}\cdot\boldsymbol{B})^2 \tag{4.16}$$

denotes the magnetic field strength squared measured in the comoving-frame.

[3]Note that this differs from the standard definition of magnetisation, σ, being the local ratio of electromagnetic to total hydrodynamic energy flux (see, e.g., Chapter 7).

Table 4.1: Injection parameters of the seven jet models. L_j is the initial injected luminosity, τ_d the exponential decay timescale of the luminosity, Σ_j the fractional magnetic contribution to the luminosity, and t_j the jet launching time after merger. The rightmost column summarises the main qualitative outcome of each model relative to the fiducial case F, discussed in detail in Sect. 4.3.1–4.3.3.

Model	L_j [erg/s]	τ_d [ms]	Σ_j [%]	t_j [ms]	Outcome
F	5.37×10^{51}	300	1.20	385	Reference model
a	8.95×10^{51}	300	1.20	385	Faster breakout, higher $\Gamma_{\max}$
b	5.37×10^{51}	500	1.20	385	Same L_j as F, higher total energy
c	5.37×10^{51}	300	1.56	385	Enhanced magnetic field preservation
d	4.12×10^{51}	300	1.56	385	Reduced L_j, less efficient drilling
e1	3.69×10^{51}	300	0.46	385	Severe jet–environment disruption
e2	3.69×10^{51}	300	0.46	185	Earlier launch, reduced disruption

4.1.4 Model Suite

To investigate how the jet's injection properties shape its subsequent evolution, a suite of seven models is then considered, summarised in Table 4.1. Each model shares the same reference BNS merger simulation, import procedure, and injection prescription as described in Sections 4.1.1–4.1.3, and differs only in the values of the jet injection luminosity L_j, decay timescale τ_d, magnetisation level Σ_j, and launching time t_j relative to merger.

Model F is the reference against which all other models are compared. Using this configuration as our baseline allows us to treat each departure from F as a controlled variation of the validated setup.

The remaining six models are organised in three pairs, each isolating one injection parameter at a time relative to F: models a and b probe the effect of jet luminosity and decay timescale (Sect. 4.3.1); models c and d probe the effect of magnetisation, at fixed and reduced luminosity respectively (Sect. 4.3.2); models *e1* and *e2* probe the effect of the launching time, at fixed low luminosity and magnetisation (Sect. 4.3.3). This choice allows us, in the analysis that follows, to attribute differences in jet dynamics and final structure to different injection conditions. We remark however that all these simulations are performed at a single resolution, which, as shown in Appendix B, can affect the jet dynamics and the resulting structure, in some cases also at a qualitative level. The comparisons among models presented here are indicative of the behaviour expected from the variations in the initial conditions, while the smallest differences between them cannot be firmly disentangled from numerical resolution effects.

4.2 Early Dynamics: Jet–Ejecta Interaction (0–3 s)

4.2.1 Jet Propagation and Cocoon Formation

As described in Section 2.5, immediately after injection the incipient jet propagates preferentially through regions of lower density in the post-merger environment. The jet front forms a shock, and a high-pressure cocoon, consisting of mixed jet–ejecta material deflected by the jet front, builds up around the jet edges, keeping it collimated. At this stage the jet is still accelerating, almost immediately converting thermal and magnetic energy into kinetic form. Part of this energy, however, is continuously lost to the environment while drilling through the ejecta (as detailed in Section 4.2.3).

The dynamical stages of this early propagation are illustrated by the 3D renderings in Fig. 4.2, reproduced from P23 and referring to their reference simulation. Shortly after injection (leftmost rendering), the jet is still deeply embedded in the ejecta and a dense shell of shocked material envelops the outflow. Near the base, this shell collimates the jet efficiently enough to sustain acceleration up to Lorentz factors of a few, while farther out the encounter with the inhomogeneous ejecta forces the front to travel along paths of lower density, breaking the axisymmetry imposed at injection and producing an irregular head morphology. Up to breakout, the jet–cocoon interface remains remarkably stable, an outcome plausibly linked to the presence of the toroidal magnetic field, which suppresses turbulent motions that develop at the interface in otherwise analogous non-magnetised setups [165].

As long as it remains confined within the dense post-merger environment, the jet stays collimated with an effective opening angle $\sim 10°$. The small-scale effects due to the intrinsically anisotropic inner environment are amplified at larger scales, ultimately shaping the final jet structure.

Across the full simulation suite, the jet is powerful enough not to be completely choked by the environment: within a time range of 200–350 ms from launching, all the outflows successfully break out from the environment. As further discussed in Section 4.3, the breakout time is sensitive to the jet injection parameters, primarily through their influence on how easily the jet traverses the environment.

Moving to model F of the parameter study as a reference, Figure 4.3 shows the distinction between the various elements of the system at breakout: the narrow jet core (1), the cocoon surrounding it and keeping it collimated (2), the inner and outer merger environment (3)–(4), composed respectively of material launched by neutron star winds after the merger, due to the accumulation of thermal and magnetic energy, and stellar material expelled or stripped away during the merger itself; and the low-density artificial atmosphere that surrogates the underlying ISM, imposed for numerical stability (5). Once the jet emerges into the outer merger environment, the exponential decay of the injected luminosity becomes dynamically relevant (rightmost renderings of Fig. 4.2): the jet base gradually detaches from the surrounding shell, while the front accelerates and expands laterally, developing a markedly asymmetric profile that reflects the imprint of the earlier interaction with the inhomogeneous ejecta. By the end of the spherical-grid evolution, the detachment is complete and

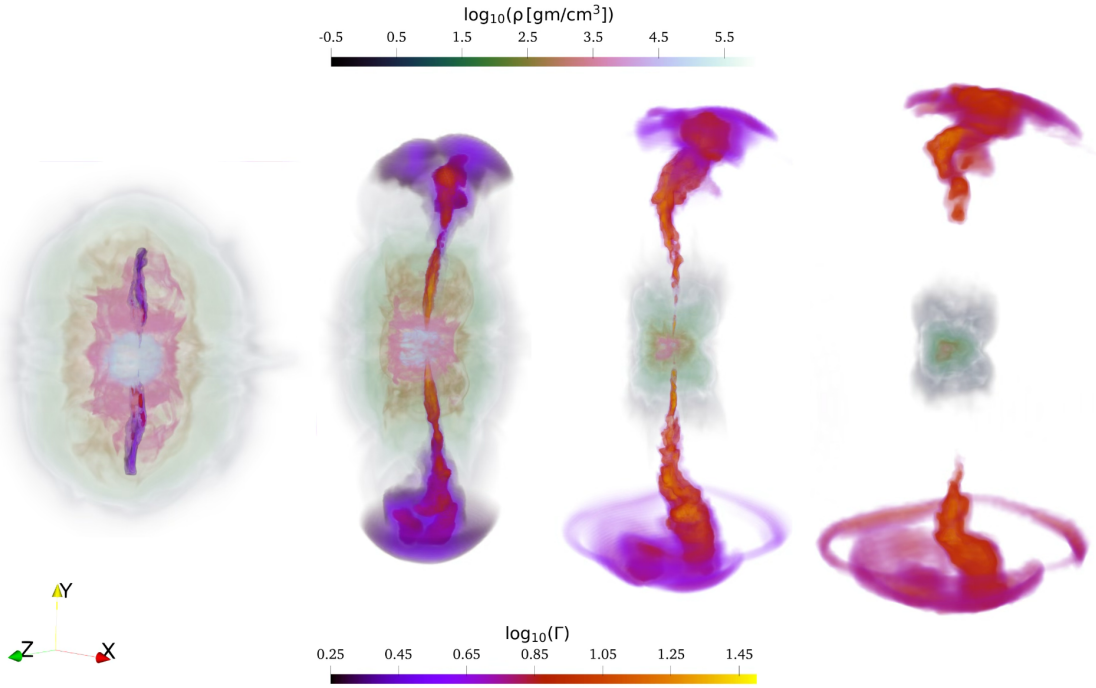


Figure 4.2: 3D renderings of the system configuration at 560 ms, 735 ms and 2.385 s after merger, in terms of rest-mass density and Lorentz factor, spatially rescaled for comparison. Adapted from P23.

the jet has organised into a more compact, ultra-relativistic structure trailed by a slower tail, providing the initial state for the extended evolution discussed in the next Section.

4.2.2 The Multi-Shock Structure at Breakout

As the jet front breaks out of the dense inner environment and propagates into the more rarefied outer region, the sharp decrease in ambient density leads to a rapid acceleration of the outermost jet material. This acceleration in turn launches a secondary forward shock into the surrounding medium, giving rise to a stratified, multi-layered shock structure. At the front, a nearly spherical outer layer moves outward into the surrounding medium, exhibiting a roughly isotropic distribution of compressed fluid whose energy content is considerably lower than the jet itself, but significant relative to the cocoon. The flow is thus divided into causally disconnected regions with different density, velocity and thermodynamic properties. Figure 4.4 illustrates this feature by superimposing contours that delineate the transition of the flow from the subsonic to the supersonic regime (white lines), thereby identifying the locations of shock fronts, on a meridional cross-section of the energy distribution at $t - t_{\rm jet} = 3$ s. This complex structure can significantly influence the subsequent radiative signatures, especially for off-axis observers (see Chapter 6). Notably, this feature is present in all our jet simulations, with slightly different properties related to the differences in jet evolution.

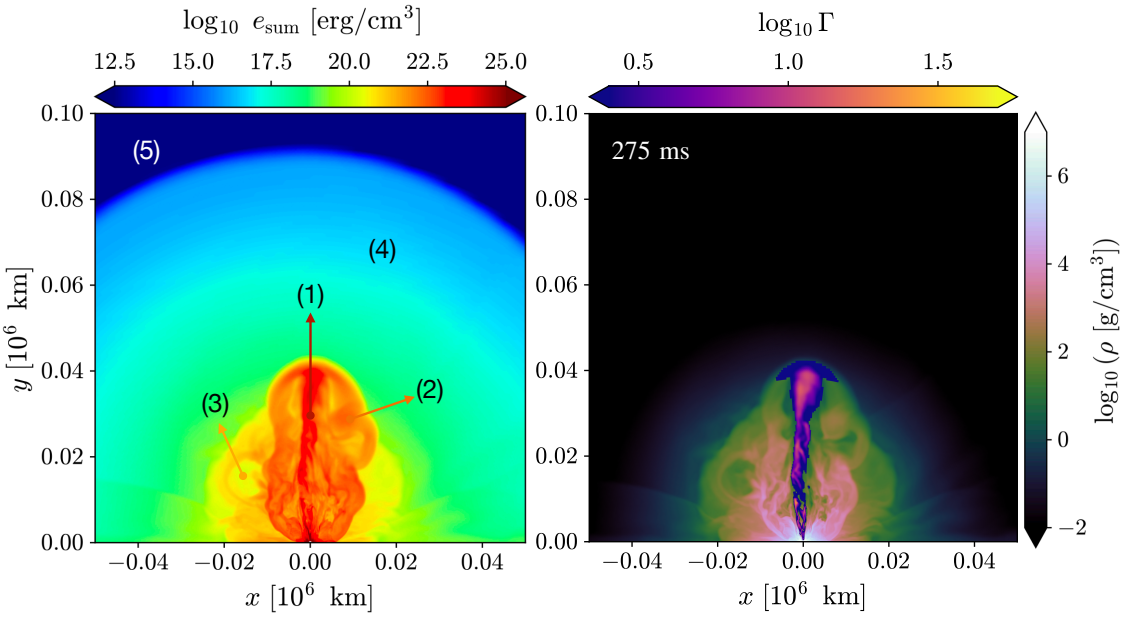


Figure 4.3: Meridional view of the energy density distribution excluding rest-mass component (left), and rest-mass density with superimposed Lorentz factor of the elements with $\Gamma > 2$ (right), at $t - t_j = 275$ ms, relative to model F (Table 4.1).

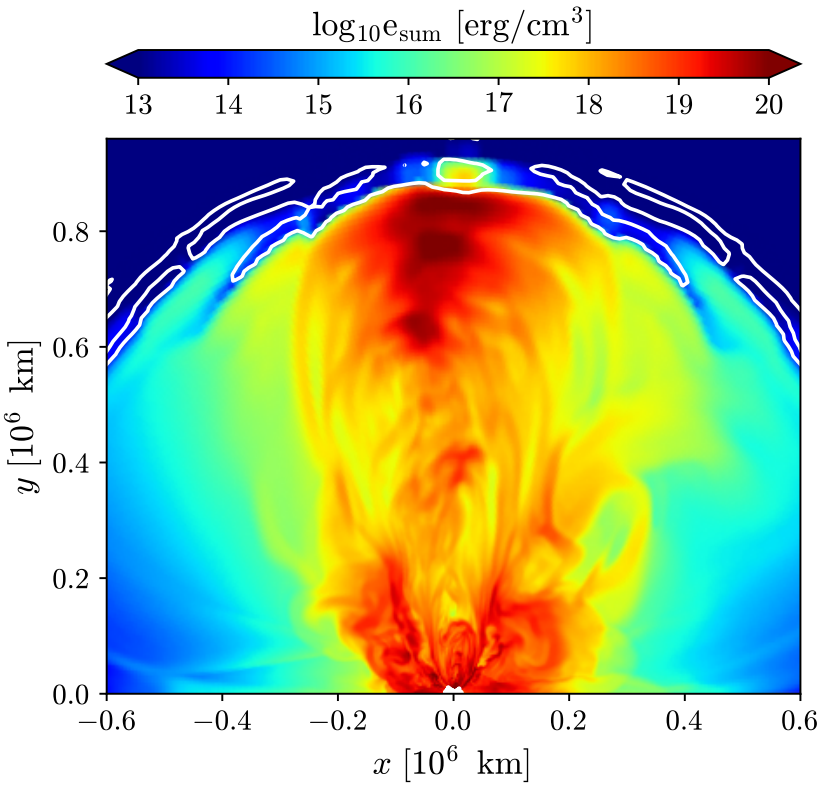


Figure 4.4: Meridional view of the energy density without rest mass component at $t - t_j = 3$ s for model F. White contours represent shocked surfaces.

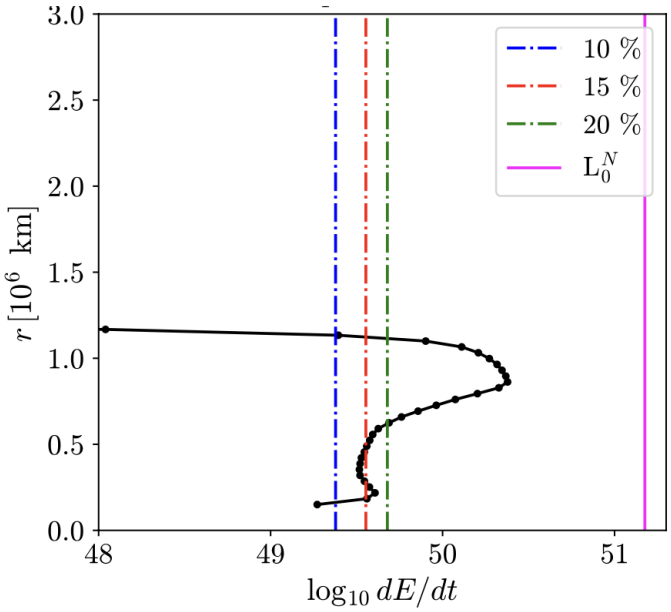


Figure 4.5: Kinetic power computed at different radii at $t - t_j = 4.49\,\mathrm{s}$ for the jet model of P23 as described in D25. The blue, red, and green dot-dashed curves indicate 10%, 15%, and 20% of the peak dE/dt, respectively. The solid magenta curve represents the initial jet luminosity at injection (north side only).

4.2.3 Energy Distribution

After breakout, the back layers injected at later times catch up with the front, and the most energetically relevant jet material gradually accumulates into a spherical-cap-shaped region that we refer to as the *jet cap*.

We define this region quantitatively here, since it is used extensively in analysing the angular distribution of the jet material, to identify a cohesive region representing the jet 'bulk'.

We compute the luminosity (dE/dt) crossing different radial surfaces of the computational domain, by evaluating the energy difference above each surface between two timesteps (i.e., for t', we compute the energy at $t' - dt$ and $t' + dt$, with dt being the time between consecutive outputs). We then identify the jet cap as the radial region where dE/dt exceeds 15% of its maximum, defined by the intersections between the luminosity curve and the line representing the fixed threshold (see Figure 4.5, referring to the simulation of P23 as described in D25). The Figure also shows that the radial limits are only weakly sensitive to moderate variations in this threshold (10%-20%). The shape of the curve reflects the behaviour of the jet injection, with an exponential decay after the initial peak. For comparison, the line labelled by L_0^N represents the maximum injected luminosity.

Following breakout, the reduced confinement from the surrounding medium allows the jet to widen significantly, reaching an aperture of $\gtrsim 20°$.

Figure 4.6, relative to the P23 simulation, shows how the energy contained in the jet cap is distributed among outflow components with different velocities. Specifically, we plot the cumulative energy,

$$E_{\mathrm{sum}}(> \beta\Gamma),$$

defined as the energy (kinetic + thermal + magnetic, without rest mass component, $E_{\rm sum}$) associated with fluid elements having $\beta\Gamma$ greater than a given threshold, as a function of that threshold. The calculation is performed by considering only the material enclosed within different angular distances from the jet injection axis. We compare four values of the half-opening angle, namely 10°, 20°, 30°, and 40°.

The profiles obtained for 10° and 20° differ by approximately 30%, whereas the discrepancy between the 20° and 30° cases is roughly one order of magnitude smaller. The profiles corresponding to 30° and 40° differ by less than 1%. These results indicate that nearly all of the jet energy is confined within 20° of the injection axis.

Most of the energy is carried by fluid elements with $\beta\Gamma \gtrsim 3$. In addition, the fastest component is entirely contained within a half-opening angle smaller than 10°, as indicated by the convergence of all profiles for $\beta\Gamma \gtrsim 10$.

Finally, we note that the Lorentz factor of the energetically dominant portion of the jet does not exceed $\Gamma \sim 20$ in this simulation, substantially below the terminal value imposed at injection, $\Gamma_\infty = 300$. The latter, however, refers to the terminal value for an individual fluid element, corresponding to the maximum Γ attainable if the element were to expand freely into vacuum and convert all of the initial internal energy into kinetic energy. It is therefore not directly comparable to the Lorentz factor characterising the jet as a whole during its propagation through the surrounding medium. The relatively low Lorentz factors reached by the energetically dominant jet material instead reflect the substantial interaction with, and energy dissipation within, the surrounding medium.

Another way of assessing the impact of this interaction is to compare the energy left in the jet cap after breakout with the total injected energy, giving an indication of how much energy has been transferred to the cocoon. For the reference simulation of P23, we find that at 3 s after launch, the jet cap contains $E_{\rm cap} \simeq 2.6 \times 10^{50}$ erg of total energy without rest–mass component, roughly 56% of the total injected energy ($E_{\rm inj} \simeq 4.6 \times 10^{50}$ erg). A further upper limit to the energy transferred during the early drilling stage is obtained by computing the energy contained within an aperture angle of 30°, which safely encloses all the jet–cocoon material, and comparing it with the injected jet energy. This comparison is only indicative, since the outflow after breakout is composed of a mixture of jet and environment material, and the gravitational attraction from the inner boundary also contributes to the energy budget. This value is found to be roughly 20% lower than the injected energy, providing additional evidence that, in this simulation, a significant fraction of the jet energy is transferred to the surroundings.

As mentioned in the previous section, part of the injected energy is also transferred to the outer shell. In section 6.1.3 we analyse the angular and radial distribution of energy for all our simulation suite and discuss the contribution of this element.

4.3 Parameter Study: Impact of Injection Conditions

Prescribing the jet at injection is less realistic than launching it self-consistently, but it allows us to vary the injection parameters one at a time and study their effect on

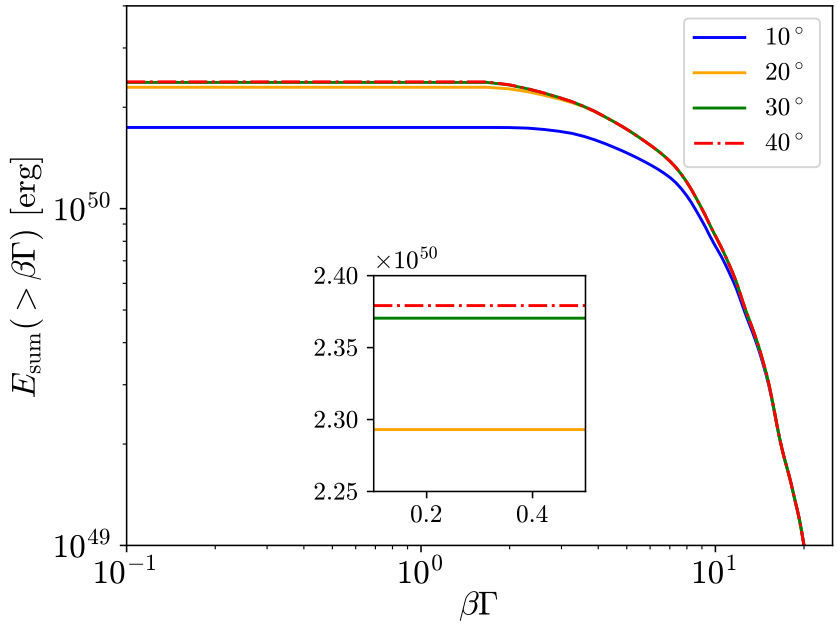


Figure 4.6: Energy $E_{\rm sum}$(without rest-mass energy contribution) of the fluid elements of the jet cap with $\beta\Gamma$ larger than a threshold, as a function of the threshold. Analysis referred to the model of P23 as presented in D25.

the jet dynamics, stability, acceleration and final structure.[4]

As described in Section 4.1.4, model F is the reference configuration against which the variations of the other models are defined. It describes a jet injected at $t_{\rm j} = 385\,{\rm ms}$ after merger with an initial two-sided luminosity $L_{\rm j} = 5.37 \times 10^{51}\,{\rm erg\,s^{-1}}$, decaying exponentially on a timescale $\tau_{\rm d} = 300\,{\rm ms}$, and a magnetic contribution to the injected luminosity of $\Sigma_{\rm j} = 1.2\%$. Under these conditions the jet drills through the post-merger ejecta and breaks out of the merger environment at $t - t_{\rm j} \simeq 275\,{\rm ms}$, after which it expands into the low-density outer atmosphere and continues to accelerate.

4.3.1 Luminosity and Decay Timescale

The jet luminosity L_j and the injection decay timescale τ_d are the two parameters most directly connected to the central engine: L_j reflects the instantaneous power extracted by the Blandford–Znajek mechanism, while τ_d is the timescale over which the accretion disc loses mass; in our case they are both fixed. These two quantities affect the jet evolution through physically distinct channels. From the jet head velocity analysis of Bromberg et al. [26], the drilling efficiency depends on the instantaneous thrust, which scales with L_j through $\tilde{L} \propto L_j/(\rho_a c^3 \theta_j^2 r^2)$. Increasing L_j at fixed τ_d therefore accelerates breakout, while increasing τ_d at fixed L_j extends the energy injection without changing significantly the pre-breakout dynamics.

[4]We refer to Appendix B for a discussion on the dependence on numerical resolution. We further note that the seven models considered probe only a restricted region of the injection-parameter space in the vicinity of the fiducial configuration, and that a comprehensive exploration of the full parameter space is beyond the scope of the present study.

Models a and b are designed to test exactly this distinction while keeping the total injected energy approximately equal. Model a increases L_j to $\approx 170\%$ of the fiducial value at fixed $\tau_d = 300$ ms; model b extends τ_d to 500 ms at the same fiducial luminosity, reaching a comparable total energy budget over the full simulation duration. The contrast between their early-time evolutions is instructive. Model a achieves breakout at $t_{\rm BO} \approx 200$ ms, roughly 75 ms earlier than the reference model F, driven by the enhanced drilling efficiency. Model b, sharing the fiducial luminosity profile in the early phase, follows a nearly identical initial trajectory and reaches breakout at a comparable time to the reference model. After breakout, however, both models retain a significantly larger fraction ($\approx 60\%$) of the injected energy in the jet cap compared to F, and reach higher local Lorentz factors, $\Gamma_{\rm max} \approx 85$ and 102 for models a and b respectively, at $t - t_j = 3$ s, vs. $\Gamma_{\rm max} \approx 56$ for model F[5].

The practical consequence for the angular structure of the jet, and therefore for the afterglow predictions of Chapter 6, is that both higher luminosity and longer injection duration produce a more energetic, better accelerated outflow, but for different reasons: model a benefits from a more efficient drilling phase, while model b accumulates more energy at late times when the environment is less dense.

4.3.2 Magnetisation

The magnetic field influences jet propagation in two main ways. First, as part of the Poynting flux, it transports energy along the jet axis. As a tension force, it resists the lateral perturbations that seed Kelvin–Helmholtz (KH) instabilities at the shear layer between the fast jet spine and the slower cocoon. The expectation from ideal RMHD is therefore that a higher magnetic field contribution Σ_j stabilises the jet–cocoon interface while marginally modifying the overall energy budget.

Models c and d both adopt $\Sigma_j = 1.56\%$, compared to $\Sigma_j = 1.20\%$ of F; model d additionally reduces L_j to $\approx 75\%$ of the fiducial value to assess whether the stabilising effect of higher magnetisation persists even when the driving power is lower. In both models the increased Σ_j is associated with a better-preserved toroidal field structure throughout propagation, consistent with the suppression of KH mixing at the jet boundary: the jet–cocoon interface remains sharper, and less baryonic material is entrained into the relativistic core. Despite this more stable interface, the global dynamics and energetics of models c and F are nearly indistinguishable in terms of breakout time and energy partitioning. The reduced luminosity in model d does result in a moderately delayed breakout ($t_{\rm BO} \approx 300$ ms) and marginally more energy deposited into the surroundings, effects attributable to the lower drilling efficiency rather than to the magnetisation itself.

The key result from this subset of models is that, over the range $\Sigma_j \approx 1.2\%$–1.56% explored here, magnetisation is a secondary parameter for the success or failure of jet breakout, which remains controlled by L_j. Its primary observable consequence is on the fine angular structure and on the stability of the jet–cocoon boundary, which have implications for the off-axis afterglow morphology discussed in Chapter 6.

[5]This quantity is distinguished from the radial and angular average over a broader region (the jet cap) reported in Table 4.2, which is lower by construction.

4.3.3 Launch Time

The jet launch time t_j (defined with respect to the BNS merger) is qualitatively different from the other injection parameters: it does not modify the jet itself but changes the environment the jet encounters. The post-merger ejecta expand with time, so a later-launched jet faces a medium that is both more massive and more spatially extended. The additional mass reflects the continued baryonic loading from magnetically-driven winds off the remnant and its accretion disk during the interval between merger and jet launch [173], while the radial expansion reflects the free coasting of the previously ejected material. Although the mean density of the ejecta actually decreases with time, the total column of baryons the jet must sweep grows: the jet has both a longer path to traverse and more inertia. Since the ejecta expand on a timescale comparable to t_j itself, even a few hundred milliseconds can substantially change the outcome.

This is demonstrated by models *e1* and *e2*, which share identical injection properties, lower L_j and lower Σ_j than *F*, and differ only in their launch time: $t_j = 385$ ms for model *e1* and $t_j = 185$ ms for model *e2*. At the earlier launch time, the surrounding medium is approximately half as massive and extends roughly four times less far in radius, providing a substantially easier path to breakout. Model *e2* achieves a successful, collimated breakout with a well-defined ultra-relativistic core and relatively modest energy losses to the environment. Model *e1*, by contrast, undergoes severe disruption: more than half of the total injected energy is deposited into the ambient medium, and the resulting angular structure at $t - t_j = 3$ s lacks a coherent energetic core.[6]

For the injection parameters of models *e1* and *e2*, a difference of just 200 ms separates a successful from an almost disrupted outflow, and both launch times are plausible given the current theoretical uncertainties on the timescales for BH formation, magnetic field organisation, and baryon-wind clearing in the immediate post-merger phase. This sensitivity connects the observable sGRB properties directly to the poorly constrained physics of the earliest post-merger evolution, a link that motivates the afterglow modelling of Chapter 6 and the discussion of model degeneracies therein.

4.4 Jet State at $t = 3$ s: A Comparative Overview

At $t - t_j = 3$ s after launch, all jet models have left the inner merger environment and entered the low-density outer atmosphere, but their configurations differ significantly as a result of the injection parameter choices explored in the previous sections. Table 4.2 compares the key dynamical properties of each model at this stage.

The jet cap boundaries $r_{\rm in}$ and $r_{\rm out}$ are identified as outlined in Section 4.2.3: within this radial range, the energy-weighted radially averaged Lorentz factor as a function

[6] See, however, Appendix B.

Table 4.2: Jet cap properties at $t - t_j = 3\,\mathrm{s}$ for all models (north side). $E_{\mathrm{iso,max}}$ is the maximum isotropic-equivalent energy; $\overline{\Gamma}_{\mathrm{max}}$ is the maximum radially averaged Lorentz factor; $E_{\mathrm{sum}}/E_{\mathrm{inj}}$ is the fraction of injected energy retained in the outflow; $E_{\mathrm{kin}}/E_{\mathrm{sum}}$ is the kinetic energy fraction; Θ_{50} is the effective half-opening angle containing 50% of E_{sum}. Data from P25.

Model	$E_{\mathrm{iso,max}}$ [10^{53} erg]	$\overline{\Gamma}_{\mathrm{max}}$	$E_{\mathrm{sum}}/E_{\mathrm{inj}}$	$E_{\mathrm{kin}}/E_{\mathrm{sum}}$	Θ_{50} [°]
F (reference)	1.52	37.2	0.50	0.92	5.81
a ($L_j \uparrow$)	4.50	62.2	0.60	0.90	5.34
b ($\tau_d \uparrow$)	4.13	75.6	0.59	0.86	4.92
c ($\Sigma_j \uparrow$)	1.82	40.8	0.51	0.89	5.80
d ($\Sigma_j \uparrow$, $L_j \downarrow$)	1.19	50.6	0.48	0.89	5.28
e1 ($\Sigma_j \downarrow$, $L_j \downarrow$)	1.10	27.8	0.44	0.90	6.58
e2 ($\Sigma_j \downarrow$, $L_j \downarrow$, $t_j \downarrow$)	2.27	70.6	0.61	0.92	4.41

of angular position is

$$\overline{\Gamma}(\Theta,\Phi) = \frac{\displaystyle\int_{r_{\mathrm{in}}}^{r_{\mathrm{out}}} \Gamma(r,\Theta,\Phi)\, e_{\mathrm{sum}}(r,\Theta,\Phi)\, r^2\, dr}{\displaystyle\int_{r_{\mathrm{in}}}^{r_{\mathrm{out}}} e_{\mathrm{sum}}(r,\Theta,\Phi)\, r^2\, dr}\,, \tag{4.17}$$

where e_{sum} is the energy density excluding the rest-mass contribution, and the weighting by e_{sum} reflects the energy actually available for conversion into radiation. The corresponding isotropic-equivalent energy per unit solid angle is

$$E_{\mathrm{iso}}(\Theta,\Phi) = 4\pi\,\frac{dE_{\mathrm{sum}}}{d\Omega} = 4\pi \int_{r_{\mathrm{in}}}^{r_{\mathrm{out}}} r^2\, e_{\mathrm{sum}}(r,\Theta,\Phi)\, dr\,. \tag{4.18}$$

$\overline{\Gamma}_{\mathrm{max}}$ and $E_{\mathrm{iso,max}}$, reported in Table 4.2, are the maxima of Eqs. (4.17) and (4.18) over (Θ,Φ). We recall that the jet injection axis in our simulation coincides with the Cartesian y-axis; the angles Θ,Φ are therefore measured from the y-axis, as opposed to the standard polar coordinate θ,ϕ measured from the z-axis. The effective half-opening angle Θ_{50} is defined as the aperture of a cone subtending the same solid angle Ω_{50} as the smallest angular contour that encloses 50% of the total energy E_{sum} and thus provides an estimate of the angular size of the jet's energetic core. Finally, $E_{\mathrm{sum}}/E_{\mathrm{inj}}$ is the ratio between the total-without-rest-mass energy contained in the jet cap at the given time and the total energy injected up to that time, while $E_{\mathrm{kin}}/E_{\mathrm{sum}}$ is the fraction of E_{sum} in kinetic form; together they quantify, respectively, how much of the injected energy survives the interaction with the environment and how successful the conversion into kinetic form is.

At $t - t_j = 3$ s, the distribution of outflowing material is stratified and non-uniform for all models, reflecting the turbulent interaction history of each jet with the surrounding medium. A few clear trends emerge from the comparison. Models with higher injected energy (either through increased luminosity, model *a*, or extended injection duration, model *b*) retain a larger fraction of that energy in the outflow ($E_{\mathrm{sum}}/E_{\mathrm{inj}} \approx 0.59$–$0.60$, vs. 0.50 for model *F*), reach higher maximum Lorentz factors, and produce a jet with a slightly more collimated energy distribution (lower Θ_{50}).

An earlier launch time (model *e2*) has a similarly strong effect: despite the same low luminosity and magnetisation as model *e1*, the substantially reduced ambient mass at the time of launch allows the jet to retain $\approx 61\%$ of the injected energy and achieve $\overline{\Gamma}_{\max} \approx 70.6$, comparable to the high-luminosity models. Models with increased magnetisation (*c*, *d*) closely reproduce model *F* in terms of global energetics, but show deviations in the detailed angular structure. Model *e1*, by contrast, loses more than half of its injected energy to the environment ($E_{\rm sum}/E_{\rm inj} = 0.44$), the lowest of all models, and its broader $\Theta_{50} = 6.58^\circ$ reflects the lack of a well-defined, collimated energetic core.

An important point shared by all models is the $E_{\rm kin}/E_{\rm sum}$ column: at $t - t_j = 3$ s, the kinetic energy fraction lies in the range 0.86–0.92. While this already represents an advanced stage of the jet evolution, it is still worth exploring whether the jet has reached the ballistic regime, with a frozen angular structure and negligible residual acceleration. Feeding the jet profiles extracted at this stage directly into an afterglow code could introduce systematic errors in the predicted light curves, particularly in the early-time rising phase, which is most sensitive to the detailed angular and velocity structure.

This is a direct consequence of the computational constraints of 3D RMHD: extending the evolution while maintaining the resolution required to track the detailed angular structure is prohibitively expensive on a static spherical grid. The strategy employed to overcome this limitation, a grid-remapping technique that transitions the outflow to a Cartesian domain and extends the evolution to $t - t_j \approx 10$ s, where $E_{\rm kin}/E_{\rm sum} \geq 0.95$ for all models, is the subject of the next chapter.

Chapter 5

Towards the Ballistic Regime

In this chapter we present the first methodological contribution of this thesis: a technique to extend our simulations, without loss of resolution, until the jet acceleration saturates and the angular structure becomes stable. This method was first presented in D25.

5.1 Why the Ballistic Regime Matters

We refer to the jet *ballistic* or *coasting* regime as the moment when all the available internal and magnetic energy has been converted into kinetic form, i.e. $E_{\rm kin}/E_{\rm sum} \to 1$ and $\partial\Gamma/\partial r \to 0$, and the structure has stabilised. This is the regime in which we consider the jet structure reliable to be extracted and used as input for the semi-analytical afterglow model presented in Chapter 6.

In 2D simulations this regime is more easily achievable, since the lower computational cost allows evolutionary times of hundreds of seconds [101, 218, 227]. In 3D the cost increases both because of the additional spatial dimension and because of the longer numerical integration time needed, making it a different problem to solve. Most published 3D simulations of GRB jets in realistic environments are therefore stopped shortly after breakout, well before the acceleration saturates.

