
Remnants of compact binary mergers and next-generation numerical relativity codes

DISSERTATION

zur Erlangung des akademischen Grades
doctor rerum naturalium (Dr. rer. nat.)



vorgelegt dem Rat der
PHYSIKALISCH-ASTRONOMISCHEN FAKULTÄT
der
FRIEDRICH-SCHILLER-UNIVERSITÄT JENA
von
Francesco ZAPPA (M.Sc.)
geboren am 30.06.1993 in Palermo (Italien)

REVIEWERS:

1. PROF. DR. SEBASTIANO BERNUZZI FRIEDRICH-SCHILLER-UNIVERSITÄT JENA
2. PROF. DR. DANIEL SIEGEL UNIVERSITÄT GREIFSWALD
3. PROF. DR. FRANCOIS FOUCART UNIVERSITY OF NEW HAMPSHIRE

DEFENSE DATE: 19 . 10 . 2023

Ehrenwörtliche Erklärung

Ich erkläre hiermit ehrenwörtlich, dass ich die vorliegende Arbeit selbständig, ohne unzulässige Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel und Literatur angefertigt habe. Die aus anderen Quellen direkt oder indirekt übernommenen Daten und Konzepte sind unter Angabe der Quelle gekennzeichnet.

Bei der Auswahl und Auswertung folgenden Materials haben mir die nachstehend aufgeführten Personen in der jeweils beschriebenen Weise unentgeltlich geholfen:

1. Prof. Sebastiano Bernuzzi
2. Prof. David Radice
3. Dr. Boris Daszuta
4. Dr. William Gregory Cook

Weitere Personen waren an der inhaltlich-materiellen Erstellung der vorliegenden Arbeit nicht beteiligt. Insbesondere habe ich hierfür nicht die entgeltliche Hilfe von Vermittlungs- bzw. Beratungsdiensten (Promotionsberater oder andere Personen) in Anspruch genommen.

Niemand hat von mir unmittelbar oder mittelbar geldwerte Leistungen für Arbeiten erhalten, die im Zusammenhang mit dem Inhalt der vorgelegten Dissertation stehen.

Die Arbeit wurde bisher weder im In- noch im Ausland in gleicher oder ähnlicher Form einer anderen Prüfungsbehörde vorgelegt.

Die geltende Promotionsordnung der Physikalisch-Astronomischen Fakultät ist mir bekannt.

Ich versichere ehrenwörtlich, dass ich nach bestem Wissen die reine Wahrheit gesagt und nichts verschwiegen habe.

Ort, Datum

Unterschrift

Abstract

Francesco ZAPPA

*Remnants of compact binary mergers and next-generation
numerical relativity codes*

Numerical relativity (NR) simulations are crucial for studying the coalescence of compact binaries. Based on NR data, we produce a model for the mass and spin of the remnant black hole (BH) for the coalescence of black hole-neutron star systems, discussing its crucial role in gravitational wave (GW) modeling and in the parameter estimation of the two signals GW200105 and GW200115. In the context of binary neutron star merger simulations, we perform the first systematic study comparing results obtained with various neutrino treatments, the presence of turbulent viscosity and different grid resolutions. We find that the time of BH formation after merger is heavily affected by grid resolution and turbulent viscosity. An early BH formation limits matter ejection from the accretion disc, as the BH swallows a significant portion of it. Our results indicate that more reliable kilonova light curves are obtained only if the various ejecta components are present. Moreover, robust r -process nucleosynthesis yields require inclusion of both neutrino emission and reabsorption in simulations. Advanced neutrino schemes and turbulent viscosity in simulations resolved beyond current standards appear necessary for reliable astrophysical predictions. To carry out computationally demanding simulations of growing complexity, next-generation NR codes that can efficiently leverage the latest pre-exascale many-core and heterogeneous infrastructures are required. To this end we develop **GR-Athena++**, a new dynamical spacetime solver built on top of **Athena++**, that shows high-order convergence properties and excellent parallel scalability up to $\mathcal{O}(10^5)$ cores in full 3D binary black hole (BBH) merger simulations. Finally we present **GR-AthenaK**, the first performance-portable spacetime solver, obtained by refactoring **GR-Athena++** with the Kokkos programming model. We demonstrate the correctness and convergence properties of **GR-AthenaK** with BBH runs on GPUs. **GR-AthenaK** shows a speedup ~ 50 on one GPU compared to **GR-Athena++** on a single CPU core.

Zusammenfassung

Francesco ZAPPA

*Überreste von Verschmelzungen kompakter Doppelsterne
und numerische Relativitätscodes der nächsten Generation*

Simulationen der numerischen Relativitätstheorie sind für die Untersuchung der Koaleszenz von kompakten Doppelsternen von entscheidender Bedeutung. Auf der Grundlage von numerischen Relativitätssimulationsdaten erstellen wir ein Modell für die Masse und den Spin des Schwarzes-Loch-Überrests für die Koaleszenz von Schwarzes-Loch-Neutronenstern-Systemen und erörtern dessen entscheidende Rolle bei der Modellierung von Gravitationswellen und bei der Parameterbestimmung der beiden Signale GW200105 und GW200115. Im Zusammenhang mit Simulationen der Verschmelzung von Neutronenstern-Doppelsternsystemen führen wir die erste systematische Studie durch, in der wir die Ergebnisse von Simulationen mit verschiedenen Neutrinobehandlungen, dem Vorhandensein von turbulenter Viskosität, und unterschiedlichen Gitterauflösungen vergleichen. Wir stellen fest, dass der Zeitpunkt der Bildung des schwarzen Lochs nach der Verschmelzung stark von der Gitterauflösung und der turbulenten Viskosität abhängt. Eine frühe Bildung des schwarzen Lochs schränkt den Materieauswurf aus der Akkretionsscheibe ein, da das schwarze Loch einen beträchtlichen Teil davon verschluckt. Unsere Ergebnisse zeigen, dass zuverlässigere Kilonova-Lichtkurven nur dann erhalten werden, wenn die verschiedenen Auswurfkomponenten vorhanden sind. Darüber hinaus erfordern robuste Ergebnisse der Nukleosynthese via des r -Prozesses die Einbeziehung sowohl der Neutrinoemission als auch der Reabsorption in die Simulationen. Simulationen mit fortschrittlichen Behandlungen der Neutrinos und der turbulenten Viskosität, die über die derzeitigen Standards hinausgehen, scheinen für zuverlässige astrophysikalische Vorhersagen notwendig zu sein. Um rechenintensive Simulationen mit wachsender Komplexität durchführen zu können, werden numerische Relativitätscodes der nächsten Generation benötigt, die die neuesten prä-Exascale Computerinfrastrukturen aus vielen und heterogenen Rechenkernen effizient nutzen können. Zu diesem Zweck entwickeln wir GR-Athena++, einen neuen dynamischen

Raumzeitlöser, der auf Athena++ aufbaut. Er zeigt Konvergenzeigenschaften hoher Ordnung und exzellente parallele Skalierbarkeit bis zu $\mathcal{O}(10^5)$ Kernen in vollständigen 3D-Simulationen der Verschmelzung von Schwarzes-Loch-Doppelsternen. Schließlich präsentieren wir GR-AthenaK, den ersten leistungsfähigen portablen Raumzeitlöser, der durch Refactoring von GR-Athena++ mit dem Kokkos-Programmiermodell entstanden ist. Wir demonstrieren die Korrektheit und Konvergenzeigenschaften von GR-AthenaK mit Simulationen von Schwarzes-Loch-Doppelsternen auf Grafikkarten. GR-AthenaK zeigt einen 50-fachen Geschwindigkeitszuwachs auf einer Grafikkarte im Vergleich zu GR-Athena++ auf einem einzelnen Prozessorkern.

Published content and contributions

Articles as leading author

For each of the following manuscripts, I substantially contributed to the writing of the entire text, the production of all figures and tables, the production of all results.

1. **Zappa F.**, Bernuzzi S., Radice D. and Perego A. “Binary neutron star merger simulations with neutrino transport and turbulent viscosity: impact of different schemes and grid resolution ”, *Mon. Not. Roy. Astron. Soc.*, vol. 520, no. 1, p. 1481-1503.
2. **Zappa F.**, Bernuzzi S., Pannarale F., Mapelli M. and Giacobbo N. “Black-Hole Remnants from Black-Hole–Neutron-Star Mergers”, *Phys. Rev. Lett.*, vol. 123, no. 4, p. 041102, 2019.

Co-authored articles

1. Gonzalez A., Gamba R., Breschi M., **Zappa F.**, Carullo G., Bernuzzi S. and Nagar A., “Numerical-Relativity-Informed Effective-One-Body model for Black-Hole-Neutron-Star Mergers with Higher Modes and Spin Precession”, *arXiv: 2212.03909*, December 2022.
Contributions: Development of the base remnant model, improved in the paper (Sec. IIA and appendix B).
2. Gonzalez A., **Zappa F.**, Breschi M., Bernuzzi S. et al., “Second release of the CoRe database of binary neutron star merger waveforms”, *preprint: 2210.16366*, October 2022.
Contributions: Development of earlier versions of some of the scripts and plots used in this work.
3. Vsevolod N., Schianchi F., Bernuzzi S., Radice D., Daszuta B., Endrizzi A., Perego A., Prakash A., **Zappa F.**, “Mapping dynamical ejecta and disk masses from numerical relativity simulations of neutron star mergers”, *Class. Quant.*

Grav., vol. 39, no. 1, p. 015008, 2022.

Contributions: Computation of the TOV sequences required to compute tidal parameters for the fitting formulas presented.

4. Daszuta B., **Zappa F.**, Cook W., Radice D, Bernuzzi S. and Morozova, V., “GR-Athena++: Puncture Evolutions on Vertex-centered Oct-tree Adaptive Mesh Refinement”, *Astrophys. J. Supp.*, vol. 257, no. 2, p. 25, 2021.
Contributions: Listed in *Chapter 3* of this *Dissertation*.
5. Agathos M., **Zappa F.**, Bernuzzi S., Perego A., Breschi M., and Radice D., “Inferring Prompt Black-Hole Formation in Neutron Star Mergers from Gravitational-Wave Data”, *Phys. Rev. D*, vol. D101, no. 4, p. 044006, 2020.
Contributions: Analysis of tidal parameter threshold and writing of corresponding text; Figures 1, 3, 7, 9, 11.
6. Bernuzzi S., Breschi M., Daszuta B., Endrizzi A., Logoteta D., Nedora V., Perego A., Schianchi F., Radice D., **Zappa F.**, Bombaci I., and Ortiz N., “Accretion-induced prompt black hole formation in asymmetric neutron star mergers, dynamical ejecta and kilonova signals”, *Mon. Not. Roy. Astron. Soc.*, Jun. 2020.
Contributions: Appendix A
7. Breschi M., Bernuzzi S., **Zappa F.**, Agathos M., Perego A., Radice D., and Nagar A., “Kilohertz gravitational waves from binary neutron star remnants: time-domain model and constraints on extreme matter”, *Phys. Rev.*, vol. D100, no. 10, p. 104029, 2019. *Contributions*: Appendix A.

Publications with LIGO-Virgo-KAGRA Collaboration

1. B. P. Abbott *et al.*, “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences,” *Astrophys. J. Lett.*, vol. 915, no.1, p. L5, 2021.
Contributions: Provide results for Sec. 4.4.
2. B. P. Abbott *et al.*, “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs,” *Phys. Rev. X*, vol. 9, no.30, p. 031040, 2019.
Contributions: Provided estimates of gravitational wave luminosity peak and radiated energy for GW170817 in table III.

Talks and seminars

1. *Talk*: “Gravitational-wave luminosity of binary neutron star mergers”, Conference “SMuK DPG meeting”, March 2019, München
2. *Talk*: “Black hole remnant of Black Hole-Neutron Star mergers”, Conference “PHAROS Annual Conference 2019: The multi-messenger physics and astrophysics of neutron stars”, April 2019, Calonge, Spain
3. *Talk*: “Black hole remnant of Black Hole-Neutron Star mergers”, Conference “GR22/Amaldi13 International Conference”, July 2019, Valencia, Spain
4. *Talk*: “GR-Athena++: an oct-tree based solver for binary black hole evolution on a vertex-centered grid”, Conference “SMuK DPG online meeting”, September 2021
5. *Talk*: “Towards Kokkos GR-Athena++: an oct-tree based spacetime solver for GPUs”, Conference “Frontiers in Numerical Relativity”, July 2022, Jena
6. *Talk*: “Impact of different input physics and grid resolution on BNS merger simulations”, Conference “SMuK DPG meeting”, March 2023, Dresden

Acknowledgements

I want to thank my supervisor, Prof. Sebastiano Bernuzzi, for all the help, support and guidance received during my Ph.D. I also explicitly thank Prof. David Radice and Prof. Albino Perego for fruitful discussions and suggestions, from which I learnt a lot.

It is my pleasure to thank all the people with whom I worked most closely during my PhD: William Cook, who I also thank for thoroughly proofreading this *Dissertation*, and Boris Daszuta. I want to thank Prof. James Stone and Hengrui Zhu for the work on AthenaK and for suggestions and bugfixing in the development of GR-AthenaK.

I extend warm greetings to the members of my group, present and past, and especially to Matteo Breschi, Vsevolod Nedora, Ssohrab Borhanian, Rossella Gamba, Gregorio Carullo and additionally to Ondřej Zelenka and Francesco Maria Fabbri for the nice time spent together.

Finally I want to thank my girlfriend, Lola, and the members of my family, Ignazia, Ignazio, Jessica, Iside, Anto and Angelica. You have always made me feel your presence with your love and encouragement throughout these years.

For my research I acknowledge support by the EU H2020 under ERC Starting Grant, no. BinGraSp-714626. Simulations were performed at the ARA cluster at Friedrich Schiller University Jena, SuperMUC_NG at the Leibniz-Rechenzentrum (LRZ) Munich and HPE Apollo Hawk at the High Performance Computing Center Stuttgart (HLRS). The ARA cluster is funded in part by DFG grants INST 275/334-1 FUGG and INST 275/363-1 FUGG, and ERC Starting Grant, grant agreement no. BinGraSp-714626. The author acknowledges the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for providing computing time on the GCS Supercomputer SuperMUC-NG at LRZ (allocation pn68wi). The author acknowledges HLRS for providing access to the supercomputer HPE Apollo Hawk under the grant number IN-TRHYGUE/44215.

Physical Constants

Name	Symbol = Value
Boltzmann constant	$k_B = 1.380649 \times 10^{-23} \text{ J K}^{-1}$
Electron mass	$m_e = 9.1093837015 \times 10^{-31} \text{ kg}$
Elementary charge	$e = 1.602176634 \times 10^{-19} \text{ C}$
Gravitational constant	$G = 6.674 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$
Neutron mass	$m_n = 1.67492749804 \times 10^{-27} \text{ kg}$
Nuclear saturation density	$\rho_{\text{nuc}} = 2.3 \times 10^{17} \text{ kg m}^{-3}$
Pi	$\pi = 3.14159265358979 \dots$
Proton mass	$m_p = 1.67262192369 \times 10^{-27} \text{ kg}$
Reduced Planck's constant	$\hbar = 1.054571817 \times 10^{-34} \text{ J s}$
Solar mass	$M_{\odot} = 1.98847 \times 10^{30} \text{ kg}$
Speed of light	$c = 299792458 \text{ m s}^{-1}$

List of Symbols

Symbol	Name	Unit
M	mass	$[M_\odot]$
r	binary separation	$[m]$, $[M_\odot]$, unit of mass $[M]$
R	radius	$[m]$, $[M_\odot]$, unit of mass $[M]$
q	mass ratio	$[\geq 1]$
ν	symmetric mass ratio	–
Λ	dimensionless tidal parameter	–
C	compactness parameter	–
a	black hole's dimensionless spin parameter	–
Q_n	convergence scaling factor of order n	–

List of Acronyms

ADM Arnowitt-Deser-Misner	21
AMR adaptive mesh refinement	21
BBH binary black hole	3
BC boundary conditions	54
BH black hole	3
BHNS black hole–neutron star	4
BNS binary neutron star	4
CB compact binary	3
CC cell-centered	54
CFL Courant-Friedrich-Lewy	21
EFE Einstein field equations	3
EOS Equation of state	9
EOB effective-one-body	5
EM electromagnetic	4
FC face-centered	54
FD finite difference	59
GPU graphical processing unit	6
GR General relativity	3
GRHD general relativistic hydrodynamics	21
GW Gravitational wave	3
HPC high-performance computing	6
ISCO innermost stable circular orbit	7
LK leakage	21

MHD magneto-hydrodynamics	6
MPI message passing interface	52
NR numerical relativity	5
NS neutron star	3
OMP OpenMP	52
RHS right-hand-side	59
RK Runge-Kutta	21
SMR static mesh refinement	92
KNEC KiloNova Explosion Code	45
SNR signal-to-noise ratio	36
SGRB short-gamma-ray burst	4
TOV Tolman-Oppenheimer-Volkoff	20
VC vertex-centered	54

Contents

Ehrenwörtliche Erklärung	i
Abstract	iii
Acknowledgements	xi
Physical Constants	xiii
List of Symbols	xv
List of Acronyms	xvi
Introduction	3
1 Model of the BH remnant of BHNS mergers	7
1.1 Physics of the remnant formation	8
1.2 Remnant’s mass and spin model	9
1.3 Applications of the model	14
1.4 GW peak luminosity	18
2 Impact of resolution and input physics on BNS merger simulations	19
2.1 Methods	20
2.2 Remnant dynamics	23
2.3 Gravitational waves	34
2.4 Ejecta mass	37
2.5 Nucleosynthesis and Kilonova light curves	43
2.6 Neutrino luminosity	47
3 Efficient BBH simulations with the GR-Athena++ spacetime solver	51
3.1 Methods	52
3.2 Mesh refinement for punctures	60
3.3 Puncture tests	64

3.4	Scaling tests	76
4	GR-AthenaK: first performance-portable dynamical spacetime solver	83
4.1	Kokkos framework	84
4.2	Kokkos version of Athena++: AthenaK	85
4.3	GR-AthenaK	86
4.4	GR-AthenaK tests	90
4.5	Towards BBH evolutions	95
	Conclusion	97
	A Numerical Relativity	103
	B Gravitational wave extraction	113
	C Geodesic spheres	115
	D Input physics in WhiskyTHC	117
	Bibliography	121

Introduction

Gravitational waves (GWs) are one of the fundamental predictions of General relativity (GR). The linearization of the Einstein field equations (EFE)

$$\mathcal{R}_{\mu\nu} - \frac{1}{2}\mathcal{R}g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \quad (1)$$

leads to the tensorial wave equations

$$\square\bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} \quad (2)$$

where $\bar{h}_{\mu\nu}$ is a small perturbation on the flat spacetime background. Considering a multipole expansion of the sources $T_{\mu\nu}$, the GW emission at leading order is proportional to the second time derivative of the source quadrupole moment, while second time derivatives of lower multipoles vanish due to the conservation of mass and total momentum of the source.

Astrophysical sources expected to produce loud GWs are those that generate large and fast-changing quadrupolar moments, such as compact binary (CB) coalescences. A CB consists of either black holes (BHs) or neutron stars (NSs), which are objects with a compactness parameter $C := \frac{M}{R} \frac{G}{c^2} \approx \mathcal{O}(10^{-1})$, where M and R represent the gravitational mass and radius of the body. The dynamics of CBs can be roughly divided into three stages. After formation, the CB quickly circularizes [1, 2] and enters a long *inspiral* phase, in which the two bodies are driven closer and closer due to the emission of energy and angular momentum via GWs. In the *merger* phase the two bodies coalesce, forming a remnant object; depending on the type of CB, this remnant object can be either a BH or a NS. The last stage is referred to as the *post-merger* phase, in which the remnant object evolves towards equilibrium. For instance, in a binary black hole (BBH) merger, the remnant object is a BH, and in the post-merger phase, it radiates away energy through its quasi normal modes (ringdown), until it reaches its fundamental state [3, 4].

The first evidence of GWs was obtained in 1974 from the observation of the pulsar

PSR J1913+16 (also known as the Hulse-Taylor binary pulsar [5]). The system consists of two NSs, one of which is the pulsar, with masses $M_1 \approx 1.44M_\odot$ and $M_2 \approx 1.39M_\odot$ respectively, orbiting the common center of mass with an orbital period that is decreasing in time. The GR prediction of the decreasing rate computed according to GW theory agrees within $\sim 0.1\%$ with the experimental value measured over a 30-years timescale observation [6].

On September 14, 2015 the first GW signal, namely GW150914, was observed by the two interferometers of the LIGO Observatory [7]. The source of the signal was identified as the coalescence of two BHs with component masses¹ of $\sim 36M_\odot$ and $\sim 29M_\odot$, which formed a remnant BH of mass $M \approx 62M_\odot$ dissipating $\sim 3M_\odot c^2$ of energy in GWs [8].

During the second observing run, the first GW signal associated with a binary neutron star (BNS) merger, namely GW170817, was detected [9] by the two LIGO interferometers and by the advanced Virgo interferometer. The GW detection was accompanied by the detection of the electromagnetic (EM) signal GRB170817A, a short-gamma-ray burst (SGRB) observed independently by the Fermi-GBM [10] and INTEGRAL telescopes [11] ~ 1.73 s after the GW detection [12]. The component masses of the BNS have been inferred as $M_1 = (1.36 - 1.60)M_\odot$ and $M_2 = (1.17 - 1.36)M_\odot$, which are consistent with the masses of known NSs [13], BNSs [14] and with theoretical predictions [15]. The joint observation of GW170817 and its associated EM counterparts represents the first direct evidence that BNS mergers are at the origin of SGRB and kilonova transients [16–34]. In particular the kilonova counterpart AT2017gfo is commonly interpreted as the UV/optical/infrared transient generated by radioactive decays of r -process elements that form in the mass ejected from the merger and the remnant [35–44].

In January 2020, during the third observing run, the LIGO-Virgo interferometers detected two GW signals consistent with the inspiral of black hole–neutron star (BHNS) binaries, named GW200105 and GW200115 [45, 46]. The masses of the primary objects constituting the two systems emitting the signals GW200105 and GW200115 were estimated as $M_1 = 8.9_{-1.5}^{+1.2} M_\odot$ and $M_1 = 5.7_{-2.1}^{+1.8} M_\odot$, respectively; these values are well above the maximum mass of NSs but compatible with BH masses. On the other hand, the masses of the secondary components were estimated as $M_2 = 1.9_{-0.2}^{+0.3} M_\odot$ and $M_2 = 1.5_{-0.3}^{+0.7} M_\odot$ respectively, thus consistent with NS masses.

¹Masses are reported as measured in the source frame.

The search for GW signals in the output data stream of interferometers involves matched filtering, which consists in testing the detectors' output against a set of waveform templates. For instance, the detection of signals from BBH mergers such as GW150914, requires about 2.5×10^5 GW templates [47]. After a signal is detected, the estimation of the source's parameters is performed with Bayesian inference techniques [48–50], which typically requires even more waveform templates covering the parameter space in a finer way. Both signal searches and parameter estimation crucially rely on accurate GW models capable of generating waveforms quickly.

The inspiral portion of GW models far from merger can be constructed analytically using e.g. post-Newtonian theory [51, 52], while other analytical treatments based on BH perturbation theory can be employed for the ringdown phase of BBH mergers to a first approximation. However, the merger phase of CB mergers and the post-merger phase for systems composed of NSs are currently inaccessible analytically. In these regimes numerical relativity (NR) simulations, in which Eqs. (1) are cast into a suitable form and solved numerically to evolve the spacetime metric [53–56], are crucial for waveform modeling. NR simulations provide the most accurate waveforms for the coalescence of CBs, but are too slow to be used alone for GW detection or parameter estimation. Accurate and fast-to-evaluate GW models are thus obtained by combining analytical and numerical methods. The three families of waveform models mainly used nowadays in GW analysis are constructed in this way, and are represented by the phenomenological approximants [57–63], NR surrogates [64–66] and models based on the effective-one-body (EOB) approach [67–69]. To model GWs in the post-merger phase of BNS and BHNS mergers it is necessary to resort to NR simulations [70–72]. The improvements of the operating sensitivity of existing detectors [73], together with incoming new experimental apparatus, such as KAGRA [74], laser interferometer space antenna (LISA) [75], Einstein Telescope [76], Cosmic Explorer [77] will enlarge the parameter space of detectable signals, requiring at the same time even more accurate GW models. This implies the need of NR simulations in extreme configurations, such as mergers of high mass ratio BBH [78, 79].