The simulations of P23 and P25 stop at $t - t_j = 3\,\mathrm{s}$. The analysis of the jet configuration at this time, presented in Section 4.4, shows that while all the models have reached an advanced stage of evolution by this point, it is still worth studying any possible further evolution.

The technique presented in the following sections was first applied to the reference simulation of P23 and presented in D25 (hereafter referred to as the *D25 model*), and subsequently extended to the parameter study detailed in Table 4.1.

5.2 The Grid Remapping Technique

After $t_{\rm beg} \equiv t - t_j = 3\,\mathrm{s}$, continuing the simulation on the original spherical grid becomes problematic. The fastest part of the outflow, which we define as the material with $\Gamma \gtrsim 2$, occupies a radial region of roughly constant width as the jet expands. On a spherical grid with logarithmic radial spacing, the number of cells covering

this shell keeps decreasing as the jet moves outward, so the resolution is lost exactly where it matters most. A spherical grid with uniform radial spacing avoids this loss, but introduces a different problem: the cell aspect ratio $r\,\Delta\theta/\Delta r$ grows with radius, and the elongated cells cause numerical diffusion across the jet boundaries. Neither choice works well beyond $t_{\rm beg}$, and resolving the angular structure properly at all times is important, since it directly affects the rising part of the afterglow light curve (see Chapter 6).

The solution we adopt is to remap the simulation onto a Cartesian grid with uniform, cubic cells, so that resolution and aspect ratio are preserved at any distance from the origin. First, at $t_{\rm beg}$, the spherical simulation data are interpolated onto a Cartesian grid. The Cartesian domain of the simulation presented in D25 has $N_x = N_z = 848$ and $N_y = 512$ points, covering $x, z \in [-2.88, 2.88] \times 10^6$ km and $y \in [0.1, 3.58] \times 10^6$ km, with a uniform cell size of 6800 km on all sides. This value is an intermediate choice between the radial spacing of the original spherical grid at the base and at the front of the fast ($\Gamma \gtrsim 2$) outflow. As a result, most of the jet is better resolved than on the original grid, and the resolution at the outer front improves by almost a factor of two, at the cost of a small loss of resolution only in the jet tail, the region closest to the origin that the fastest material has already passed through. The simulation is evolved in this domain until $t_{\rm end} \equiv t - t_j = 9.5$ s. For the parameter study we have explored a different grid, maintaining the same resolution, but with a larger $y_{\rm max}$, slightly decreasing the x, z range and overall increasing the number of points. The aim was pushing the simulations even later in time, up to $t = 20$ s and assess any possible further evolution, at the cost of losing part of the material at larger angles, exiting the grid sideways. We find no significant differences in the jet properties after $t \approx 10$ s, therefore for consistency between the simulations we consider $t_{\rm end}$ as the final time. Only the northern hemisphere ($y > 0$) of the original domain is imported into the Cartesian grid. At the time of remapping the northern and southern jets are causally disconnected, so this choice does not introduce any error and simply reduces the computational cost by about a factor of two. An analogous procedure could in principle be applied to the southern hemisphere; however, this extension is not pursued in the present study. Figure 5.1 presents a schematic (not to scale) representation of the remapping procedure.

As boundary conditions in the Cartesian domain we apply standard outflow (zero-gradient) conditions on all sides of the box except on the inner boundary in the y direction (see below). The $y_{\rm min}$ face, however, is not placed at $y = 0$ but at a value chosen to exclude most of the post-merger ejecta that, by the time of the import, no longer affects the jet evolution, while still including all the fast ($\Gamma > 2$) outflow of interest. On this face, the outflow condition is modified: instead of copying the initial boundary value $v_{y,0}$ into the domain at all times, which would cause a continuous and unphysical inflow of material, we impose an exponentially decaying velocity $v_y(t) = v_{y,0} \exp(-t/\tau)$ with $\tau = 50$ ms. This does not affect the jet dynamics, since the jet is already far from this boundary at the time of remapping, but it allows a much cleaner check of energy conservation. Regarding the numerical methods, we use the HLL Riemann solver, PPM reconstruction, third-order Runge-Kutta time stepping, and hyperbolic divergence cleaning (Chapter 3).

We assess the accuracy of the remapping using as a reference the D25 model,

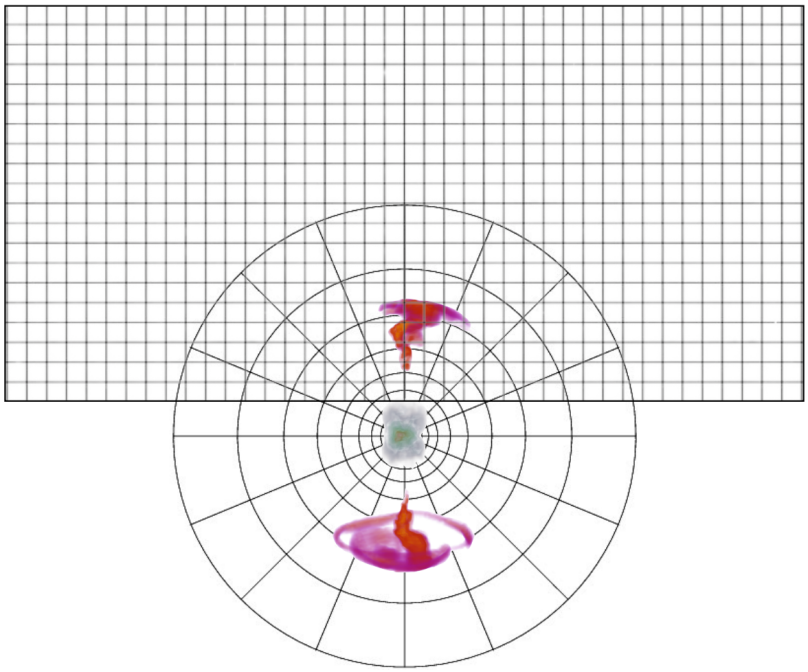


Figure 5.1: Schematic representation of the grid remapping technique (not to scale).

starting with a check on energy conservation. We find that $E_{\rm sum} \equiv E_{\rm kin} + E_{\rm int} + E_{\rm mag}$, computed over the whole Cartesian domain, increases by about 0.7% in the first second after remapping, due to the spurious material still entering through the $y_{\rm min}$ boundary before the exponential damping takes full effect. After this initial transient, $E_{\rm sum}$ is conserved to within 0.08% for the rest of the simulation, which suggests that no significant numerical errors are introduced by the remapping or by the evolution that follows.

A comparison of the rest-mass density profile along the y-axis before and after remapping (see Figure 5.2 where the two profiles are shown as blue lines of different opacity) shows almost precise agreement up to $r \sim 2.25 \times 10^6$ km, corresponding to the maximum radius of the original interpolated data. The slight difference at larger distances is likely due to an evolution of the atmosphere close to the outer boundary of the initial simulation, that is instead re-initialised in the extended one following the same prescription, as described in Chapter 4. We also check the consistency of the jet structure by comparing the angular profiles of the azimuthally averaged Lorentz factor and isotropic-equivalent energy of the jet cap, computed as in Section 4.4, just before and just after remapping. To compute these profiles from the Cartesian output, the data are remapped a second time onto a spherical grid whose resolution is at least as fine as the Cartesian cell size, so no further information is lost in this second interpolation. We find good overall agreement, with a maximum relative difference below 10%, consistent with the resolution change introduced by this double-remapping procedure.

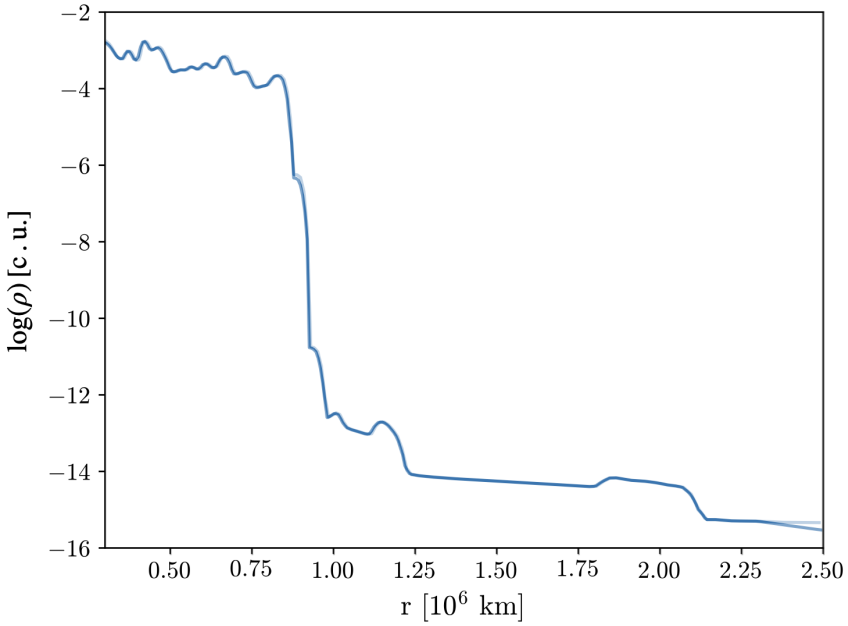


Figure 5.2: Rest-mass density profile along the jet axis before and after the remap to the Cartesian domain at $t_{\rm beg}$, in different shades of blue.

5.3 Assessment of the Magnetic Field Divergence

As described in Section 3.2.3, the magnetic field must satisfy $\nabla \cdot \boldsymbol{B} = 0$ at all times, and this constraint has to be assessed for the remapping procedure as much as for the rest of the evolution. Analytically, the coordinate change from spherical to Cartesian preserves the divergence-free condition exactly, since it is only a change of representation of the same physical field. In practice, however, the interpolation of a discretely sampled field from one grid to another does not preserve $\nabla \cdot \boldsymbol{B} = 0$ to machine precision: each component of $\boldsymbol{B}$ is interpolated independently, and the discrete divergence computed on the new grid can develop a small but non-zero value, effectively introducing spurious magnetic monopoles at the level of numerical truncation error.

We control this residual error using the same hyperbolic divergence cleaning scheme already employed throughout the spherical phase of the simulation (Chapter 3), following the approach of Dedner et al. [42] as implemented in PLUTO [129]. This method couples an auxiliary scalar field ψ to the induction equation, so that any divergence error is transported and damped as a combination of an advected and a decaying wave, rather than being left to accumulate locally.

An alternative strategy, commonly employed in RMHD codes, is to evolve the magnetic vector potential $\boldsymbol{A}$ instead of $\boldsymbol{B}$ directly, recovering the magnetic field as $\boldsymbol{B} = \nabla \times \boldsymbol{A}$ and using a constrained-transport scheme for its update. Since the divergence of a curl vanishes identically, this approach guarantees $\nabla \cdot \boldsymbol{B} = 0$ to round-off precision by construction [51]. In the context of our remapping procedure, however, adopting this formulation would require first constructing a vector potential for the magnetic field in the original domain, including a choice of gauge and appropriate boundary conditions, followed by remapping $\boldsymbol{A}$ and reconstructing $\boldsymbol{B}$ from its curl on

the new grid. This introduces an additional reconstruction step and, at the discrete level, the remapping operation does not in general commute with the curl, so the resulting magnetic field need not be identical to that obtained by directly remapping $\boldsymbol{B}$. We therefore retain the direct evolution and remapping of $\boldsymbol{B}$ together with hyperbolic divergence cleaning. Since this approach already maintains $\nabla \cdot \boldsymbol{B}$ at a very low level, comparable to that obtained in the original spherical simulation, we consider it sufficient for the purposes of this work without introducing the additional complexity associated with the vector-potential formulation.

5.4 Energy Conversion and Approach to Saturation

5.4.1 Evolution of Energy Components

A first, direct way to check how close the jet is to the ballistic regime is to look at how the conversion of thermal and magnetic energy into kinetic form saturates over time. We track the energy budget over the whole computational domain from jet launch to $t_{\rm end}$. For the D25 model, $E_{\rm sum} \approx 4.1 \times 10^{50}$ erg throughout the extended evolution, and at $t_{\rm beg} = 3$ s the kinetic energy fraction is $E_{\rm kin}/E_{\rm sum} \simeq 94\%$, so some further acceleration is still expected at later times. By $t_{\rm end}$, this fraction reaches 98%, and the trend is clearly saturating.

The remaining 2% could in principle still convert on much longer timescales, or stay locked in the shocked cocoon. The polynomial extrapolation of the $E_{\rm kin}/E_{\rm sum}$ trend for the fiducial model gives an idea of how much further evolution would be needed to reach a full conversion. Extrapolating the trend shown in the previous subsection, a 99.9% conversion would require more than 200 s of evolution, well beyond the timescale reachable with this setup, and consistent with the estimate of Xie et al. [227]. Figure 5.3 shows the energetic evolution discussed here.

Assuming full conversion to kinetic energy already at $t_{\rm end}$ introduces an error of only about 2% on the jet kinetic energy. This translates into an underestimate of the Lorentz factor by about 1% with respect to the saturation value, since $\sqrt{E_{\rm sum}/E_{\rm kin}} = 1/\sqrt{0.98} \approx 1.01$ (we refer to the next section for a more detailed analysis on this topic). We consider this error acceptable and adopt this as the practical ballistic limit throughout this thesis.

The same picture holds across the parameter study: thermal and magnetic energy are progressively converted into kinetic form during the propagation, sustaining the jet acceleration until it enters the ballistic regime. At $t_{\rm beg}$, the kinetic energy fraction at the jet cap is between 86% and 92% for the different models (Table 4.2). By $t - t_j = 6$ s, this fraction rises to about 95% for model *a*, 93% for model *b*, 96% for model *F*, 94% for models *c*, *d*, and *e1*, and 96% for model *e2*. By $t_{\rm end}$, all models reach a saturation stage, with kinetic energy fractions between 95% and 97%. The relative increase in kinetic energy fraction between 3 and 6 s is much larger than between 6 and 9 s, which shows that the residual acceleration becomes less efficient as the jet approaches its final configuration.

Table 5.1 shows the jet cap properties at $t_{\rm end}$ for all seven models, extending the

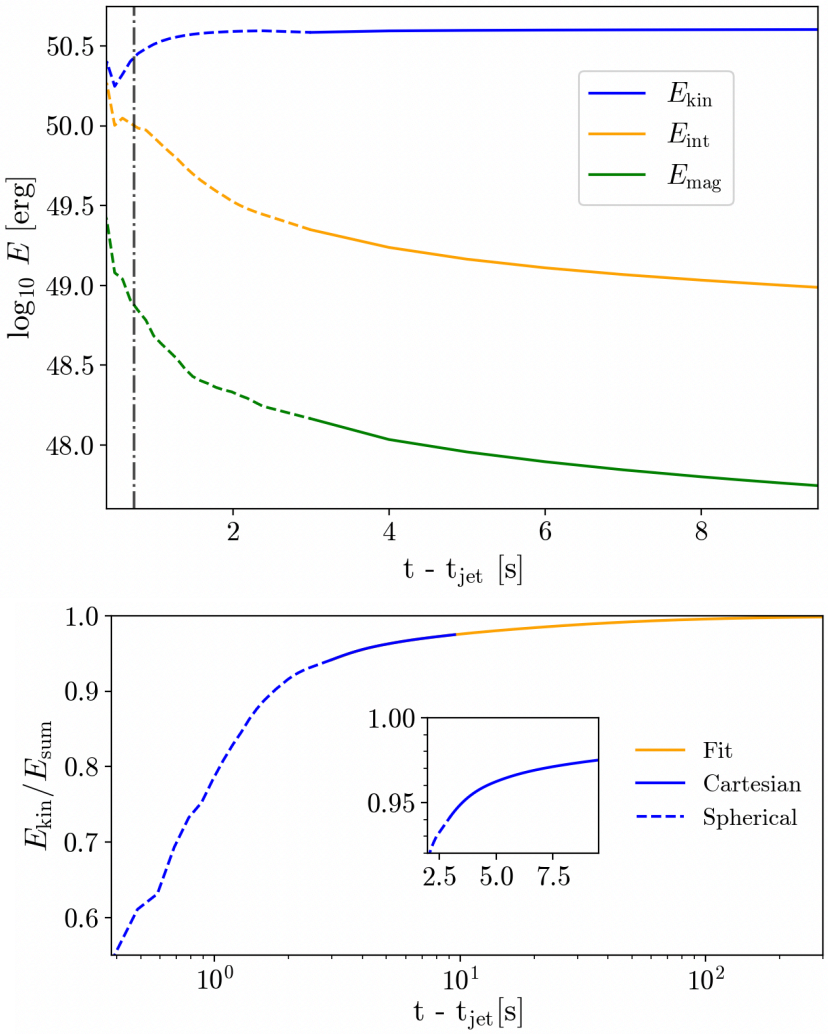


Figure 5.3: Temporal evolution of the energy component of the simulation of D25. The kinetic-to-total ratio is extrapolated at later times with a polynomial fit.

comparison already given at $t_{\rm beg}$ in Table 4.2. The discussion on the features related to the angular structure is left to the next chapter.

Table 5.1: Jet cap properties at $t_{\rm end} = 9.5\,$s for all models. The jet cap properties at $t_{\rm end}$ are: kinetic energy fraction $E_{\rm kin}/E_{\rm sum}$, effective half-opening angle Θ_{50}, and offset of the jet barycentre from the injection axis $\Theta_{\rm off}$. Data from D26.

Model	$E_{\rm kin}/E_{\rm sum}$	Θ_{50} [°]	$\Theta_{\rm off}$ [°]
F	0.97	5.6	1.4
a	0.97	5.1	1.6
b	0.95	4.9	1.0
c	0.97	5.0	2.7
d	0.96	5.2	1.5
e1	0.96	6.1	3.6
e2	0.97	4.6	0.9

Taken together, these results demonstrate that substantial changes in the jet structure can occur well after breakout, at times extending up to nearly 10 s after jet launching. In a configuration such as the one considered here, namely a GRB jet injected into a confined post-merger environment, this is significantly later than the engine shut-off and corresponds to the phase of essentially free jet expansion. We therefore conclude that, for our simulation suite, $t_{\rm end}$ represents a robust epoch at which to extract stable jet-structure configurations to be used in afterglow modelling.

5.4.2 The Saturation of the Jet Structure

The energy budget alone does not tell us how close the jet is to its final structure, since the same amount of residual energy can be distributed very differently across the outflow depending on the local conditions. To look at this more directly, we employ a *saturation Lorentz factor* Γ_s: the value of Lorentz factor that a fluid element would reach if all its remaining internal and magnetic energy were converted into kinetic form. For a representative element of the outflow, the lab-frame kinetic energy density is $e_{\rm kin} = \Gamma(\Gamma-1)\,\rho_0 c^2$, where ρ_0 is a representative comoving rest mass density. If we substitute $e_{\rm kin}$ with the available energy for conversion $e_{\rm sum}$, and approximate $\Gamma(\Gamma-1) \approx \Gamma^2$ in the ultra-relativistic limit, we get

$$\Gamma_s \approx \sqrt{e_{\rm sum}/e_{\rm kin}}\,\Gamma\,, \tag{5.1}$$

so the ratio $\sqrt{e_{\rm sum}/e_{\rm kin}}$ tells us, at any given time, how far a fluid element still is from its terminal velocity. We apply this analysis to the reference D25 simulation and average over the radial extent of the jet cap at fixed (Θ, Φ) to get the saturation profile $\overline{\Gamma}_s(\Theta, \Phi)$, and its azimuthal average

$$\overline{\Gamma}_s(\Theta) = \frac{\int_0^{2\pi} \overline{\Gamma}_s(\Theta, \Phi) d\Phi}{2\pi}. \tag{5.2}$$

Comparing $\overline{\Gamma}(\Theta)$ to $\overline{\Gamma}_s(\Theta)$ at different times is therefore a more direct test than the global energy fraction. The resulting profiles are shown in Figure 5.4.

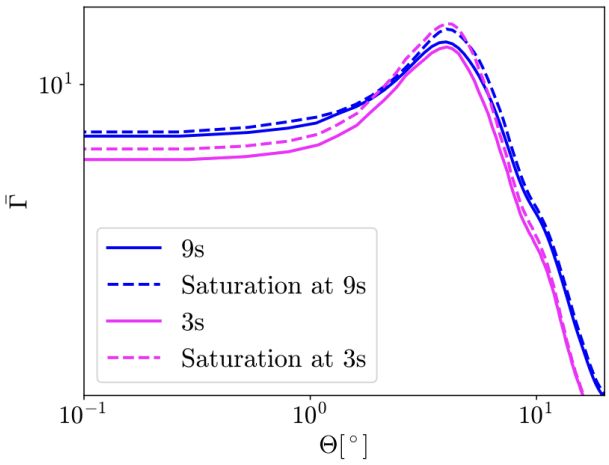


Figure 5.4: Angular profile of the radially and azimuthally averaged Lorentz factor, together with its saturation value, at $t_{\rm beg}$ and $t_{\rm end}$.

At $t_{\rm beg}$ the two profiles differ by up to about 10%; at the latest time this gap shrinks to 2–5% within $\Theta \lesssim 10^\circ$. Around the peak the difference remains visible for both times, while being slightly reduced at 9 s and remains small at larger angles, so the jet cap is close to its terminal velocity structure over most of the angular range of interest. It is also useful to compare the saturation profiles at the two times directly: $\overline{\Gamma}_s(\Theta)$ at $t_{\rm beg}$ and at $t_{\rm end}$ agree to within about 2% near the peak, but differ by 8% at small Θ and by as much as 35% in the wings. Since the afterglow signal is quite sensitive to the input Lorentz factor profile, this residual difference means a light curve predicted from the $t_{\rm beg}$ structure and one predicted from the $t_{\rm end}$ structure would not agree in general, a point we quantify in Chapter 6.

For the D25 model, we compare the azimuthally averaged angular profiles of Lorentz factor, defined as in Eq. 5.2 for its saturation value, and isotropic-equivalent energy

$$E_{\rm iso}^{\Phi_{\rm avg}}(\Theta) = \frac{\int_0^{2\pi} E_{\rm iso}(\Theta, \Phi) d\Phi}{2\pi} \tag{5.3}$$

at four representative times: break-out ($t_{\rm BO} \approx 350\,$ms) in the original spherical domain, right after the beginning of the Cartesian evolution ($t_1 = 3.5\,$s), an intermediate time ($t_2 = 6.5\,$s), and close to the end at $t_3 = 9\,$s. The comparison, shown in Figure 5.5, shows clear acceleration and a significant evolution of the energy distribution following breakout. This behaviour persists beyond that point: between t_1 and t_2, the average Lorentz factor profile retains an approximately self-similar shape while overall increasing in intensity, whereas the isotropic-equivalent energy profile gradually flattens between 1° and 4° and steepens at larger angles, moving towards a slope $\propto \Theta^{-4}$ beyond about 10°. The differences between t_2 and t_3 are much smaller, which confirms that by the later time the angular distributions are nearly frozen.

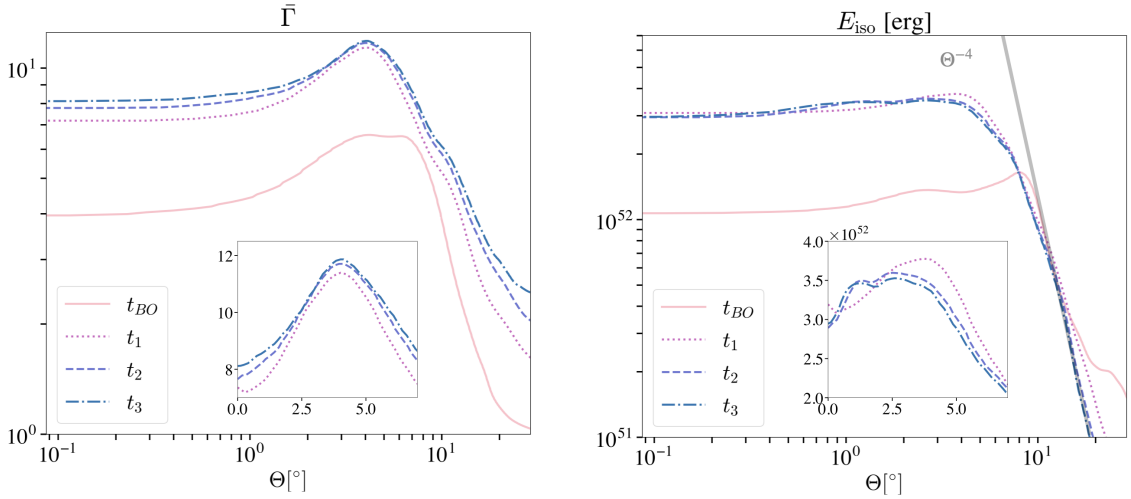


Figure 5.5: Angular profiles of the azimuthally averaged jet cap Lorentz factor (left) and isotropic equivalent energy (right) at different evolutionary times for the simulation of D25.

5.5 Jet Dynamics

Figure 5.6 displays the Lorentz factor, rest-mass density, pressure, and magnetic field strength on the xy-plane at three representative epochs of the Cartesian evolution, corresponding to $t - t_j \simeq 3$, 6.5, and 9.5 s. In the Lorentz factor maps, the high-velocity outflow appears composed of an outer front expanding radially, a short trailing tail behind it, and a quasi-spherical shell of material accelerated during the earlier injection phase which envelops the tail. The radial stratification of Γ across the tail traces the fading behaviour of the injected luminosity: fluid elements launched earlier, when L_j was higher, now lead the outflow, while progressively slower material trails behind. Between 3 and 9.5 s the high-velocity region ($\Gamma \gtrsim 2$) does not spread appreciably in the radial direction, whereas its lateral extent grows steadily, gradually approaching the geometry of a spherical-shell sector with an approximately constant half-opening angle of $\sim 30°$. By this stage the jet has decoupled from the merger environment, and its subsequent evolution across the artificial floor medium is limited to the mild residual acceleration and some morphological adjustments described before, with no dynamically significant back-reaction from the ambient. The rest-mass density distribution (second column) shows the entrained environmental material inherited from the spherical grid, only partially transferred to the Cartesian domain, expanding slowly and preserving the low-density funnel excavated along the y-axis during the earlier jet passage. Near the lower boundary at $y = 10^5$ km, the exponentially decaying v_y imposed to suppress further mass and energy injection (Section 5.2) leaves a clear imprint on the density field, with analogous features appearing in the magnetic field strength map (third column).

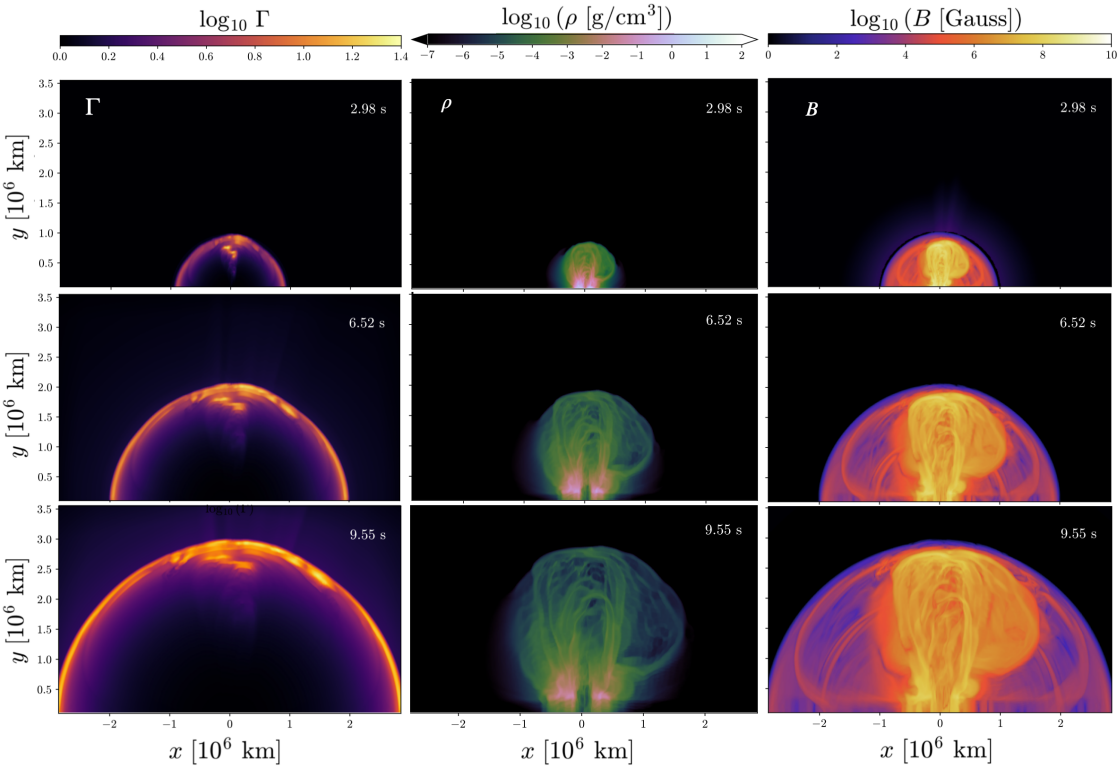


Figure 5.6: Meridional view of Lorentz factor, rest-mass density and magnetic field intensity at three representative times for the simulation of D25.

Chapter 6

Connecting the Jet Structure to Electromagnetic Emission Across Models

The previous chapters established the numerical framework for propagating relativistic jets through a realistic BNS post-merger environment and for extending the simulation to the ballistic regime via grid remapping, and presented the main results on the early-time dynamical interaction between the jet and the environment. We now turn to the central analysis of this thesis: characterising the angular and velocity structure of the resulting outflows across the jet models introduced in Sections 4.2 and 4.3, and quantifying the imprint of this structure on synthetic afterglow light curves.

6.1 Final Angular and Radial Structure

The analysis of the angular and velocity distributions of the jet material at $t_{\rm end}$ serves as direct input for the afterglow calculations. In this section, we first examine the two-dimensional front-view maps of isotropic-equivalent energy and average Lorentz factor, their axisymmetised profiles, and then turn to the radial velocity profiles that reflect the layered internal structure of the outflow.

6.1.1 Front-View Angular Distributions

The angular distribution of the jet material can be characterised, as a first diagnostic, through the front-view maps of isotropic-equivalent energy $E_{\rm iso}(\Theta, \Phi)$ and radially averaged Lorentz factor $\bar{\Gamma}(\Theta, \Phi)$ at the jet cap (as defined in Section 4.4), evaluated at $t_{\rm end}$.

The analysis of the simulation presented in D25, illustrated in Figure 6.1, already reveals pronounced departures from axisymmetry. A first key result is a clear misalignment between the actual location of the jet's energetic core and the nominal injection axis, evidenced by a local maximum in $E_{\rm iso}$ that is offset by approximately 4°–6° from the centre. A second result is the presence of significant azimuthal variations in both the energy and the Lorentz factor around the maximum of the distribution, with fluctuations of order ~ 3, and dispersions of up to an order of

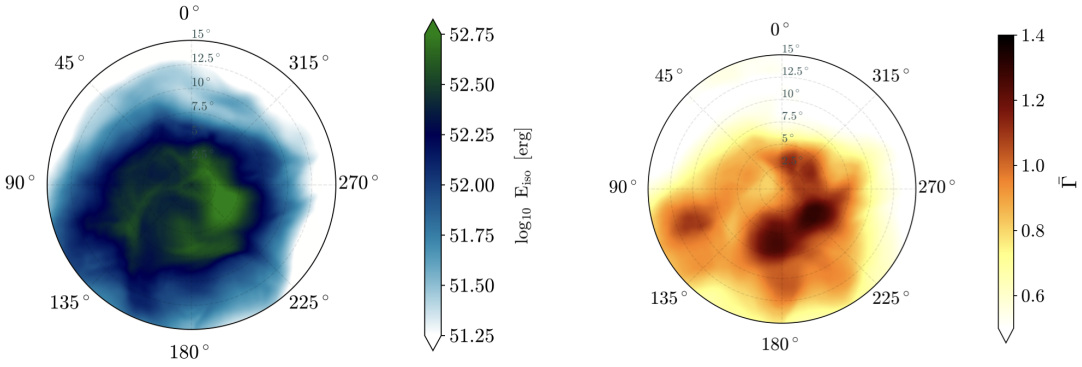


Figure 6.1: Front-view angular distribution of isotropic equivalent energy and average Lorentz factor of the jet cap, at $t_{\rm end}$, for the simulation of D25.

magnitude in the wings at larger angles. A more comprehensive representation of the structural characteristics is provided by the three-dimensional rendering of the kinetic energy density shown in Figure 6.2, in which the same asymmetric features can be recognised.

We extend this analysis here to the full parameter study; the resulting maps are displayed in Figure 6.3. We note that a different colour scale is adopted with respect to the D25 model analysis, in order to properly compare the features across all models.

In each panel, the red contour encloses the smallest region containing 50% of the total energy, providing a measure of the angular extent of the jet's energetic core. A white cross marks the energy-weighted barycentre of the distribution in angular space. The displacement of this point from the injection axis, $\Theta_{\rm off}$, quantifies the degree of asymmetry imparted by the interaction with the non-uniform merger environment. None of the simulated jets exhibits an axisymmetric energy distribution: the barycentre is offset from the injection axis by 0.9°–3.6° across all models (see Table 5.1).

A clear trend emerges from the comparison. The most energetic or highly magnetised models (*a, b, c, d*) display a more centred, tightly collimated core, consistent with a reduced interaction with the environment during the early propagation phase. Model *e2*, despite its lower energy and magnetisation, also produces a compact, well-centred core, a consequence of the shorter extent of the environment it traverses owing to its earlier launch time. In contrast, model *F* shows a more pronounced imprint of the environmental interaction, while model *e1* is the most affected: its energy distribution is broad, displaced, and lacks a well-defined peak, consistent with the marginal breakout discussed in Section 4.3.3.[1]

The effective half-opening angle Θ_{50} (already defined in Section 4.4 as the opening angle of a cone subtending the same solid angle as the region enclosed by the red contour) is a proxy of how much collimation the jet retains at large scales. Across all models, Θ_{50} lies close to 5°, with model *e2* producing the narrowest cone (4.6°)

[1]See, however, Appendix B.

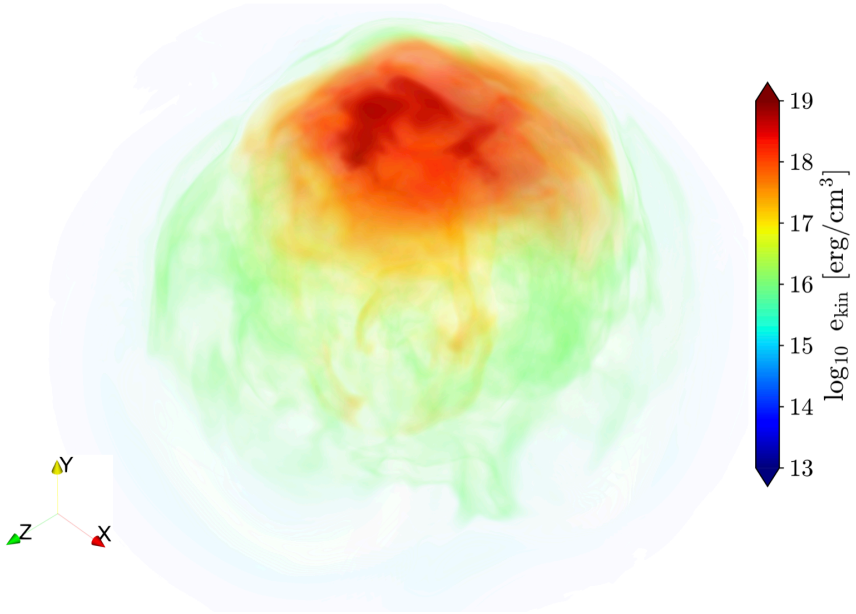


Figure 6.2: 3D rendering of the kinetic energy density at $t_{\rm end}$ for the simulation of D25.