GW modeling aside, NR simulations are fundamental to study the dynamics of BHNS and BNS mergers [80–83]. An open question regards the understanding of the fate of the remnant in BNS mergers, as well as the conditions that lead to a BH prompt collapse [84–93]. The use of NR simulations is crucial in the quantitative analysis of matter ejection from BNS mergers, including the determination of its composition [94, 95]. This ejection is influenced in part by the properties of the accretion disc that forms

around the remnant object. For instance, the features of the ejected material regulate the EM emission and the r -process nucleosynthesis yields from BNS mergers, which can also be estimated with numerical methods [96–100]. Complete models of the mass ejecta are necessary to understand the still poorly known connection with the EM counterparts. These require detailed *ab-initio* simulations involving extreme gravity, relativistic magneto-hydrodynamics (MHD), and advanced microphysics models for the NS matter, including neutrino interactions and transport. In particular, long post-merger simulations with advanced input physics are needed to study the matter ejection mechanisms developing on seconds timescales [81, 101–109].

Such simulations can be computationally expensive and their feasibility is constrained by the availability of high-performance computing (HPC) infrastructures [110]. Thus, accurate NR codes that remain performant as HPC resources are scaled up and simultaneously allow for a simple extension of input physics, are crucial. Finally, the recent technological advances in terms of heterogeneous HPC platforms, push towards the production of portable NR codes, capable of efficiently making use of the new architectures which include GPU accelerators.

In this *Dissertation*, the remnants of BHNS and BNS mergers are studied utilizing NR simulations. Moreover, the *Dissertation* showcases the advancements of two new next-generation NR codes, which are crucial for detailed simulations of CB coalescences with enhanced computational efficiency.

Summary of content In *Chapter 1*, NR data are used to construct models to characterize the remnant BH of BHNS mergers; applications of such models in GW modeling and in BHNS detection are discussed, as well as consequences on the expected EM counterparts from synthetic population analysis of BHNS systems. *Chapter 2* is devoted to the systematic study of different neutrino treatments, turbulent viscosity and grid resolution effects on BNS merger simulations; the dynamics and thermodynamics of the system are discussed, together with the matter ejection and the implications on GW observations, EM counterparts and r -process nucleosynthesis yields. In *Chapter 3* the new NR code **GR-Athena++** is introduced, demonstrating its convergence and scalability properties with benchmark BBH simulations. In *Chapter 4* the developments of the first performance-portable NR code **GR-AthenaK** are presented, as well as the first numerical experiments showing convergence and performance properties with spacetime runs on graphical processing units (GPUs). Lastly, a summary of the results and future work and perspectives is reported in the Conclusion.

Chapter 1

Model of the BH remnant of BHNS mergers

The detections of the GW signals GW200105 and GW200115 [45] confirm that mergers of a stellar-mass BH and a NS are sources of GWs, detectable by ground-based laser-interferometers and possibly accompanied by EM counterparts [111–117]. Based on these observations, the merger rate of BHNS has been estimated as $\sim 45 \text{ Gpc}^{-3} \text{ yr}^{-1}$, while larger rates $\sim 130 \text{ Gpc}^{-3} \text{ yr}^{-1}$ are obtained for broader mass distributions of the component objects [45].

NR simulations are the only means to study BHNS mergers [112–114, 118–127]. The aftermath of such events is a remnant BH surrounded by an accretion disc of NS matter. Simulations indicate that the NS tidal disruption is a characteristic feature of the dynamics of quasi-circular BHNS mergers. On the contrary, quasi-circular BNS mergers with mass ratio up to ~ 2 do not present significant tidal disruption, [e.g., 128, 129].

Accurate models for the mass and dimensionless spin of the remnant BH are necessary to construct GW models to be used for matched filtering searches and parameter estimation; for instance, they enter the computation of the GW quasi-normal modes in the post-merger phase [72, 130]. An analytical formula for the BH remnant’s mass and dimensionless spin can be found using mass and angular momentum conservation arguments [131, 132] (see also [133]). This approach builds on estimates of the radiated energy and the binary orbital angular momentum based on the expressions for test particles on Kerr background at the innermost stable circular orbit (ISCO), and on the disc mass fits of [134]. Results are accurate to a few percent, which is comparable to the energy radiated in GWs, with the largest uncertainty coming from the disc mass estimates from simulations [e.g., 96, 135].

In this *Chapter*, an analytical model for the BH remnant of BHNS coalescences is

constructed directly from NR simulations data. In §1.1 the physics of the remnant formation is reviewed, motivating the choice of the parameters used in the model. In §1.2 the construction of the model for the remnant BH's mass and spin is described, discussing its behaviour in the parameter space; additionally, the model is extended for the case of precessing binaries. §1.3 contains applications of the model, demonstrating its importance in astrophysical predictions and observations. Ultimately, in §1.4 a model for the peak luminosity of the GW (2, 2) mode is presented.

Throughout this *Chapter*, geometric units $c = G = 1$ are employed and physical quantities are expressed in units of M_{\odot} , unless otherwise specified.

1.1 Physics of the remnant formation

For a physical understanding of the remnant's formation it is useful to resort to the concept of ISCO for the motion of test particles on equatorial trajectories around a Kerr BH. For a Kerr BH of mass M_{BH} , spin S_{BH} and dimensionless spin parameter $a_{\text{BH}} := S_{\text{BH}}/M_{\text{BH}}^2$, $r_{\text{ISCO}} = M_{\text{BH}}f(a_{\text{BH}})$, where [136]

$$\begin{aligned} f(a_{\text{BH}}) &:= 3 + Z_2 \mp \sqrt{(3 - Z_1)(3 + Z_1 + 2Z_2)}; \\ Z_1 &:= 1 + (1 - a_{\text{BH}}^2)^{1/3} [(1 + a_{\text{BH}})^{1/3} + (1 - a_{\text{BH}})^{1/3}]; \\ Z_2 &:= \sqrt{3a_{\text{BH}}^2 + Z_1^2}. \end{aligned} \quad (1.1)$$

$f(a_{\text{BH}}) \in [1, 9]$ is a monotonically decreasing function of a_{BH} , with range $a_{\text{BH}} \in [-1, 1]$. Here, positive(negative) values of a_{BH} correspond to BH spins that are aligned(antialigned) to the orbital angular momentum of the test particle. Particles that orbit the BH at a distance smaller than r_{ISCO} cannot have a stable circular motion and plunge into the BH.

The NS in a merging BHNS system undergoes tidal disruption if it reaches a characteristic distance r_{TD} from the BH larger than r_{ISCO} . r_{TD} is expected to scale in the same way as r_{MS} , which is the distance of the NS from the BH at which mass shedding occurs. r_{MS} is determined by the condition that the tidal forces exerted by the BH on its companion overcomes the NS self-gravity at the stellar surface. Hereafter, M_{BH} and a_{BH} denote the mass and dimensionless spin of the initial BH, respectively. Defining the binary mass ratio $q := M_{\text{BH}}/M_{\text{NS}}$, indicating with R_{NS} the NS radius and using keplerian relations, this condition implies $r_{\text{TD}} \lesssim r_{\text{MS}} \propto q^{1/3}R_{\text{NS}}$, with an approximately weak dependence on the BH spin [80]. Since $R_{\text{NS}}/M_{\text{BH}} = (qC_{\text{NS}})^{-1}$,

where $C_{\text{NS}} := M_{\text{NS}}/R_{\text{NS}}$ is the NS compactness, the parameter that regulates the onset of tidal disruption is the ratio $\xi = r_{\text{TD}}/r_{\text{ISCO}} \propto C_{\text{NS}}^{-1} q^{-2/3} f(a_{\text{BH}})^{-1}$. This argument shows that ultimately NS' tidal disruption depends on three physical parameters: q , C_{NS} , a_{BH} . Despite r_{ISCO} being formally defined for a test particle, this conclusion is expected to hold also for the NS, which has a finite size. Tidal disruption in this case is determined by the binary's tidal and spin-orbit interactions.

Simulations have shown that if the BH is non-spinning, or its spin is antialigned with the orbital angular momentum, tidal disruption occurs in BHNSs with $q \lesssim 3$. Generally, large, aligned BH spins $a_{\text{BH}} \gtrsim +0.5$ favour tidal disruption because spin-orbit interactions push r_{ISCO} to smaller values. As an example, $r_{\text{ISCO}} = 1 M_{\text{BH}}$ for a Kerr BH with $a_{\text{BH}} = +1$, as opposed to $r_{\text{ISCO}} = 6 M_{\text{BH}}$ for a non-spinning BH. Disruption is also favored by low values of C_{NS} , which are obtained for stiff Equations of state (EOSs) and imply large NS tidal deformabilities¹ [137, 138]. For a fixed M_{NS} , larger deformabilities imply larger R_{NS} , while smaller q correspond to smaller M_{BH} .

Tidal disruption leads to the formation of an accretion disc in the merger remnant. Simulations predict remnant discs with baryon (rest) masses as large as $M_b^{\text{disc}} \gtrsim 0.1 M_{\odot}$ [134, 135], thus creating the conditions to ignite a SGRB [139–141].

Kyutoku et al. [112] tentatively classifies the phenomenology of BHNS mergers into three classes, based on the ratio ξ . For Type-I and Type-III mergers, tidal disruption occurs far from and close to r_{ISCO} , respectively. In Type-II mergers instead, typically $q \gtrsim 3$, $a_{\text{BH}} \lesssim 0$, which implies $r_{\text{TD}} < r_{\text{ISCO}}$ and therefore the NS plunges onto the BH before tidal disruption occurs. These three classes differ by their GW spectra and M_b^{disc} . For instance, Type-II mergers have a GW spectrum very similar to BBH ones [see e.g. 142–146].

1.2 Remnant's mass and spin model

All the models in this work are constructed from the data of NR simulations of quasi-circular BHNS mergers reported in Kyutoku et al. [112], Kyutoku, Shibata, and Taniguchi [118], and Kyutoku et al. [120] and performed with the SACRA code [147]. A comprehensive collection of all data used in this work can be found in the appendix of [148]. These simulations adopt different NS matter EOSs and (anti)aligned BH spin values. The simulations used for this work are performed with non-spinning NSs and

¹See §1.2 for the definition of NS tidal deformability.




FIGURE 1.1: Relative differences between the data points and the fit. Left panel: X_{\bullet} , the maximum difference is of the order of 1%. Right panel: a_{\bullet} , the maximum difference is below 3%. For both cases the coefficient of determination of the fit is $R^2 \approx 0.92$.

therefore the NS spin is neglected in the models of this work; however, this is expected to be a good approximation for realistic systems [149, 150].

The total gravitational mass of the binary is denoted as $M := M_{\text{BH}} + M_{\text{NS}}$, and the mass and spin of the remnant BH as M_{\bullet} and S_{\bullet} , respectively. The re-scaled remnant BH mass and spin are defined as $X_{\bullet} := M_{\bullet}/M$ and $a_{\bullet} := S_{\bullet}/M_{\bullet}^2$, respectively. Motivated by the discussion of §1.1, the parameters chosen to map X_{\bullet} and a_{\bullet} are the symmetric mass ratio $\nu := q/(1+q)^2 \in [0, 1/4]$, a_{BH} and the dimensionless NS quadrupolar tidal polarizability parameter $\Lambda := 2k_2/(3C_{\text{NS}}^5)$ [137]. In the latter, k_2 is the gravito-electric quadrupolar Love number, a monotonically decreasing function of C_{NS} [138]. Λ encodes² the tidal interactions at leading order in the post-Newtonian dynamics. Typically, $\Lambda \approx 100 - 2500$ for NSs in BHNS systems, depending on M_{NS} and EOS. In the model, no assumptions on the EOS are made for the NS, which is characterized solely by Λ . The parameter space spanned by the data used in this work is $\Lambda \in [100, 2500]$, $q \in [2, 7]$ (i.e. $\nu \in [0.109, 0.222]$), and $a_{\text{BH}} \in [-0.5, 0.75]$.