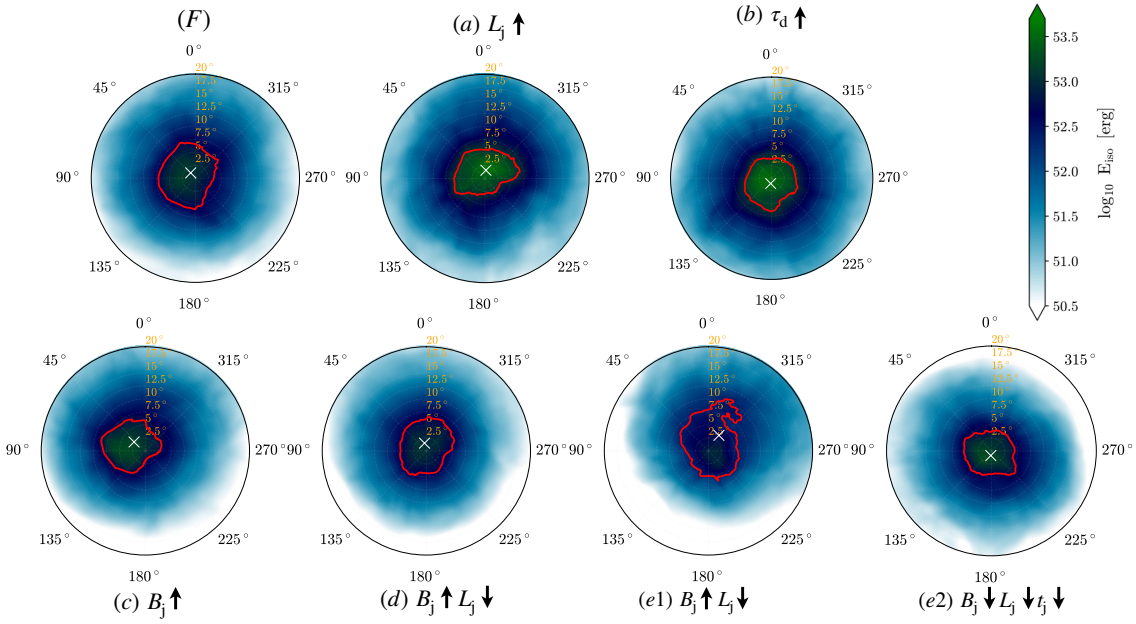


Figure 6.3: Front-view distributions of the isotropic-equivalent energy of the jet cap for the jet models of the parameter study at $t_{\rm end}$. The red line shows the smallest contour containing 50% of the energy; the white cross points at the jet's energy barycentre.

and model *e1* the widest (6.1°; see Table 5.1).

Models *a* and *b*, which carry the largest injected energies, also exhibit the highest peak values of the isotropic-equivalent energy (see Eq. 5.3), $E_{\rm iso,max} = 4.57 \times 10^{53}$ and 3.93×10^{53} erg respectively, against 1.46×10^{53} erg for the reference model *F*, roughly a factor of three higher. The maximum of the average Lorentz factor (Eq. 5.2) follows the same trend, reaching $\overline{\Gamma}_{\rm max}|_{\Phi\rm avg} \simeq 80$ for *b*, close to the values typically inferred for GRB jets, and 41 for *a*, against 28 for *F*.

Model *e2* shares the same behaviour despite its lower injected luminosity: launched earlier into a less obstructed environment, it retains more of the injected energy than *F* and reaches $E_{\rm iso,max} = 2.33 \times 10^{53}$ erg and $\overline{\Gamma}_{\rm max}|_{\Phi\rm avg} \simeq 54$. Model *e1*, launched later with otherwise identical injection parameters, provides the direct contrast: its peak energy falls below the reference, 0.97×10^{53} erg, with $\overline{\Gamma}_{\rm max}|_{\Phi\rm avg} \simeq 24$, consistent with its marginal breakout and with the absence of a well-defined core.

Models *c* and *d*, which vary the initial magnetisation, show a milder response. Model *c* is within one unit of *F* in the average Lorentz factor (27 against 28) while reaching a somewhat higher peak energy (1.94×10^{53} erg). Model *d* injects less energy than *F*, reaches a lower $E_{\rm iso,max} = 1.26 \times 10^{53}$ erg, but ends with $\overline{\Gamma}_{\rm max}|_{\Phi\rm avg} \simeq 34$. Both behaviours are consistent with the conclusion of Section 4.3.2, namely that magnetisation, over the narrow range explored here, has a limited impact on the global energetics. The dependence of these parameters on the resolution is discussed in Appendix B.

6.1.2 Axisymmetrised profiles

To facilitate a comparison with analytical jet profiles, we compute the azimuthally averaged isotropic-equivalent energy $E_{\rm iso}^{\Phi_{\rm avg}}(\Theta)$ (Eq. 5.3) at $t_{\rm end}$ for all models.

The resulting profiles are shown in Figure 6.4. All models feature a flat energetic core extending to $\Theta \lesssim 2°$–$3°$, followed by wings of progressively decreasing energy. The wing profiles are not smooth: they display multiple segments with different slopes, reflecting the layered interaction history of the outflow with the environment. This portion of the structure has a direct influence on the rising part of the off-axis afterglow light curves, as discussed in Section 6.4. To allow for a direct visual comparison, we show the profiles alongside a reference power-law decay given by $E_{\rm iso} \propto \Theta^{-\alpha}$. Immediately outside the core aperture, the majority of models exhibit a slope of $\alpha \approx 3$. At larger angles, $\Theta > 10°$, most models instead follow a steeper decay with an index $\alpha \approx 6$, while model *e1* shows a shallower profile with $\alpha \approx 4$, consistent with its broader and less collimated structure. A Gaussian profile, $E_{\rm iso}(\Theta) = E_{\rm iso}(0) \exp\left[-\frac{\Theta^2}{2\theta_c^2}\right]$, with $\theta_c = 4°$, is also shown in comparison with model *a*. We conclude that the jet structures produced by our simulations are complex and articulated, and do not generally follow a simple power-law or Gaussian functional form.

The main structural parameters at $t_{\rm end}$ for all models are summarised in Table 5.1.

6.1.3 Radial Velocity Stratification

Beyond the angular distribution, the internal velocity structure of the outflow carries essential information for afterglow modelling. In the simplest jet models, all the

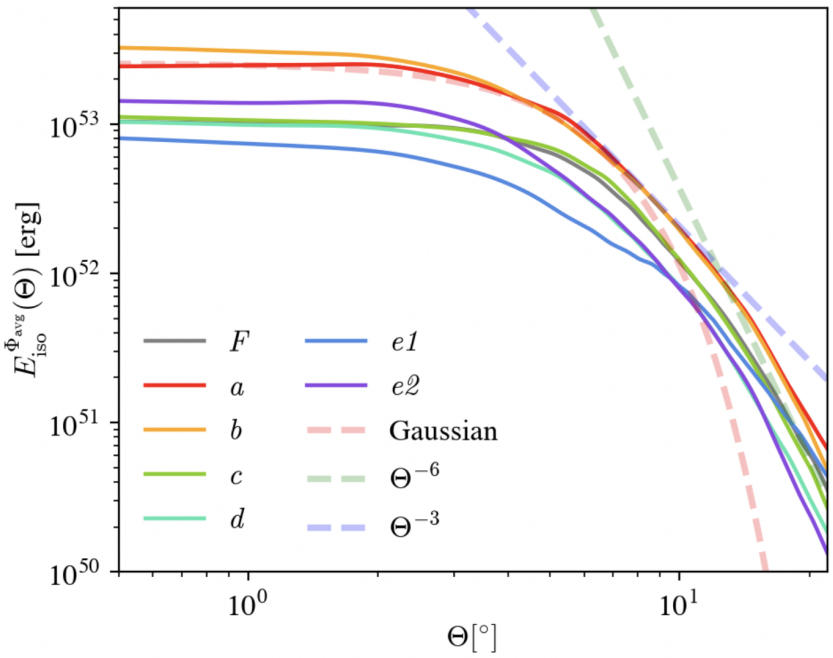


Figure 6.4: Isotropic equivalent energy profiles of the azimuthally averaged jet cap for all models of the parameter study at $t_{\rm end}$.

energy at a given angle is carried by material moving at a single, well-defined Lorentz factor. The outflows produced by our simulations are far from this idealisation: at each angle, energy is distributed over a broad range of four-velocities, reflecting the layered structure that develops during the jet propagation through the merger environment.

To quantify this stratification, we extract from the simulation output at $t_{\rm end}$ the differential energy distribution $dE_{\rm iso}/du\,(\Theta)$, which describes the amount of isotropic-equivalent energy carried by fluid elements with four-velocity $u \equiv \Gamma\beta$ (in units of the speed of light) at each polar angle Θ. The profiles are computed from the azimuthally averaged total energy distribution (including the rest-mass contribution), binning the simulation cells in $(\Theta,\, u)$ space and summing the energy content in each bin.

Physically, $dE_{\rm iso}/du$ gives the energy carried by shells of different speed. A narrow profile concentrated around a single value of u would correspond to the standard single-shell approximation used in most afterglow models. A broad profile, on the other hand, implies that slower shells carry a non-negligible fraction of the energy and will progressively catch up with the decelerating forward shock, injecting additional energy at later observer times, following the mechanism discussed in Section 2.8. The velocity profiles therefore constitute the key input for the extended afterglow model of Section 2.8.1.

Figure 6.5 shows the differential velocity profiles $dE_{\rm iso}/d\log(u)$, which allows for a better understanding of the energy distribution in the u range, at several representative polar angles Θ for the reference model F at $t_{\rm end}$.

The jet core ($\Theta \lesssim 5^\circ$) carries the highest energies, concentrated in a relatively narrow range of $u \sim 20$–60. A component with energy comparable to the cocoon at

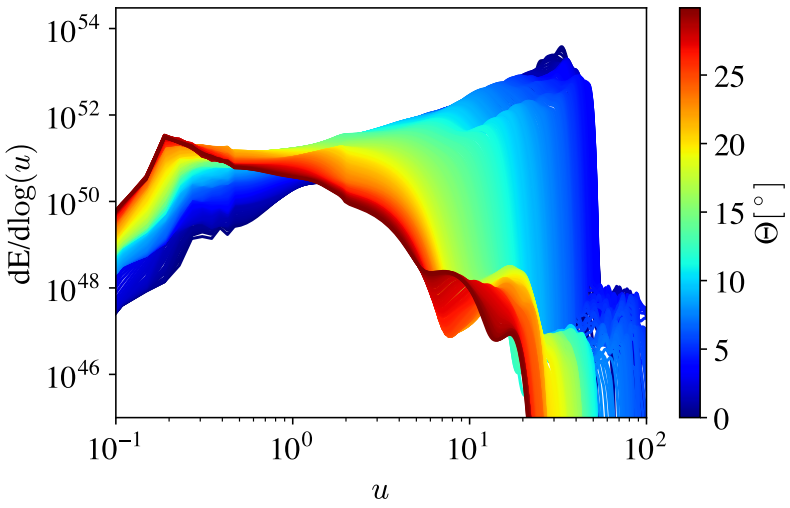


Figure 6.5: Velocity profile of model *F*. The colorbar indicates different polar angles.

larger angles, carried by slower elements in the range $u \sim 1 - 10$, is present also for small angles from the injection axis. This broad distribution arises directly from the jet's complex interaction history: internal shocks, mixing at the jet–cocoon boundary, and the progressive conversion of magnetic energy into kinetic form all contribute to spreading the energy over a range of velocities.

At intermediate angles ($\Theta \sim 5°$–$15°$), the profiles shift to lower values of u and the total energy per unit four-velocity decreases by several orders of magnitude. However, a high-velocity tail persists: material with $u > 10$ is present even at $\Theta \sim 15°$, despite carrying a small fraction of the total energy. This fast component outside the core is associated with the multi-shock structure described in Section 4.2.2 and, as discussed in Section 6.4.1, can produce observable signatures for off-axis observers.

At large angles ($\Theta \gtrsim 20°$), the profiles become dominated by sub-relativistic material ($u \lesssim 1$), corresponding to the cocoon. The energy content at these angles is comparatively small, but it is not negligible in the context of afterglow modelling at large viewing angles.

To compare the velocity structure across all models simultaneously, we construct the joint energy distribution in the Θ–u plane. For each polar angle, we compute the percentiles of the energy distribution as a function of u from the cumulative azimuthally averaged velocity profiles. The results are presented in Figure 6.6: the dark shaded region corresponds to the interquartile range (25th–75th percentile), and the lighter region extends between the 5th and 95th percentiles. We denote the n-th percentile of the four-velocity as u_n.

This representation captures at a glance both the characteristic speed at each angle and the range of speeds present. At small angles ($\Theta \lesssim 4°$), several models show relatively flat profiles or mild peaks around $\Theta \approx 2°$, likely reflecting the misalignment between the jet axis and the line of sight to the energy-weighted barycentre. The median four-velocity in this angular range typically lies between 30 and 60, confirming the ultra-relativistic nature of the core. The substantial width of the interquartile

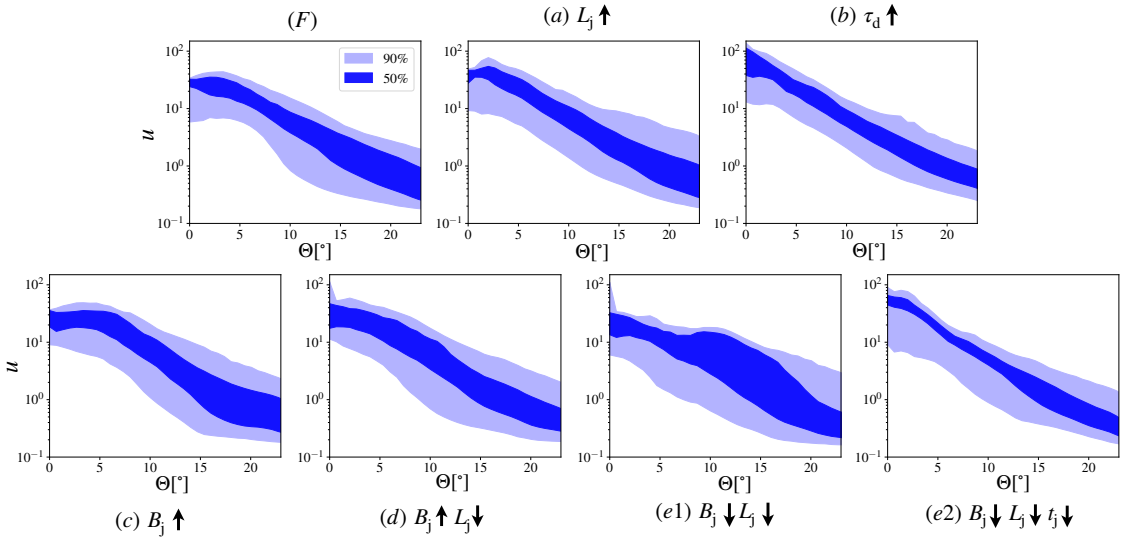


Figure 6.6: Jet energy distribution in the $u - \Theta$ plane. The dark blue shaded region identifies the interquartile range (25th–75th percentiles) while the range between the 5th and 95th percentiles is presented in light blue.

range, however, reinforces what was already visible in the differential profiles: even the fastest part of the outflow is not a homogeneous shell but a stratified structure with a broad spread in Lorentz factor. Consistently with the previous discussion, we observe that for $\Theta \approx 4°$ and beyond, the median and upper percentile curves decrease gradually. The upper limit of the 95th percentile u_{95} remains above 10 out to $\Theta \approx 15°$, indicating that a small fraction of rapidly moving material persists at relatively large angular offsets from the jet axis. At large angles ($\Theta \gtrsim 20°$), the outflow is predominantly sub-relativistic ($u \lesssim 1$) and the velocity distribution is concentrated at low values, though some faster components survive carrying a minor fraction of the total energy budget.

These results demonstrate that our simulated jets exhibit strong angular and velocity stratification, with a fast, narrowly collimated core surrounded by slower, less energetic wings. The transition between these regions is smooth rather than abrupt, in contrast with the predictions of top-hat or two-component jet models. The broad velocity distributions at every angle imply a sustained energy injection into the external shock over an extended period, as progressively slower shells catch up with the decelerating blast wave. This highlights the need for afterglow models that incorporate the full distribution $dE_{\rm iso}/du$ rather than a single representative Lorentz factor at each angle, which is precisely what the extended model of Section 2.8.1 provides.

Both the angular asymmetry and the radial stratification are persistent features that arise in all models regardless of the injection parameters, and represent a robust prediction of our 3D RMHD jet simulations in realistic BNS environments. As discussed before and as shown in the resolution study of Appendix B, the simulations presented here have not yet reached full numerical convergence, and the fine details of the jet structure retain some dependence on the grid resolution. The qualitative features described in this section are, however, robust and hold across all models.

6.2 From Simulation to Observable: the Modelling Pipeline

Given the three-dimensional jet structures extracted in the previous section, the next step is to determine what an observer at a given viewing angle would see. This section describes the pipeline connecting the RMHD simulation output to synthetic multi-wavelength afterglow light curves, and establishes the reference framework for the analysis that follows.

The afterglow emission is computed using the semi-analytical model of Salafia et al. [193], with the two extensions developed by O. S. Salafia and described in Section 2.8: (i) the inclusion of the full radial velocity profile $dE_{\rm iso}/du\,(\Theta, \Phi)$ extracted from the simulation at $t_{\rm end}$, and (ii) the relaxation of the assumption of axial symmetry, allowing the computation of light curves from the full 3D energy and velocity distributions including azimuthal variations. A key feature of this pipeline is that the angular and velocity structures entering the afterglow calculation are not parametric assumptions but direct outputs of the RMHD simulations. This distinguishes our approach from many previous structured-jet afterglow studies that relied on analytical profiles.

The input to the afterglow code is the total energy distribution $E_{\rm iso}(\Theta, \Phi, u)$, including the rest-mass component, at $t_{\rm end}$. The free parameters of the model are the electron energy fraction ε_e, the magnetic energy fraction ε_B, the electron power-law index p, the ISM number density n, the luminosity distance d_L, the redshift z, and the viewing angles (θ_v, ϕ_v). The output is the flux density $F_\nu(t)$ at a chosen observing frequency ν and viewing angle.

To isolate the effects of different structural features gradually, we perform most of the analysis using the azimuthally averaged profiles. The full 3D structure is employed when studying the impact of azimuthal asymmetries and when performing the extended parameter estimation described in Section 6.4.3.

As a first illustration, Figure 6.7 presents the light curves for model F in the optical, radio, and X-ray bands, for both an on-axis and an off-axis observer. The purpose of this comparison is to show that the structural features discussed in this chapter leave qualitatively similar signatures across all wavelengths. Given the superior data coverage in the radio band, all subsequent analysis is presented at 3 GHz frequency for clarity. We note that the MCMC fitting procedure described in Section 6.3.1 takes into account the full multi-wavelength dataset.

6.3 Application to GW170817

6.3.1 Dataset and Fitting Procedure

A central element of the approach presented in this thesis is the direct comparison with observational data, which also determines the reference afterglow parameters adopted throughout the following analysis. We perform this comparison by fitting synthetic light curves to the multi-wavelength afterglow data of GRB 170817A, the short GRB associated with the gravitational wave event GW170817.

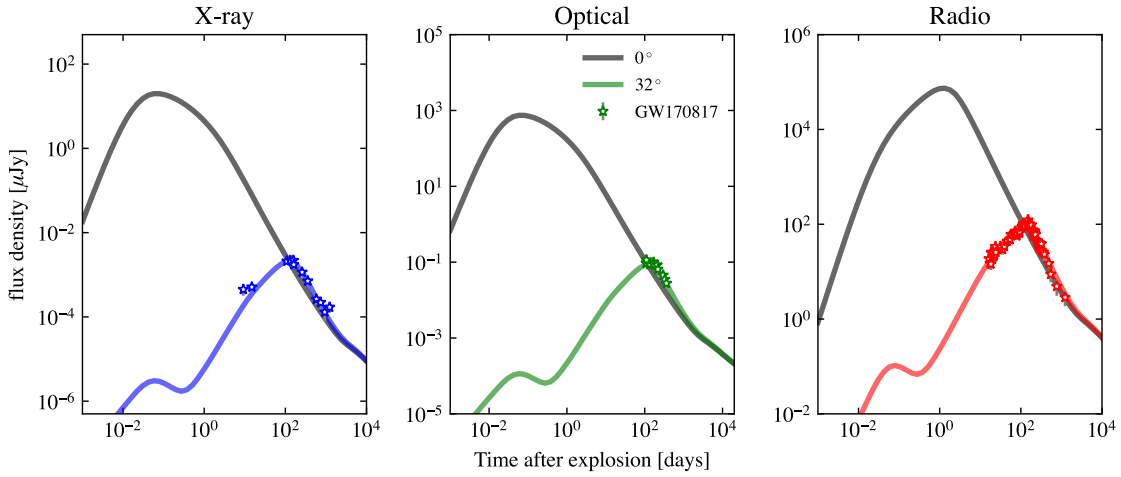


Figure 6.7: Multi-wavelength afterglow light curves for model F, at viewing angles $\theta_{\rm v} = 0°, 32°$.

We employ the observational dataset compiled by Makhathini et al. [111], which includes radio, optical, and X-ray data spanning the interval from 0.5 to 940 days post-merger. The radio flux densities are drawn from the compilation of Mooley et al. [138] and references therein. The optical detections (HST/F606W) are from Fong et al. [57]. The radio data further include reprocessed observations originally reported by Resmi et al. [185], Margutti et al. [114], and Alexander et al. [7], and the X-ray and additional optical (HST/F814W) data have been reanalysed by Makhathini et al. [111] to ensure a consistent flux determination across all bands.

The fitting is performed through a Markov Chain Monte Carlo technique using the `emcee` Python package [58]. The source is placed at the redshift $z = 0.0096$ and luminosity distance $d_L = 40\,\mathrm{Mpc}$ [32]. We fix the electron energy fraction to $\varepsilon_e = 0.1$, based on theoretical expectations from particle-in-cell simulations of relativistic shocks [204], and the electron power-law index to $p = 2.166$, as constrained by multi-wavelength spectral observations of the GW170817 synchrotron emission [57]. This leaves three free parameters for the axisymmetrised fits: n, ε_B, and θ_v; or four when the azimuthal angle is included in the analysis (n, ε_B, θ_v, ϕ_v).

We adopt the following priors. The ISM density is assigned a log-uniform prior in the range $n \in (10^{-5}, 10^{-1})\,\mathrm{cm}^{-3}$, where the lower limit reflects the smallest densities expected in galactic outskirts [56] and the upper limit is motivated by observational constraints on the GW170817 environment [80]. The magnetic energy fraction has a log-uniform prior in $\varepsilon_B \in (10^{-8}, 10^{-1})$, bounded below by a loose requirement that the post-shock magnetic field exceeds typical intergalactic values and above by PIC simulation expectations [204]. The viewing angle prior is isotropic, i.e. uniform in $\cos\theta_v$ between 0 and 1.

The MCMC chains are evolved for 5000 iterations. This number was determined by progressively increasing the chain length up to 10 000 iterations and verifying that the posterior distributions had already stabilised at 5000 steps, thus providing a good balance between accuracy and computational cost.

Table 6.1: Results of the MCMC fitting of each model to GW170817 afterglow data. All fits use the axisymmetrised jet structure at $t_{\rm end}$. Values are the median and 90% credible interval of the marginalised posterior. Fixed parameters: $\varepsilon_e = 0.1$, $p = 2.166$, $z = 0.0096$, $d_L = 40\,\mathrm{Mpc}$.

Model	$\log(n/\mathrm{cm}^{-3})$	$\log \varepsilon_B$	θ_v [deg]
F	$-1.99^{+0.3}_{-0.39}$	-3.49 ± 0.12	$32.0^{+3.3}_{-3.4}$
a	$-1.53^{+0.40}_{-0.39}$	-4.25 ± 0.12	$32.9^{+3.4}_{-3.5}$
b	$-1.69^{+0.48}_{-0.46}$	$-4.24^{+0.13}_{-0.15}$	$30.9^{+3.8}_{-3.7}$
c	-2.05 ± 0.39	$-3.53^{+0.12}_{-0.13}$	$31.3^{+3.3}_{-3.1}$
d	-2.41 ± 0.35	-3.16 ± 0.12	$28.7^{+2.9}_{-2.5}$
e1	$-1.69^{+0.29}_{-0.27}$	$-3.05^{+0.10}_{-0.09}$	$38.4^{+2.8}_{-2.7}$
e2	$-2.46^{+0.37}_{-0.39}$	-3.33 ± 0.12	$27.0^{+2.8}_{-2.7}$

6.3.2 Best-Fit Results and Parameter Constraints

We first present the results of the MCMC fits performed using the axisymmetrised jet structures of all seven models. The posterior probability distributions are summarised in Table 6.1, and the corresponding corner plots are shown in Figures 6.8 and 6.9.

A general trend emerges across the parameter space: the fits favour values of $\log \varepsilon_B \sim [-4, -3,]$ combined with relatively high ISM densities (some models show $n > 0.01\,\mathrm{cm}^{-3}$). This combination is in some tension with observational indications of a low-density environment around GW170817 [80]. The best-fit viewing angles fall in the range 27°–33°, for most models, with the exception of model *e1*, which yields $\theta_v \simeq 38°$, reflecting the qualitatively different angular structure of this marginal jet.

We note that a more refined estimate of θ_v should also account for the constraint provided by the apparent superluminal motion of the afterglow centroid, as measured with radio very long baseline interferometry [68, 69, 77, 138, 139, 145]. For GW170817, the viewing-angle range found here is generally inconsistent with the centroid-motion constraints ($\sim$ 19°–25°). In the context of a comprehensive parameter study, modelling the centroid displacement alongside the light curve could change these results, potentially ruling out some jet structures that would otherwise reproduce the observed emission.

An important limitation of the current analysis is the well-known degeneracy between afterglow parameters. In particular, n and ε_B can be strongly correlated, as the peak flux depends on the combination $n^{1/2}\varepsilon_B^{1/2}$: increasing n, and hence the number of emitting electrons, can therefore be partially compensated by decreasing ε_B, which reduces the magnetic-field strength, and vice versa. The electron energy fraction ε_e can likewise be correlated with other model parameters through its effect on the characteristic synchrotron frequency, while p is primarily constrained by the observed spectral slope. We therefore fix ε_e and p using independent observational and theoretical constraints. These degeneracies allow all jet structures produced by our simulations to fit the GW170817 afterglow data within the current uncertainties and priors used. Therefore, afterglow fitting alone cannot readily distinguish between the different injection models studied here.

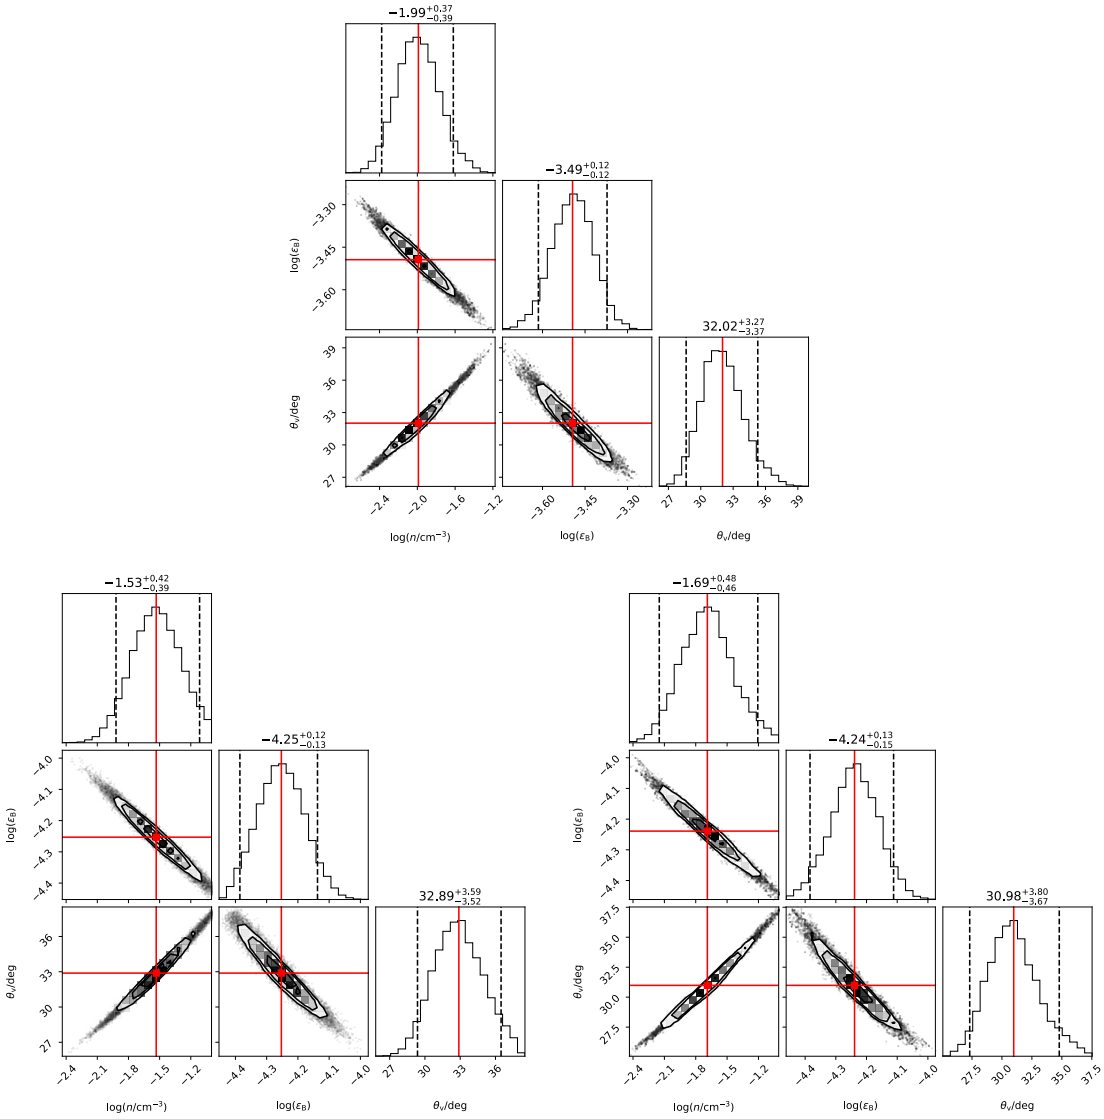


Figure 6.8: Corner plots of MCMC fit results for models F, a, and b to the GW170817 data. Each panel shows marginalized one-dimensional histograms and two-dimensional contour posteriors of the afterglow jet-model parameters (F: top; a: bottom left; b: bottom right). Vertical black dashed lines mark the 5th and 95th percentiles, and red solid lines mark the medians.

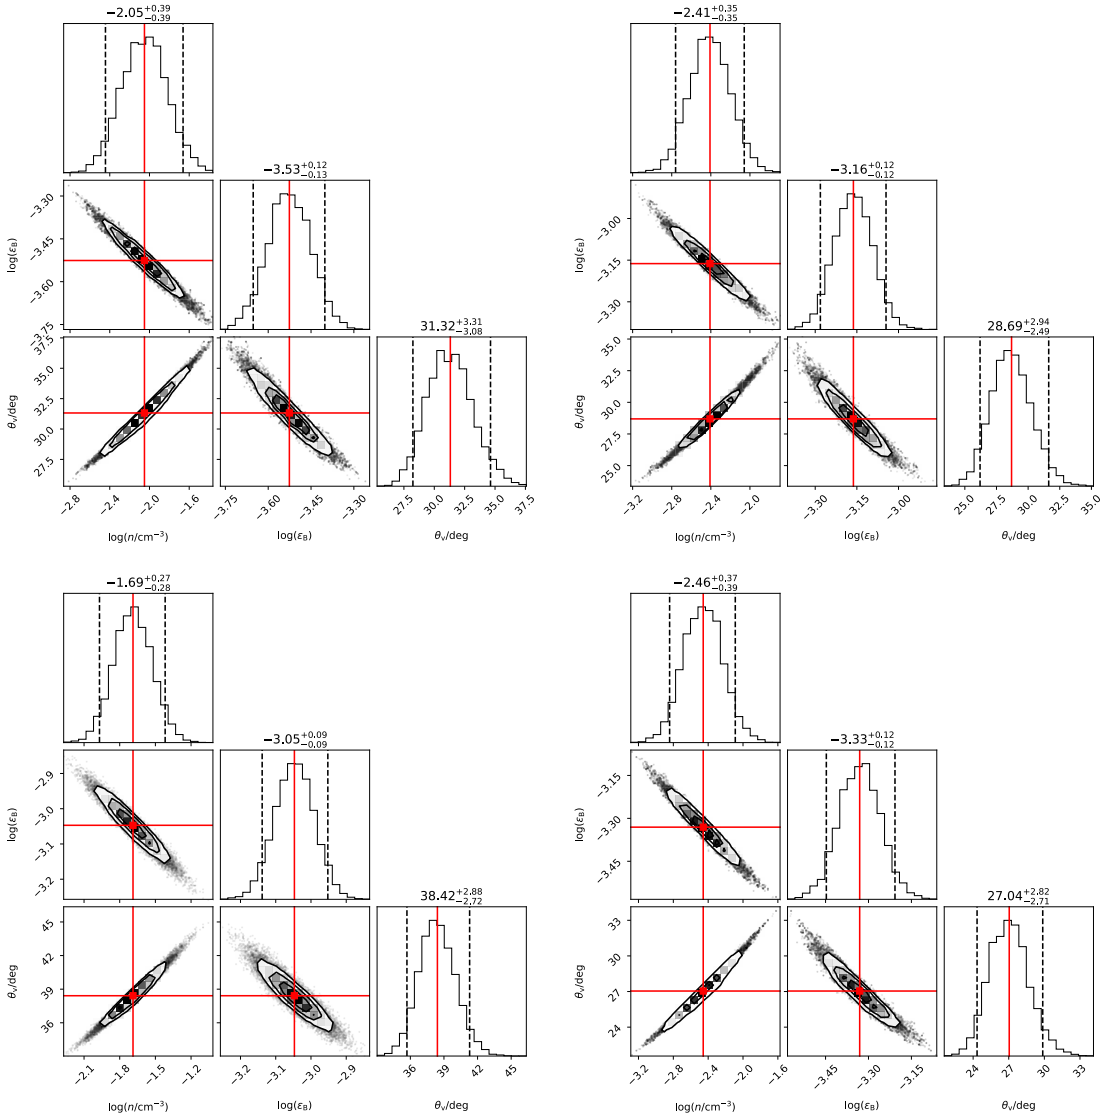


Figure 6.9: Corner plots of MCMC fit results of models *c, d, e1* and *e2* to the GW170817 data.