X_{\bullet} , a_{\bullet} of BHNS mergers are mapped as

$$F : (\nu, a_{\text{BH}}, \Lambda) \rightarrow (X_{\bullet}, a_{\bullet}) = F_{\text{BBH}}(\nu, a_{\text{BH}}) \frac{1 + p_1(\nu, a_{\text{BH}})\Lambda + p_2(\nu, a_{\text{BH}})\Lambda^2}{(1 + [p_3(\nu, a_{\text{BH}})]^2\Lambda)^2}. \quad (1.2)$$

²For compact binaries, tidal interactions at leading order are described by $\tilde{\Lambda} := \frac{16}{13} \frac{(m_1+12m_2)m_1^4\Lambda_1}{(m_1+m_2)^2} + 1 \leftrightarrow 2$. In the case of BHNS systems with $m_1 = M_{\text{NS}}$, $m_2 = M_{\text{BH}}$, $\Lambda_2 = 0$ and $\tilde{\Lambda} = \frac{16}{13} \frac{(1+12q)\Lambda}{(1+q)^5}$

TABLE 1.1: Best fit parameters for the remnant's mass and spin and for the GW peak luminosity, with respective determination coefficient R^2 of the fits.

F	k	p_{k10}	p_{k11}	p_{k20}	p_{k21}	R^2
X_\bullet	1	-1.83×10^{-3}	2.39×10^{-3}	4.29×10^{-3}	9.8×10^{-3}	0.921
	2	2.34×10^{-7}	-8.28×10^{-7}	-1.64×10^{-6}	8.08×10^{-6}	
	3	-2.01×10^{-2}	1.32×10^{-1}	6.51×10^{-2}	-1.43×10^{-1}	
a_\bullet	1	-5.44×10^{-3}	7.91×10^{-3}	2.33×10^{-2}	2.48×10^{-2}	0.92
	2	-8.57×10^{-7}	-2.82×10^{-6}	6.61×10^{-6}	4.29×10^{-5}	
	3	-3.04×10^{-2}	2.55×10^{-1}	1.48×10^{-1}	-4.28×10^{-1}	
L_{peak}	1	3.08×10^{-2}	-4.18×10^{-2}	-5.17×10^{-2}	3.19×10^{-1}	0.98
	2	-1.23×10^{-5}	8.84×10^{-6}	1.05×10^{-4}	-3.88×10^{-5}	
	3	3.30×10^{-1}	-3.76×10^{-2}	-9.2×10^{-1}	1.44	

where F_{BBH} are the models for the mass and spin of the remnant BH for BBH mergers developed in [151] and $p_k(\nu, a_{\text{BH}})$ are polynomials of the form

$$p_k(\nu, a_{\text{BH}}) = p_{k1}(a_{\text{BH}}) \nu + p_{k2}(a_{\text{BH}}) \nu^2 \quad (1.3)$$

$$p_{kj}(a_{\text{BH}}) = p_{kj0} a_{\text{BH}} + p_{kj1} . \quad (1.4)$$

Eq. (1.2) includes by construction the BBH limit for $\Lambda \rightarrow 0$ (no tidal effects) and the test-mass limit for $\nu \rightarrow 0$. Note that the dependence on a_{BH} is linear. The coefficient $p_3(\nu, a_{\text{BH}})$ is squared in order to avoid poles in the denominator.

The values of the best-fit parameters p_{kjl} are reported in Tab. 1.1. Relative differences with respect to the fit are shown in Fig. 1.1. The fits have determination coefficient $R^2 \approx 0.92$ and the residuals are normally distributed with mean ~ 0 and standard deviations $\sim 0.25 \times 10^{-2}$ and ~ 0.01 , respectively. The maximum relative differences are below 1% for the remnant mass fit and below 3% for the the remnant spin, which are smaller than the differences due to the EOS [118]. The NR data do not extend to $\nu \lesssim 0.1$ and $a_{\text{BH}} \lesssim -0.5 \vee a_{\text{BH}} \gtrsim 0.75$, thus the model effectively extrapolates into those regions. The extrapolation leads in some cases to unphysical values of $X_\bullet > 1$ and $a_\bullet < -1$. This behaviour is fixed by forcing the model to agree with the BBH case, i.e., imposing $X_\bullet/F_{\text{BBH}} = 1$ and $|a_\bullet| \leq 1$. The model has been recently updated in Gonzalez et al. [72] with 19 additional NR simulations from [152], which correspond to mergers of non-spinning BHNS systems with $q \lesssim 3$. Moreover, the fitting formula for X_\bullet has been slightly modified to improve the behaviour of the model in the region of the parameter space of large Λ and small ν , where no NR data




FIGURE 1.2: Contour plots of the mass-rescaled remnant BH mass $X_{\bullet} := M_{\bullet}/M$ (top row) and of the dimensionless spin parameter $a_{\bullet} := S_{\bullet}/M_{\bullet}^2$ (bottom row) as a function of the symmetric mass ratio ν and of the NS tidal polarizability parameter Λ . Each column corresponds to fixed values of the initial BH spin parameter a_{BH} reported on top, corresponding to those of the NR simulations employed. White markers indicate the NR data used to construct the model.

are available. Details about the new fitting formulas are in Gonzalez et al. [72].

In the simulations, X_{\bullet} and a_{\bullet} are measured from the apparent horizon. The former should also correspond to

$$X_{\bullet} = 1 - \frac{E_{\text{GW}}}{M} - \frac{M_{\text{g}}^{\text{disc}}}{M} - \frac{M_{\text{g}}^{\text{ejecta}}}{M}, \quad (1.5)$$

where E_{GW} is the total energy radiated in GWs during the coalescence and $M_{\text{g}}^{\text{disc}}$ and $M_{\text{g}}^{\text{ejecta}} < M_{\text{g}}^{\text{disc}}$ are respectively the disc and ejecta mass contributions to the gravitational energy, which cannot be directly measured in the simulations [120]. Eq. (1.5) is key to understand the physics of the behaviour of X_{\bullet} . In BBH mergers finite mass-ratio effects are repulsive, implying that the GW emission is more efficient for larger ν and the same effect is seen in the BHNS dynamics.

The model of Eq. 1.2 is represented in Fig. 1.2, which clearly shows that the smallest values of X_{\bullet} are obtained for larger values $\nu \rightarrow 1/4$. The precise behaviour of X_{\bullet} , however, depends on the competition between the energy emitted in GWs and the effect of tidal disruption, as per Eq. (1.5). The second column of Fig. 1.2 shows that

for non-spinning BHNS binaries, the value of X_\bullet slightly increases with respect to the BBH case as $\Lambda > 0^+$ and for a given ν . Tidal disruption does not occur for small values of Λ (e.g. $\Lambda \lesssim 500$ for $\nu = 0.22$ and $a_{\text{BH}} = 0$), so this effect is solely due to the fact that tidal interactions are attractive and reduce the GW emission with respect to the $\Lambda = 0$ case (i.e., E_{GW} decreases so X_\bullet grows, with $M_b^{\text{disc}} \approx 0$). As Λ becomes sufficiently large (and $\nu \rightarrow 1/4$), tidal disruption occurs and only part of the remnant mass contributes to the final BH mass. Consequently, as Λ increases beyond a certain critical value, X_\bullet starts to decrease because part of the NS mass is accreted in the disc instead of being swallowed by the BH. Note that the peak mass is more pronounced for $\nu \rightarrow 1/4$ and disappears for sufficiently small ν , for which Type-II mergers occur, i.e. the whole NS plunges into the BH. Focusing on spin effects, for a given ν , the remnant mass decreases for increasing $a_{\text{BH}} > 0$ because the ratio ξ increases. This is a consequence of the repulsive character of the spin-orbit interaction for aligned positive spins. Notably, the peak at small Λ is no longer present for sufficiently large values of a_{BH} . On the other hand, for $a_{\text{BH}} < 0$, the spin-orbit interactions are attractive and have the same sign as the tidal interactions. As a consequence, for smaller a_{BH} , X_\bullet increases and the peak at small Λ is more pronounced.

For non-spinning BBHs, S_\bullet is expected to decrease for increasing ν , due to the same finite mass-ratio effect described above. Due to the M_\bullet^2 normalization, however, a_\bullet shows the opposite behaviour. In the BHNS case, the remnant BH has a larger dimensionless mass-rescaled spin with respect to the BBH case and it increases with Λ , for small $\Lambda > 0$. This happens because the NS is less compact and therefore less angular momentum is dissipated via GWs. Above a peak value, however, tidal disruption occurs and the angular momentum redistributes into the disc that forms around the remnant BH. For $|a_{\text{BH}}| \lesssim 0.5$ and a given value of ν , the final a_\bullet is roughly linear in a_{BH} (see Eq. (1.2) and [112]).

Extension to precessing BHNS systems Although the models of this work are developed from non-precessing BHNS data, they can be extended to the case of generic BH spins [132, 153, 154]. For this purpose, the initial spin is mapped as follows

$$a_{\text{BH}} \rightarrow a_{\text{BH}}^z = a_{\text{BH}} \cos \beta , \quad (1.6)$$

where β is the angle between a_{BH} and the initial total orbital angular momentum \mathbf{L} of the system. In this case the model will yield a_\bullet^z instead of a_\bullet . This prescription assumes that the direction of the total angular momentum $\mathbf{J} := \mathbf{L} + \mathbf{S}$ is approximately




FIGURE 1.3: Comparison of the model with the prescription of Eq. (1.6) against NR data of binaries with spin precession of [122]. β is the angle between the initial orbital angular momentum and the initial BH spin. All available data correspond to fixed values of $(a_{\text{BH}}, q) = (0.75, 5)$ and the errorbars are conservatively taken as 1% of the values.

preserved during the merger process, so that the direction θ of the final spin is given by the projection $\cos \theta = \hat{\mathbf{J}} \cdot \hat{\mathbf{L}}$. This prescription is validated against simulations of binaries with precessing spins. Fig. 1.3 shows that the model predictions for X_{\bullet} and α_{\bullet} using Eq. (1.6) are in good agreement with the NR data of Kawaguchi et al. [122], which are not used to determine the model.

1.3 Applications of the model

1.3.1 Population of BHNS remnants

As a first application, the model Eq. (1.5) is used to predict the distribution of BH remnant's masses and spins for given synthetic populations of BHNS systems merging at redshift $z \leq 1$. This study can give insights on how frequently tidal disruption events and associated EM counterparts are expected.

The latter are constructed by convolving the binary population-synthesis from the MOBSE code [155–157] with the Illustris cosmological simulation [158–160], see [161–163]. In particular, the run CC15 α 5 of [163] is adopted, where the common-envelope parameter is chosen as $\alpha = 5$ and natal kicks are drawn from a Maxwellian

distribution with a single root-mean square velocity $v_\sigma = 15 \text{ km s}^{-1}$ for both electron-capture and core-collapse supernovae. The minimum(maximum) mass of a BH(NS) is set to $5(2) M_\odot$. This assumption enforces the existence of a mass gap between BHs and NSs, which is suggested by dynamical mass measurements of compact objects in X-ray binaries [164, 165]. BH spins are added by randomly drawing spin magnitudes $|a_{\text{BH}}| \in [0, 1]$ from a truncated Maxwellian distribution with root mean square σ . Here, spins are considered isotropically oriented with respect to the binary orbital plane with $\langle a_{\text{BH}} \rangle = 0.2$ as fiducial distribution or aligned spin distributions with $\sigma = \{0.1, 0.35, 0.5, 0.7\}$, corresponding to average values $\langle a_{\text{BH}} \rangle = \{0.2, 0.5, 0.75, 0.95\}$, respectively.

The population synthesis predicts BH component masses below $10M_\odot$, distributed narrowly about $M_{\text{BH}} \approx 5M_\odot$ and $M_{\text{BH}} \approx 8M_\odot$ [162]. The metallicity Z of a star is defined as the mass fraction of the composing elements heavier than helium. The population depends very weakly on progenitors' metallicities for $Z \leq 0.002$, but for $Z \geq 0.003$ the smallest BHs are suppressed and only BH with $M_{\text{BH}} \approx 8M_\odot$ are obtained. This is a consequence of the dependence of the delay time on the progenitor's metallicity. The delay time is defined as the time elapsed from the formation of the progenitor stars to the BHNS merger. Metal-rich progenitors have longer delay times than metal-poor ones and thus they do not merge within the Hubble time, especially for small BH masses. Additionally, NS masses $M_{\text{NS}} \gtrsim 1.3M_\odot$ are favoured.