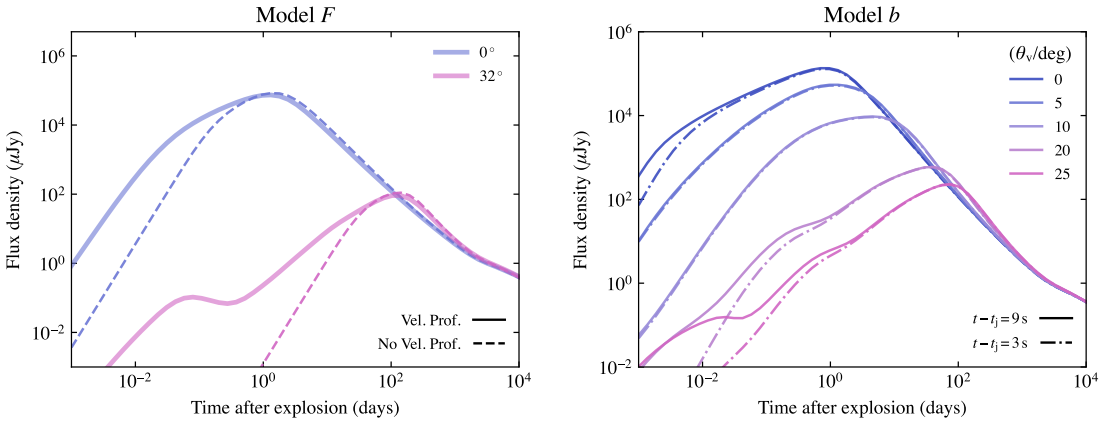


Figure 6.10: Afterglow light curves at 3GHz for model F (left) with or without including the velocity profile, and for model b (right) at $t_{\rm beg}$ and towards $t_{\rm end}$, for different observing angles.

6.4 Observational Signatures of the 3D Structure

6.4.1 Impact of Velocity Stratification

We begin by assessing the effect of including the radial velocity profile on the predicted afterglow emission. The left panel of Figure 6.10 shows the comparison, for the reference model F, between light curves computed with and without the velocity stratification described in Section 2.8.1. The comparison is shown for both an on-axis ($\theta_v = 0°$) and an off-axis observer at the GW170817 best-fit viewing angle, adopting the reference values of the microphysical parameters obtained from the MCMC fit to GW170817 data (Table 6.1, Section 6.3.2). This choice provides a physically motivated reference for a consistent comparison across models and viewing angles.

The light curves without the velocity profile correspond to a single-shell approximation, where all the outflow material at a given angle is assigned a single, average Lorentz factor computed from the total energy of the outflow.

The most prominent difference between the two prescriptions appears in the early, rising phase of the off-axis light curve. In the single-shell case, the flux rises steeply as the relativistic beaming cone of the decelerating blast wave opens towards the observer. When the full velocity distribution is included, the rise is significantly more gradual. The physical interpretation is that the energy is distributed over a range of four-velocities at every angle. Slower inner shells, which lag behind the fastest material, progressively catch up with the decelerating forward shock and deposit additional energy into it. This produces a contribution to the emission that is received at distinct observer times, such that the measured flux increases as a superposition of contributions arriving over an extended interval.

Differences also appear in the on-axis light curve, although they are smaller. Even within the jet core, the Θ–u distributions (Section 6.1.3) show a non-negligible amount of energy carried by slower material ($\Gamma \sim 3$–10). The presence of these inner, slower shells arriving at the forward shock at progressively later times produces a smoother and more extended rise also for the on-axis case.

For the off-axis observer, an additional feature becomes visible: a faint, early peak preceding the main afterglow maximum. This emission originates from the quasi-spherical outer shell described in Section 4.2.2, which propagates at mildly relativistic velocities. Along the jet axis, this component is entirely outshone by the core emission. For viewing angles $\theta_v \gtrsim \theta_j$, where the observer has no direct line of sight to the ultra-relativistic core, the emission from the outer shell provides the leading contribution at very early times. This precursor peak is present in all models and its presence is confirmed by the resolution study of Appendix B, however a detailed characterisation of its structure and energy content would require complete numerical convergence.

While its low flux and early occurrence make this component challenging to observe, it represents an interesting prediction of structured jet models that may become accessible with future observatories operating at higher sensitivity. Appendix C is dedicated to a discussion on the detectability of this feature.

At late times ($t \gg t_{\rm peak}$), the two prescriptions converge: all shells have contributed to the blast wave by then, and the asymptotic decay is governed by the standard Blandford-McKee dynamics regardless of the initial velocity distribution.

6.4.2 Impact of Residual Evolution

As discussed in Chapter 5, the evolution between $t_{\rm beg}$ and $t_{\rm end}$ is characterised by a progressive approach to the ballistic regime, during which the remaining thermomagnetic energy is converted into kinetic form and the angular structure stabilises. By $t_{\rm end}$, the kinetic energy fraction has reached 95%–97% in all models, and the core structure is largely frozen. The residual changes during this evolution are mostly related to the redistribution of material at large polar angles ($\Theta \gtrsim 10°$), where the outer shell and cocoon interface are still settling, and to the remaining acceleration process in the slowest-converging models.

To verify the impact of these residual modifications on the predicted afterglow, we compare light curves computed from the axisymmetrised profile of model b at $t_{\rm beg}$ and at $t_{\rm end}$ at different viewing angles (right panel of Figure 6.10). We choose model b because it is the slowest in approaching the ballistic regime, owing to its longer injection decay timescale τ_d, and therefore provides a conservative upper bound on the expected differences within our model suite.

The post-peak decay is effectively identical between the two snapshots, being insensitive to the details of the input structure. The early rising phase, instead, shows visible differences for off-axis observers. As stated before, the rising part of an off-axis light curve is shaped by the energy distribution at the angles that first enter the observer's beaming cone as the blast wave decelerates. These are precisely the intermediate-to-large angles ($\Theta \gtrsim 10°$) where the outflow structure is still evolving between $t_{\rm beg}$ and $t_{\rm end}$, as the cocoon material and the outer shell continue to rearrange. The jet core, by contrast, is already mostly stable at $t_{\rm beg}$, which is why the on-axis curve is much less affected.

The magnitude of this effect is moderate: the differences are smaller than those produced by, for example, the inclusion of velocity stratification or the choice of azimuthal viewing angle (a dedicated analysis is present in the next Section).

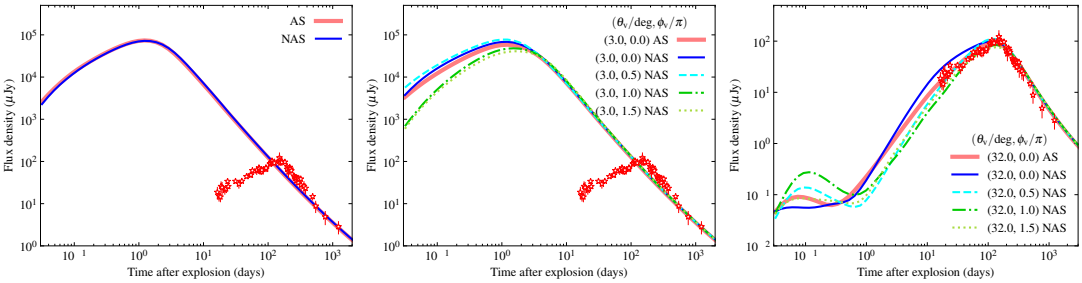


Figure 6.11: On-axis, slightly off-axis and off-axis (left to right) afterglow light curves at 3GHz for model F, for different polar and azimuthal observing angles.

Nevertheless, this comparison confirms that extracting the jet structure shortly after break-out can affect the predicted rising phase.

6.4.3 Impact of Azimuthal Asymmetries

We now turn to the observational consequences of the azimuthal asymmetries documented in Section 6.1.1, using model F as the reference case.

For an on-axis observer ($\theta_v = 0°$), the comparison between the axisymmetrised (AS) and non-axisymmetrised (NAS) light curves shows negligible differences (left panel of Figure 6.11). An on-axis observer initially receives emission primarily from the nearly uniform jet core, where the angular energy distribution is approximately flat and symmetric. As the blast wave decelerates, the relativistic beaming cone widens and progressively includes emission from the more asymmetric wings. Although these outer regions progressively contribute, their effect on the total flux is typically modest and at later times it is increasingly governed by the overall blast-wave dynamics rather than the detailed angular structure.

For a slightly off-axis observer ($\theta_v = 3°$), the situation changes substantially. The middle panel in Figure 6.11 compares the AS and NAS light curves at four azimuthal viewing angles $\phi_v = 0, \pi/2, \pi, 3\pi/2$. Order-of-magnitude differences appear in the rising phase, reducing towards the peak flux. This reflects the strongly non-uniform azimuthal distribution of energy already at these intermediate angles.

At a larger viewing angle of $\theta_v = 32°$, corresponding to the best-fit value for model F with respect to GW170817 data, the azimuthal dependence remains significant (right panel of Figure 6.11). The early-time precursor peak associated with the outer shell (Section 6.4.1) undergoes substantial variations with ϕ_v, demonstrating that while this component is quasi-spherical in morphology, its energy distribution retains considerable azimuthal structure.

Figure 6.12 extends this comparison to all models, confirming that the different structural features produced by the various injection parameters leave distinct imprints on the light curves at different azimuthal angles. The level of asymmetry and the energy content of the outer shell vary across the parameter space, translating into model-dependent azimuthal variations in the predicted emission. A model-by-model comparison, at the chosen numerical resolution, is presented in the next section.

As a further quantification of the impact of asymmetries on parameter estimation,

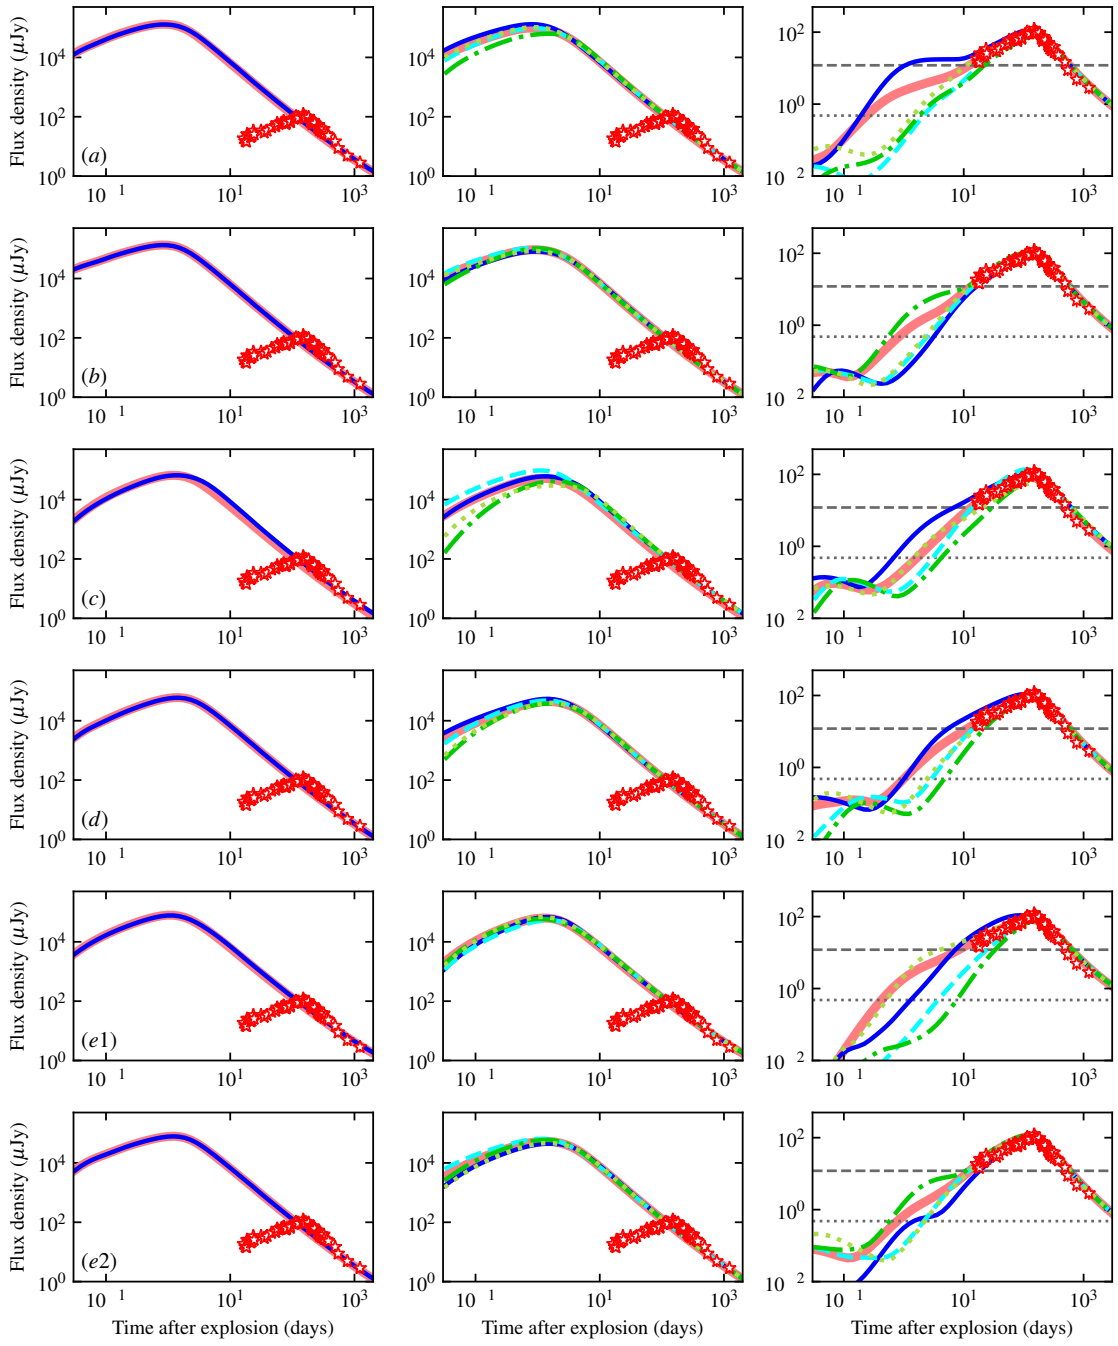


Figure 6.12: Same as figure 6.11, for all models of the parameter study. Horizontal lines on the off-axis light curve mark the 3σ sensitivities of the VLA (dashed) and ngVLA (dotted) at 3 GHz for a fiducial 1 hr on-source integration (see Appendix C).

we performed MCMC fits to the GW170817 data using the full 3D (NAS) structure of model F at three azimuthal viewing angles ($\phi_v = 0, \pi/2, \pi$). The setup follows the procedure described in Section 6.3.1. The resulting posterior distributions for the viewing angle are $\theta_v \simeq 41.6° \pm 3°, 26.3° \pm 3°, 23.9° \pm 3°$ for the three values of ϕ_v, respectively (Figure 6.13). All are incompatible at the 90% credibility level with the value obtained from the axisymmetrised fit ($\theta_v \simeq 32° \pm 3°$). Similar shifts appear in the estimates of ε_B and n. The NAS fits yield systematically larger values of $\log \varepsilon_B$ (ranging from -2.94 ± 0.12 to -2.41 ± 0.17, compared to -3.49 ± 0.12 in the AS case) and comparable or smaller values of $\log n$ (as low as -3.36 ± 0.50).

We emphasise that the intrinsic asymmetry described here is a robust feature that appears across the tested models and resolutions. We conclude that azimuthal asymmetry is a commonly under-explored feature that introduces an additional degree of freedom in the interpretation of afterglow data from structured jets.

6.4.4 Impact of Injection Parameters

Figure 6.14 presents the comparison of the afterglow emission across all jet models. To limit the number of simultaneous degrees of freedom, this comparison is based on the axisymmetrised jet structures; the additional variations arising from the full 3D structure were presented in the previous section. Both on-axis and off-axis light curves are shown, computed with a fixed set of afterglow parameters corresponding to the best-fit values of model F (Table 6.1), to enable a direct comparison of the structural effects.

The on-axis light curves show that more successful models (*b, a, e2*) produce a brighter peak at a later time with respect to the others. This scaling is a direct consequence of the dependence of the characteristic peak time on the jet energy through the Sedov length R_S, which sets the fundamental length scale of the blast wave dynamics [23]:

$$t_{\rm peak} \propto \frac{R_S}{c} \propto \left[\frac{E_{\rm iso}(\Theta = 0)}{n}\right]^{1/3} . \tag{6.1}$$

A larger R_S corresponds to a greater deceleration radius, meaning that the forward shock sweeps up more ambient material before reaching the deceleration point, resulting in both a later and a brighter peak.

Jet structures with steeper energy decay outside the core produce a correspondingly steeper rise of the light curve, as more energy enters the observer's line of sight when the decelerating outflow widens its beaming cone. The multi-segment structure of the azimuthally averaged energy profiles (Figure 6.4), characterised by different power-law slopes at different angular ranges, translates directly into changes of slope and breaks in the rising part of the light curves.

Models *c* and *d* differ very little from the reference case *F*, consistently with the minor variations in the injection parameters. The differences are most visible in the properties of the outer shell emission, indicating that the initial magnetisation, over the range explored here, has a limited impact on the overall afterglow morphology while still influencing the details of the outermost shock structure.

Models *e1* and *e2* provide an informative contrast. Both feature lower-luminosity jets, but differ in their launch time. Model *e2*, launched earlier, traverses a less

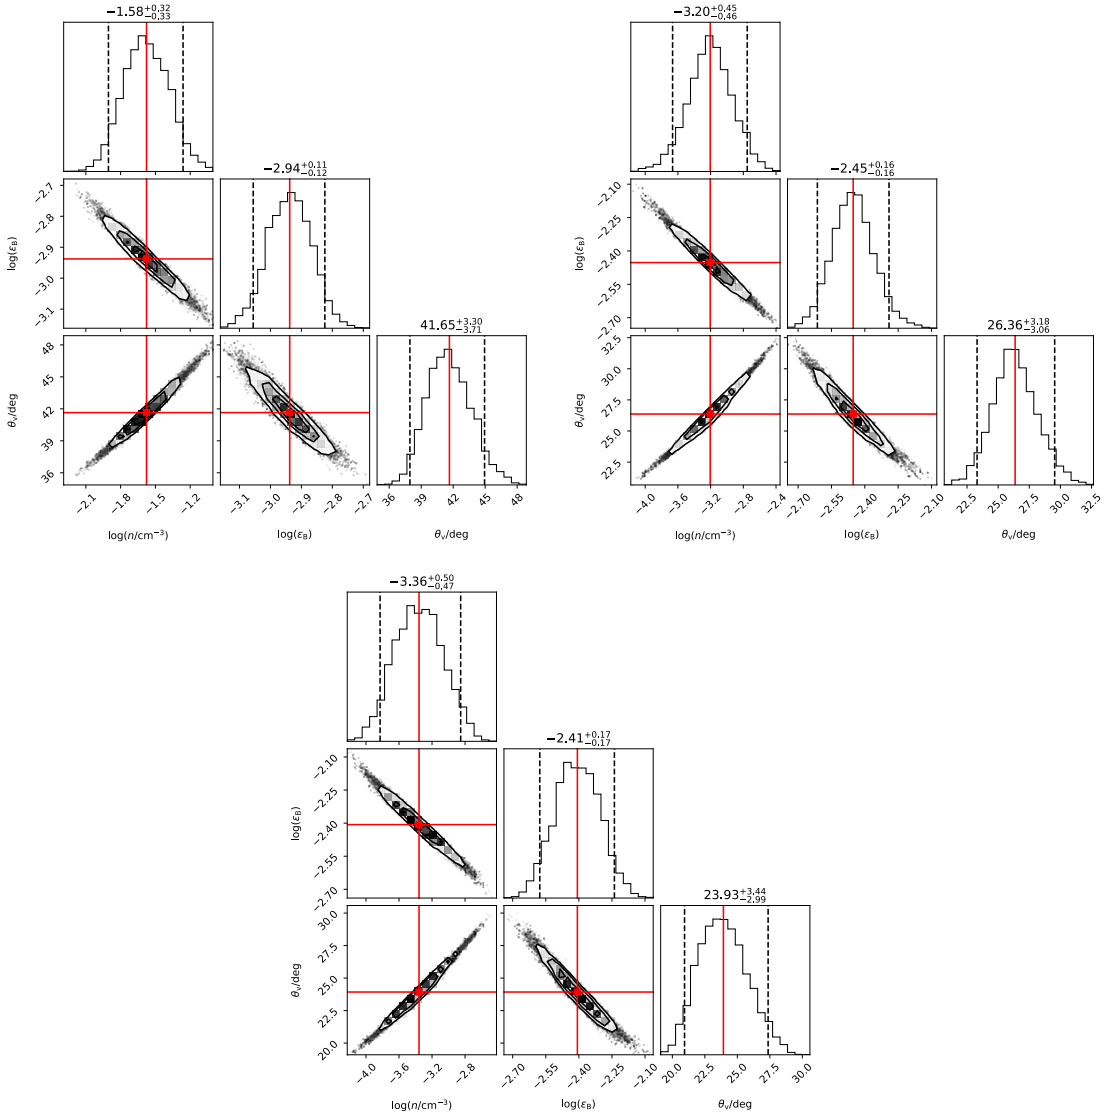


Figure 6.13: Corner plots of MCMC fit results of model F to the GW170817 data, considering the non-axisymmetrised jet structure and different azimuthal viewing angles, i.e. $\phi_{\rm v} = 0$ (upper left), $\phi_{\rm v} = \pi/2$ (upper right), $\phi_{\rm v} = \pi$ (bottom).

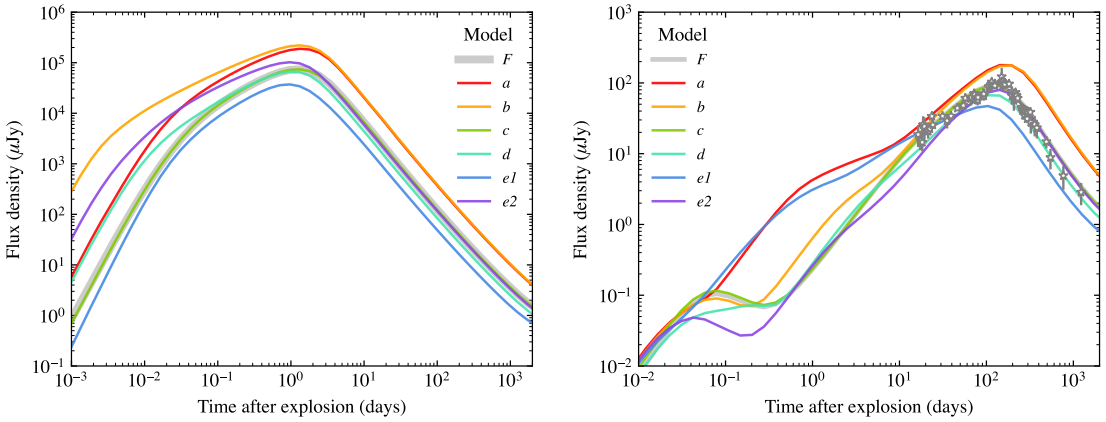


Figure 6.14: On-axis and off-axis ($\theta_{\rm v} = 32°$) light curves for all models of the parameter study, for the fixed set of afterglow parameters corresponding to the best-fit of model *F*.

extended ejecta environment and achieves a more complete breakout. The resulting afterglow is comparable to those of models *a* and *b*, with a well-defined peak and a standard evolution. Model *e1*, launched later and experiencing a much stronger interaction with the environment, lacks a well-defined energetic core and shows a qualitatively different rising phase: an extended, dimmer plateau replaces the typical peak, reflecting the broad and shallow angular energy distribution[2].

We conclude that the qualitative trends described above are consistent with the physical expectations for the chosen parameter variations, and the overall comparison across models provides useful insight into the sensitivity of the afterglow emission to the jet initial conditions.

[2]See, however, Appendix B.

Chapter 7

Gamma-ray Burst Jets in Core-collapse Supernova Environments

7.1 Introduction and Scientific Context

The core collapse of a massive star is another progenitor channel for GRB jets, introduced in Section 2.2.1 and mostly linked to long-duration bursts. In the collapsar scenario, the iron core of the progenitor can collapse to a black hole surrounded by an accretion disc, and the same central-engine physics discussed in Section 2.3 for the merger case, namely high angular momentum and strong magnetic fields powering a jet through the Blandford-Znajek mechanism or neutrino annihilation, applies here as well. The jet must then drill through the remaining stellar envelope before it can break out and produce the prompt γ-ray emission at large radii, following the same breakout physics introduced in Section 2.5.

The propagation of relativistic jets through stellar envelopes has been studied numerically by many authors, starting from the pioneering hydrodynamic simulations of MacFadyen and Woosley [109] and Aloy et al. [8], and extended over the years to magnetised jets and to the dynamical ejecta of binary neutron star mergers in the short GRB context [e.g. 66, 166]. Almost all of these simulations share a common simplification: the stellar environment is taken either directly from a one-dimensional pre-collapse stellar model in hydrostatic equilibrium, or from simple analytic density profiles. The role of the accompanying supernova has subsequently been addressed in a few studies. For example, De Colle et al. [41] used 2D axisymmetric relativistic hydrodynamic simulations in which a parametrised supernova, injected using a piston model, is followed by a relativistic jet. They showed that the outcome depends on the ordering of the cocoon breakout time, the supernova breakout time, and the time required for the cocoon to engulf the star. Urrutia et al. [219] extended this study to 3D, finding that a jet launched 1 s after the supernova initially propagates through the low-density cavity produced by the expanding supernova ejecta, before encountering the supernova shock front and being strongly decelerated. In both studies, the simulations are purely hydrodynamic and the supernova is injected parametrically into a pre-existing stellar progenitor profile. However, if the central

engine forms after a successful (or partially successful) magnetorotational explosion, the jet is launched into a medium that has already been profoundly reshaped by the supernova shock. The innermost envelope has been partially unbound and accelerated, strong anisotropies have developed along the rotation axis, and the density and pressure profiles can differ substantially from hydrostatic ones. This is the same conceptual gap identified for the merger channel in Chapter 1: idealised environments miss the asymmetries that a realistic progenitor imposes on the jet, and the same three-dimensional treatment developed throughout this thesis is applied here to a different progenitor.

Here, we assume that after the collapse of the proto-neutron star to a black hole, the accreting black hole launches a relativistic jet that is evolved within the three-dimensional post-explosion environment produced by a neutrino-magnetohydrodynamic core-collapse simulation, until (possible) breakout from the stellar envelope. The work presented here is still in a preliminary stage. The relativistic jet simulation, and its jet-off counterpart for comparison, should be considered as first working tests in this context. They are not yet performed in numerical convergence regime, and the jet simulation is subject to some numerical errors internally handled by the code, which will be tackled in future, more resolved, productions. The results presented should therefore be taken as indicative behaviours, more than fully reliable diagnostics. Future developments also include an exploration of the impact of different injection parameters, from canonical powerful jets to choked ones, with different possible interactions between the jet and the supernova shock. Moreover, the afterglow emission will eventually be predicted from the simulation output, following the same philosophy of the semi-analytical framework used in Chapter 6. The novelty of this study consists in having a GRB jet that evolves in a stellar environment that is being actively crossed by a supernova shock, in full 3D, allowing us to assess the relative effect on the supernova shock evolution, and on the jet dynamics, of their interaction.

This chapter presents the physical and numerical setup and the main results obtained: the progenitor and the core-collapse simulation that provide the environment (Sections 7.2–7.3), the analysis of the initial data (Section 7.4), the remapping into the PLUTO code and the preliminary simulations (Section 7.5). We caution that no publication describes the specific snapshot used here as initial data at the moment. The characterisation presented here is part of the work of this thesis, and a detailed analysis performed by the authors of the core-collapse simulation is in preparation at the time of writing.

7.2 Progenitor Model

The progenitor considered is a $26\,M_\odot$ rapidly rotating Wolf-Rayet star from the grid of Aguilera-Dena et al. [4]. In the reference paper, the initial stellar models begin with uniform chemical composition and evolve quasi-chemically homogeneously throughout core hydrogen burning. Upon transition to core helium burning, they contract and experience significant rotationally enhanced mass loss. At the pre-supernova stage, hydrogen is essentially absent across all models, regardless of stellar mass. This depletion results from the combined effects of efficient internal mixing,

which transports most of the hydrogen into the core, and rotationally driven mass loss, which removes the remaining hydrogen from the surface.

For the series to which our progenitor belongs, the stars retain very little helium (less than about $0.06\,M_{\odot}$) regardless of their initial mass. The helium remains well mixed during helium burning, and the models never develop a distinct core-envelope structure or significant helium shell burning. Following core hydrogen exhaustion, all models undergo overall contraction. In standard massive-star evolution, a hydrogen-burning shell develops around the helium core at this stage and drives the envelope to expand into a supergiant. Here, the rapidly rotating, chemically homogeneous models retain essentially no hydrogen in the envelope by the time core H burning ends, so no strong hydrogen-shell energy source can develop and no envelope expansion follows. The absence of a hydrogen envelope drives continued contraction on a neutrino-mediated Kelvin-Helmholtz timescale, leading to spin-up and substantial centrifugal mass loss in the final evolutionary stages. The resulting pre-supernova progenitors are rapidly rotating, nearly pure carbon-oxygen stars, in line with previous studies of homogeneous evolution, and are plausible progenitors of Type Ic superluminous supernovae, broad-lined Type Ic supernovae, and long gamma-ray bursts [4].

This progenitor was evolved with the one-dimensional stellar evolution code MESA [168–170], which assumes hydrostatic equilibrium. The calculation is stopped just before the final collapse phase, when the hydrostatic assumption becomes inapplicable.

7.3 Core-collapse Simulation and Initial Data

The model is then evolved through core bounce and the launching of a magnetorotational explosion by a dedicated core-collapse simulation. The methodological framework and code closures follow those of Obergaulinger and Aloy [157], applied to the AguileraDena+2018 progenitor with parameters within the range explored in that series and its follow-ups [155, 156, 158]. The reference papers contain the detailed physics of core collapse in rapidly rotating, magnetised progenitors of this class; what follows is limited to the elements needed to interpret the snapshot we import into PLUTO, together with the numerical values we read directly from the snapshot files provided by the authors.

The simulation is performed with the Aenus-ALCAR code [88, 155], which solves the equations of special relativistic magnetohydrodynamics coupled to a spectral two-moment (M1) neutrino transport scheme. Gravity is included through an effective relativistic potential of the type described by Marek et al. [112]. The equation of state combines the hybrid EoS at low densities with a tabulated nuclear EoS at high densities (see Section 7.4.6). The pre-collapse model carries a large-scale poloidal magnetic field of the strength and geometry discussed in the reference paper, subsequently amplified during collapse and by MRI-driven turbulence in the post-bounce phase.

The magnetorotational explosion mechanism relies on a specific pre-collapse configuration, and the progenitor series adopted here meets its requirements in a coherent way. Pre-collapse differential rotation is not strictly demanded for post-collapse magnetic-field amplification, since collapse itself generates differential

rotation even from an initially rigid profile. It is, however, needed if the field is already strong by the time of collapse: sustained differential rotation through the late evolutionary stages is what allows magnetic-field amplification, through winding and through the magnetorotational instability, to act before bounce rather than only after it. Chemical homogeneity has a consequence on the magnetic field: efficient mixing suppresses the strong internal shear layers that would otherwise act as dynamo sites during the main-sequence and post-main-sequence phases, limiting how much the field can grow before collapse. The pre-collapse model therefore arrives at core collapse with sustained fast rotation and a magnetic field that is present but not yet saturated, both of which are needed for the MR mechanism to become effective in the post-bounce phase.

After bounce, the magnetic field is compressed and wrapped around the newly formed proto-neutron star, and is further amplified by MRI-driven turbulence and by neutrino-driven convection. The amplified field builds an ordered, funnel-shaped topology anchored on the PNS, which channels post-shock material towards the rotation axis and delivers energy to the stalled shock preferentially in that direction. Neutrino heating, magnetic stresses and centrifugal support therefore act constructively to revive the shock, with the last two enhanced by the fast rotation inherited from the progenitor. The resulting explosion is intrinsically anisotropic: the compact remnant is oblate rather than spherical, and the ejecta come out in a wide but far from isotropic cone, typically extending 30–35° around the rotation axis. Once the field grows strong enough to enforce approximate rigid rotation in the affected shells, further amplification through winding saturates.

For the specific run adopted here, magnetic stresses dominate the momentum budget of the explosion, but neutrino heating is important throughout the pre-explosion phase and remains non-negligible during shock revival. The stalled shock is revived approximately 400 ms after core bounce and expands slowly outward. The explosion energy at the time of black hole formation is $\sim 0.9 \times 10^{51}$ erg. This value is not the asymptotic explosion energy that would be reached in a run without black hole formation: the collapse of the proto-neutron star truncates the accretion power available to sustain the outflow, and an uninterrupted evolution would accumulate somewhat more energy.

The reference simulation is followed for approximately 1.7 s from core bounce, at which point the snapshot used as initial condition for PLUTO is extracted, a few milliseconds before black hole formation. Immediately after the central engine has switched from proto-neutron star to black hole plus (potential) accretion disc, the stellar envelope has not yet had time to react to the switch. This is the physical moment at which the assumption of a newly launched relativistic jet, imposed as an inner boundary condition in the following simulations, is best justified. The subsequent evolution of the stellar envelope, with and without an injected jet, is followed in PLUTO as described in the following sections.

7.4 Analysis of the Initial Data

We first characterise the environment through its density (Section 7.4.1), its velocity and rotation (Section 7.4.2), and its magnetic field configuration (Section 7.4.3); we

then identify the supernova shock (Section 7.4.4) and use these results to define the excision radius and the properties of the central engine that will be imposed at the inner boundary of the PLUTO runs (Section 7.4.5).

7.4.1 Mass and Density Distribution

Figure 7.1 shows the logarithmic mass density on the equatorial and meridional planes, at different spatial scales. The distribution at radii smaller than 1.2×10^9 cm is strongly anisotropic: the equatorial region is filled by material still bound to the proto-compact object, while along the polar direction the low-density, high-entropy channel opened by the magnetorotational explosion is clearly visible. The equatorial-to-polar density contrast at $r \sim 10^9$ cm exceeds one order of magnitude.

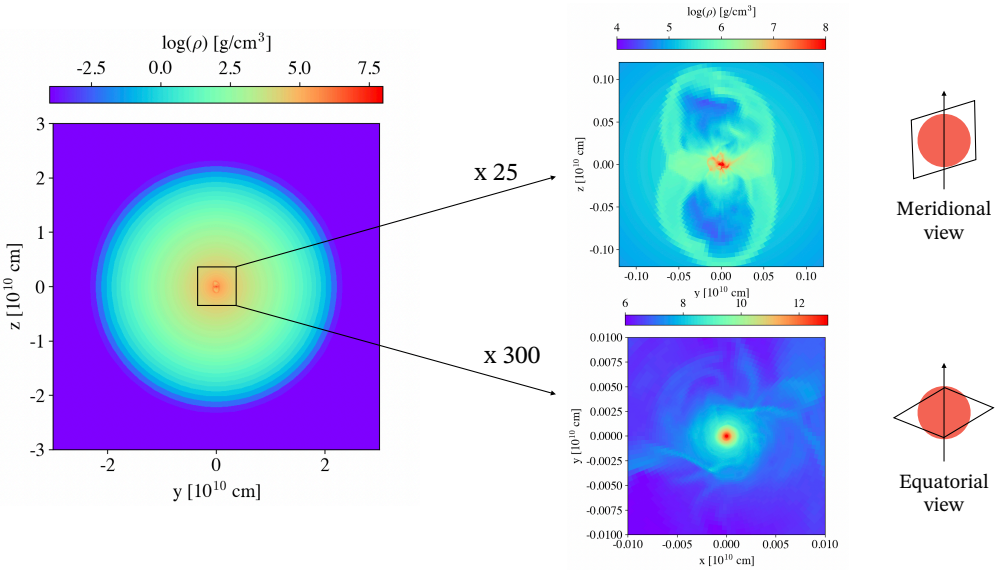


Figure 7.1: Slices of the logarithmic mass density from the Aenus-ALCAR snapshot used as initial condition, on the meridional and equatorial planes, at different levels of zoom.