From the distribution of M_{BH} and a_{BH} of the synthetic population, the distributions of X_\bullet and a_\bullet are computed as follows. A representative set of EOSs is chosen, for each of which the distribution of Λ on the NS population is calculated. The remnant properties are then determined according to Eq. (1.2) with the prescription of Eq. (1.6). The top row of Fig. 1.4 shows that the distributions of X_\bullet and a_\bullet depend weakly on the EOS employed. The metallicity of the progenitor stars instead affects the final result, as can be seen in the bottom row. For metallicities $Z \leq 0.002$ bimodal distributions peaked around $M_\bullet \approx 7M_\odot$ and $M_\bullet \approx 9M_\odot$ are found, independently from the EOS employed. BHNS populations generated with large metallicities of the progenitor stars only produce the more massive remnants. The remnant spins inferred from Eq. (1.2) and the isotropic/aligned spin population with $\langle a_{\text{BH}} \rangle \approx 0.2$ are distributed around $a_\bullet^z \approx 0.4$ with standard deviation ~ 0.1 . Overall, the distributions obtained track the populations generated by the population synthesis code. The values $\gtrsim 7 M_\odot$ obtained for the final masses regardless the EOS employed point to the plunge scenario as the most likely, and therefore from these populations it is unlikely that large ejecta masses




FIGURE 1.4: Top row: X_{\bullet} and a_{\bullet} distributions for different EOSs. Bottom row: X_{\bullet} and a_{\bullet} distributions for different values of the metallicity Z of the progenitor stars. The latter plot is relative to the SLy EOS; the fiducial isotropic spin distribution for the synthetic BHNS population for the plots in this figure is peaked around $\langle a_{\text{BH}} \rangle = 0.2$.

and bright EM counterparts will be emitted.

Using the model of [135] it is possible to estimate the distribution of M_b^{disc} associated to the remnant BH. Fig. 1.5 shows that aligned, low-spin distributions result in $\gtrsim 99\%$ of the remnants having M_b^{disc} smaller than $M_{\text{SGRB}}^b = 0.075 \frac{M_{\text{NS}}^b}{1.5} M_{\odot}$, independently from the adopted EOS. Disc masses above M_{SGRB}^b are necessary to produce SGRBs of 1 s duration [166, 167]. Remnants with significant disc masses are found for aligned spin distributions with $\langle a_{\text{BH}} \rangle \gtrsim 0.5$. In these cases, the largest discs are found for the stiffest EOSs, corresponding to $\Lambda \gtrsim 1700$. Softer EOSs, corresponding to $\Lambda \lesssim 400$, give massive discs only for $\lesssim 20\%$ of the binaries which have $\langle a_{\text{BH}} \rangle \gtrsim 0.75$. In this context, the results obtained from aligned spin distributions can be considered *upper limits* to the ones from isotropic spin distributions.

1.3.2 BH remnants of GW200105 and GW200115

The models discussed in this *Chapter* have been used to estimate the final mass and spin of the remnant BH for the systems that produced the two signals named GW200105 and GW200115 [45]. In particular, under the hypothesis of BHNS coalescence for the two events, posterior samples for the binary parameters have been generated with the two models IMRPhenomNSBH [168] and




FIGURE 1.5: Remnant disc baryonic mass distribution for different EOSs and for low (left) and high (right) aligned BH spin distributions. The mass threshold represents the minimum mass of the disc that allows the production of SGRBs with 1 s duration. The percentage of binaries with disc mass larger than this threshold is provided in the legend for each EOS employed.

SEOBNRv4_ROM_NRTidalv2_NSBH [130]. The latter crucially rely on the a_{\bullet} and X_{\bullet} fits presented here to calculate the GW amplitude in the merger and ringdown phase. Applying Eqs. (1.2) on the BHNS posterior samples, it was possible to compute the distributions of a_{\bullet} and M_{\bullet} for the two events. For GW200105, the remnant mass and spin have been estimated as $M_{\bullet} = 10.4_{-2.0}^{+2.7}$ and $a_{\bullet} = 0.43_{-0.03}^{+0.04}$, while for GW200115, $M_{\bullet} = 7.8_{-1.6}^{+1.4}$ and $a_{\bullet} = 0.38_{-0.02}^{+0.04}$.

1.3.3 Extension of TEOBResumS for BHNS

TEOBResumS is a GW model based on the EOB formalism [67–69, 169–171] capable of producing accurate waveforms for the coalescence of CBs including spin, higher modes and tidal effects [146, 172–176]. In Gonzalez et al. [72], TEOBResumS is augmented with an inspiral-merger-ringdown GW model for BHNS mergers constructed by computing the deviations from the BBH GW model with NR fitting formulas. Notably, TEOBResumS is the first GW model for BHNS coalescences including higher modes and precession effects. Updated version of the remnant models presented in this *Chapter* are used in Gonzalez et al. [72] to discriminate among different NS tidal disruption cases and to construct the ringdown sector.




FIGURE 1.6: Distribution of the GW peak luminosity as a function of the NS Λ for all the values of a_{BH} and for all the different values of ν in the available dataset.

1.4 GW peak luminosity

The approach used in this work can be used to estimate the GW peak luminosity, thus complementing the results derived for BBHs and BNS systems in [88, 177] (see [178] for an application of those results). The GW peak luminosity is computed from the (2, 2) mode of the GW strain,

$$L_{\text{peak}} \approx \max_t \frac{1}{16\pi} \left(\left| \frac{dh_{22}(t)}{dt} \right|^2 \right) \quad (1.7)$$

where $h_{\ell m}$ is defined through Eq. B.10.

The expression used to fit L_{peak} is slightly different from Eq. (1.2) and reads:

$$L_{\text{peak}}(\nu, a_{\text{BH}}, \Lambda) = L_{\text{peak}}^{\text{BBH}}(\nu, a_{\text{BH}}) \frac{(1 + p_1(\nu, a_{\text{BH}})\Lambda + p_2(\nu, a_{\text{BH}})\Lambda^2)^2}{(1 + [p_3(\nu, a_{\text{BH}})]^2\Lambda)^4}. \quad (1.8)$$

where the numerator is squared in order to enforce positive values for the luminosity. The denominator instead is raised to the 4th power in order to get a finite value for the asymptotic limit of $\Lambda \rightarrow \infty$. The best fit parameters p_{kjl} are reported in Tab. 1.1. The model delivers results with accuracy at the 20% level, but note that, since only the (2, 2) mode data are used, these results represent a *lower limit* for the GW peak luminosity from BHNS mergers.

Fig. 1.6 shows the peak luminosity model as a function of Λ and for the values of ν and a_{BH} sampled by the NR dataset. The behaviour of L_{peak} is analogous to the one X_{\bullet} detailed in §1.2.

Chapter 2

Impact of resolution and input physics on BNS merger simulations

NR simulations represent a fundamental approach for the prediction of astrophysical observables from the mergers of BNS systems and their aftermath [83, 179], allowing one to identify the different mechanisms for mass ejection and the kinematical and thermodynamical properties of the unbound material.

The inclusion of weak interactions mediated by neutrinos is key for the physical understanding of such processes. Neutrinos are produced in the high-temperature regions forming due to the collision of the two NSs and, on longer timescales, from the hottest regions of the NS remnant and accretion torus [82, 115, 140, 180–185]; neutrino emission is the dominant process for NS matter cooling and, together with neutrino-matter interactions, it determines the outcome of the r -process nucleosynthesis and the colour of the kilonova [186–188], as well as the geometry and the mass of the dynamical ejecta [39, 189, 190]. Neutrino-matter interactions may also impact the high-density regions of the remnant through out-of-equilibrium effects [82]; however, while some NR works report damping of the density oscillations in the remnant NS [191] and possible signature in the post-merger GW signal [192, 193], others employing more sophisticated neutrino transport schemes and higher grid resolutions do not find significant out-of-thermodynamic equilibrium effects on the post-merger dynamics or GW emission [194]. Multi-resolution studies employing a consistent neutrino transport scheme appear necessary to assess the impact of out-of-equilibrium effects.

MHD instabilities and turbulence can affect the matter flow after merger [195–199] and provide crucial processes for the SGRB jet-launching mechanism [200–203]. The possible development of global large-scale magnetic stresses can enhance mass ejecta [42, 103, 198, 199, 204]. One of the main open issues is to adequately resolve the amplification of magnetic fields with realistic strengths and self-consistent turbulent flow

[197]. Sub-grid models have been proposed as a lightweight way to include magnetic field effects in the simulations [205–208].

This *Chapter* reports the first systematic study of the impact of different neutrino transport schemes, turbulent viscosity and of the role of finite grid resolution on the outcome of a BNS simulation and associated predictions of astrophysical observables.

The rest of this *Chapter* is organised as follows. In section §2.1 the NR code used and the different input physics and the simulation setups considered in this work are described. In section §2.2 the dynamics and thermodynamical evolution of the system and the properties of the remnant object are discussed. Section §2.3 is devoted to the study of the effects on the GW emission. Section §2.4 studies the dynamical ejecta mass and its composition. In section §2.5 the r -process nucleosynthesis yields and kilonova light curves associated to the ejecta computed in post-processing from the simulations are compared. Section §2.6 examines the variations in neutrino luminosities and average energies between the two neutrino schemes including neutrino reabsorption.

Throughout the text masses are expressed in units of solar masses M_{\odot} , while temperature and energy are reported in MeV. The other quantities are expressed in SI or cgs units.

2.1 Methods

Matter model, initial data and evolution methods In the simulations of this work, the NS matter is modelled using the SLy4-SOR EOS (hereafter SLy), a finite-temperature, composition-dependent EOS based on a Skyrme potential for the nucleonic interaction [209, 210]. This EOS includes baryons (both free and bound in nuclei), electrons, positrons and photons as the relevant degrees of freedom. The SLy EOS predicts a maximum Tolman-Oppenheimer-Volkoff (TOV) gravitational mass of $M_{\text{max}}^{\text{TOV}} \approx 2.05 M_{\odot}$ and a radius for a $1.4 M_{\odot}$ NS of $R_{1.4} \approx 11.9$ km. Both these values are compatible with the observations of extremely massive millisecond pulsars [211, 212], with results obtained by the NICER collaboration [213, 214], and with LIGO-Virgo detections [215]; see also Breschi et al. [44] for a multi-messenger analysis based on NR data.

Irrotational initial data in quasi-circular orbit are produced with the pseudo-spectral multi-domain code *Lorene* [216]. Initial data are constructed with the minimum temperature slice $T = 0.01$ MeV of the SLy EOS and neutrino-less β -equilibrium

inside the two component NSs is assumed at the beginning of the evolution.

The system is evolved using the 3+1 Z4c free evolution scheme described in §A.2, coupled with the general relativistic hydrodynamics (GRHD) equations (see §A.3). The simulations are performed with the `WhiskyTHC` code [217–221], which is built on top of the `Cactus` framework [222, 223]. In particular, the spacetime is evolved with the `CTGamma` code [224] which is part of the `Einstein Toolkit` [225, 226]. The time evolution is performed with the method of lines, using 4th-order finite-differencing spatial derivatives for the metric and the strongly-stability preserving 3rd-order Runge-Kutta (RK) scheme [227] as the time integrator. The time step is set according to the Courant-Friedrich-Lewy (CFL) criterion and the CFL factor is set to $\alpha_{\text{CFL}} = 0.15$. Berger-Oliger conservative adaptive mesh refinement (AMR) [228] with sub-cycling in time and refluxing is employed [229, 230], as provided by the `Carpet` module of the `Einstein Toolkit` [231]. `WhiskyTHC` implements a number of microphysics schemes to include neutrino effects and turbulent viscosity in the fluid in the simulations. In this work the effects of these schemes on BNS simulations are systematically compared. The neutrino treatments considered in this work are the leakage (LK), M0 and M1 schemes, while the general relativistic large-Eddy-simulations (GRLES) scheme is adopted for turbulent viscosity. All these schemes are described in some detail in App. D.

The simulation domain consists of a square cuboid $\Omega = [-1512, 1512] \times [-1512, 1512] \times [0, 1512]$ km³. Reflection symmetry about the xy -plane is used for $z < 0$. The grid setup consists of 7 refinement levels centred on the two NSs or in the merger remnant, with the finest level covering entirely each star. The simulations in this work are performed at low resolution (LR), standard resolution (SR) and high resolution (HR), for which the minimum spacings in the finest refinement level are $\delta x_{\text{LR}} \approx 247$ m, $\delta x_{\text{SR}} \approx 185$ m, $\delta x_{\text{HR}} \approx 123$ m.

Simulation sample In this work, the NS component masses of the BNS system and the EOS are chosen in such a way that the merger results in a remnant NS close to BH collapse; the idea is to see how much employing different neutrino treatments, turbulent viscosity and finite grid resolution can affect the outcome of such border-line case system. The two NS of the system have gravitational mass $M_1 = M_2 = 1.30 M_\odot$, corresponding to baryonic masses $M_{1b} = M_{2b} = 1.42 M_\odot$ and symmetric mass ratio $\nu := M_1 M_2 / (M_1 + M_2)^2 = 0.25$. The initial separation is set to ~ 45 km. This setup results in a BNS system with initial total gravitational mass, Arnowitt-Deser-Misner (ADM) mass and ADM angular momentum respectively $M := M_1 + M_2 \approx 2.60 M_\odot$, $M_{\text{ADM}} \approx 2.57 M_\odot$, $J_{\text{ADM}} \approx 6.82 M_\odot^2$.

This study is based on a total of 15 evolutions of the same initial data. A first set of simulations are performed solving only spacetime and hydrodynamics equations, and are labeled hereafter as HY. Preliminary results about these simulations are presented in Appendix B of Breschi et al. [70]. The effect of neutrinos is included using the three different schemes implemented in `WhiskyTHC` and described in App. D: three independent sets of simulations are performed employing the LK scheme, the LK scheme coupled with the M0 scheme, and the more advanced M1 scheme. The three different types of simulations are labelled as LK, M0 and M1, respectively. Finally, an additional set of simulations includes both M0 and viscosity and is referred to as VM0. Throughout the text the expression *physics scheme* is used to refer to the different realizations employed in the runs. Each of these 5 realizations is run at different resolutions, namely LR, SR, and HR.

To ensure stable runs with the M1 scheme, the following choices are made. Firstly, the relative tolerance parameter that is used to solve the implicit time step in the source term (see [194]) is set to 10^{-10} . Secondly, local thermodynamical equilibrium depending on the equilibration timescale in a specific cell is additionally enforced. In particular, if for a given cell the corresponding time step contains more than X e-foldings of the equilibration time, the neutrinos average energies are forced to be at equilibrium for the evolution of the neutrinos number densities in the selected cell. This prevents failures of the runs and the development of spurious features in regions of high density and low electron fraction Y_e ¹ in the first few ms after collision. The parameter X has been set to 20 for LR and HR run and to 10 for the SR run.

In the following, a particular simulation is named by indicating first the scheme used and then the resolution; for instance M0-SR is the simulation with M0 scheme run at standard resolution. A complete list of all the simulations is reported in the first column of Tab. 2.1. The simulations are performed for a minimum of 35 ms (LK-HR) up to a maximum of 155 ms (LK-LR). Some runs in which a BH forms are affected at later time by the numerical instability described in Radice et al. [194] and thus were not continued. Simulation data are analyzed to a safe evolution time reported in the second column of Tab. 2.1; no spurious effects are observed before this time.

¹See beginning of App. D for the definition of the electron fraction Y_e .




FIGURE 2.1: Time evolution of the main quantities describing the dynamics of the system. Top row: maximum rest-mass density ρ_{\max} expressed in units of the nuclear saturation density $\rho_{\text{nuc}} \approx 2.3 \times 10^{14} \text{ g cm}^{-3}$. Bottom row: (minus) reduced binding energy $-e_b$ (solid lines) and reduced angular momentum of the system j_{rem} (dashed lines). Time is shifted by the time of merger.

2.2 Remnant dynamics

The two NSs revolve around the common center of mass for about 6 orbits before colliding within ~ 14 ms from the beginning of the simulation. The moment of merger is conventionally defined as the peak amplitude of the (2,2) GW mode², and is denoted as t_{merg} hereafter. The evolution before this moment is referred as the *inspiral-merger* phase, while the evolution after t_{merg} is called *post-merger* phase. After merger a remnant NS forms, which survives for at least a few tens of ms. In six of the simulations the remnant NS collapses to a BH at $t - t_{\text{merg}} \gtrsim 18.3$ ms.

2.2.1 Remnant evolution

The overall remnant evolution is well described in terms of the maximum rest-mass density, ρ_{\max} , (minus) the reduced binding energy, $-e_b$, and the reduced angular momentum, j_{rem} , of the system. The latter two quantities are defined as

$$e_b := \frac{M_{\text{ADM}} - E_{\text{GW}} - M}{M\nu}, \quad j_{\text{rem}} := \frac{J_{\text{ADM}} - J_{\text{GW}}}{M^2\nu}, \quad (2.1)$$

²The GW amplitude is defined in Eq. (B.11).

where E_{GW} , J_{GW} are the radiated energy and angular momentum calculated from the multipolar GWs [232, 233]. The evolution of these quantities is reported in Fig. 2.1, comparing the different schemes in each panel and resolution effects across the three columns.

For $t < t_{\text{merg}}$ the evolution is qualitatively and quantitatively very similar for all the runs. As it can be clearly seen at negative times in Fig. 2.1, both ρ_{max} (top row), $-e_b$ and j_{rem} (bottom row) curves do not display any significant differences across the runs. This is expected, since neutrino production and viscosity effects are negligible in the two NSs before collision. In this regime, increasing the resolution has the only effect of accelerating the merger process, [see, e.g., 233, 234]. However, this effect is not visible in Fig. 2.1, because all curves are shifted by t_{merg} .

For $t > t_{\text{merg}}$, ρ_{max} rapidly increases as the NS cores merge reaching $\gtrsim 6 \rho_{\text{nuc}}$ within 10 ms; the damped oscillations in ρ_{max} are caused by the bounces of the two cores in the process. At about 10 ms post-merger, the outcome of the GW-dominated (early) post-merger phase is a remnant NS, formed by a core that is slowly rotating surrounded by a rapidly rotating envelope. The absolute value of the binding energy after t_{merg} measures the compactness of the remnant NS and it increases in time due to the emission of gravitational energy. The panels in the bottom row of Fig. 2.1 shows that most of E_{GW} and J_{GW} are radiated within $t - t_{\text{merg}} \approx 10$ ms [88, 235]. Comparing to the top row, this period coincides with the time in which the large oscillations of ρ_{max} are strongly dampened and the remnant NS stabilises or collapses. The physical explanation is that the remnant NS has a large and rapidly evolving quadrupole moment and is therefore an efficient emitter of gravitational radiation. The emission increases the remnant's compactness and reduces its angular momentum, thus driving the remnant NS towards axisymmetry and eventually stationarity. Overall, the gravitational energy and angular momentum emission show qualitatively a similar evolution for all the runs. In all cases about the same values of $-e_b \approx 0.12$ and $j_{\text{rem}} \approx 2.9$ are reached at $t - t_{\text{merg}} \approx 5$ ms, and only after this time some differences start developing among the runs.

During the GW-dominated phase, turbulent viscosity has a dominant impact on the remnant's core dynamics with respect to the inclusion of neutrinos. In particular, ρ_{max} and $-e_b$ in VM0-LR and VM0-HR runs are comparably smaller with respect to the other runs at the same resolution, especially at later times. This effect is explained by viscosity enabling angular momentum transport from the rapidly rotating envelope to the slowly rotating remnant core, which gains more rotational support. This causes

TABLE 2.1: Main properties of the remnant disc and ejecta for all simulations. The end time of the simulation t_{end} and of BH collapse t_{BH} are measured with respect to merger. The baryonic masses of disc and ejecta, respectively M_b^{disc} and M_{ejecta} , are expressed in solar masses. M_b^{disc} is computed at the latest available time before BH collapse occurs, or at t_{end} if a BH does not form. The ejecta quantities are extracted with the Bernoulli criterion on a spherical surface at 443 km. The electron fraction Y_e and the specific entropy s are reported as mass-weighted averages. The emission angle is calculated as the mass-weighted root mean square (RMS) of the emission latitudes. The ejecta analysis is performed until $t - t_{\text{merg}} = 20.3$ ms, corresponding to the earliest t_{end} of the entire set of simulations, i.e. to the run LK-HR.

Simulation	t_{end} [ms]	t_{BH} [ms]	M_b^{disc} [M_{\odot}]	M_{ej} [M_{\odot}]	$M_{\text{ej}}^{v > 0.6c}$ [M_{\odot}]	$\langle Y_e \rangle$	$\theta_{\text{ej}}^{\text{RMS}}$ [$^{\circ}$]	$v_{\text{ej}, \infty}/c$	$\langle s \rangle$ [k_B/bar]
HY-LR	109	-	1.85×10^{-1}	1.10×10^{-2}	1.11×10^{-5}	0.05	34	0.16	16
LK-LR	140	-	1.76×10^{-1}	2.41×10^{-3}	8.60×10^{-6}	0.13	28	0.18	13
M0-LR	94	-	1.57×10^{-1}	6.70×10^{-3}	1.34×10^{-5}	0.23	34	0.16	17
VM0-LR	104	-	1.80×10^{-1}	6.44×10^{-3}	1.48×10^{-5}	0.23	34	0.15	17
M1-LR	35.8	-	2.42×10^{-1}	6.59×10^{-3}	2.02×10^{-6}	0.24	36	0.17	16
HY-SR	109	-	1.64×10^{-1}	8.43×10^{-3}	2.73×10^{-5}	0.049	33	0.19	17
LK-SR	114	-	8.14×10^{-2}	2.35×10^{-3}	1.23×10^{-5}	0.16	30	0.21	14
M0-SR	64.3	64	7.55×10^{-2}	5.85×10^{-3}	3.92×10^{-5}	0.22	32	0.18	16
VM0-SR	35.8	21	7.58×10^{-2}	4.02×10^{-3}	3.09×10^{-5}	0.23	33	0.19	18
M1-SR	41.8	-	1.51×10^{-1}	4.13×10^{-3}	1.29×10^{-5}	0.24	37	0.19	18
HY-HR	27.2	25.6	1.10×10^{-1}	7.20×10^{-3}	2.44×10^{-5}	0.044	34	0.19	18
LK-HR	20.3	19.9	6.77×10^{-2}	1.92×10^{-3}	1.47×10^{-6}	0.17	29	0.2	16
M0-HR	28.6	20.2	8.98×10^{-2}	5.11×10^{-3}	7.96×10^{-6}	0.26	34	0.16	18
VM0-HR	61.3	60.9	9.46×10^{-2}	6.14×10^{-3}	2.80×10^{-5}	0.24	34	0.16	18
M1-HR	28.4	18.3	8.87×10^{-2}	4.82×10^{-3}	4.43×10^{-6}	0.29	35	0.21	18

a decrease of the central density of the remnant star, which makes it more stable [205, 236].

The grid resolution has a significant impact on the fate of the remnant. LR simulations present the smallest GW emission, which leads to a less compact and more rotationally supported remnant NS core. At LR, gravitational collapse is never observed within the simulated time. At higher resolution, overall binding energies are larger and remnant angular momenta are smaller for all the runs comparing one by one to the LR simulations. For the M0-SR case BH collapse happens at ~ 64 ms post-merger (see third column of Tab. 2.1), while BH formation is not observed for HY-SR, LK-SR and M1-SR runs within the end of the simulations. The HR runs show the largest binding energy magnitudes, i.e. the largest compactnesses for the remnant NS. As a consequence, BH collapse occurs as early as $t - t_{\text{merg}} \approx 18.3$ ms for the M1 simulation, $t - t_{\text{merg}} \approx 20$ ms for LK and M0 runs and $t - t_{\text{merg}} \approx 26$ ms for the HY case. In VM0-HR the BH collapse is delayed by about 40 ms with respect to the other HR runs, due to the effect of viscosity, as described above. The runs that employ the




FIGURE 2.2: Correlation between the binding energy of the system and the maximum rest-mass density. The latter is rescaled by the central density of the maximum-mass TOV star predicted by the SLy EOS.

M1 transport scheme show a monotonic increase of ρ_{\max} and $-e_b$ and a monotonic decrease of j_{rem} with resolution.