Complementing the 2D visualisation, the left panel in Figure 7.2 shows the radial density profile along the z axis across the full radial extent of the AENUS-ALCAR simulation. A flat plateau at small radii ($r \lesssim 10^6$ cm) corresponds to the proto-neutron star, whose interior is not modelled in the RMHD approximation adopted here. Outside this plateau, the profile steepens and traces the envelope material relevant for the jet-propagation problem. To avoid contamination by the PNS interior, all radial profiles shown in the following are truncated at $r_{\rm min} = 10^6$ cm, and the same lower bound is applied consistently throughout the analysis.

The cumulative mass enclosed within radius r,

$$M_r \equiv M(<r) = \int_{r_{\rm min}}^{r} r'^2\, dr' \int_0^{\pi} \sin\theta\, d\theta \int_0^{2\pi} d\phi\; \rho(r', \theta, \phi), \tag{7.1}$$

gives a first estimate of the mass enclosed by any future excision boundary. The middle panel of Figure 7.2 shows that M_r increases smoothly and monotonically

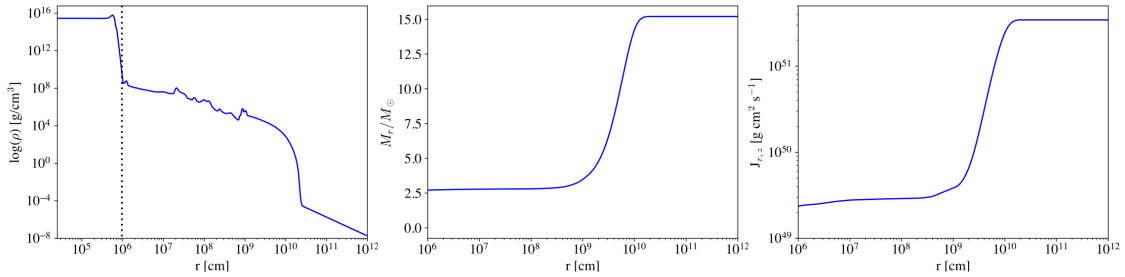


Figure 7.2: Left: Radial mass density profile along the polar axis. The flat inner plateau at $r \lesssim 10^6$ cm corresponds to the proto-neutron star, which is not part of the physical problem addressed here. Centre: cumulative mass profile M_r as a function of radius. Right: Cumulative angular momentum J_r as a function of radius.

between $r \sim 10^9$ and 10^{10} cm, so the enclosed mass is a well-defined function of the excision choice in this range.

7.4.2 Angular Momentum and Rotation

The velocity field is visually quantified by the distribution of the radial velocity of the shocked ejecta, shown in the middle panel of Figure 7.4. It shows that the outer polar layers are already expanding, while a large equatorial region is still infalling and feeds the central object.

To quantify the angular-momentum content we compute the cumulative axial angular momentum enclosed within radius r (aligned with the rotation axis),

$$J_{r,z} \equiv J_z(<r) = \int_{r_{\min}}^{r} r'^2\, dr' \int_0^{\pi} \sin\theta\, d\theta \int_0^{2\pi} d\phi\; \rho\, v_\phi\, r' \sin\theta. \tag{7.2}$$

The right panel of Figure 7.2 shows the cumulative angular-momentum profile, which follows the behaviour of the cumulative mass.

The dimensionless spin parameter of the material enclosed within r, evaluated using the axial component of the angular momentum (aligned with the assumed rotation axis of the eventual black hole),

$$a_*(r) = \frac{c\, J_{r,z}}{G M_r^2}, \tag{7.3}$$

stays close to 0.4 throughout the stellar interior. Evaluating Eq. (7.3) at $r = r_{\min} = 10^6$ cm, near the surface of the proto-neutron star, gives $a_* \simeq 0.36$; at $r = 10^8$ cm the value has risen to $a_* \simeq 0.4$, and it remains close to this value out to $r \sim 10^9$ cm. Equation (7.3) is a proxy for the spin of the enclosed material, however not a precise prediction of the final black-hole spin: only a fraction of that material will actually be accreted (part is expected to remain in a disc or to be ejected by disc winds during the launching phase), and angular momentum can be redistributed during accretion. Under the simplifying assumption that the material within $r_{\rm exc}$ is accreted with its angular momentum conserved, and given the near-constancy of $a_*(r)$ across the stellar interior, we adopt $a_* \simeq 0.4$ as an indicative spin estimate for the central engine in Section 7.4.5.

7.4.3 Magnetic Field Configuration

Figure 7.3 shows a two-dimensional slice of the magnetic field intensity together with a 3D rendering of the field streamlines close to the PNS. The field intensity is enhanced along the rotation axis and in the equatorial belt, the two regions where winding and magnetorotational amplification are most efficient. The ratio $p_{\rm mag}/p_{\rm th}$ stays below unity throughout the domain of interest: the environment is thermally supported and the field is dynamically subdominant in the bulk of the ejecta, though magnetic effects on the jet-medium interaction near the jet boundary will not be negligible. The 3D streamlines show a large-scale helicoidal component coexisting with a substantial toroidal wind-up in the equatorial belt.

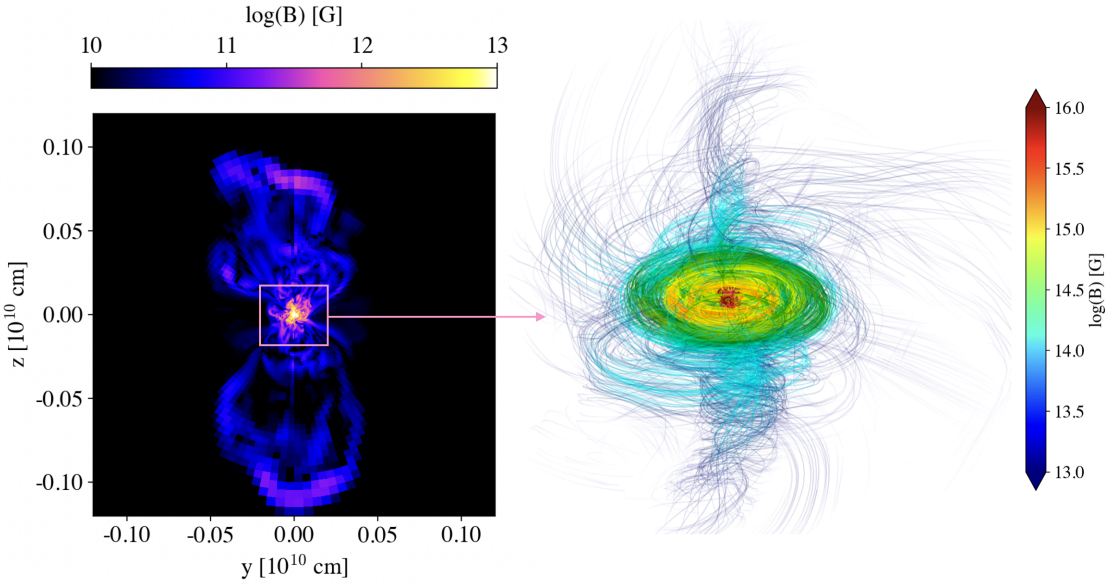


Figure 7.3: Two-dimensional slice in the y–z plane of the magnetic field intensity (left), and 3D rendering of magnetic field streamlines around the PNS (right).

For a divergence-free field, the signed integral of $\mathbf{B}\cdot\hat{\mathbf{n}}$ over a closed sphere vanishes identically, so the appropriate quantity for a Blandford-Znajek estimate is the flux through a single hemisphere. We compute the radial magnetic flux threading the northern hemisphere of a spherical surface of radius r as

$$\Phi(r) = \int_{\theta\le\pi/2} B_r(r,\theta,\phi)\, r^2 \sin\theta \, d\theta \, d\phi, \tag{7.4}$$

which matches the convention of the BZ formula recalled below. Given the black-hole spin estimate of the previous section, the associated Blandford-Znajek luminosity can be written as [212]

$$L_{\rm BZ} \sim \frac{\kappa}{4\pi c}\,\Phi^2\,\Omega_H^2, \qquad \Omega_H = \frac{a_*\, c}{2\, r_H(a_*)}, \tag{7.5}$$

where $r_H = (GM/c^2)\left[1+\sqrt{1-a_*^2}\right]$ is the horizon radius and $\kappa \sim 0.04$ is a numerical constant that depends on the field geometry. This form refines the original Blandford

and Znajek [24] split-monopole scaling, and is valid for moderate values of the dimensionless spin. The value of $\Phi(r)$ measured on a sphere well outside the excised volume is only a rough diagnostic of the flux that will eventually thread the horizon: near the black hole the flux can accumulate through MRI-driven amplification, but it can also reconnect, escape or fail to be accreted altogether. We therefore treat the value we obtain from Eq. (7.4) at $r_{\rm exc}$, and the corresponding $L_{\rm BZ}$, as illustrative estimates of the engine power that is not resolved in these simulations. Numerical values of Φ and $L_{\rm BZ}$ at the excision radius, together with the alternative luminosity estimate used to fix the fiducial jet power, are given in Section 7.4.5.

7.4.4 Shock Finding

The supernova shock is the dominant feature of the post-explosion environment.We identify it on the discretised grid by combining five proxies applied along every angular direction, each sensitive to a physical signature of a shock: entropy dropping steeply with radius, a non-negligible radial velocity, a compressive sign of its radial derivative, a steep radial drop of the pressure, and a decrease of the radial velocity across a spatial window wider than a single-cell finite difference. In terms of the dimensionless logarithmic slopes $d\ln X/d\ln r$, after smoothing the radial profiles over a few cells, the conditions read

$$\frac{d\ln S}{d\ln r} < -0.25, \qquad \frac{d\ln v_r}{d\ln r} < -1, \qquad \frac{d\ln p}{d\ln r} < -1, \tag{7.6}$$

$$v_r(i-3) > v_r(i+3), \tag{7.7}$$

where i is the radial cell index. The four logarithmic-slope thresholds are chosen operatively, at values consistent with the choices adopted in previous works using the same code. A cell is flagged as shocked when all conditions are simultaneously satisfied, and the supernova shock is defined as the outermost flagged cell along each angular direction. This multi-criterion definition rejects the numerous weaker features present at smaller radii and returns a well-defined supernova-shock surface even where the shock is oblique.

Extending the procedure to three dimensions confirms that the shock is highly anisotropic, spanning radii from $r \sim 6\times10^8$ cm along the equatorial direction, to $r \gtrsim 10^9$ cm in the polar one (see the right panel of Figure 7.4). This asymmetry is a signature of the magnetorotational explosion, which propagates more efficiently along the polar direction where magnetic stresses and neutrino heating are strongest.

7.4.5 Excision Radius and Central-engine Parameters

To follow the jet evolution for orders of magnitude in space by maintaining a sufficient resolution and a reasonable computational cost, as well as to exclude the effects of general relativity that are involved in the remnant collapse, but not included in the PLUTO code, we need to define an internal excision radius.

The diagnostics collected in the previous subsections bracket the excision radius from above and from below. The upper limit follows from the requirement that most of the supernova-shocked structure lies inside the computational domain, so

that the jet-shock interaction can be followed self-consistently; this bounds $r_{\rm exc}$ at $\lesssim 6 \times 10^8$ cm, the innermost point of the anisotropic shock front identified in Section 7.4.4. The lower limit follows from the region likely to be accreted onto the central compact object on a dynamical timescale, identified by the joint conditions on the sign and magnitude of the infall velocity discussed below.

Figure 7.4 combines three complementary diagnostics on the meridional plane: the shock indicator applied to the pressure distribution (right), the radial velocity in units of c (centre) identifying regions with negative velocity, and the sonic condition $|v_r|/c_s > 1$ (left). Where the infall is supersonic ($|v_r|/c_s > 1$), pressure gradients cannot decelerate the flow before it reaches the compact object, and the material is effectively accreted. Where the infall is subsonic ($|v_r|/c_s < 1$), the flow is redirected outward by pressure and advection rather than reaching the centre. The sonic condition therefore selects the region of the post-explosion environment that will contribute to the mass and angular momentum of the central object, as opposed to the region that will be advected outward with the ejecta. In the meridional plane this supersonically infalling region extends out to $r \sim 2$–3×10^8 cm along the equator, and defines the lower bound on $r_{\rm exc}$.

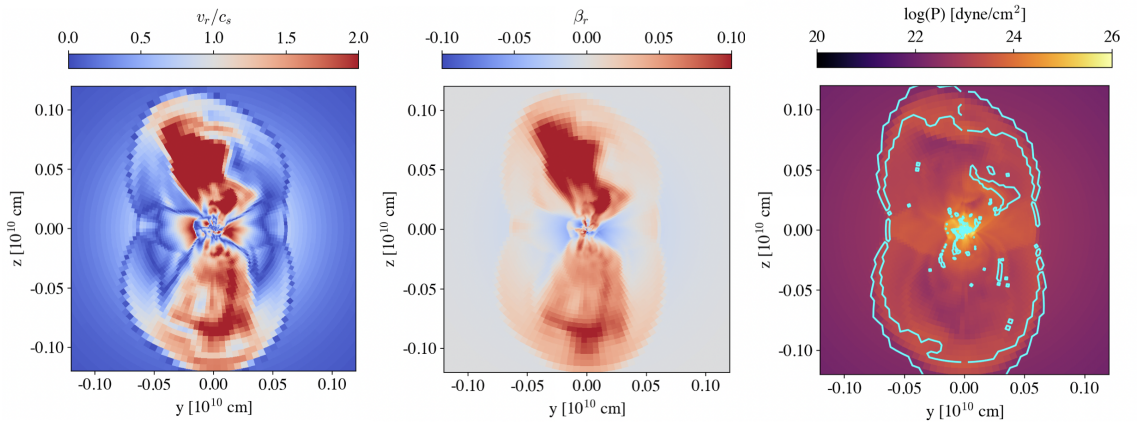


Figure 7.4: Diagnostics used to bracket the excision radius on the meridional plane. Left: sonic condition $|v_r|/c_s > 1$, tracing material likely to be accreted onto the central compact object. Centre: radial velocity v_r/c. Right: shock indicator overplotted on the pressure distribution.

We adopt $r_{\rm exc} = 3 \times 10^8$ cm. This value sits inside the upper limit by a comfortable margin, includes the equatorial region likely to accrete onto the central object (which we replace by an inner boundary carrying the mass of the enclosed material, see below), and excludes only material that is either bound to the proto-compact object or dominated by physics not captured by RMHD, in particular the innermost strongly gravitating region.

Central-engine parameters at $r_{\rm exc}$. The material enclosed within the excision boundary carries the mass, angular momentum and magnetic flux that will end up (partly) in the newly formed black hole and its accretion disc. Its properties determine both the mass and spin of the central engine and, together with the

magnetic flux available at that radius, an estimate of the Blandford-Znajek power. Evaluating Eq. (7.1) at $r_{\rm exc}$ gives an enclosed mass $M_{\rm exc} \simeq 2.3\,M_\odot$, adopted as the mass of the central point source used to set the Newtonian gravitational potential in the PLUTO runs (Section 7.5.1), and as the mass of the newly formed black hole for the purposes of the horizon-radius estimate[1]. From Section 7.4.2 we take $a_* \simeq 0.4$ as an indicative spin estimate for the material likely to be accreted (see the caveats there), well within the physical range and large enough to support a Blandford-Znajek engine. Evaluating Eq. (7.4) at the same radius gives $\Phi(r_{\rm exc}) \sim 10^{26}\,{\rm G\,cm^2}$; substituting into Eq. (7.5) yields $L_{\rm BZ} \sim 10^{47}\,{\rm erg\,s^{-1}}$. This value of luminosity is a lower-bound diagnostic rather than a prediction: since $L_{\rm BZ} \propto \Phi_{\rm BH}^2$, closing the gap with a canonical engine luminosity $L_j \sim 10^{50}\,{\rm erg\,s^{-1}}$ requires only a factor ~ 30 of amplification in the magnetic flux between $r_{\rm exc}$ and the horizon, plausibly achieved through the combination of geometric concentration in the polar funnel, flux advection, and MRI-driven amplification at the disc scale, none of which is resolved by this setup.

Threshold luminosity and adopted jet power. Rather than deriving L_j from $L_{\rm BZ}$, we set the target jet luminosity through an independent physical constraint. Following Aloy et al. [10], a jet injected into an ambient medium must exceed a characteristic luminosity in order for its head to propagate supersonically with respect to the surrounding material. For a jet of half-opening angle θ_j injected at radius r_{exc}, this threshold can be written, up to factors of order unity, as

$$L_{\rm thresh} \sim \pi r_{exc}^2 \theta_j^2 \langle\rho\rangle c_s^2 c, \tag{7.8}$$

that is, the product of an effective jet cross-section $\pi r_{exc}^2\theta_j^2$, an envelope enthalpy-density scale $\langle\rho\rangle c_s^2$, and the speed of light. Physically, it is the jet power at which the ram pressure delivered by the jet at its working surface matches the thermal pressure of the surrounding envelope. $\langle\rho\rangle$ and c_s characterize the ambient medium at the injection location. We therefore use $L_{\rm thresh}$ to set the characteristic initial jet luminosity.

It is important, however, that exceeding this luminosity does not by itself guarantee successful breakout. The engine must also remain active for a time comparable to or longer than the jet propagation time through the envelope. Thus, a successful breakout requires both

$$L_j \gtrsim L_{\rm thresh} \qquad \text{and} \qquad t_{\rm eng} \gtrsim t_{\rm b}, \tag{7.9}$$

where $t_{\rm b}$ is the jet breakout time, which itself depends on the jet luminosity and the density structure of the surrounding envelope. Consequently, $L_{\rm thresh}$ specifies

[1] This is a rough estimate, providing mostly an upper limit. Integrating the rest-mass density over the volume gives the baryonic mass M_b, while the gravitational mass M_g that sets the Newtonian potential and horizon radius is smaller by the fraction of rest-mass energy lost to gravitational binding. Part of this energy is lost during collapse to a black hole. In addition, not all the material enclosed in the selected radius is necessarily accreted. We expect however this discrepancy to be smaller than other uncertainties in the central engine, such as the accretion history inside the excised region and the magnetic flux at the horizon.

the minimum instantaneous power for supersonic propagation, while $t_{\rm eng}$ determines whether this power is sustained for long enough for the jet to break out.

Evaluating Eq. (7.8) on the snapshot analysed here, with $\theta_j = 10^\circ$ as adopted for the jet-injection setup (Section 7.5.3) and $\langle\rho\rangle$ computed as the angle-averaged density at the shock radius, we obtain $L_{\rm thresh} \sim 10^{50}\,{\rm erg\,s^{-1}}$. We adopt this value as the fiducial jet luminosity for the production run, so that the run probes the marginal-breakout regime. A study around this threshold, exploring the range $L_j \in [10^{48}, 10^{51}]\,{\rm erg\,s^{-1}}$ and its consequences for jet propagation, breakout and structure at large radii, is left as future work.

7.4.6 Equation-of-state Matching

The equation of state used in Aenus-ALCAR is hybrid, in the sense of combining two density regimes on separate tables [181]: the SFHo nuclear equation of state [83, 206] at high densities, and a low-density treatment accounting for baryonic, electronic and photonic pressure components. In PLUTO the immediately available options are the Taub-Matthews EoS [128] or an ideal gas with constant adiabatic index $\hat{\gamma}$. Some mismatch between the thermodynamic variables of the two codes is unavoidable, and needs to be characterised and minimised before the environment can be evolved on the PLUTO grid.

A first attempt with the Taub-Matthews EoS, initialised directly from the original primitive variables, produces an effective adiabatic index distribution centred around $\hat{\gamma} \approx 5/3$ rather than the $\sim 4/3$ of the input data. The two closures are therefore not consistent on the original primitive variables. An attempted workaround is to solve for the pressure that conserves energy and momentum between the two codes, keeping ρ, Γ and $\mathbf{B}$ fixed, by inverting

$$E_{\rm HD} = \rho c^2 \Gamma^2 h(P) - P, \qquad h(P) = \frac{5}{2}\frac{P}{\rho c^2} + \sqrt{\frac{9}{4}\left(\frac{P}{\rho c^2}\right)^2 + 1}, \tag{7.10}$$

where $h(P)$ is the Taub-Matthews specific enthalpy. This procedure raises the pressure but leaves the effective adiabatic index still concentrated around 5/3, and is therefore not able to reproduce the radiation-dominated regime of the input data.

A more direct alternative adopts an ideal-gas EoS with fixed $\hat{\gamma} = 4/3$ and solves the analogous energy-conservation inversion,

$$E_{\rm HD} = \rho c^2 \Gamma^2 h(P) - P, \qquad h(P) = 1 + \frac{\hat{\gamma}}{\hat{\gamma} - 1}\frac{P}{\rho c^2}. \tag{7.11}$$

The recovered pressure distribution matches the original one over the spatial region of interest. Visible differences appear only at radii excluded from the PLUTO domain, plausibly where the original data switch between EoS regimes and where the ideal-gas assumption breaks down. Moreover, imposing the same ideal-gas closure everywhere without applying any energy-momentum matching leaves the specific-enthalpy and total-energy-density differences with respect to the original data confined to $r < 10$ km, well inside the excised zone and outside the PLUTO domain. We therefore adopt an ideal-gas EoS with $\hat{\gamma} = 4/3$ for all subsequent runs, using the original primitive variables directly.

7.5 PLUTO simulations

7.5.1 Grid Remapping and Setup

The collapsed star and the supernova environment are imported into the PLUTO code [131] after excising the central region at $r_{\rm exc} = 3 \times 10^8$ cm (Section 7.4.5). Only the northern hemisphere ($z > 0$) is imported, and the data are rotated by 90° so that the jet is injected along the Cartesian y-axis rather than along the polar axis of the spherical grid, which is a coordinate singularity; the same strategy was adopted for the merger environment (Section 4.1.1). A gravitational potential from a central point mass of $M_{\rm exc} = 2.3\,M_\odot$ (Section 7.4.5) is imposed as a source term in the momentum and energy equations. The radial extent of the grid spans $r \in [3 \times 10^8, 5 \times 10^{12}]$ cm, discretised with 364 logarithmically spaced cells, and the two angular directions use 216 uniformly spaced cells each, the latter covering slightly less than π from the injection axis. The radial range is smaller than the extent of the original data, but it still contains all the stellar material of interest, and preserves resolution in the region where the jet-shock interaction will occur without inflating the total number of points beyond what is affordable given the timestep constraint. We perform two separate simulations, one including the injection of a relativistic jet as detailed in the next section, and a jet-off evolution of the stellar material only, as a control. The target injection time is 20 s, however the jet simulation was terminated shortly after $t \simeq 10$ s due to a numerical failure affecting the reconstruction of the primitive variables along the jet axis in the current setup, while the jet-off case was able to evolve smoothly for the full simulation time. Any comparison between the two is therefore restricted to their common evolution interval, and no conclusion about the asymptotic outcome of the jet is drawn from either simulation. Future work will extend this framework to track the subsequent expansion phase, possibly employing a configuration analogous to the grid-remapping technique developed for the merger scenario discussed in Chapter 5.

The rotation used to bring the jet axis onto the Cartesian y direction is a 90° rotation about the x axis. In Cartesian components, the new coordinates (x', y', z') are related to the original ones (x, y, z) by

$$(x', y', z') = (x,\ z,\ -y), \tag{7.12}$$

and the analogous transformation is applied cell by cell to any vector field, in particular the fluid velocity and the magnetic field:

$$(V^{x'}, V^{y'}, V^{z'}) = (V^x,\ V^z,\ -V^y). \tag{7.13}$$

Interpolation onto the target spherical grid then follows in the new coordinate system.

The magnetic field is the only remap ingredient for which cell-by-cell interpolation is not automatically consistent, because there is no guarantee that the discrete divergence of the interpolated field vanishes on the new grid. Two remap strategies are in principle available: a direct interpolation of the Cartesian components of $\mathbf{B}$ (faster, but subject to the introduction of monopoles during the remap), and a reconstruction based on a vector potential $\mathbf{A}$ such that $\mathbf{B} = \nabla \times \mathbf{A}$ on the target grid (solenoidal by construction, but computationally heavier and sensitive to the

interpolation of $\mathbf{A}$ itself). The runs presented here use the direct method. The residual divergence introduced by the interpolation is controlled during the subsequent evolution by PLUTO's hyperbolic divergence-cleaning scheme [42], which extends the induction equation by including a scalar potential ψ that is transported and damped on the fluid timescale, propagating any $\nabla \cdot \mathbf{B}$ error out of the domain at the fastest local signal speed.

The status of the magnetic divergence is monitored throughout the run. A small number of cells with strongly monopolar signatures are present already at the first output and persist during the evolution, mainly along the outer boundary and, more sparsely, in the ambient region outside the jet cone; these are the expected artefacts of the boundary treatment and of the initial direct interpolation on a discrete grid. The hyperbolic cleaning approach keeps the error under control and monitored, but its dynamical insignificance has not been demonstrated yet.

The details of the numerical setup are mostly identical to those used in the previous simulations (Chapter 4). Outside the jet nozzle at the excision boundary ($\alpha > \theta_j$; see Section 7.5.3 for the definition of α), and over the entire boundary of the jet-off reference run of Section 7.5.2, the ghost zones use a diode-type condition: outflowing material ($v_r > 0$) is damped exponentially,

$$v_r(t) = v_{r_0}\, e^{-t/\tau}, \qquad \tau = 0.1 \text{ s}, \tag{7.14}$$

modelling the post-explosion rarefaction of the boundary and preventing it from indefinitely including outward-moving material into the domain, while infalling material ($v_r < 0$) is left free to cross the boundary, in order to allow accretion onto the compact remnant. Density and pressure in the ghost zones are floored at 2% of their own angular average at $t = 0$, which keeps the boundary from rarefying into a numerical vacuum without fighting the physical decline during accretion, and the ghost-zone magnetisation is additionally capped, $\rho \geq B^2/\sigma_{\max}$ with $\sigma_{\max} = 100$, to keep the primitive-variable inversion well conditioned where the field locally dominates.

To handle the steep initial gradients, the first few radial cells are evolved with a diffusive setup, i.e., linear reconstruction with a MINMOD limiter, which smooths the initial data without dissipating the physical structure at larger radii. At larger scales the run continues with the less diffusive parabolic reconstruction.

7.5.2 Jet-off Evolution

Before injecting any jet, the initial data are evolved under the boundary conditions of Section 7.5.1 with no jet source: the whole inner boundary is treated as diode ambient, at the same resolution and with the same numerics as the jetted run described in Section 7.5.3. This jet-off run provides a reference against which any jet-induced structure can later be measured.

7.5.3 Jet Evolution

The jet builds on the same framework used for the merger case described in Chapter 4 [66, 115].

Given the coordinate change, it is convenient to introduce a second set of spherical coordinates (r, α, β), obtained from the original coordinates (r, θ, ϕ) by the rigid rotation that maps $\hat{z}$ onto $\hat{y}$, with

$$\cos\alpha \;=\; \hat{e}_r \cdot \hat{y} \;=\; \sin\theta\,\sin\phi, \tag{7.15}$$

so that α is the polar angle measured from the jet axis and β the azimuth around it. We impose transverse pressure balance through the jet edge, and derive the jet pressure profile as follows. The equilibrium problem can be posed directly in α, which is the frame in which the injected outflow is axisymmetric ($\partial_\beta = 0$).

At the injection surface we assume a stationary, axisymmetric and purely radial flow, $\mathbf{v} = v^r \hat{e}_r$, threaded by a poloidal and a toroidal field component, $\mathbf{B} = B^r \hat{e}_r + B^\beta \hat{e}_\beta$. With $b^\mu = \big[\Gamma(\mathbf{v}\cdot\mathbf{B}),\, B^i/\Gamma + \Gamma(\mathbf{v}\cdot\mathbf{B})v^i\big]$, the field component parallel to the flow is boost invariant while the perpendicular one is reduced by the Lorentz factor,

$$b_\mathrm{p} = B^r, \qquad b_\varphi = \frac{B^\beta}{\Gamma}, \qquad b^2 \equiv b^\mu b_\mu = b_\mathrm{p}^2 + b_\varphi^2. \tag{7.16}$$

Following Geng et al. [66], and unlike the single parameter Σ_j adopted in Chapter 4, the poloidal and toroidal contributions are prescribed independently through two magnetisation parameters,

$$\sigma_\mathrm{r} \equiv \frac{b_\mathrm{p}^2}{\rho\, h_\mathrm{g}}, \qquad \sigma_\varphi \equiv \frac{b_{\varphi,\mathrm{m}}^2}{\rho\, h_\mathrm{g}}, \qquad h_\mathrm{g} \equiv \frac{h^*}{1 + \sigma_\mathrm{r} + \sigma_\varphi}, \tag{7.17}$$

where $b_{\varphi,\mathrm{m}}$ is the peak comoving toroidal field over the aperture and h_g is the reference gas enthalpy consistent with the prescribed total specific energy h^* (defined below). The poloidal field is uniform across the aperture, while the toroidal component follows the angular profile of Geng et al. [66, their Eq. 6],

$$B^\beta(\alpha) = \frac{2\, B_{\varphi,\mathrm{m}}\, x}{1 + x^2}, \qquad x \equiv \frac{\alpha}{\theta_\mathrm{m}}, \qquad \theta_\mathrm{m} = 0.4\,\theta_j, \tag{7.18}$$

which vanishes linearly on the axis, peaks at $\alpha = \theta_\mathrm{m}$ with $B^\beta = B_{\varphi,\mathrm{m}} = \gamma\, b_{\varphi,\mathrm{m}}$, and decreases outwards to the jet edge. For the run described here $\sigma_\mathrm{r} = 0.1$ and $\sigma_\varphi = 0.4$, for a nominal total magnetisation $\sigma_\mathrm{r} + \sigma_\varphi = 0.5$ at $\alpha = \theta_\mathrm{m}$. This is a free parameter, following a motivated choice: the magneto-rotational engine invoked for this progenitor (Section 7.3) extracts rotational energy through a strongly wound-up toroidal field, and a weakly magnetised jet would not be representative of that mechanism.

We impose transverse pressure balance through the jet edge, and derive the jet pressure profile as follows. The spatial part of the ideal–RMHD stress tensor reads

$$T^{ij} = w_\mathrm{tot}\,\Gamma^2 v^i v^j - b^i b^j + p_\mathrm{tot}\,\delta^{ij}, \qquad w_\mathrm{tot} = \rho h + b^2, \qquad p_\mathrm{tot} = p + \tfrac{1}{2}b^2, \tag{7.19}$$

and its α component in the orthonormal jet-aligned basis is

$$\frac{1}{r^3}\,\partial_r\!\left(r^3 T_{r\alpha}\right) + \frac{1}{r\sin\alpha}\,\partial_\alpha(\sin\alpha\, T_{\alpha\alpha}) + \frac{1}{r\sin\alpha}\,\partial_\beta T_{\beta\alpha} - \frac{\cot\alpha}{r}\,T_{\beta\beta} = 0. \tag{7.20}$$

For the assumed configuration $T_{r\alpha} = T_{\beta\alpha} = 0$ and $\partial_\beta = 0$, while $T_{\alpha\alpha} = p_{\rm tot}$ and $T_{\beta\beta} = w_{\rm tot}\Gamma^2 v_\beta^2 + p_{\rm tot} - b_\varphi^2$. The isotropic contributions cancel identically and the radial dependence factors out, leaving the transverse balance in the general form

$$\frac{\mathrm{d}}{\mathrm{d}\alpha}\left(p + \frac{b^2}{2}\right) = \cot\alpha\left[\left(\rho h + b^2\right)\Gamma^2 v_\beta^2 - b_\varphi^2\right], \tag{7.21}$$

which is the condition adopted by Martí [115] and Geng et al. [66]. Since the injected outflow is non-rotating ($v_\beta = 0$) and the poloidal field is uniform across the aperture, Eq. (7.21) reduces to the ordinary differential equation actually solved at injection,

$$\boxed{\frac{\mathrm{d}p}{\mathrm{d}\alpha} = -\,b_\varphi^2(\alpha)\cot\alpha - \frac{1}{2}\frac{\mathrm{d}b_\varphi^2}{\mathrm{d}\alpha}} \tag{7.22}$$

The two terms on the right-hand side are the hoop stress of the closed toroidal field lines and the gradient of the toroidal magnetic pressure, respectively. The profile (7.18) guarantees that the equation is regular everywhere: since $b_\varphi \propto \alpha$ as $\alpha \to 0$, the $\cot\alpha$ source vanishes linearly on the axis and no numerical regularisation is required.

Following Geng et al. [66], we prescribe top-hat profiles for the density and the radial fields, and let the gas pressure be the quantity that adjusts to enforce the balance: Eq. (7.22) is integrated inwards from the jet edge to the axis on a fixed table in α, with the boundary value at $\alpha = \theta_j$ fixed by the equation of state of the injected plasma. Here we adopt an ideal gas with adiabatic index $\hat\gamma = 4/3$, so that $h = 1 + 4p/\rho$ and, denoting by $h^* \equiv (\rho h + b^2)/\rho$ the total specific energy of the injected material,

$$p(\theta_j) = \frac{\rho\,(h^* - 1) - b^2(\theta_j)}{4}, \qquad h^*(\theta_j) = \frac{\Gamma_\infty}{\Gamma_j}. \tag{7.23}$$

It should be stressed that, with ρ held uniform across the aperture, the condition $h^* = \Gamma_\infty/\gamma_j$ can be imposed at the edge only. The same equilibrium that determines $p(\alpha)$ also fixes the run of the total specific energy,

$$\rho\,\frac{\mathrm{d}h^*}{\mathrm{d}\alpha} = 4\frac{\mathrm{d}p}{\mathrm{d}\alpha} + \frac{\mathrm{d}b_\varphi^2}{\mathrm{d}\alpha} = -\left[4b_\varphi^2\cot\alpha + \frac{\mathrm{d}b_\varphi^2}{\mathrm{d}\alpha}\right], \tag{7.24}$$

which is manifestly non-zero, so that the injected outflow is intrinsically structured in terminal Lorentz factor. Consistently, the density normalisation is obtained by equating the injected energy flux to the nominal engine power using the aperture-averaged value $\langle h^*\rangle$ rather than its edge value,

$$L_j = \rho\,\Gamma_j^2 v^r\left[\langle h^*\rangle - b_{\rm p}^2/\rho\right]A_j, \qquad \langle h^*\rangle = \frac{\int_0^{\theta_j} h^*(\alpha)\sin\alpha\,\mathrm{d}\alpha}{\int_0^{\theta_j}\sin\alpha\,\mathrm{d}\alpha}, \tag{7.25}$$

with $A_j = 2\pi R_{\rm in}^2(1 - \cos\theta_j)$. By the same token, the magnetisations defined in Eq. (7.17) are attained only approximately. The alternative closure, in which h^* is held uniform and the density absorbs the equilibrium through $\rho(\alpha) = [4p(\alpha) + b^2(\alpha)]/(h^* - 1)$, removes both effects at the cost of a non-uniform injected density; it will be tested against the present setup in a forthcoming set of runs.

Because both the source term of Eq. (7.22) and the boundary condition (7.23) are linear in the density normalisation, the solution $p(\alpha)$ scales exactly linearly with ρ. This property is used to ramp the engine on and off without ever leaving transverse equilibrium.

Kinematically, we adopt the same terminal Lorentz factor, initial Lorentz factor and half-opening angle used throughout the BNS model suite, $\Gamma_j = 3$, $\Gamma_\infty = 300$ ($h^* = 100$ at the jet edge), $\theta_j = 10°$. Given L_j, Γ_∞, Γ_j, $\sigma_{\rm r}$ and σ_φ, the jet density follows directly from the luminosity integral (7.25).

The engine timing also differs from the merger case. Rather than the exponentially decaying luminosity $L(t) = L_j\, e^{-t/\tau_d}$ used for the BNS models, the engine here is switched on smoothly over $t_{\rm on} = 50$ ms after a delay $t_{\rm delay} = 0.3$ s, and held at constant luminosity for a plateau of $t_{\rm engine} = 20$ s (however the current run has only reached 10 s of simulation time). The delay is set to match the free-fall time at the excision radius for the enclosed mass,

$$t_{\rm ff} = \sqrt{\frac{\pi^2 r_{\rm exc}^3}{8GM_{\rm exc}}} \simeq 0.33 \text{ s}, \tag{7.26}$$

so it represents, approximately, the time needed for the innermost material to fall in and settle into a disc able to power the jet, a process this setup does not resolve directly.

7.5.4 Diagnostics and Results

What follows is a first look at the jet run and its jet-off control, evolved to $t = 10$ s. This is a single pair of simulations at one resolution, and the setup still has a few numerical aspects under refinement. The radial grid is logarithmically spaced, so it resolves the region close to injection well, but the same cell count covers a much larger physical extent once the flow reaches large radii; resolution is therefore comfortably adequate only over the smaller-scale part of the domain. The results below are meant as a first, qualitative test of the pipeline developed for the merger channel now applied to a core-collapse environment; a resolution study and the remaining numerical refinements are left for a follow-up.

Morphology and dynamics. Figures 7.5–7.7 show meridional slices of the mass density, radial velocity and magnetic field intensity, for the jet run and the jet-off control, at $t = 0$ s, $t = 5$ s and $t = 10$ s. The jet-off simulation shows the expansion of the SN shock, that with time acquires a more spherical morphology (see below), rarefies and remains at mild velocities. The jet simulation instead presents an outflow that is initially well collimated and considerably faster than the rest of the environment, but at later times appears to be disrupted, with strong mixing taking place away from the injection nozzle, and the fluid remaining at non-relativistic velocities. At the current resolution and up to the last reliable output, the jet decelerates and mixes strongly, showing behaviour consistent with a jet that may ultimately be choked, rather than with a successful outflow of the type usually associated with observed GRBs. This conclusion however cannot be considered final, since the engine was intended to operate for 20 s, a fact that is expected to influence

the following dynamics. The presence of the jet nevertheless affects the environment, by inducing strong turbulence and modifying the propagation of the shock. Given the resolution caveat above, this comparison is only qualitative.

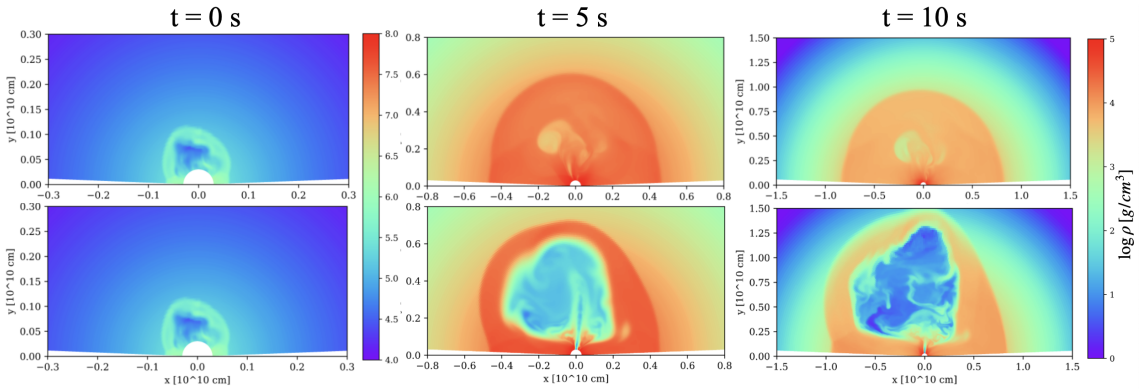


Figure 7.5: Meridional slices of the logarithmic rest-mass density for the jet-off control (top) and the jet run (bottom), at $t = 0$ s (left), $t = 5$ s (middle) and $t = 10$ s (right).

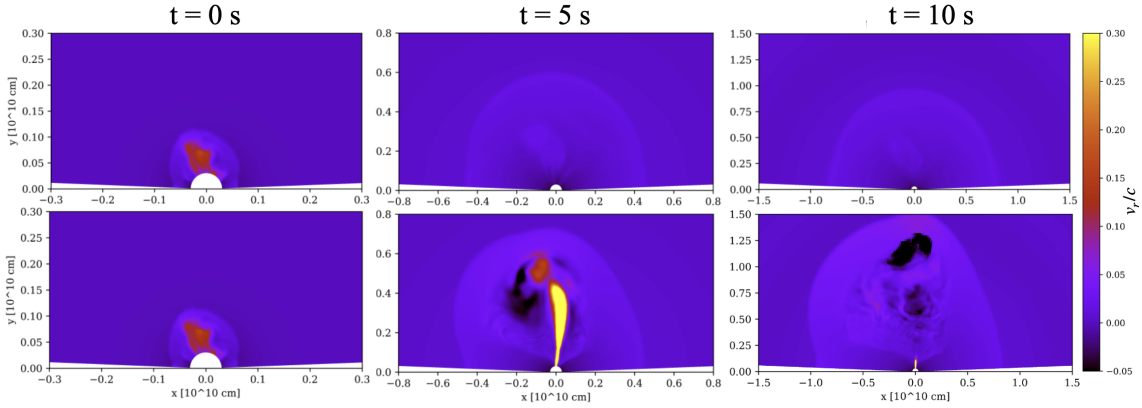


Figure 7.6: Meridional slices of the radial velocity for the jet-off control (top) and the jet run (bottom), at $t = 0$ s (left), $t = 5$ s (middle) and $t = 10$ s (right).

Mass and energy budget. Analytically, the integrated injected energy over the 10 s of jet evolution is $E_{\rm inj} = \int_0^{10\,\rm s} L_j\,dt \simeq 9.6 \times 10^{50}$ erg for the adopted $L_j \sim 10^{50}\,\mathrm{erg\,s^{-1}}$, assuming the luminosity is uniform across the jet cross-section. The energy actually injected is higher, of the order of 1.2×10^{51} erg ($\sim 20\%$ above the uniform estimate). The discrepancy follows from the top-hat density profile: transverse pressure balance (Section 7.5.3) requires the pressure, and with it the local magnetisation and terminal Lorentz factor, to rise toward the jet axis relative to the edge, where the injection is calibrated; locally, this amounts to nearly a factor of three, though most of the cone's solid angle sits far closer to the edge value, limiting the angle-integrated excess to about 20%. Over the same 10s interval the domain-integrated mass declines by about 6% in both runs, and this decline is

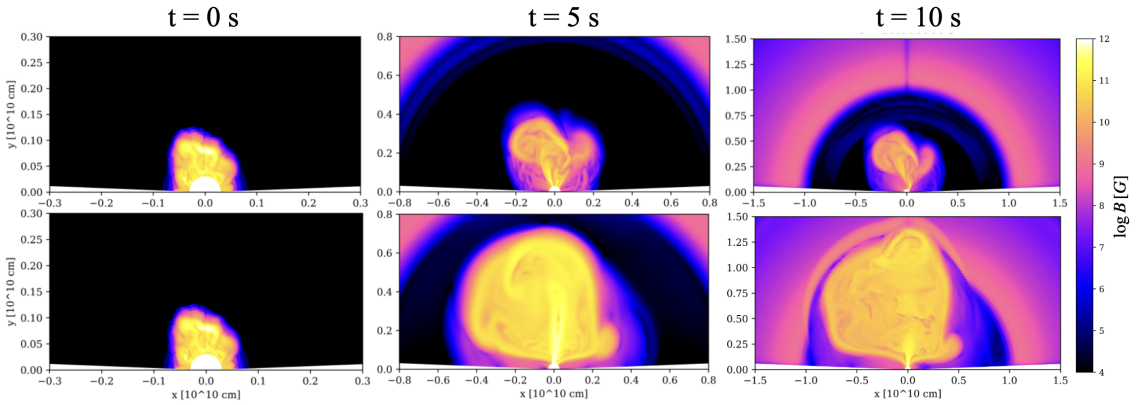


Figure 7.7: Meridional slices of the magnetic field intensity for the jet-off control (top) and the jet run (bottom), at $t = 0$ s (left), $t = 5$ s (middle) and $t = 10$ s (right).

consistent, to a few percent, with the accretion rate estimated independently from the mass flux through the inner boundary, obtained from the surface integral of the mass flux Dv_r[2] over the innermost active radial shell; together with the outward flux from the outer boundary. In the jet run, the total energy (defined here as the sum of the lab-frame kinetic $\rho\Gamma(\Gamma - 1)c^2$, thermal $[p\Gamma^2/(\hat{\gamma} - 1) + p(\Gamma^2 - 1)]$ and magnetic $[b^2(\Gamma^2 - 1/2) - b_t^2]$ components, integrated over the domain volume and excluding the gravitational and rest-mass contributions) grows by roughly 86% over the same interval. This growth is comparable to $E_{\rm inj}$ minus the estimated energy lost through the boundaries. Almost all of the increase is stored in the kinetic and magnetic components, while the thermal component grows modestly in the first second and is immediately converted into kinetic form. In the jet-off run, thermal and magnetic energy are continuously converted into kinetic energy as the ejecta continue to expand, and the mass loss to accretion leads the total energy to diminish by about a quarter. The accretion rate at the inner boundary is mildly enhanced in the jet run at early times, and becomes slower than the jet-off case once the cocoon has had time to inflate, consistent with the jet's outward pressure eventually suppressing infall.

Shock front. Figure 7.8 compares the shock radius, taken as the radius of steepest pressure decline, at 0°, 45° and 70° from the jet axis between the two runs at $t = 10$ s. This is intended as a quick diagnostic, less precise than the shock finding method presented in Section 7.4.4, at a fixed azimuthal angle ($\phi = 90°$). In the jet-off run at $t = 10$ s the ratio of the maximum to the minimum shock radius across the three sampled angles is $R_{\rm max}/R_{\rm min} \simeq 1.2$, indicating that the shock has evolved towards a more spherical shape (a value of 1 would correspond to full sphericity). In the jet run the same ratio is $R_{\rm max}/R_{\rm min} \simeq 2$: the on-axis shock runs well ahead of the shock at 70°, as expected, since the jet locally outruns the ambient shock along its own direction of propagation, while away from the axis the shock continues to trace the slower, SN-driven expansion.

Taken together, these diagnostics show the setup behaving as physically expected

[2]With $D = \Gamma\rho$ being the laboratory-frame mass density.

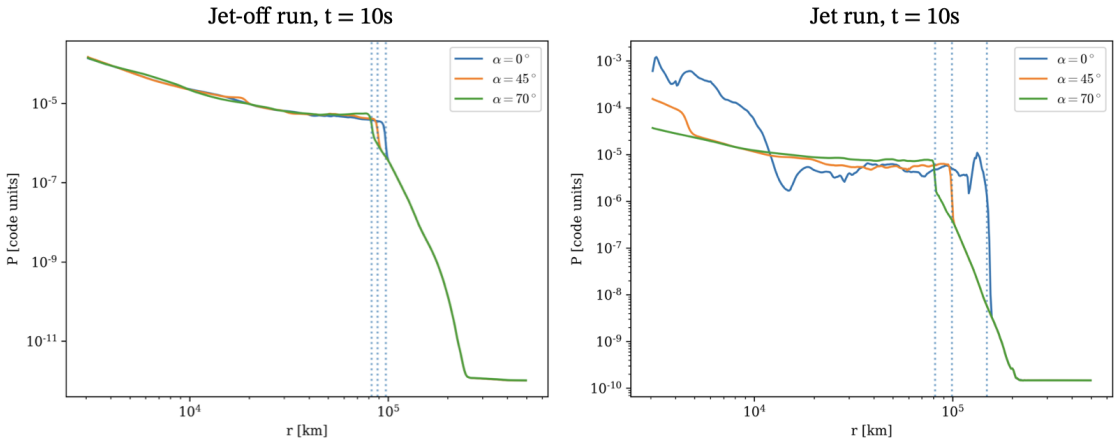


Figure 7.8: Radius of the steepest pressure feature at 0°, 45° and 70° from the jet axis, for the jet run and the jet-off control, at $t = 10$ s. In the jet-off run this feature is the supernova shock at all angles; in the jet run it is the jet-driven forward shock at 0° and the supernova shock at 45° and 70°. Vertical dashed lines mark the identified shock radii for each angle.

in every case checked so far: comparable mass loss between the two runs, energy growth in the jet run, an accretion response to the jet consistent with its outward pressure, and a shock morphology that becomes markedly more anisotropic once the jet is active. None of this should yet be read as a quantitative result; it is a first confirmation that the collapsar extension of the pipeline produces sensible dynamics, ahead of a resolution study and the further numerical work referenced in Section 7.5.1.

A completed production run is expected to show a physically meaningful departure from spherical symmetry in the stellar environment, driven by the still-active supernova shock. This evolving environment should, in turn, leave an imprint on the jet's own structure and propagation. Because the adopted luminosity sits at the estimated breakout threshold (Section 7.4.5), whether the jet actually breaks out over the full duration of the engine is itself an open question the run is designed to address. Should breakout occur, the resulting jet properties will be characterised and compared against what is typically inferred for long duration GRBs, then used as input to the jet-to-afterglow framework of Chapters 5 and 6 to compute the associated electromagnetic emission, which can in turn be compared with observed SN-driven GRB light curves and spectra.

Chapter 8

Conclusions and Future Perspectives

8.1 Summary and Discussion

The main aim of this Thesis was to build a pipeline that connects a realistic progenitor environment, the simulation of the jet propagating through it, and the prediction of the resulting afterglow signal.

The choice of a three-dimensional and realistic progenitor is a defining feature of the pipeline, and its consequences on the jet structure and on the observable signal, which are commonly overlooked when analytical and axisymmetric jet structures are assumed in the interpretation of observational data, are a central outcome of this work. Quantifying these effects, and providing an effective connection between progenitor physics, jet dynamics, and observations, are the two intertwined objectives around which the results of this Thesis are organised.

The objectives set out in the introductory Section 1.3 are addressed as follows. In the context of binary neutron star mergers, a first step in this connection was provided by the work of Pavan et al. [165, 166], which constitutes the starting point of the present Thesis.

Reaching a nearly ballistic regime in 3D. The first contribution presented here is devoted to obtaining a jet structure that can be used as input for the afterglow prediction with the required degree of confidence. Since the afterglow scales cannot be reached directly within a three-dimensional simulation, our approach relies on a semi-analytical prediction, and for this prediction to be meaningful the jet structure fed into it must be as definitive as possible, in the sense of being in a stage advanced enough that no significant residual evolution is expected. To this aim, we developed a method that pushes the simulations inherited from P23 up to a saturating regime that approaches the ballistic phase, by switching to a Cartesian domain at the end of the original simulation and evolving the jet further in time while preserving the resolution needed to resolve its small-scale structure.

We first presented the dynamical evolution of the fiducial simulation inherited from Pavan et al. [166], and we identified features that confirmed the expected behaviour of a relativistic jet propagating in a confining medium, and also gained

new understanding of this interaction. Among the latter, the most impactful is the formation of a multi-shock structure that gives rise to an outer shell of material carrying non-negligible energy, together with the evolution of the jet material into a section of spherical cap that displays strong angular asymmetries and radial stratification. Both features can leave an imprint on the observable emission.

Exploring jet injection conditions and charactering three-dimensional jet structure. These results are confirmed by an exploration of the jet injection parameter space around the fiducial configuration, in which we vary the initial magnetisation, luminosity, total injected energy, and launching time. The comparison across models is carried out with the caveat that the simulations are not at full numerical convergence, and therefore a strict one-to-one mapping between injection conditions and the resulting jet structure is not guaranteed; the refined model-to-model comparison should be interpreted with the corresponding caution. Nonetheless, the qualitative trends are robust: more luminous and more energetic jets have a more successful evolution, breaking out earlier, retaining a larger fraction of the injected energy, and reaching higher average Lorentz factors and stronger collimation, while weaker jets are strongly affected by their interaction with the environment and, in the most marginal cases, fail to develop a coherent ultra-relativistic core. The angular structure analysis quantifies the level of asymmetry and stratification produced by the interaction with a realistic environment, which is one of the main strengths of the initial data used in this work.

We then assessed whether the jet undergoes significant further evolution up to the final time reached in our extended simulations, namely 10 s after jet launching. In most cases, except for the more energetic models that reach the saturating regime later, at 3 s after launching the jets already display an advanced evolutionary stage, and the residual changes are moderate. We nonetheless identify some residual rearrangement and acceleration of the jet material, which motivates the extension of the simulation up to a clearer converging trend in the angular distribution and Lorentz factor of the outflow. Only at that point do we consider the jet structure sufficiently settled to be used as input for the afterglow prediction.

Predicting multi-wavelength afterglow light curves. The afterglow modelling relies on the semi-analytical framework of Salafia et al. [193], in the extended version developed by O. S. Salafia that admits an arbitrary, non-analytical jet structure with azimuthal asymmetries and radial stratification. Applying this framework to our full set of jet models, we find that: first, including the velocity profile inherited from the radial stratification broadens the afterglow light curve, most notably in the rising phase; second, the residual dynamical changes captured by the extended simulation leave a visible imprint on the light curve at large viewing angles, again in the rising phase, which is the one more sensitive to the details of the jet structure, while the on-axis emission is dominated by the already settled symmetric core and shows negligible differences.

The impact of the azimuthal variation is assessed by rotating the observer around the jet symmetry axis for a single jet structure, and reveals up to order-of-magnitude changes, mostly in the early, rising part of the light curve. The detectability analysis

addresses the observable counterpart of the outer shell identified in the dynamical study, which imprints itself on the afterglow as an early bump seen off axis, preceding the main peak of several hours. In some configurations this feature reaches flux levels within the sensitivity range of future instruments, but the signal depends strongly on many parameters, such as the interstellar medium density and on the viewing angle, so we refrain from any strong quantitative claim on its detectability.

Furthermore, we performed a fit of our full set of jet models to the afterglow data of GRB 170817A with two purposes: first, to obtain physically informed parameter values with which to compute the light curves and carry out the analysis just described; second, to test how our jet models compare with real observational data, mainly as a proof of concept of the pipeline. Given the number of free parameters, essentially all of our jet structures, with the caveats presented, provide fits in agreement with the data of this specific event. A more discriminating analysis, that could exclude some of the structures or constrain some of the models by comparison with observations, would require both a more computationally converged mapping between injection conditions and final jet structure, and, from the afterglow side, tighter constraints on the microphysical parameters to allow narrower priors in the fit. The radio centroid motion contribution was intentionally not included, as reproducing the specific observed event was not the scope of this work; its inclusion in future analyses may allow some of the jet structures to be further constrained. Moreover, fitting a specific event is not the target of this work: the application to GRB 170817A serves as a proof of concept for the pipeline, not as a claim on the physical parameters of that event. However, a broader statistical analysis exploiting multiple observed events, could provide useful insights on the connection between the jet structure and the emission pattern.

Finally, the jet structure at the end of the extended simulations, in addition to being the input for the afterglow prediction, can provide a natural input for internal shock models of the prompt γ-ray emission. Exploring this connection is beyond the scope of the present Thesis, but represents a further natural line of development that would extend the reach of this framework from the afterglow signal back to the prompt emission phase.

Extending the framework to different progenitor environments. A natural extension of this work is the application of the same pipeline to the second progenitor channel of gamma-ray bursts, namely the core-collapse of a rapidly rotating, strongly magnetised massive star. The progenitor environment used in this Thesis is provided by the magnetorotational core-collapse supernova simulations of M. A. Aloy and M. Obergualinger.

This Thesis addresses a preliminary stage of work in this direction, in which the output of a magnetorotational core-collapse supernova simulation is used as an input for injecting a relativistic jet, along the same lines followed for the merger-driven case. The goal is analogous: to build a consistent connection between the progenitor physics and the jet structure, from a realistic intrinsically inhomogeneous progenitor to the observable emission of a gamma-ray burst that evolves in such an environment, as opposed to the one-dimensional homogeneous configurations typically assumed. The results shown in Chapter 7 are preliminary, and represent the first step of a

longer programme that will follow the same methodological line established in the merger-driven case.

8.2 Limitations of the Current Framework and Outlook

Several limitations of the present analysis have already been discussed in the relevant chapters, and are summarised here.

On the simulation side, the main limitation concerns numerical resolution; this is discussed in detail in the following subsection. On the afterglow modelling side, the current version of the code does not allow for a fully self-consistent treatment of a stratified velocity profile together with lateral expansion of the outflow. Lateral expansion is expected to affect the light curves at late times, in particular for the post-peak decline, as illustrated in Appendix D. The combined effect of including both a velocity profile and lateral expansion is however not straightforward to anticipate, since the velocity profile tends to broaden the light curve while lateral expansion tends to steepen it, and the two contributions may not combine linearly. A proper quantification of the effect of neglecting lateral expansion, together with a comparison with existing afterglow codes, is left to future work and to a possible future version of the code that includes lateral expansion in a self-consistent way.

The parameter study presented here is not exhaustive: a broader exploration of the jet injection parameter space would give a more comprehensive picture of the variety of merger-driven jets that can develop in the BNS environment and of their afterglow signatures, and is left to future work.

The jet is imposed manually as an injection into the merger environment, and is therefore physically decoupled from the progenitor evolution. Ongoing efforts aim at extending the BNS simulations up to the launching of an outflow that could then be used as a self-consistent jet input. Along the same line, the BNS simulation adopted here does not include neutrino transport, that can modify the environment through which the jet propagates, and whose inclusion will be addressed in future work.

8.2.1 Limitations due to resolution

The resolution study presented in Appendix B covers a factor 1.2 refinement along each spatial direction and is best read as a resolution comparison rather than a formal convergence test. Its coverage is uneven across the model suite. Models *e1* and *e2*, the two lowest-luminosity and lowest-magnetisation configurations, completed at the higher resolution and provide the cleanest comparison; the remaining five models could not be evaluated at the same resolution on equal terms. Models *b*, *c* and *d* failed before reaching the target simulation time in their standard high-resolution setup. Model *d* was rerun with a more diffusive scheme near the injection nozzle and completed, but with an energy budget lower than *e1*, which is inconsistent with its injection parameters and indicates that the patched run cannot be regarded as reliable. Model *a* completed the higher-resolution run, but the early-time evolution shows clear numerical difficulties whose imprint is visible in the final structure. This

indicates that the current numerical configuration is not adequate to handle these simulations at higher resolution, because of the stronger gradients involved, and that a dedicated study of a configuration that is at the same time robust and accurate will be required before a systematic high-resolution parameter study can be carried out, which is left to future work.

Within these limits, two consequences for the results presented in this Thesis should be kept in mind. First, the qualitative features on which the conclusions rest, in particular the departure from axisymmetry, the presence of an outer shell and the resulting early-time afterglow signature, are confirmed at the higher resolution wherever the comparison could be reliably carried out; the corresponding statements should be read as robust. Second, the fine-scale model-to-model differences, in particular the separation between models within the *c*, *d* and *F* triplet, cannot be firmly disentangled from resolution effects at the fiducial setup, and comparisons at that level are indicative. In the case of *e1*, the model closest to marginal breakout, the higher-resolution simulation is measurably more successful, with a larger fraction of the injected energy retained by the jet and a higher terminal Lorentz factor. Thus, its qualitative behaviour is not an artefact of resolution; however, the exact degree of marginality remains resolution-dependent.

Appendices

Appendix A

Publications Related to This Thesis

The results presented in this thesis are based on, or directly extend, the following publications.

- E. Dreas, A. Pavan, R. Ciolfi, A. Celotti, *Approaching ballistic motion in 3D simulations of GRB jets in realistic BNS merger environments*, A&A, 694, A200 (2025). This work forms the basis of Chapter 5.
- A. Pavan, R. Ciolfi, E. Dreas, J. V. Kalinani, *Role of injection parameters in jet propagation through realistic BNS merger environments*, MNRAS, 540, 1345–1362 (2025). This work forms the basis of the parameter study presented in s 4 and 6 .
- E. Dreas, O. S. Salafia, A. Pavan, R. Ciolfi, A. Celotti, *Evolution and afterglow emission of gamma-ray burst jets from binary neutron star mergers*, submitted to A&A (2026). This work forms the basis of the afterglow modelling presented in Chapter 6.
- E. Dreas, A. Pavan, R. Ciolfi, A. Celotti, *Simulating the quasi-ballistic regime of a short Gamma-Ray Burst jet*, submitted to Mem. Soc. Astron. Italiana (2026).

Appendix B

Resolution Study

This chapter documents a resolution study performed on the simulations analysed in Chapters 4 and 6 with the aim of assessing the impact of the resolution on the results presented in this thesis.

B.1 Motivation and Setup

All models presented in Section 4.3 were re-run at a resolution increased by a factor of 1.2 along each spatial direction with respect to the fiducial setup described in Section 4.1. Due to limited computational resources, the present analysis should be understood as a resolution comparison rather than a formal convergence study.

B.2 Outcome Across Models

The outcome of the higher-resolution runs was strongly dependent on the injection luminosity and magnetisation of the model considered. The least energetic and least magnetised models, *e1* and *e2*, completed successfully. Models *b, c* and *d* failed before reaching the target simulation time. Model *d* has been successfully re-stabilized through the implementation of a more diffusive numerical scheme in the vicinity of the injection zone, thereby preventing the numerical failure previously encountered in this region around $t - t_{\rm j} \approx 2$ s; while this patch allowed the simulation to complete, its final energy budget is the lowest among all models, lower even than that of model *e1*, which is inconsistent with the injection parameters of the two models and suggests that the reliability of this run remains compromised even after the patch. Model *a* completed the full evolution, but inspection of the energy density evolution reveals clear signs of numerical difficulties (and thus possible artifacts) at early times, resulting in a final structure that is considerably more stratified, and less energetic, than that of model *F*, again without a physical justification based on the respective injection conditions.

A plausible interpretation is that the numerical scheme becomes increasingly stressed, at higher resolution, in configurations characterised by higher injection energy and magnetisation. Such configurations produce steeper spatial gradients in the vicinity of the jet injection nozzle; at higher resolution these gradients are

Table B.1: Jet cap properties of models *e1* and *e2* at the fiducial resolution (from Table 2 of P25, north side) and at increased resolution (*e1*$_{\mathrm{HR}}$, *e2*$_{\mathrm{HR}}$) at $t - t_{\mathrm{j}} = 3\,\mathrm{s}$.

	e1	*e1*$_{\mathrm{HR}}$	*e2*	*e2*$_{\mathrm{HR}}$
$E_{\mathrm{iso,max}}$ $[10^{53}\,\mathrm{erg}]$	1.10	1.21	2.27	2.58
$\overline{\Gamma}_{\mathrm{max}}$	27.8	40.7	70.6	76.2
$E_{\mathrm{sum}}/E_{\mathrm{inj}}$	0.44	0.47	0.61	0.63
$E_{\mathrm{kin}}/E_{\mathrm{sum}}$	0.90	0.94	0.92	0.92
Θ_{50} [°]	6.58	5.89	4.41	4.06

increasingly resolved rather than smoothed by numerical diffusion, which appears to trigger localised numerical inconsistencies that can propagate to later times and larger scales.

B.3 Case Study: Models *e1* and *e2*

As discussed in Section 4.3 (see also P25), model *e1* is characterised by relatively low injection luminosity and magnetisation, and its breakout is only marginally successful: at the fiducial resolution, the jet loses more than half of its injected energy to the surrounding environment, and the resulting angular structure lacks the high-$\overline{\Gamma}$ core found in all the other models. This makes *e1* a particularly sensitive probe of numerical effects, since small changes in how efficiently the jet drills through the environment can have a strong impact on its final state. Directly comparable is model *e2*, different only for the launching time.

Table B.1 reports the jet cap properties of model *e1* and *e2* at the fiducial resolution (from Table 2 of P25, north side) and at the increased resolution considered here (*e1*$_{\mathrm{HR}}$, *e2*$_{\mathrm{HR}}$). All quantities point in the same direction: at higher resolution, the jet retains a larger fraction of the injected energy ($E_{\mathrm{sum}}/E_{\mathrm{inj}}$: 0.44 → 0.47 (*e1*), 0.61 → 0.63 (*e2*)), channels a larger (or equal) part of it into kinetic form ($E_{\mathrm{kin}}/E_{\mathrm{sum}}$: 0.90 → 0.94 (*e1*), 0.92 → 0.92 (*e2*)), and reaches a maximum Lorentz factor higher by 45 per cent for *e1* and 8 per cent for *e2* ($\overline{\Gamma}_{\mathrm{max}}$: 27.8 → 40.7 (*e1*), 70.6 → 76.2 (*e2*)). The jet cap is also somewhat more collimated (Θ_{50}: 6.6° → 5.9° (*e1*), 4.41° → 4.06° (*e2*)). The differences, throughout all the analysis, are considerably milder for the case of *e2*.

Taken together, these changes describe a jet that is consistently more successful at higher resolution: less energy is lost during propagation, more of what remains is converted into directed kinetic motion, and the resulting outflow is both faster and more collimated. This is qualitatively what one would expect if the higher resolution allows the code to better resolve the shocks and the mixing processes responsible for energy losses during the drilling phase, though this simulation is not yet in convergence regime. We observe that the discrepancies between the two resolutions in the northern hemisphere are of the same order of magnitude as the discrepancies between the northern and southern hemispheres at the fiducial resolution. This indicates that the detailed mapping between injection parameters and jet structure is highly sensitive to small-scale fluctuations arising in the early stages of the dynamics.

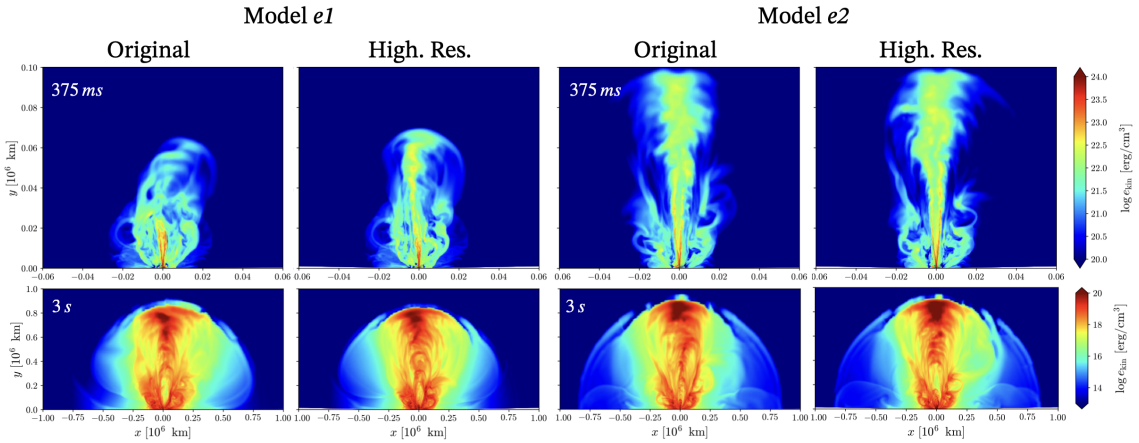


Figure B.1: Meridional view of the energy density without rest-mass component at 375 ms (top) and 3 s (bottom) after jet launching, for model *e1* (left) and *e2* (right) at the two resolutions.

Figure B.1 compares meridional maps of the energy density (excluding the rest-mass contribution) for models *e1* and *e2* at two times, $t - t_j = 0.375\,\mathrm{s}$ and $t - t_j = 3\,\mathrm{s}$, run at both the fiducial and the enhanced spatial resolutions. For model *e1*, the higher-resolution simulation yields a visibly more successful outflow, whereas for model *e2* the two resolutions produce nearly indistinguishable configurations. In both cases, the presence of the outer shell and an asymmetric angular distribution are confirmed.

The associated on-axis afterglow light curves, presented in Fig. B.2, remain qualitatively unchanged across resolutions despite these structural modifications, suggesting that near-axis predictions are relatively insensitive to the numerical resolution adopted. The picture is different off-axis: for model *e1*, the extended pre-peak plateau that distinguishes it from the more energetic models at fiducial resolution is substantially reduced at higher resolution, and the rising branch converges toward those of the other models. The peak time is essentially preserved, while its flux increases marginally. Model *e2*, consistently with the structural analysis, exhibits no appreciable change.

The main features identified in our simulations that influence the light curve (velocity and energy stratification, complex angular structure) are instead qualitatively robust to the choice of resolution. However, the resolution-dependent variations in the jet structure can still influence the detailed shape of the light curve, in a way that is hard to quantify and possibly comparable to some of the differences between models. The results presented in this thesis should be interpreted with this caveat in mind.

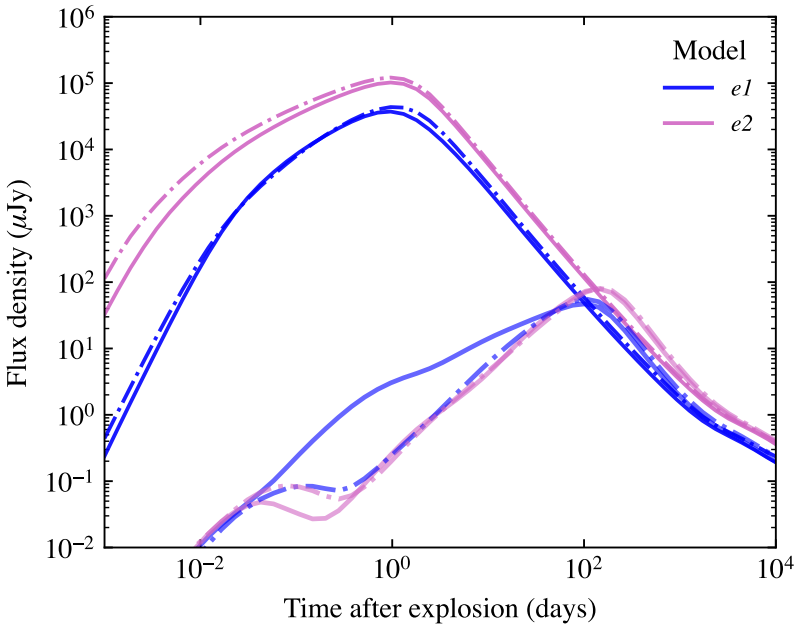


Figure B.2: Afterglow light curves at 3 GHz for model *e1* and *e2* (blue and magenta respectively) at the original resolution (full line) and increased resolution (dashed line), on-axis and with $\theta_{\rm v} = 32°$, computed for the best-fit values of model F for both resolutions, for comparison.

Appendix C

Detectability of the Outer-shell Feature

In this appendix we assess the detectability of the observational feature associated with the outer shell, introduced in Sect. 4.2.2 and further discussed in Sect. 6.1.3 in the context of the afterglow light curves.

Drawing a firm conclusion on this point is not straightforward, since the flux density of the early peak depends on several parameters at once. The source distance has the strongest influence, but the flux also depends on the ISM density and on the observing angle, both in the axisymmetrised case and when different azimuthal viewing angles are considered. Different jet structures lead to considerable differences in this feature as well (see Figures 6.11 and 6.12).