The run VM0-SR presents an unexpected behaviour. The remnant NS collapses quite early, around 20 ms post-merger. Comparing to VM0-LR and VM0-HR runs, the density oscillations at 8 – 10 ms appear less dampened and ρ_{\max} keeps increasing until the NS eventually collapses. This behaviour has never been observed in previous works where viscosity was included in the same way [90, 208]. It can be speculated that this result is related to the specific simulation setup that, for this particular BNS, is not yet in a convergent regime at SR. Higher resolution simulations would be required to explore the possibility of obtaining consistent results.

The analysis of the merger dynamics in terms of ρ and energetics is weakly dependent on the particular setup of the simulations and thus it robustly captures the merger dynamics. This is summarised considering the gauge invariant $\rho_{\max}(-e_b)$ curves in Fig. 2.2. The plot shows that the two quantities are clearly correlated, which implies that ρ_{\max} can be, in principle, estimated from a measurement of the total GW radiated energy [237]. The robustness of the correlation showed in Fig. 2.2 indicates also that the simulations analysed here are internally self consistent with each other. The figure also highlights the fact that in all simulations in which BH collapse occurs, this happens for values of ρ_{\max} below the central density of the maximum-mass TOV star, in particular with $\rho_{\max} \gtrsim 70\% \rho_{\max}^{\text{TOV}}$ [93]. This result indicates that gravitational




FIGURE 2.3: Comparison of the first 20 ms post-merger evolution of the remnant NS for all LR runs. Each row represents a different simulation, while each column corresponds to snapshots taken at the time expressed on top, which refers to $t - t_{\text{merg}}$. The left and right half of each subplot show respectively the temperature profile in linear scale and the density profile in logarithmic scale, both on the equatorial plane. The black contour levels represent iso-density curves. Moving away from the centre, they correspond to decreasing densities of 10^{15} , 10^{14} , 10^{13} , 10^{12} ... g cm^{-3} . The thickest black line corresponding to $\rho = 10^{13} \text{ g cm}^{-3}$ conventionally denotes the interface between the remnant NS core and the disc.

collapse is mainly determined by the remnant core, which is slowly rotating and cold.

2.2.2 Thermodynamic evolution of the remnant

To discuss the impact of different neutrino schemes and viscosity on the thermodynamics of the remnant NS core, in Figs. 2.3 and 2.4 the rest-mass density and temperature profiles on the equatorial plane for LR and HR runs, respectively, are reported. The remnant NS is conventionally considered as the region enclosed within the iso-density surface $\rho = 10^{13} \text{ g cm}^{-3}$, indicated with thick black curves in the plots. Comparing the




FIGURE 2.4: Same as Fig. 2.3 but for HR runs. In the frame corresponding to M1-HR at $t - t_{\text{merg}} = 20$ ms the black circle represents the apparent horizon of the BH that forms at 18.3 ms post-merger.

profiles at LR and HR, the most evident effect is that remnants corresponding to the HR runs are more compact; this is in agreement with the binding energy analysis of the system in §2.2.1. The snapshots $t - t_{\text{merg}} = 0$ ms (first column) show the moment in which the two NSs touch and the cores start to fuse, causing the matter at the collisional interface to warm up because part of the kinetic energy is converted into thermal energy. At 5 and 10 ms post-merger (second and third column respectively) the hot matter produced at the collisional interface forms two hotspots at peak temperature $T \approx 70\text{--}80$ MeV that revolve around the colder core [82, 238, 239]. At later time $t - t_{\text{merg}} = 15$ ms, the hot matter is concentrated in an annulus with a more uniform temperature $T \approx 60\text{--}70$ MeV.

The structure of the remnant NS after the GW-phase is almost axisymmetric. The density profile decreases monotonically with the radial coordinates, while the temperature profile does not. In particular, the central densest region of $\rho \gtrsim 10^{15}$ g cm $^{-3}$ is characterised by $T \lesssim 20$ MeV. In the region of densities $\rho \in [10^{14}, 10^{15}]$ g cm $^{-3}$ the

temperature first increases up to $T \approx 60 - 70$ MeV and then it decreases down to $T \approx 20$ MeV. The layer of density $\rho \in [10^{13}, 10^{14}]$ g cm $^{-3}$ is colder, with temperatures $T \lesssim 20$ MeV.

With the exception of the M1 runs, which are discussed below, no significant differences are seen in the remnant density and temperature profiles comparing runs with different physics schemes at the same resolution, as expected. The inclusion of neutrino emission with the LK scheme does not impact significantly the thermodynamics of the remnant's core, where matter is at high density. Including neutrino reabsorption with the M0 scheme also does not affect the remnant appreciably, because the component of trapped neutrinos is neglected and because free-streaming neutrinos mostly interact with the lower-density material around the remnant NS. Turbulent viscosity is also not expected to have a strong impact on the thermodynamics of the remnant core, because the effects of the viscosity model implemented here are by construction small at densities higher than 10^{13} g cm $^{-3}$ [208]. In particular, no increase in the core temperature due to kinetic energy being converted into thermal energy due to viscosity is observed.

A comparison of the internal temperature of the remnant star between the runs with LK and M1 at 15 ms post-merger reveals an effect due to neutrino radiation in optically thick conditions. The hot annulus at densities $\rho \in [10^{14}, 10^{15}]$ g cm $^{-3}$ shows lower temperatures in the M1 run compared to the LK case, with $T_{\text{M1}}^{\text{peak}} \approx 88\% T_{\text{LK}}^{\text{peak}}$. This difference is a physical effect due to the emergence of a neutrino trapped gas that converts fluid thermal energy into radiation energy [82]. The left plot of Fig. 2.5 shows the consequent difference in the matter composition of the remnant's core in the region $\rho \in [10^{14}, 10^{15}]$ g cm $^{-3}$, corresponding to the hot annulus of matter. While in LK runs the remnant core retains its pristine Y_e with peaks of $Y_e \approx 0.058 - 0.059$, in M1 runs Y_e can be locally 40% larger than these values. These variations are consistent for both LR and HR resolutions and with Fig. 9 of Perego, Bernuzzi, and Radice [82]. The analysis of [82] was performed in postprocessing from simulations where the neutrino trapped component was not considered, and the conclusion was that the presence of a neutrino gas would cause a $\sim 33\%$ increase in Y_e . Here this effect is confirmed in simulations that do include the neutrino trapped component inside the remnant (see also [194]).

The right plot of Fig. 2.5 shows that the thermodynamical conditions inside the remnant are such that locally, in the high-temperature region, the neutrino fractions follow the hierarchy $Y_{\nu_e} < Y_{\nu_x} < Y_{\bar{\nu}_e}$ [82, 115, 194]. This is confirmed for all resolutions




FIGURE 2.5: Left plot: comparison of Y_e inside the remnant between LK and M1 in a 2D snapshot at $t - t_{\text{merg}} = 15$ ms on the equatorial plane. The black contour levels represent iso-density curves. Moving away from the centre, they correspond to decreasing densities of 10^{15} , 10^{14} , 10^{13} g cm $^{-3}$. The remnant of the runs with M1 shows an annulus of higher Y_e with respect to the LK runs at densities $\rho \in [10^{14}, 10^{15}]$ g cm $^{-3}$, coinciding with the hot annuli of matter in Figs. 2.3 and 2.4 Right plot: neutrino fraction in the remnant core of the M1 simulations at $t - t_{\text{merg}} = 15$ ms. Each column corresponds to the fraction of one of the three species simulated. Neutrino production inside the remnant core is favoured for the species $\bar{\nu}_e$ and disfavoured for ν_e . This finding is robust with resolution. In both plots, the thickest black line corresponding to $\rho = 10^{13}$ g cm $^{-3}$ conventionally denotes the interface between the remnant NS core and the disc.

and explained as follows. The matter constituting the hot annulus is characterised by densities $\rho \gtrsim 10^{14}$ g cm $^{-3}$ and temperatures of a few tens of MeV. This is matter initially in cold, neutrino-less weak equilibrium coming from the collisional interface of the fusing NS cores that both decompresses and heats up. Electrons in these conditions are highly degenerate and relativistic, and their chemical potential μ_e is weakly sensitive to density and temperature variations. On the other hand, neutrons and even more protons are non-degenerate, since their Fermi temperature T_F is such that $T \gtrsim T_F$ and $Y_p \sim 0.1Y_n$ due to the initial neutron richness. The chemical potentials of protons (μ_p) and of neutrons (μ_n) are negative, but the magnitude of the former increases faster than the magnitude of the latter. Then, the chemical potential of neutrinos at equilibrium, $\mu_{\nu_e, \text{eq}} = \mu_p - \mu_n + \mu_e$, becomes negative and in particular $-\mu_{\nu_e, \text{eq}} \approx 120$ MeV. For thermalised neutrinos in weak equilibrium, $\mu_{\bar{\nu}_e} = -\mu_{\nu_e}$ and $Y_\nu \propto TF_2(\mu_\nu/T)$, where $F_2(x)$ is the Fermi function of order 2, so that $Y_{\nu_e} < Y_{\bar{\nu}_e}$. Electron antineutrinos form a mildly degenerate Fermi gas, because the temperature




FIGURE 2.6: Time evolution of the disc mass M_b^{disc} comparing all runs. The vertical dashed lines mark the time of BH collapse. Time is shifted by the time of merger.

is high and the degeneracy parameter $\eta_{\bar{\nu}_e} = \mu_{\bar{\nu}_e}/T \approx 2.5 - 2.7$. Therefore, while the production of electron neutrinos is suppressed due to the higher neutron degeneracy, the production of electron antineutrinos is not and a gas of $\bar{\nu}_e$ forms, with $Y_{\bar{\nu}_e}$ reaching peaks of ~ 0.04 . In comparison, the maximum of Y_{ν_e} is of the order of 10^{-3} , while $\max(Y_{\nu_x}) \approx 0.035 - 0.039$ is found depending on the resolution. This means that locally each neutrino species constituting the effective species x can be, on average, a factor 4 less abundant than electron antineutrinos.

2.2.3 Disc evolution

After merger, part of the matter expelled during the collision accretes around the remnant object, forming a disc. The baryonic mass of the disc M_b^{disc} is computed from the simulations according to Eq. A.28. The disc is conventionally defined as the baryon matter with density lower than $10^{13} \text{ g cm}^{-3}$, as in Shibata et al. [240] and therefore the integration domain V is restricted to the region of $\rho < 10^{13} \text{ g cm}^{-3}$ if a massive NS is present. If a BH forms, the domain is instead restricted by excluding the points inside the apparent horizon using the minimum lapse criterion, i.e. retaining only points for which $\min \alpha \geq 0.3$ (see the discussion in appendix of [90] for this choice. In this definition there is no distinction between bound and unbound material, which are both considered part of the disc. The error introduced by not subtracting the unbound material in the computation of the disc mass is lower than NR uncertainties, which can be estimated by comparing M_b^{disc} at different resolutions. M_b^{disc} for all runs, computed for each simulation at the last available time prior BH collapse, are listed in the fourth column of Tab. 2.1.




FIGURE 2.7: 2D snapshots of the xy -plane (top plot) and of the xz -plane (bottom plot) showing the properties of the accretion disc around the NS remnant at $t - t_{\text{merg}} = 20$ ms. In each plot all runs at LR (top row) and HR (bottom row) are compared. The right and left half of each frame show the matter specific entropy s and electron fraction Y_e , respectively. The thickest black curve is the isodensity contour $\rho = 10^{13}$ g cm $^{-3}$ delimiting the NS remnant core, while the other thinner curves moving outwards are isodensity curves corresponding to $\rho = 10^{12}, 10^{11}, 10^{10}, \dots$ g cm $^{-3}$. In the M1-HR frame the black circle represents the apparent horizon of the BH, that forms at 18.3 ms post-merger.

In Fig. 2.6 the time evolution of M_b^{disc} during the first 40 ms post-merger is reported. The largest increase in M_b^{disc} happens within ~ 10 ms post-merger, as a result of the collision and of the successive bounces of the two merging NS cores. On this timescale M_b^{disc} reaches values of the order of $\sim 10^{-1} M_\odot$, and then it stays constant for a few tens of ms, if the remnant NS does not collapse. When a BH forms, the disc mass drastically drops because a large fraction of the disc gets swallowed by the BH.