To place the predicted signal on a quantitative footing, we compare it to the sensitivity of current and forthcoming radio interferometers at 3 GHz. The theoretical 1σ rms noise of an interferometric observation is set by the radiometer equation,

$$\Delta I_m = \frac{\mathrm{SEFD}}{\eta_c \sqrt{n_{\mathrm{pol}}\, N(N-1)\, t_{\mathrm{int}}\, \Delta\nu}}\,, \tag{C.1}$$

where the system equivalent flux density (SEFD) encapsulates the effective sensitivity of a single antenna including receiver noise and aperture efficiency, η_c is the correlator efficiency, n_{pol} the number of polarisation products ($n_{\mathrm{pol}} = 2$ for Stokes I), N the number of antennas, t_{int} the on-source integration time, and $\Delta\nu$ the observing bandwidth. The $\sqrt{N(N-1)}$ scaling reflects the number of independent baselines that contribute coherently to the image plane. Since Eq. (C.1) depends on the effective (rather than nominal) values of the SEFD and bandwidth, for both facilities considered here we adopt the sensitivity returned by the corresponding official exposure calculators, which internally account for band edges, RFI-affected sub-bands, and other calibration losses that reduce the usable bandwidth relative to the nominal value.

For the Karl G. Jansky Very Large Array (VLA), the Exposure Calculator[1] returns, for S-band (2–4 GHz), dual polarisation, natural weighting, and one hour of on-source integration, a naturally-weighted rms of $\Delta I_m \simeq 4\ \mu$Jy/beam, corresponding

[1] https://obs.vla.nrao.edu/ect/

to a 3σ detection threshold of ~ 12 μJy. This is consistent with the early radio non-detections of GW170817, for which upper limits of comparable magnitude were reported over the first ~ 15 days after merger [6, 81], and with a direct evaluation of Eq. (C.1) using the standard VLA S-band parameters ($N = 27$, SEFD $\simeq 370$ Jy, $\eta_c = 0.93$, $\Delta\nu = 2$ GHz; Perley et al. 174).

For the next-generation VLA (ngVLA; Selina et al. 202), we adopt the sensitivity provided by the official ngVLA Exposure Calculator[2] for the current reference design. In the main+long subarray, Band 1, with dual polarisation, natural weighting, the full instantaneous bandwidth $\Delta\nu = 2.30$ GHz centred at 3 GHz, and one hour of on-source integration, the calculator returns $\Delta I_m \simeq 0.16$ μJy/beam, corresponding to a 3σ detection threshold of ~ 0.5 μJy. This represents more than an order-of-magnitude improvement over the VLA.

The 3σ VLA and ngVLA thresholds adopted above are overlaid on the off-axis light curves obtained from the axisymmetrised structure of model F, for a set of viewing angles and ISM densities, in Fig. C.1. For reference, we also show the early GW170817 upper limits and non-detections from Alexander et al. [6] and Hallinan et al. [81], as compiled in the uniform dataset of Makhathini et al. [111]. Under the most favourable combination of parameters explored here (lowest ISM density, intermediate viewing angle), the peak of the shell feature approaches, and in some cases crosses, the ngVLA detection threshold. In general, however, the predicted signal remains well below the GW170817 upper limits, indicating that these do not yet provide a meaningful constraint on this specific signature.

A further layer of complexity arises from azimuthal asymmetries and from differences among the various jet models. To give a qualitative sense of this additional variability, in Figure 6.12 we overlay the same sensitivity lines on the off-axis light curves of all models, including the spread introduced by different azimuthal viewing angles. While several of these curves appear promising for future facilities, a more fundamental limitation is set by time. For model F, the outer-shell peak occurs at $t \sim 10^{-2}$–10^{-1} d after merger, i.e. between roughly a quarter of an hour and two and a half hours. This is representative of the general behaviour: across all models the feature peaks at comparably early times, although the precise value varies considerably from model to model, reflecting the dependence on the jet injection parameters. In every case, this timescale is comparable to, or shorter than, the 1 h integration adopted to build the sensitivity curves.

Even the most rapid radio follow-up campaigns obtained to date for a BNS merger fall short of this requirement. The electromagnetic counterpart of GW170817 was itself identified remarkably early, within ~ 11 h of the merger [38], and the earliest available VLA observation followed at $t \simeq 0.57$ d [6], only a few hours later. This response, on a timescale of hours rather than days, was exceptionally fast by the standards of the time, and yet it still trails the predicted peak of model F by roughly eleven hours. Reaching such early timescales would require sub-hour localisation-to-pointing response, a capability that has continued to improve in recent years with faster alert distribution and progressively more automated follow-up pipelines. Moreover, even under such favourable circumstances, if the intrinsic duration of the feature is comparable to or shorter than the required integration time, the resulting

[2] https://ngvla.nrao.edu/page/performance

flux estimate would effectively average over a rapidly evolving source, further diluting the sensitivity to this specific signature.

The detectability of the outer-shell feature is therefore limited jointly by flux threshold and by timing. Given the number of parameters that concurrently shape the predicted signal, we refrain from drawing strong quantitative claims on its detectability. The feature nonetheless remains a potentially valuable diagnostic for structured jets, one that could become more readily accessible as higher-sensitivity facilities, faster and more precise localisation, and closer events become available.

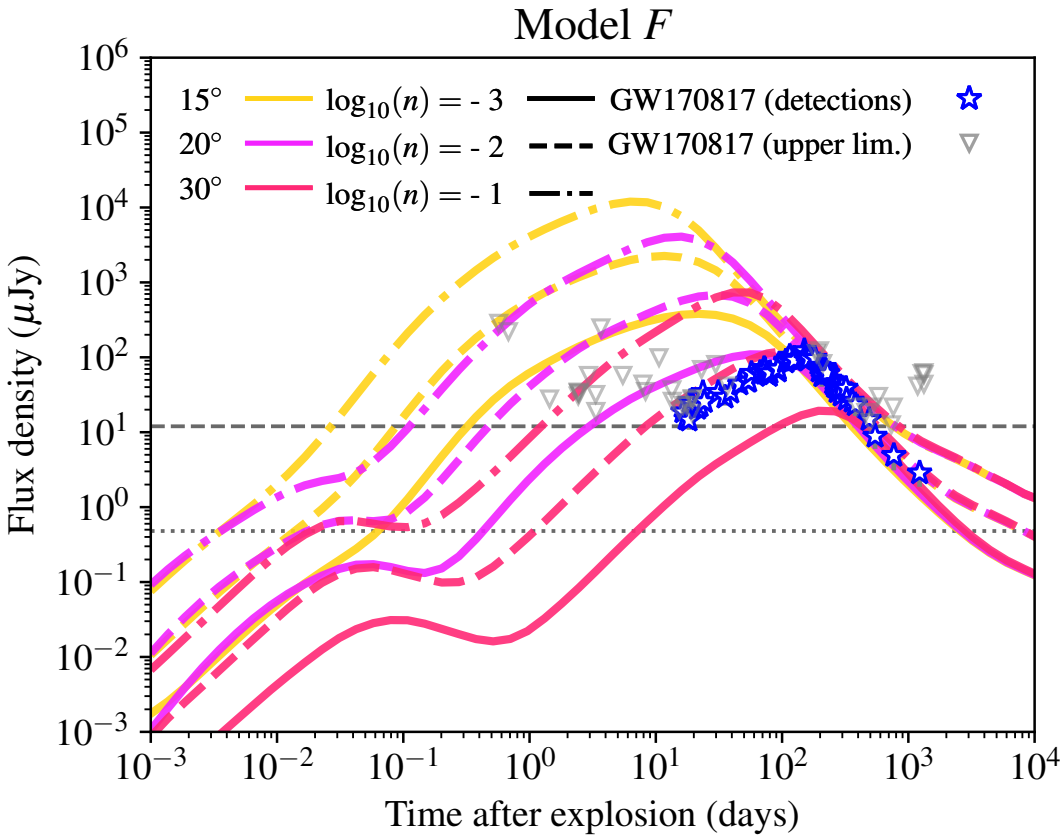


Figure C.1: Off-axis afterglow light curves at 3 GHz for model F (axisymmetrised jet structure), for a set of viewing angles $\theta_v = 15°, 20°, 30°$ (colours) and ISM densities $n = 10^{-3}, 10^{-2}, 10^{-1}$ cm^{-3} (line styles), illustrating the dependence of the early shell-associated peak on these parameters. Horizontal dashed and dotted grey lines mark the 3σ VLA and ngVLA sensitivities, respectively, for a fiducial 1 h integration at 3 GHz, obtained from the corresponding official exposure calculators (see text). Grey stars with error bars and downward-pointing arrows show, respectively, GW170817 radio detections and 3σ upper limits at frequencies between 1 and 10 GHz, rescaled to 3 GHz, as compiled by Makhathini et al. [111] from Alexander et al. [6] and Hallinan et al. [81].

Appendix D

On the Neglect of Lateral Expansion

A limitation of the afterglow model employed in this work is the absence of the treatment of the lateral expansion together with the inclusion of the velocity profile. In the current version, the two can only be included in a mutually exclusive manner. We therefore compare their impact separately, and attempt at estimating what could be the combined effect, and the consequent influence on the work presented.

Figure D.1 shows the on-axis and off-axis light curves for model F for the two cases, compared with baseline of not including any of the two. As already discussed in Section 6.1.3, the velocity profile leads to a broadening of the light curve, allowing not only the angular structure, but also the radial stratification of the outflow to shape the rising light curve. The inclusion of the lateral expansion instead has an overall steepening effect, influencing also the post-peak decay, and slightly rising the peak itself, while having a negligible impact on the rising part.

Consequently, we can expect a combination of the two to steepen the declining phase and slightly affect the peak emission, with respect to the current prediction. However, if the velocity profile is quite broad, the effects of the lateral expansion can be considerably mitigated, as lateral gradients within the shock are diluted at the speed of sound. The contribution of the velocity profiles could therefore be preserved. We leave to a future work a detailed quantification of this effect, when a joint implementation of the two mechanisms will be available.

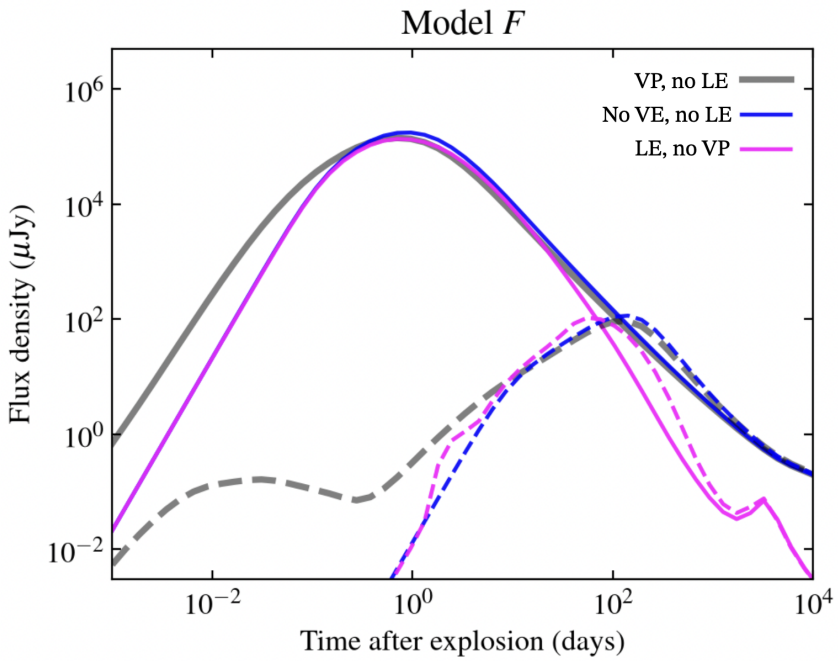


Figure D.1: Afterglow light curves at 3 GHz for model F, on-axis (solid) and off-axis, (dashed), isolating the effect of the velocity profile (VP) and of lateral expansion (LE) against the common baseline of neither ("no VP, no LE", blue). Grey: VP only. Magenta: LE only.

Bibliography

[1] B. P. Abbott et al. Multi-messenger observations of a binary neutron star merger. *ApJ Letters*, 848:L12, 2017. doi: 10.3847/2041-8213/aa91c9.

[2] B. P. Abbott et al. GW170817: Observation of gravitational waves from a binary neutron star inspiral. *Physical Review Letters*, 119:161101, 2017. doi: 10.1103/PhysRevLett.119.161101.

[3] B. P. Abbott et al. Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817A. *ApJ Letters*, 848:L13, 2017. doi: 10.3847/2041-8213/aa920c.

[4] D. R. Aguilera-Dena, N. Langer, T. J. Moriya, and A. Schootemeijer. Related progenitors for superluminous supernovae, broad-lined Type Ic supernovae, and long gamma-ray bursts. *The Astrophysical Journal*, 858:115, 2018. doi: 10.3847/1538-4357/aabfc1.

[5] A. Akmal, V. R. Pandharipande, and D. G. Ravenhall. Equation of state of nucleon matter and neutron star structure. *Phys. Rev. C.*, 58:1804–1828, 1998. doi: 10.1103/PhysRevC.58.1804.

[6] K. D. Alexander, E. Berger, W. Fong, P. K. G. Williams, C. Guidorzi, R. Margutti, B. D. Metzger, J. Annis, P. K. Blanchard, D. Brout, D. A. Brown, H.-Y. Chen, R. Chornock, P. S. Cowperthwaite, M. Drout, T. Eftekhari, J. Frieman, D. E. Holz, M. Nicholl, A. Rest, M. Sako, M. Soares-Santos, and V. A. Villar. The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. VI. Radio Constraints on a Relativistic Jet and Predictions for Late-time Emission from the Kilonova Ejecta. *ApJ Letters*, 848(2):L21, 2017. doi: 10.3847/2041-8213/aa905d.

[7] K. D. Alexander, R. Margutti, P. K. Blanchard, W. Fong, E. Berger, A. Hajela, T. Eftekhari, R. Chornock, P. S. Cowperthwaite, D. Giannios, et al. A decline in the x-ray through radio emission from GW170817 continues to support an off-axis structured jet. *ApJ Letters*, 863:L18, 2018. doi: 10.3847/2041-8213/aad637.

[8] M. A. Aloy, E. Müller, J. M. Ibáñez, J. M. Martí, and A. MacFadyen. Relativistic jets from collapsars. *ApJ Letters*, 531:L119–L122, 2000. doi: 10.1086/312537.

[9] M. A. Aloy, H.-T. Janka, and E. Müller. Relativistic outflows from remnants of compact object mergers and their viability for short gamma-ray bursts.

Astronomy and Astrophysics, 436:273–311, 2005. doi: 10.1051/0004-6361:20041865.

[10] M. A. Aloy, C. Cuesta-Martínez, and M. Obergaulinger. On the existence of a luminosity threshold of GRB jets in massive stars. *Monthly Notices of the Royal Astronomical Society*, 478:3576–3589, 2018. doi: 10.1093/mnras/sty1212.

[11] L. Amati et al. Intrinsic spectra and energetics of BeppoSAX gamma-ray bursts with known redshifts. *Astronomy and Astrophysics*, 390:81–89, 2002. doi: 10.1051/0004-6361:20020722.

[12] A. M. Anile. *Relativistic Fluids and Magneto-fluids: With Applications in Astrophysics and Plasma Physics*. Cambridge University Press, Cambridge, 1989.

[13] S. Ascenzi, G. Oganesyan, M. Branchesi, S. Dall'Osso, O. S. Salafia, and G. Ghirlanda. Structured jets and the environment in GRB afterglows. *Monthly Notices of the Royal Astronomical Society*, 499:2519–2534, 2020. doi: 10.1093/mnras/staa2983.

[14] A. Balasubramanian et al. Continued evolution of the outflow of GW170817 as revealed by high-resolution VLBI observations. *ApJ Letters*, 914:L20, 2021. doi: 10.3847/2041-8213/ac0140.

[15] S. A. Balbus and J. F. Hawley. A powerful local shear instability in weakly magnetized disks. I. linear analysis. *The Astrophysical Journal*, 376:214–222, 1991. doi: 10.1086/170270.

[16] D. Band et al. Batse observations of gamma-ray burst spectra. I. spectral diversity. *The Astrophysical Journal*, 413:281–292, 1993. doi: 10.1086/173000.

[17] T. W. Baumgarte and S. L. Shapiro. *Numerical Relativity: Solving Einstein's Equations on the Computer*. Cambridge University Press, 2010. doi: 10.1017/CBO9781139193344.

[18] A. Bauswein et al. Equation of state constraints from the threshold binary mass for prompt collapse of neutron star mergers. *ApJ Letters*, 908:L1, 2021. doi: 10.3847/2041-8213/abda4a.

[19] M. C. Begelman and D. F. Cioffi. Overpressured cocoons in extragalactic radio sources. *ApJ Letters*, 341:L41–L44, 1989. doi: 10.1086/185449.

[20] P. Beniamini, J. Granot, and R. Gill. Afterglow light curves from structured jets with a off-axis components. *Monthly Notices of the Royal Astronomical Society*, 493:3521–3534, 2020. doi: 10.1093/mnras/staa538.

[21] E. Berger. Short-duration gamma-ray bursts. *Annual Review of Astronomy and Astrophysics*, 52:43–105, 2014. doi: 10.1146/annurev-astro-081913-035926.

[22] G. S. Bisnovatyi-Kogan. The explosion of a rotating star as a supernova mechanism. *Sov. Astron.*, 14:652–655, 1971.

[23] R. D. Blandford and C. F. McKee. Fluid dynamics of relativistic blast waves. *Physics of Fluids*, 19:1130–1138, 1976. doi: 10.1063/1.861619.

[24] R. D. Blandford and R. L. Znajek. Electromagnetic extraction of energy from Kerr black holes. *Monthly Notices of the Royal Astronomical Society*, 179: 433–456, 1977. doi: 10.1093/mnras/179.3.433.

[25] O. Bromberg and A. Tchekhovskoy. Relativistic MHD simulations of core-collapse GRB jets: 3D instabilities and magnetic dissipation. *Monthly Notices of the Royal Astronomical Society*, 456:1739–1760, 2016. doi: 10.1093/mnras/stv2591.

[26] O. Bromberg, E. Nakar, T. Piran, and R. Sari. The propagation of relativistic jets in external media. *The Astrophysical Journal*, 740:100, 2011. doi: 10.1088/0004-637X/740/2/100.

[27] O. Bromberg, E. Nakar, and T. Piran. Are low-luminosity gamma-ray bursts generated by relativistic jets? *ApJ Letters*, 739:L55, 2012. doi: 10.1088/2041-8205/739/2/L55.

[28] N. Bucciantini. Relativistic hydrodynamics and magneto-hydrodynamics, 2013. Lecture notes, INAF–Arcetri Observatory.

[29] A. Burrows, L. Dessart, E. Livne, C. D. Ott, and J. Murphy. Simulations of magnetically driven supernova and hypernova explosions in the context of rapid rotation. *The Astrophysical Journal*, 664:416–434, 2007. doi: 10.1086/519161.

[30] D. N. Burrows and J. Racusin. Swift X-ray afterglows: Where are the X-ray jet breaks? *Nuovo Cimento B Serie*, 121(10):1273–1287, October 2006. doi: 10.1393/ncb/i2007-10252-7.

[31] Z. Cano, S.-Q. Wang, Z.-G. Dai, and X.-F. Wu. The observer's guide to the gamma-ray burst supernova connection. *Adv. Astron.*, 2017:8929054, 2017. doi: 10.1155/2017/8929054.

[32] M. Cantiello, J. B. Jensen, J. P. Blakeslee, et al. A precise distance to the host galaxy of the binary neutron star merger GW170817 using surface brightness fluctuations. *ApJ Letters*, 854:L31, 2018. doi: 10.3847/2041-8213/aaaa64.

[33] Ho-Sang Chan, Ore Gottlieb, Jonatan Jacquemin-Ide, Matteo Cantiello, and Mathieu Renzo. Jets from Scratch: Dynamo-generated Poloidal Magnetic Fields in Three-dimensional Collapsar Simulations. *ApJ Letters*, 1007(2):L41, August 2026. doi: 10.3847/2041-8213/ae9156.

[34] R. Ciolfi. Collimated outflows from long-lived binary neutron star merger remnants. *MNRAS Lett.*, 495:L66–L70, 2020. doi: 10.1093/mnrasl/slaa048.

[35] R. Ciolfi. The key role of magnetic fields in binary neutron star mergers. *Gen. Rel. Grav.*, 52:59, 2020. doi: 10.1007/s10714-020-02714-x.

[36] P. Colella and P. R. Woodward. The piecewise parabolic method (PPM) for gas-dynamical simulations. *J. Comput. Phys.*, 54:174–201, 1984. doi: 10.1016/0021-9991(84)90143-8.

[37] E. Costa et al. Discovery of an X-ray afterglow associated with the γ-ray burst of 28 february 1997. *Nature*, 387:783–785, 1997. doi: 10.1038/42885.

[38] D. A. Coulter, R. J. Foley, C. D. Kilpatrick, M. R. Drout, A. L. Piro, B. J. Shappee, M. R. Siebert, J. D. Simon, N. Ulloa, D. Kasen, B. F. Madore, A. Murguia-Berthier, Y.-C. Pan, J. X. Prochaska, E. Ramirez-Ruiz, A. Rest, and C. Rojas-Bravo. Swope Supernova Survey 2017a (SSS17a), the optical counterpart to a gravitational wave source. *Science*, 358(6370):1556–1558, December 2017. doi: 10.1126/science.aap9811.

[39] P. S. Cowperthwaite et al. The electromagnetic counterpart of the binary neutron star merger LIGO/Virgo GW170817. *ApJ Letters*, 848:L17, 2017. doi: 10.3847/2041-8213/aa8fc7.

[40] P. D'Avanzo et al. The evolution of the x-ray afterglow emission of GW 170817/GRB 170817A in XMM-Newton observations. *Astronomy and Astrophysics*, 613:L1, 2018. doi: 10.1051/0004-6361/201832664.

[41] F. De Colle, P. Kumar, and P. Hoeflich. The large landscape of supernova, GRB, and cocoon interactions. *Monthly Notices of the Royal Astronomical Society*, 512(3):3627–3637, May 2022. doi: 10.1093/mnras/stac742.

[42] A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, and M. Wesenberg. Hyperbolic divergence cleaning for the MHD equations. *J. Comput. Phys.*, 175: 645–673, 2002. doi: 10.1006/jcph.2001.6961.

[43] M. Della Valle et al. An enigmatic long-lasting γ-ray burst not accompanied by a bright supernova. *Nature*, 444:1050–1052, 2006. doi: 10.1038/nature05374.

[44] L. Dessart, C. D. Ott, A. Burrows, S. Rosswog, and E. Livne. Neutrino signatures and the neutrino-driven wind in the aftermath of a neutron star coalition. *The Astrophysical Journal*, 690:1681–1705, 2009. doi: 10.1088/0004-637X/690/2/1681.

[45] E. Dreas, A. Pavan, R. Ciolfi, and A. Celotti. Approaching ballistic motion in 3D simulations of gamma-ray burst jets in realistic binary neutron star merger environments. *Astronomy and Astrophysics*, 694:A200, 2025. doi: 10.1051/0004-6361/202451578.

[46] E. Dreas, O. S. Salafia, A. Pavan, R. Ciolfi, and A. Celotti. Evolution and afterglow emission of gamma-ray burst jets from binary neutron star mergers. *Astronomy and Astrophysics*, 2026. submitted.

[47] P. C. Duffell, E. Quataert, D. Kasen, and H. Klion. Jet cocoon energetics and the origin of short gamma-ray burst afterglows. *The Astrophysical Journal*, 866:3, 2018. doi: 10.3847/1538-4357/aadf8e.

[48] G. C. Duncan and P. A. Hughes. Simulations of relativistic extragalactic jets. *ApJ Letters*, 427:L97–L100, 1994. doi: 10.1086/187372.

[49] R. Duque, P. Beniamini, F. Daigne, and R. Mochkovitch. Cosmological population of GRBs and afterglows with structured jets. *Astronomy and Astrophysics*, 658:A51, 2022. doi: 10.1051/0004-6361/202141926.

[50] D. Eichler, M. Livio, T. Piran, and D. N. Schramm. Nucleosynthesis, neutrino bursts and gamma-rays from coalescing neutron stars. *Nature*, 340:126–128, 1989. doi: 10.1038/340126a0.

[51] Z. B. Etienne, V. Paschalidis, Y. T. Liu, and S. L. Shapiro. Relativistic magnetohydrodynamics in dynamical spacetimes: improved electromagnetic gauge condition for adaptive mesh refinement grids. *Phys. Rev. D*, 85:024013, 2012. doi: 10.1103/PhysRevD.85.024013.

[52] C. R. Evans and J. F. Hawley. Simulation of magnetohydrodynamic flows: A constrained transport method. *The Astrophysical Journal*, 332:659–677, 1988. doi: 10.1086/166684.

[53] R. Fernández and B. D. Metzger. Delayed outflows from black hole accretion tori following neutron star binary coalescence. *Monthly Notices of the Royal Astronomical Society*, 435:502–517, 2013. doi: 10.1093/mnras/stt1312.

[54] D. Finstad, S. De, D. A. Brown, E. Berger, and C. M. Biwer. Measuring the viewing angle of GW170817 with electromagnetic and gravitational waves. *ApJ Letters*, 860:L2, 2018. doi: 10.3847/2041-8213/aac6c1.

[55] W. Fong and E. Berger. The locations of short gamma-ray bursts as evidence for compact object binary progenitors. *The Astrophysical Journal*, 776:18, 2013. doi: 10.1088/0004-637X/776/1/18.

[56] W. Fong, E. Berger, R. Margutti, and B. A. Zauderer. A decade of short-duration gamma-ray burst broadband afterglows. *The Astrophysical Journal*, 815:102, 2015. doi: 10.1088/0004-637X/815/2/102.

[57] W. Fong, P. K. Blanchard, K. D. Alexander, R. Strausbaugh, A. Cucchiara, P. S. Cowperthwaite, C. R. Angus, et al. The optical afterglow of GW170817: An off-axis structured jet and deep constraints on a globular cluster origin. *ApJ Letters*, 883:L1, 2019. doi: 10.3847/2041-8213/ab3d9e.

[58] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC hammer. *PASP*, 125:306–312, 2013. doi: 10.1086/670067.

[59] D. A. Frail et al. Beaming in gamma-ray bursts: Evidence for a standard energy reservoir. *ApJ Letters*, 562:L55–L58, 2001. doi: 10.1086/338119.

[60] A. S. Fruchter et al. Long γ-ray bursts and core-collapse supernovae have different environments. *Nature*, 441:463–468, 2006. doi: 10.1038/nature04787.

[61] J. P. U. Fynbo et al. No supernovae associated with two long-duration γ-ray bursts. *Nature*, 444:1047–1049, 2006. doi: 10.1038/nature05375.

[62] A. Gal-Yam, D. B. Fox, P. A. Price, et al. A novel explosive process is required for the γ-ray burst GRB 060614. *Nature*, 444:1053–1055, 2006. doi: 10.1038/nature05373.

[63] T. J. Galama et al. An unusual supernova in the error box of the γ-ray burst of 25 april 1998. *Nature*, 395:670–672, 1998. doi: 10.1038/27150.

[64] Y. A. Gallant and A. Achterberg. Ultra-high-energy cosmic ray acceleration at relativistic shock fronts. *Monthly Notices of the Royal Astronomical Society*, 305:L6–L10, 1999. doi: 10.1046/j.1365-8711.1999.02566.x.

[65] N. Gehrels, J. P. Norris, S. D. Barthelmy, et al. A new gamma-ray burst classification scheme from GRB 060614. *Nature*, 444:1044–1046, 2006. doi: 10.1038/nature05376.

[66] J.-J. Geng, B. Zhang, A. Kölligan, R. Kuiper, and Y.-F. Huang. Propagation of a short GRB jet in the ejecta: Jet launching delay time, jet structure, and GW170817/GRB 170817A. *ApJ Letters*, 877:L40, 2019. doi: 10.3847/2041-8213/ab224b.

[67] G. Ghirlanda, L. Nava, G. Ghisellini, A. Celotti, and C. Firmani. Short versus long gamma-ray bursts: spectra, energetics, and luminosities. *Astronomy and Astrophysics*, 496(3):585–595, March 2009. doi: 10.1051/0004-6361/200811209.

[68] G. Ghirlanda, O. S. Salafia, Z. Paragi, et al. Compact radio emission indicates a structured jet was produced by a binary neutron star merger. *Science*, 363:968–971, 2019. doi: 10.1126/science.aau8815.

[69] G. Gianfagna, L. Piro, F. Pannarale, H. Van Eerten, F. Ricci, and G. Ryan. Potential biases and prospects for the hubble constant estimation via electromagnetic and gravitational-wave joint analyses. *Monthly Notices of the Royal Astronomical Society*, 528:2600–2613, 2024. doi: 10.1093/mnras/stae198.

[70] J. H. Gillanders et al. Modelling the optical emission of GRB 230307A. *arXiv*, 2023. arXiv:2308.00633.

[71] A. Goldstein et al. An ordinary short gamma-ray burst with extraordinary implications: Fermi-GBM detection of GRB 170817A. *ApJ Letters*, 848:L14, 2017. doi: 10.3847/2041-8213/aa8f41.

[72] B. P. Gompertz et al. The case for a minute-long merger-driven gamma-ray burst from fast-cooling synchrotron emission. *Nature Astron.*, 7:67–79, 2023. doi: 10.1038/s41550-022-01819-4.

[73] O. Gottlieb, E. Nakar, T. Piran, and K. Hotokezaka. A cocoon shock breakout as the origin of the γ-ray emission in GW170817. *Monthly Notices of the Royal Astronomical Society*, 479:588–600, 2018. doi: 10.1093/mnras/sty1462.

[74] O. Gottlieb, O. Bromberg, C. B. Singh, and E. Nakar. Black hole to photosphere: 3d GRMHD simulations of collapsars. *Monthly Notices of the Royal Astronomical Society*, 498:3320–3333, 2020. doi: 10.1093/mnras/staa2567.

[75] O. Gottlieb, B. D. Metzger, E. Quataert, et al. Black hole to photosphere: GRMHD simulations of collapsars. *ApJ Letters*, 958:L33, 2023. doi: 10.3847/2041-8213/ad096e.

[76] Ore Gottlieb, Mathieu Renzo, Brian D. Metzger, Jared A. Goldberg, and Matteo Cantiello. She's Got Her Mother's Hair: Unveiling the Origin of Black Hole Magnetic Fields through Stellar to Collapsar Simulations. *ApJ Letters*, 976(1):L13, November 2024. doi: 10.3847/2041-8213/ad8563.

[77] T. Govreen-Segal and E. Nakar. Analytic model for off-axis GRB afterglow images: geometry measurement and implications for measuring h_0. *Monthly Notices of the Royal Astronomical Society*, 524:403–425, 2023. doi: 10.1093/mnras/stad1628.

[78] J. Granot and P. Kumar. Constraining the structure of Gamma-Ray Burst jets through the afterglow light curves. *The Astrophysical Journal*, 591:1086–1096, 2003. doi: 10.1086/375489.

[79] J. Granot and R. Sari. The shape of spectral breaks in gamma-ray burst afterglows. *The Astrophysical Journal*, 568:820–829, 2002. doi: 10.1086/338966.

[80] A. Hajela, R. Margutti, K. D. Alexander, et al. Two years of nonthermal emission from the binary neutron star merger GW170817: Rapid fading of the jet afterglow and first constraints on the kilonova fastest ejecta. *ApJ Letters*, 886:L17, 2019. doi: 10.3847/2041-8213/ab5226.

[81] G. Hallinan, A. Corsi, K. P. Mooley, K. Hotokezaka, E. Nakar, M. M. Kasliwal, D. L. Kaplan, D. A. Frail, S. T. Myers, T. Murphy, K. De, D. Dobie, J. R. Allison, K. W. Bannister, V. Bhalerao, P. Chandra, T. E. Clarke, S. Giacintucci, A. Y. Q. Ho, A. Horesh, N. E. Kassim, S. R. Kulkarni, E. Lenc, F. J. Lockman, C. Lynch, D. Nichols, S. Nissanke, N. Palliyaguru, W. M. Peters, T. Piran, J. Rana, E. M. Sadler, and L. P. Singer. A radio counterpart to a neutron star merger. *Science*, 358(6370):1579–1583, 2017. doi: 10.1126/science.aap9855.

[82] A. Harten, P. D. Lax, and B. van Leer. On upstream differencing and Godunov-type schemes for hyperbolic conservation laws. *SIAM Rev.*, 25:35–61, 1983. doi: 10.1137/1025002.

[83] M. Hempel and J. Schaffner-Bielich. A statistical model for a complete supernova equation of state. *Nucl. Phys. A*, 837:210–254, 2010. doi: 10.1016/j.nuclphysa.2010.02.010.

[84] J. Hjorth et al. A very energetic supernova associated with the γ-ray burst of 29 march 2003. *Nature*, 423:847–850, 2003. doi: 10.1038/nature01750.

[85] B. James, A. Janiuk, and F. H. Nouri. Modeling the Gamma-Ray Burst Jet Properties with 3D General Relativistic Simulations of Magnetically Arrested Accretion Flows. *The Astrophysical Journal*, 935(2):176, August 2022. doi: 10.3847/1538-4357/ac81b7.

[86] A. Janiuk, B. James, and I. Palit. Variability of Magnetically Dominated Jets from Accreting Black Holes. *The Astrophysical Journal*, 917(2):102, August 2021. doi: 10.3847/1538-4357/ac0624.

[87] Z.-P. Jin, X. Li, Z. Cano, S. Covino, Y.-Z. Fan, and D.-M. Wei. The light curve of the macronova associated with the long-short burst GRB 060614. *ApJ Letters*, 811:L22, 2015. doi: 10.1088/2041-8205/811/2/L22.

[88] O. Just, M. Obergaulinger, and H.-T. Janka. A new multidimensional, energy-dependent two-moment transport code for neutrino-hydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 453:3386–3413, 2015. doi: 10.1093/mnras/stv1892.

[89] O. Just, M. Obergaulinger, H.-T. Janka, A. Bauswein, and N. Schwarz. Neutron-star merger ejecta as obstacles to neutrino-powered jets and the importance of nucleonic pair annihilation for gamma-ray burst engines. *ApJ Letters*, 816:L30, 2016. doi: 10.3847/2041-8205/816/2/L30.

[90] D. Kasen, B. Metzger, J. Barnes, E. Quataert, and E. Ramirez-Ruiz. Origin of the heavy elements in binary neutron-star mergers from a gravitational-wave event. *Nature*, 551:80–84, 2017. doi: 10.1038/nature24453.

[91] A. Kathirgamaraju, A. Tchekhovskoy, D. Giannios, and R. Barniol Duran. Em counterparts of structured jets from 3D GRMHD simulations. *Monthly Notices of the Royal Astronomical Society*, 484:L98–L103, 2019. doi: 10.1093/mnrasl/slz012.

[92] K. Kiuchi, P. Cerdá-Durán, K. Kyutoku, Y. Sekiguchi, and M. Shibata. Efficient magnetic-field amplification due to the Kelvin-Helmholtz instability in binary neutron star mergers. *Phys. Rev. D.*, 92:064034, 2015. doi: 10.1103/PhysRevD.92.064034.

[93] R. W. Klebesadel, I. B. Strong, and R. A. Olson. Observations of gamma-ray bursts of cosmic origin. *ApJ Letters*, 182:L85–L88, 1973. doi: 10.1086/181225.

[94] S. S. Komissarov and S. A. E. G. Falle. The large-scale structure of FR-II radio sources. *Monthly Notices of the Royal Astronomical Society*, 297:1087–1108, 1998. doi: 10.1046/j.1365-8711.1998.01547.x.

[95] C. Kouveliotou, C. A. Meegan, G. J. Fishman, N. P. Bhat, M. S. Briggs, T. M. Koshut, W. S. Paciesas, and G. N. Pendleton. Identification of two classes of gamma-ray bursts. *ApJ Letters*, 413:L101–L104, 1993. doi: 10.1086/186969.

[96] P. Kumar and J. Granot. The evolution of a structured relativistic jet and gamma-ray burst afterglow light curves. *The Astrophysical Journal*, 591:1075–1085, 2003. doi: 10.1086/375186.

[97] P. Kumar and B. Zhang. The physics of gamma-ray bursts & relativistic jets. *Phys. Rep.*, 561:1–109, 2015. doi: 10.1016/j.physrep.2014.09.008.

[98] E. Laplace, S. Justham, M. Renzo, Y. Götberg, R. Farmer, D. Vartanyan, and S. E. de Mink. Different to the core: The pre-supernova structures of massive single and binary-stripped stars. *Astronomy and Astrophysics*, 656: A58, December 2021. doi: 10.1051/0004-6361/202140506.

[99] J. M. Lattimer and D. N. Schramm. Black-hole-neutron-star collisions. *ApJ Letters*, 192:L145–L147, 1974. doi: 10.1086/181612.

[100] D. Lazzati, B. J. Morsony, and M. C. Begelman. Very high efficiency photospheric emission in long-duration gamma-ray bursts. *The Astrophysical Journal*, 700:47–55, 2009. doi: 10.1088/0004-637X/700/1/47.

[101] D. Lazzati, R. Perna, B. J. Morsony, D. Lopez-Camara, M. Cantiello, R. Ciolfi, B. Giacomazzo, and J. C. Workman. Late time afterglow observations reveal a collimated relativistic jet in the ejecta of the binary neutron star merger GW170817. *Physical Review Letters*, 120:241103, 2018. doi: 10.1103/PhysRevLett.120.241103.

[102] J. M. LeBlanc and J. R. Wilson. A numerical example of the collapse of a rotating magnetized star. *The Astrophysical Journal*, 161:541–551, 1970. doi: 10.1086/150558.

[103] A. J. Levan, B. P. Gompertz, O. S. Salafia, et al. Heavy element production in a compact object merger observed by JWST. *Nature*, 626:737–741, 2024. doi: 10.1038/s41586-023-06759-1.

[104] E. M. Levesque, L. J. Kewley, E. Berger, and H. Jabran Zahid. The host galaxies of gamma-ray bursts. I. interstellar medium properties of ten nearby long-duration gamma-ray burst hosts. *The Astronomical Journal*, 139:694–703, 2010. doi: 10.1088/0004-6256/139/2/694.

[105] L.-X. Li and B. Paczyński. Transient events from neutron star mergers. *ApJ Letters*, 507:L59–L62, 1998. doi: 10.1086/311680.

[106] X. J. Li, Z. B. Zhang, and K. Zhang. Overall spectral properties of prompt emissions with diverse segments in Swift/BAT short gamma-ray bursts. *Astronomy and Astrophysics*, 657:A124, January 2022. doi: 10.1051/0004-6361/202140747.

[107] Y. Lithwick and R. Sari. Lower limits on Lorentz factors in gamma-ray bursts. *The Astrophysical Journal*, 555:540–545, 2001. doi: 10.1086/321455.

[108] B. Lowell, J. Jacquemin-Ide, A. Tchekhovskoy, and A. Duncan. Rapid Black Hole Spin-down by Thick Magnetically Arrested Disks. *The Astrophysical Journal*, 960(1):82, January 2024. doi: 10.3847/1538-4357/ad09af.

[109] A. I. MacFadyen and S. E. Woosley. Collapsars: Gamma-ray bursts and the explosion of massive stars. *The Astrophysical Journal*, 524:262–289, 1999. doi: 10.1086/307790.

[110] A. I. MacFadyen, S. E. Woosley, and A. Heger. Supernovae, jets, and collapsars. *The Astrophysical Journal*, 550:410–425, 2001. doi: 10.1086/319698.

[111] S. Makhathini, K. P. Mooley, M. Brightman, K. Hotokezaka, A. J. Nayana, H. T. Intema, D. Dobie, et al. The panchromatic afterglow of GW170817: The full uniform data set, modeling, comparison with previous results, and implications. *The Astrophysical Journal*, 922:154, 2021. doi: 10.3847/1538-4357/ac1ffc.

[112] A. Marek, H. Dimmelmeier, H.-T. Janka, E. Müller, and R. Buras. Exploring the relativistic regime with Newtonian hydrodynamics: an improved effective gravitational potential for supernova simulations. *Astronomy and Astrophysics*, 445:273–289, 2006. doi: 10.1051/0004-6361:20052840. [Key is Marek2005 but published 2006; keep the key as used in the chapter or rename consistently.].

[113] B. Margalit and B. D. Metzger. The multi-messenger matrix: The future of neutron star merger constraints on the nuclear equation of state. *ApJ Letters*, 880:L15, 2019. doi: 10.3847/2041-8213/ab2ae2.

[114] R. Margutti, K. D. Alexander, X. Xie, L. Sironi, B. D. Metzger, A. Kathirgamaraju, et al. The binary neutron star event LIGO/Virgo GW170817 160 days after merger: Synchrotron emission across the electromagnetic spectrum. *ApJ Letters*, 856:L18, 2018. doi: 10.3847/2041-8213/aab2ad.

[115] J. M. Martí. Numerical simulations of jets from active galactic nuclei. *Galaxies*, 3:199–217, 2015. doi: 10.3390/galaxies3040199.

[116] J. M. Martí, E. Müller, and J. M. Ibáñez. Numerical relativistic hydrodynamics: local characteristic approach. *Astronomy and Astrophysics*, 281:L9–L12, 1994.

[117] J. M. Martí, E. Müller, J. A. Font, J. M. Ibáñez, and A. Marquina. Morphology and dynamics of relativistic jets. *The Astrophysical Journal*, 479:151–163, 1997. doi: 10.1086/303842.

[118] J. C. McKinney, A. Tchekhovskoy, and R. D. Blandford. General relativistic magnetohydrodynamic simulations of magnetically choked accretion flows around black holes. *Monthly Notices of the Royal Astronomical Society*, 423 (4):3083–3117, July 2012. doi: 10.1111/j.1365-2966.2012.21074.x.

[119] C. A. Meegan et al. Spatial distribution of γ-ray burst sources. *Nature*, 355: 143–145, 1992. doi: 10.1038/355143a0.

[120] A. Mei et al. Gigaelectronvolt emission from a compact binary merger. *Nature*, 612:236–239, 2022. doi: 10.1038/s41586-022-05404-x.

[121] P. Mészáros and M. J. Rees. Relativistic fireballs and their impact on external matter. *The Astrophysical Journal*, 405:278–284, 1993. doi: 10.1086/172360.

[122] P. Mészáros and M. J. Rees. Optical and long-wavelength afterglow from gamma-ray bursts. *The Astrophysical Journal*, 476:232–237, 1997. doi: 10.1086/303625.

[123] B. D. Metzger, A. Arcones, E. Quataert, and G. Martínez-Pinedo. The effects of r-process heating on fallback accretion in compact object mergers. *Monthly Notices of the Royal Astronomical Society*, 402:2771–2784, 2010. doi: 10.1111/j.1365-2966.2009.16107.x.

[124] B. D. Metzger, G. Martínez-Pinedo, S. Darbha, E. Quataert, A. Arcones, D. Kasen, R. Thomas, P. Nugent, I. V. Panov, and N. T. Zinner. Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei. *Monthly Notices of the Royal Astronomical Society*, 406: 2650–2662, 2010. doi: 10.1111/j.1365-2966.2010.16864.x.

[125] B. D. Metzger, D. Giannios, T. A. Thompson, N. Bucciantini, and E. Quataert. The protomagnetar model for gamma-ray bursts. *Monthly Notices of the Royal Astronomical Society*, 413:2031–2056, 2011. doi: 10.1111/j.1365-2966.2011.18280.x.

[126] M. R. Metzger, S. G. Djorgovski, S. R. Kulkarni, C. C. Steidel, K. L. Adelberger, D. A. Frail, E. Costa, and F. Frontera. Spectral constraints on the redshift of the optical counterpart to the γ-ray burst of 8 may 1997. *Nature*, 387:878–880, 1997. doi: 10.1038/43132.

[127] A. Mignone and G. Bodo. An HLLC riemann solver for relativistic flows – II. magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 368:1040–1054, 2006. doi: 10.1111/j.1365-2966.2006.10162.x.

[128] A. Mignone and J. C. McKinney. Equation of state in relativistic magnetohydrodynamics: variable versus constant adiabatic index. *Monthly Notices of the Royal Astronomical Society*, 378:1118–1130, 2007. doi: 10.1111/j.1365-2966.2007.11849.x.

[129] A. Mignone and P. Tzeferacos. A second-order unsplit godunov scheme for cell-centered MHD: the CTU-GLM scheme. *J. Comput. Phys.*, 229:2117–2138, 2010. doi: 10.1016/j.jcp.2009.11.026.

[130] A. Mignone, T. Plewa, and G. Bodo. The piecewise parabolic method for multidimensional relativistic fluid dynamics. *ApJ Supplements*, 160:199–219, 2005. doi: 10.1086/430905.

[131] A. Mignone, G. Bodo, S. Massaglia, T. Matsakos, O. Tesileanu, C. Zanni, and A. Ferrari. PLUTO: a numerical code for computational astrophysics. *ApJ Supplements*, 170:228–242, 2007. doi: 10.1086/513316.

[132] A. Mignone, M. Ugliano, and G. Bodo. A five-wave Harten-Lax-van Leer Riemann solver for relativistic magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 393:1141–1156, 2009. doi: 10.1111/j.1365-2966.2008.14221.x.

[133] A. Mignone, P. Tzeferacos, and G. Bodo. High-order conservative finite difference GLM–MHD schemes for cell-centered MHD. *J. Comput. Phys.*, 229: 5896–5920, 2010. doi: 10.1016/j.jcp.2010.04.013.

[134] A. Mizuta and M. A. Aloy. Angular energy distribution of collapsar-jets. *The Astrophysical Journal*, 699:1261–1273, 2009. doi: 10.1088/0004-637X/699/2/1261.

[135] A. Mizuta, T. Yamasaki, S. Nagataki, and S. Mineshige. Collimated jets from collapsars and their interaction with the stellar matter. *The Astrophysical Journal*, 651:960–978, 2006. doi: 10.1086/507861.

[136] K. P. Mooley, A. T. Deller, O. Gottlieb, et al. Superluminal motion of a relativistic jet in the neutron-star merger GW170817. *Nature*, 561:355–359, 2018. doi: 10.1038/s41586-018-0486-3.

[137] K. P. Mooley, A. T. Deller, O. Gottlieb, et al. Superluminal motion of a relativistic jet in the neutron-star merger GW170817. *Nature*, 561:355–359, 2018. doi: 10.1038/s41586-018-0486-3.

[138] K. P. Mooley, D. A. Frail, D. Dobie, et al. A strong jet signature in the late-time light curve of GW170817. *ApJ Letters*, 868:L11, 2018. doi: 10.3847/2041-8213/aaeda7.

[139] K. P. Mooley, J. Anderson, and W. Lu. Optical superluminal motion measurement in the neutron-star merger GW170817. *Nature*, 610:273–276, 2022. doi: 10.1038/s41586-022-05347-3.

[140] B. J. Morsony, D. Lazzati, and M. C. Begelman. Temporal and angular properties of gamma-ray burst jets emerging from massive stars. *The Astrophysical Journal*, 665:569–598, 2007. doi: 10.1086/519483.

[141] P. Mösta, C. D. Ott, D. Radice, L. F. Roberts, E. Schnetter, and R. Haas. A large-scale dynamo and magnetoturbulence in rapidly rotating core-collapse supernovae. *Nature*, 528:376–379, 2015. doi: 10.1038/nature15755.

[142] H. Nagakura, H. Ito, K. Kiuchi, and S. Yamada. Jet propagation in collapsing massive stars with realistic equations of state. *The Astrophysical Journal*, 731:80, 2011. doi: 10.1088/0004-637X/731/2/80.

[143] H. Nagakura, K. Hotokezaka, Y. Sekiguchi, M. Shibata, and K. Ioka. Jet collimation in the ejecta of double neutron star mergers. *ApJ Letters*, 784:L28, 2014. doi: 10.1088/2041-8205/784/2/L28.

[144] E. Nakar and T. Piran. The observable signatures of GRB cocoons. *The Astrophysical Journal*, 834:28, 2017. doi: 10.3847/1538-4357/834/1/28.

[145] E. Nakar and T. Piran. Afterglow constraints on the viewing angle of binary neutron star mergers and determination of the hubble constant. *The Astrophysical Journal*, 909:114, 2021. doi: 10.3847/1538-4357/abd6cd.

[146] R. Narayan, B. Paczyński, and T. Piran. Gamma-ray bursts as the death throes of massive binary stars. *ApJ Letters*, 395:L83–L86, 1992. doi: 10.1086/186493.

[147] R. Narayan, I. V. Igumenshchev, and M. A. Abramowicz. Magnetically arrested disk: an energetically efficient accretion flow. *PASJ*, 55:L69–L72, 2003. doi: 10.1093/pasj/55.6.L69.

[148] A. Nathanail, R. Gill, O. Porth, C. M. Fromm, and L. Rezzolla. 3d magnetized jet breakout from neutron star merger ejecta: after-glow emission from GW170817. *Monthly Notices of the Royal Astronomical Society*, 502:1843–1855, 2021. doi: 10.1093/mnras/stab115.

[149] L. Nativi, G. P. Lamb, S. Rosswog, C. Lundman, and G. Kowal. Are the jets of short gamma-ray bursts affected by the viewing angle of the binary orbital plane? *Monthly Notices of the Royal Astronomical Society*, 509:903–913, 2022. doi: 10.1093/mnras/stab2982.

[150] L. Nava, G. Ghirlanda, G. Ghisellini, and A. Celotti. Spectral properties of 438 GRBs detected by fermi/GBM. *Astronomy and Astrophysics*, 530:A21, 2011. doi: 10.1051/0004-6361/201016270.

[151] L. Nava, L. Sironi, G. Ghisellini, A. Celotti, and G. Ghirlanda. Afterglow emission in gamma-ray bursts – I. physical distinctions between spectral regimes. *Monthly Notices of the Royal Astronomical Society*, 433:2107–2121, 2013. doi: 10.1093/mnras/stt872.

[152] J. P. Norris and J. T. Bonnell. Short gamma-ray bursts with extended emission. *The Astrophysical Journal*, 643:266–275, 2006. doi: 10.1086/502796.

[153] J. P. Norris, R. J. Nemiroff, J. T. Bonnell, J. D. Scargle, C. Kouveliotou, W. S. Paciesas, C. A. Meegan, and G. J. Fishman. Attributes of pulses in long bright gamma-ray bursts. *The Astrophysical Journal*, 459:393–412, 1996. doi: 10.1086/176902.

[154] J. A. Nousek et al. Evidence for a canonical gamma-ray burst afterglow light curve in the swift XRT data. *The Astrophysical Journal*, 642:389–400, 2006. doi: 10.1086/500724.

[155] M. Obergaulinger and M. A. Aloy. Protomagnetar and black hole formation in high-mass stars. *Monthly Notices of the Royal Astronomical Society*, 469: L43–L47, 2017. doi: 10.1093/mnrasl/slx046.

[156] M. Obergaulinger and M. A. Aloy. Magnetorotational core collapse of possible GRB progenitors. I. explosion mechanisms. *Monthly Notices of the Royal Astronomical Society*, 492:4613–4634, 2020. doi: 10.1093/mnras/staa096.

[157] M. Obergaulinger and M. A. Aloy. Magnetorotational core collapse of possible GRB progenitors – I. explosion mechanisms. *Monthly Notices of the Royal Astronomical Society*, 503:4942–4971, 2021. doi: 10.1093/mnras/stab295.

[158] M. Obergaulinger and M. A. Aloy. Magnetorotational core collapse of possible GRB progenitors. IV. a wider range of progenitors. *Monthly Notices of the Royal Astronomical Society*, 512:2489–2507, 2022. doi: 10.1093/mnras/stac613.

[159] E. O'Connor and C. D. Ott. Black Hole Formation in Failing Core-Collapse Supernovae. *The Astrophysical Journal*, 730(2):70, April 2011. doi: 10.1088/0004-637X/730/2/70.

[160] G. Oganesyan, S. Ascenzi, M. Branchesi, O. S. Salafia, S. Dall'Osso, and G. Ghirlanda. Structured jets and x-ray plateaus from gyration of jets around the magnetic axis of a millisecond pulsar in gamma-ray bursts. *ApJ Letters*, 893:L31, 2020. doi: 10.3847/2041-8213/ab8218.

[161] F. Özel and P. Freire. Masses, radii, and the equation of state of neutron stars. *Annual Review of Astronomy and Astrophysics*, 54:401–440, 2016. doi: 10.1146/annurev-astro-081915-023322.

[162] B. Paczyński. Gamma-ray bursters at cosmological distances. *ApJ Letters*, 308:L43–L46, 1986. doi: 10.1086/184740.

[163] A. Palmese et al. Evidence for dynamically driven formation of the GW170817 neutron star binary in NGC 4993. *ApJ Letters*, 849:L34, 2017. doi: 10.3847/2041-8213/aa9660.

[164] A. Panaitescu and P. Kumar. Properties of relativistic jets in gamma-ray burst afterglows. *The Astrophysical Journal*, 571:779–789, 2002. doi: 10.1086/340094.

[165] A. Pavan, R. Ciolfi, J. V. Kalinani, and A. Mignone. Short gamma-ray burst jet propagation in BNS merger environments. *Monthly Notices of the Royal Astronomical Society*, 506:3483–3498, 2021. doi: 10.1093/mnras/stab1959.

[166] A. Pavan, R. Ciolfi, J. V. Kalinani, and A. Mignone. 3d simulations of magnetized short gamma-ray burst jets in BNS merger environments. *Monthly Notices of the Royal Astronomical Society*, 524:260–285, 2023. doi: 10.1093/mnras/stad1910.

[167] A. Pavan, R. Ciolfi, and J. V. Dreas E., Kalinani. Role of injection parameters in jet propagation through realistic binary neutron star merger environments. *Monthly Notices of the Royal Astronomical Society*, 540(1):1345–1362, June 2025. doi: 10.1093/mnras/staf795.

[168] B. Paxton, L. Bildsten, A. Dotter, F. Herwig, P. Lesaffre, and F. Timmes. Modules for experiments in stellar astrophysics (MESA). *ApJ Supplements*, 192:3, 2011. doi: 10.1088/0067-0049/192/1/3.

[169] B. Paxton, M. Cantiello, P. Arras, L. Bildsten, E. F. Brown, A. Dotter, C. Mankovich, M. H. Montgomery, D. Stello, F. X. Timmes, and R. Townsend. Modules for experiments in stellar astrophysics (MESA): Planets, oscillations, rotation, and massive stars. *ApJ Supplements*, 208:4, 2013. doi: 10.1088/0067-0049/208/1/4.

[170] B. Paxton, P. Marchant, J. Schwab, E. B. Bauer, L. Bildsten, M. Cantiello, L. Dessart, R. Farmer, H. Hu, N. Langer, R. H. D. Townsend, D. M. Townsley, and F. X. Timmes. Modules for experiments in stellar astrophysics (MESA):

Binaries, pulsations, and explosions. *ApJ Supplements*, 220:15, 2015. doi: 10.1088/0067-0049/220/1/15.

[171] A. Pe'er. Dynamical model of an expanding shell. *ApJ Letters*, 752:L8, 2012. doi: 10.1088/2041-8205/752/1/L8.

[172] Asaf Pe'er. Gamma-Ray Bursts: What Do We Know Today That We Did Not Know 10 Years Ago? *Galaxies*, 13(1):2, December 2024. doi: 10.3390/galaxies13010002.

[173] A. Perego, S. Rosswog, R. M. Cabezón, O. Korobkin, R. Käppeli, A. Arcones, and M. Liebendörfer. Neutrino-driven winds from neutron star merger remnants. *Monthly Notices of the Royal Astronomical Society*, 443(4):3134–3156, October 2014. doi: 10.1093/mnras/stu1352.

[174] R. A. Perley, C. J. Chandler, B. J. Butler, and J. M. Wrobel. The Expanded Very Large Array: A New Telescope for New Science. *ApJ Letters*, 739(1):L1, June 2011. doi: 10.1088/2041-8205/739/1/L1.

[175] P. C. Peters. Gravitational radiation and the motion of two point masses. *Phys. Rev.*, 136:B1224–B1232, 1964. doi: 10.1103/PhysRev.136.B1224.

[176] T. Piran. Gamma-ray bursts and the fireball model. *Phys. Rep.*, 314:575–667, 1999. doi: 10.1016/S0370-1573(98)00127-6.

[177] K. A. Postnov and L. R. Yungelson. The evolution of compact binary star systems. *Living Rev. Relativ.*, 17:3, 2014. doi: 10.12942/lrr-2014-3.

[178] D. J. Price and S. Rosswog. Producing ultrastrong magnetic fields in neutron star mergers. *Science*, 312:719–722, 2006. doi: 10.1126/science.1125201.

[179] J. L. Racusin, E. W. Liang, D. N. Burrows, A. Falcone, T. Sakamoto, B. B. Zhang, B. Zhang, P. Evans, and J. Osborne. Jet Breaks and Energetics of Swift Gamma-Ray Burst X-Ray Afterglows. *The Astrophysical Journal*, 698 (1):43–74, June 2009. doi: 10.1088/0004-637X/698/1/43.

[180] D. Radice, A. Perego, K. Hotokezaka, S. Bernuzzi, S. A. Fromm, and L. F. Roberts. Binary neutron star mergers: Mass ejection, electromagnetic counterparts, and nucleosynthesis. *The Astrophysical Journal*, 869:130, 2018. doi: 10.3847/1538-4357/aaf054.

[181] M. Rampp and H.-T. Janka. Radiation hydrodynamics with neutrinos: Variable Eddington factor method for core-collapse supernova simulations. *Astronomy and Astrophysics*, 396:361–392, 2002. doi: 10.1051/0004-6361:20021398.

[182] J. C. Rastinejad, B. P. Gompertz, A. J. Levan, et al. A kilonova following a long-duration gamma-ray burst at 350 Mpc. *Nature*, 612:223–227, 2022. doi: 10.1038/s41586-022-05390-w.

[183] M. J. Rees and P. Mészáros. Relativistic fireballs: energy conversion and time-scales. *Monthly Notices of the Royal Astronomical Society*, 258:41P–43P, 1992. doi: 10.1093/mnras/258.1.41P.

[184] M. J. Rees and P. Mészáros. Refreshed shocks and afterglow longevity in gamma-ray bursts. *ApJ Letters*, 496:L1–L4, 1998. doi: 10.1086/311244.

[185] L. Resmi, S. Schulze, C. H. Ishwara-Chandra, K. Misra, J. Buchner, M. De Pasquale, et al. Low-frequency view of GW170817/GRB 170817a with the giant metrewave radio telescope. *The Astrophysical Journal*, 867:57, 2018. doi: 10.3847/1538-4357/aae1a6.

[186] J. E. Rhoads. The dynamics and light curves of beamed gamma-ray burst afterglows. *The Astrophysical Journal*, 525:737–749, 1999. doi: 10.1086/307907.

[187] E. Rossi, D. Lazzati, and M. J. Rees. Afterglow light curves, viewing angle and the jet structure of gamma-ray bursts. *Monthly Notices of the Royal Astronomical Society*, 332:945–950, 2002. doi: 10.1046/j.1365-8711.2002.05363.x.

[188] S. Rosswog. Fallback accretion in the aftermath of a compact binary merger. *MNRAS Lett.*, 376:L48–L51, 2007. doi: 10.1111/j.1745-3933.2007.00284.x.

[189] Stephan Rosswog and Marcus Bruggen. *Introduction to High-Energy Astrophysics*. Cambridge University Press, 2003.

[190] G. Ryan, H. van Eerten, L. Piro, and E. Troja. Exploring the diversity of short gamma-ray burst afterglows with BOXFIT. *The Astrophysical Journal*, 896: 166, 2020. doi: 10.3847/1538-4357/ab93cf.

[191] G. Sadeh and E. Waxman. The Nonthermal Emission Following GW170817 is Consistent with a Conical Radially Stratified Outflow with Initial Lorentz Factor $\lesssim$10. *The Astrophysical Journal*, 987(2):178, July 2025. doi: 10.3847/1538-4357/ade150.

[192] O. S. Salafia and G. Ghirlanda. The structure of gamma-ray burst jets. *Galaxies*, 10:93, 2022. doi: 10.3390/galaxies10050093.

[193] O. S. Salafia, G. Ghirlanda, S. Ascenzi, and G. Ghisellini. On the jet structure and the Lorentz factor distribution of GRB 170817A. *Astronomy and Astrophysics*, 628:A18, 2019. doi: 10.1051/0004-6361/201935831.

[194] R. Sari and P. Mészáros. Impulsive and varying injection in gamma-ray burst afterglows. *ApJ Letters*, 535:L33–L37, 2000. doi: 10.1086/312689.

[195] R. Sari, T. Piran, and R. Narayan. Spectra and light curves of gamma-ray burst afterglows. *ApJ Letters*, 497:L17–L20, 1998. doi: 10.1086/311269.

[196] R. Sari, T. Piran, and R. Narayan. Spectra and light curves of gamma-ray burst afterglows. *ApJ Letters*, 497:L17–L20, 1998. doi: 10.1086/311269.

[197] R. Sari, T. Piran, and J. P. Halpern. Jets in gamma-ray bursts. *ApJ Letters*, 519:L17–L20, 1999. doi: 10.1086/312109.

[198] Nikhil Sarin and Paul D. Lasky. The evolution of binary neutron star post-merger remnants: a review. *General Relativity and Gravitation*, 53(6):59, June 2021. doi: 10.1007/s10714-021-02831-1.

[199] V. Savchenko et al. INTEGRAL detection of the first prompt gamma-ray signal coincident with the gravitational-wave event GW170817. *ApJ Letters*, 848:L15, 2017. doi: 10.3847/2041-8213/aa8f94.

[200] Hidetomo Sawai and Shoichi Yamada. The evolution and impacts of magnetorotational instability in magnetized core-collapse supernovae. *The Astrophysical Journal*, 817(2):153, 2016. doi: 10.3847/0004-637X/817/2/153.

[201] F. R. N. Schneider, Ph. Podsiadlowski, and B. Müller. Pre-supernova evolution, compact-object masses, and explosion properties of stripped binary stars. *Astronomy and Astrophysics*, 645:A5, January 2021. doi: 10.1051/0004-6361/202039219.

[202] R. J. Selina, E. J. Murphy, M. McKinnon, A. Beasley, B. Butler, C. Carilli, B. Clark, S. Durand, A. Erickson, W. Grammer, R. Hiriart, J. Jackson, B. Kent, B. Mason, M. Morgan, O. Y. Ojeda, V. Rosero, W. Shillue, S. Sturgis, and D. Urbain. The ngVLA Reference Design. In E. J. Murphy, editor, *Science with a Next Generation Very Large Array*, volume 517 of *ASP Conf. Ser.*, page 15, 2018.

[203] M. Shibata and K. Hotokezaka. Remnant massive neutron stars of binary neutron star mergers: evolution process and gravitational waveform. *Annu. Rev. Nucl. Part. Sci.*, 69:477–507, 2019. doi: 10.1146/annurev-nucl-101918-023625.

[204] L. Sironi, A. Spitkovsky, and J. Arons. The maximum energy of accelerated particles in relativistic collisionless shocks. *The Astrophysical Journal*, 771:54, 2013. doi: 10.1088/0004-637X/771/1/54.

[205] K. Z. Stanek et al. Spectroscopic discovery of the supernova 2003dh associated with GRB 030329. *ApJ Letters*, 591:L17–L20, 2003. doi: 10.1086/376976.

[206] A. W. Steiner, M. Hempel, and T. Fischer. Core-collapse supernova equations of state based on neutron star observations. *The Astrophysical Journal*, 774:17, 2013. doi: 10.1088/0004-637X/774/1/17.

[207] T. Sukhbold, T. Ertl, S. E. Woosley, J. M. Brown, and H.-T. Janka. Core-collapse Supernovae from 9 to 120 Solar Masses Based on Neutrino-powered Explosions. *The Astrophysical Journal*, 821(1):38, April 2016. doi: 10.3847/0004-637X/821/1/38.

[208] J. L. Synge. *The Relativistic Gas*. North-Holland, Amsterdam, 1957.

[209] N. R. Tanvir et al. The emergence of a lanthanide-rich kilonova following the merger of two neutron stars. *ApJ Letters*, 848:L27, 2017. doi: 10.3847/2041-8213/aa90b6.

[210] A. H. Taub. Relativistic rankine-hugoniot equations. *Phys. Rev.*, 74:328–334, 1948. doi: 10.1103/PhysRev.74.328.

[211] T. M. Tauris et al. Formation of double neutron star systems. *The Astrophysical Journal*, 846:170, 2017. doi: 10.3847/1538-4357/aa7e89.

[212] A. Tchekhovskoy, R. Narayan, and J. C. McKinney. Prograde and retrograde black holes: Whose jet is more powerful? *The Astrophysical Journal*, 711: 50–63, 2010. doi: 10.1088/0004-637X/711/1/50.

[213] A. Tchekhovskoy, R. Narayan, and J. C. McKinney. Efficient jet formation by a strongly magnetized black hole. *Monthly Notices of the Royal Astronomical Society*, 418:L79–L83, 2011. doi: 10.1111/j.1745-3933.2011.01147.x.

[214] E. F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*. Springer, Berlin, Heidelberg, 3rd edition, 2009. doi: 10.1007/b79761.

[215] E. F. Toro and V. A. Titarev. MUSTA fluxes for systems of conservation laws. *Journal of Computational Physics*, 216:403–429, 2006. doi: 10.1016/j.jcp.2005.12.012.

[216] E. Troja, C. L. Fryer, B. O'Connor, et al. A nearby long gamma-ray burst from a merger of compact objects. *Nature*, 612:228–232, 2022. doi: 10.1038/s41586-022-05327-3.

[217] E. Troja et al. A thousand days after the merger: continued x-ray emission from GW170817. *Monthly Notices of the Royal Astronomical Society*, 498: 5643–5651, 2020. doi: 10.1093/mnras/staa2626.

[218] G. Urrutia, F. De Colle, A. Murguia-Berthier, and E. Ramirez-Ruiz. Jet propagation and sGRB light curves in neutron star merger environments. *Monthly Notices of the Royal Astronomical Society*, 503:4363–4371, 2021. doi: 10.1093/mnras/stab723.

[219] G. Urrutia, F. De Colle, and D. López-Cámara. Three-dimensional numerical simulations of structured gamma-ray burst jets. *Monthly Notices of the Royal Astronomical Society*, 518(4):5145–5153, February 2023. doi: 10.1093/mnras/stac3401.

[220] V. A. Villar et al. The combined ultraviolet, optical, and near-infrared light curves of the kilonova associated with the binary neutron star merger GW170817. *ApJ Letters*, 851:L21, 2017. doi: 10.3847/2041-8213/aa9c84.

[221] Yi-Ying Wang, Shao-Peng Tang, Zhi-Ping Jin, and Yi-Zhong Fan. The Late Afterglow of GW170817/GRB 170817A: A Large Viewing Angle and the

Shift of the Hubble Constant to a Value More Consistent with the Local Measurements. *The Astrophysical Journal*, 943(1):13, January 2023. doi: 10.3847/1538-4357/aca96c.

[222] S. E. Woosley. Gamma-ray bursts from stellar mass accretion disks around black holes. *The Astrophysical Journal*, 405:273–277, 1993. doi: 10.1086/172359.

[223] S. E. Woosley and J. S. Bloom. The supernova gamma-ray burst connection. *Annual Review of Astronomy and Astrophysics*, 44:507–556, 2006. doi: 10.1146/annurev.astro.43.072103.150558.

[224] S. E. Woosley and A. Heger. The progenitor stars of gamma-ray bursts. *The Astrophysical Journal*, 637:914–921, 2006. doi: 10.1086/498500.

[225] S. E. Woosley and Alexander Heger. Long Gamma-Ray Transients from Collapsars. *The Astrophysical Journal*, 752(1):32, June 2012. doi: 10.1088/0004-637X/752/1/32.

[226] Zhao-Feng Wu, Michail Damoulakis, Paz Beniamini, and Dimitrios Giannios. Maximal jet energy of gamma-ray bursts through the blandford–znajek mechanism. *The Astrophysical Journal Letters*, 980(2):L28, feb 2025. doi: 10.3847/2041-8213/adaeb8. URL https://doi.org/10.3847/2041-8213/adaeb8.

[227] X. Xie, J. Zrake, and A. MacFadyen. Numerical simulations of the jet dynamics and synchrotron radiation of binary neutron star merger event GW170817/GRB 170817A. *The Astrophysical Journal*, 863:58, 2018. doi: 10.3847/1538-4357/aacf9a.

[228] B. Yang, Z.-P. Jin, X. Li, et al. A possible macronova in the late afterglow of the long-short burst GRB 060614. *Nature Communications*, 6:7323, 2015. doi: 10.1038/ncomms8323.

[229] J. Yang, S. Ai, B.-B. Zhang, et al. A long-duration gamma-ray burst with a peculiar origin. *Nature*, 626:742–745, 2024. doi: 10.1038/s41586-023-06916-6.

[230] J. Yang et al. A long-duration gamma-ray burst with a peculiar origin. *Nature*, 612:232–235, 2022. doi: 10.1038/s41586-022-05403-y.

[231] S.-C. Yoon and N. Langer. On the evolution of rapidly rotating massive white dwarfs towards supernovae or collapses. *Astronomy and Astrophysics*, 435:967–985, 2005. doi: 10.1051/0004-6361:20042542.

[232] Y.-W. Yu, B. Zhang, and H. Gao. Bright "merger-nova" from the remnant of a neutron star binary merger: a signature of a newly born, rapidly rotating millisecond pulsar. *ApJ Letters*, 776:L40, 2013. doi: 10.1088/2041-8205/776/2/L40.

[233] I. Zalamea and A. M. Beloborodov. Neutrino-antineutrino annihilation above rotating black holes. *Monthly Notices of the Royal Astronomical Society*, 410:2302–2308, 2011. doi: 10.1111/j.1365-2966.2010.17600.x.

[234] Y. Zenati, H. B. Perets, and S. Toonen. Neutron star–white dwarf mergers: early evolution, physical properties and outcomes. *Monthly Notices of the Royal Astronomical Society*, 499:4097–4113, 2020. doi: 10.1093/mnras/staa3008.

[235] B. Zhang. *The Physics of Gamma-Ray Bursts*. Cambridge University Press, 2018. doi: 10.1017/9781139226530.

[236] B. Zhang and P. Mészáros. Gamma-ray burst afterglow with continuous energy injection: Homogeneous and wind environments. *The Astrophysical Journal*, 566:712–722, 2002. doi: 10.1086/338247.

[237] B. Zhang et al. Physical processes shaping gamma-ray burst X-Ray afterglow light curves: Theoretical implications from the swift XRT observations. *The Astrophysical Journal*, 642:354–370, 2006. doi: 10.1086/500723.

[238] W. Zhang, S. E. Woosley, and A. I. MacFadyen. Relativistic jets in collapsars. *The Astrophysical Journal*, 586:356–371, 2003. doi: 10.1086/367609.

[239] Z.-B. Zhang and C.-S. Choi. An analysis of the durations of Swift gamma-ray bursts. *Astronomy and Astrophysics*, 484(2):293–297, June 2008. doi: 10.1051/0004-6361:20079210.