In the following, the differences arising by using different physics schemes are discussed for HR runs. In HY-HR, when neutrinos are not included, M_b^{disc} is the largest,

being almost double the LK-HR one. Even before BH formation, the LK run exhibits the smallest M_b^{disc} among all the runs, with $M_b^{\text{disc}} \approx 0.06 M_\odot$. This is explained by the fact that neutrino leakage cools down the lower-density matter around the NS core, causing the outer shells of the remnant NS to be less inflated and to expel less matter. When neutrino reabsorption is present (M0-HR, M1-HR) the M_b^{disc} increases, reaching $M_b^{\text{disc}} \approx 0.09 M_\odot$ with very similar values comparing the two runs. For VM0-HR, angular momentum and matter transport enhanced by viscosity has the effect of increasing the disc mass with respect to M0 only. Eventually M_b^{disc} reaches an intermediate value between HY-HR and M0-HR ones.

M_b^{disc} depends sensitively on resolution. LR runs present the largest M_b^{disc} for all simulations. Here, the minimum mass is found for LK-LR run, with $\sim 0.12 M_\odot$, while in M0-LR, VM0-LR and HY-LR runs M_b^{disc} reaches similar masses $\sim 0.15 M_\odot$. The largest M_b^{disc} is obtained for M1-LR simulation, with almost $\sim 0.25 M_\odot$. For this resolution a stable rotating NS forms, and on timescales longer than the ones shown in the plot M_b^{disc} slowly increases with time. This is due to the fact that at later time some matter is expelled from the outer shell of the remnant NS and becomes part of the disc [81]. For increasing resolution the disc mass decreases comparing each run with its lower resolution counterparts, except for HY-SR. The decrease can be as large as 44% (M0-LR vs. M0-SR). HR runs show systematically the smallest values of M_b^{disc} .

Finite resolution also affects M_b^{disc} indirectly by determining different times of BH collapse. Higher resolution simulations can predict final disc masses that are much smaller than lower resolution ones, because when a BH forms it swallows part of the disc. This is apparent in the run M1-HR, where the final disc mass drastically drops with respect to the mass before BH collapse. The presence of such lighter discs can have a large impact on the emission of gravitationally unbound material from the disc at secular timescales [see, e.g., 96, 241]. However, as long as gravitational collapse does not occur, the simulations show that the spread of M_b^{disc} due to different physics schemes gets smaller as the resolution increases.

Fig. 2.7 shows a comparison of the geometric properties and matter composition of the disc among the LR and HR runs as 2D snapshots of the xy -plane (top plot) and xz -plane (bottom plot) at $t - t_{\text{merg}} = 20$ ms. The geometry of the disc can be analysed referring to the black iso-density contours in the figure. The high-density portion of the disc $\rho \in [10^{12}, 10^{13}] \text{ g cm}^{-3}$ extends to ~ 20 km in the equatorial plane and ~ 10 km in the xz -plane. The region $\rho \in [10^{11}, 10^{12}] \text{ g cm}^{-3}$ is more inflated when neutrinos are present, compared to the HY case, in both xy - and xz - planes. The low-density

$\rho \approx 10^{10} \text{ g cm}^{-3}$ tails of the disc extends up to tens of km from the central object on the equatorial plane.

The most evident difference among resolutions is that discs are geometrically smaller for higher resolutions. If we consider the iso-density curve $\rho = 10^{10} \text{ g cm}^{-3}$ on the orbital plane, it extends to ~ 90 km for LK-LR and ~ 65 km for LK-HR. Similar numbers are found for M0 runs, while in VM0 runs the difference between LR and HR is smaller, ~ 10 km. The largest difference is found in M1 runs, for which the curve extends to $\gtrsim 100$ km for LR and to ~ 65 km at HR.

The composition of the disc is described in terms of the entropy and electron fraction profiles in Fig. 2.7. The high-density matter $\rho \in [10^{12}, 10^{13}] \text{ g cm}^{-3}$ is characterised by low electron fraction and low entropy because it is composed of fresh matter expelled from the remnant NS core. In the HY runs the electron fraction is frozen at $Y_e = 0.05$ because neutrinos are not simulated. Comparing HY (left column) in the bottom plot with the others, it can be seen that the presence of neutrinos clears the polar regions right above the remnant NS [221, 242]. In the runs with neutrinos, in these regions, Y_e increases up to values $\gtrsim 0.2$, indicating that matter protonises. In LK runs, for decreasing density and increasing distance from the remnant NS, the Y_e first increases as mentioned above, then decreases to ~ 0.1 at $\rho \in [10^{10}, 10^{11}] \text{ g cm}^{-3}$. At lower densities and high latitude $Y_e \lesssim 0.25$. In the region right above the remnant $Y_e \approx 0.4$, at LR. At HR the remnant is close to BH collapse and this causes an increase of temperature and consequent increase of electron fraction in the low-density matter above the remnant. The M0 and VM0 runs show different disc composition with respect to LK but similar between each others. Here, a fraction of neutrinos streaming out of the NS remnant is absorbed by lower-density material, increasing its Y_e . Y_e in the shell $\rho \in [10^{10}, 10^{11}] \text{ g cm}^{-3}$ is larger than in the LK case at the same density. For increasing latitudes (and decreasing density) Y_e increases, reaching values up to $Y_e \approx 0.35$. In the M1 runs the electron fraction has larger values when comparing shells of same density to M0 or VM0. In particular matter at high latitude and low density reaches $Y_e \approx 0.5$, and is thus quantitatively different from M0 runs. Such features are robust against changes of resolution.

2.3 Gravitational waves

In this section, the GWs extracted from the simulations are compared. The modes of the gravitational wave strain $h_{\ell m}$ are computed according to Eq. (B.10) in App. B.

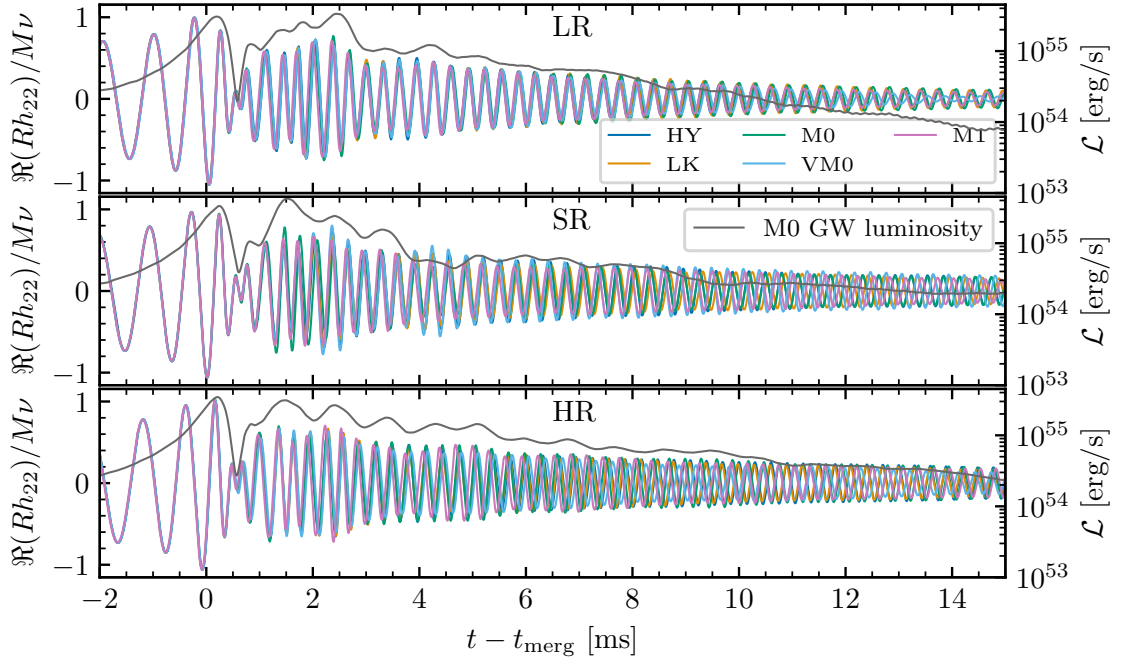


FIGURE 2.8: Real part of the $(\ell, m) = (2, 2)$ mode of the GW strain normalised by the total mass of the system M , the symmetric mass ratio ν and the extraction radius R . The grey curve corresponds to one representative GW luminosity curve (namely the M0 runs). Time is shifted by the merger time.

In Fig. 2.8 the $(2, 2)$ – mode of the GWs is compared among the runs up to 16 ms post-merger. Additionally, the GW luminosity $\mathcal{L}_{\text{GW}} := dE_{\text{GW}}/dt$ is reported for one representative run (M0 for each resolution). Up to merger, no significant differences among the waveforms are observed. The GW amplitude (see Eq. B.11) peaks at $RA_{22}^{\text{merg}}/M\nu \approx 1.06$ with a merger frequency of $f_{22}^{\text{merg}} \approx 1.9$ kHz. The post-merger spectrum peak frequency is $f_2 \approx 3.2$ kHz. These three quantities are measured quite robustly from the simulations studied here. At LR, the maximum variations of A_{22}^{merg} , f_{22}^{merg} and f_2 among all the runs are respectively $\sim 0.3\%$, $\sim 1.3\%$, $\sim 2.1\%$. At SR the maximum variations of these quantities are below 0.7%. Lastly, for HR runs the maximum variations of A_{22}^{merg} , f_{22}^{merg} and f_2 among all the runs are respectively $\sim 0.38\%$, $\sim 1.1\%$, $\sim 2.1\%$. The differences due to finite resolution are instead generally larger. Maximum differences between SR and HR of $\sim 1.3\%$ for A_{22}^{merg} , $\sim 4.1\%$ for f_{22}^{merg} and $\sim 2\%$ for f_2 are found. The GW luminosity peaks shortly after merger at $\mathcal{L}_{\text{GW}}^{\text{peak}} \approx 3.5 \times 10^{55}$ erg s $^{-1}$, consistently with Zappa et al. [88]. A maximum variation of the peak value of about $\sim 34\%$ is found among the runs.

In the post-merger waveforms, significant differences in the amplitude, frequency and phase evolution develop among the runs. Analyzing the phase convergence among different resolutions for fixed physics prescription, approximately first order convergence is obtained. To quantitatively compare the impact of resolution and of physics input on the waveforms, a faithfulness analysis is performed. The faithfulness between two waveforms $h_1(t)$ and $h_2(t)$ is defined as

$$\mathcal{F} := \max_{t_c, \phi_c} \frac{(h_1|h_2)}{\sqrt{(h_1|h_1)(h_2|h_2)}}, \quad (h_1|h_2) := 4\Re \int \frac{\tilde{h}_1(f)\tilde{h}_2^*(f)}{S_n(f)} df, \quad (2.2)$$

where t_c, ϕ_c are the time and phase of the waveforms at a reference time, and $(h_1|h_2)$ is the Wiener inner product; in the latter, the symbol \sim denotes the fourier transform, and $S_n(f)$ is the power spectral density, which here is chosen as the one of the Einstein Telescope. The unfaithfulness is defined as $\overline{\mathcal{F}} := 1 - \mathcal{F}$. In the context of GW parameter estimation, two waveforms are distinguishable if their faithfulness satisfies the necessary criterion [243]

$$\mathcal{F} > 1 - \frac{\epsilon^2}{2\varrho^2} \implies \varrho \approx \sqrt{\frac{N}{2\overline{\mathcal{F}}}}, \quad (2.3)$$

where ϱ is the matched-filtered signal-to-noise ratio (SNR) and $\epsilon^2 = N$ is chosen, with N being the number of intrinsic parameters of the system [244]. The latter condition on ϱ represents the minimum SNR that allows to detect the differences between two waveforms.

The runs are compared in pairs, in such a way that the two runs in a pair have either the same resolution (e.g. HY-LR and LK-LR) or are simulated with the same physics scheme (e.g. HY-LR and HY-SR), excluding comparisons of the kind HY-LR and LK-SR. At LR, the maximum mismatch is found between HY-LR and M0-LR runs, with $\overline{\mathcal{F}} \approx 0.087$. At SR the mismatches are generally larger and reach a maximum value of $\overline{\mathcal{F}} \approx 0.2$ between LK-SR and M1-SR and also between M0-SR and VM0-SR runs. At HR the mismatches are the largest, with maximum $\overline{\mathcal{F}} \approx 0.37$ between the HY-HR and M1-HR runs. Comparing runs at different resolutions, the mismatches are of the order of few times 10^{-1} in almost all comparisons, except for HY-LR vs. HY-SR and LK-LR vs. LK-SR, for which $\overline{\mathcal{F}}$ is few times 10^{-2} .

This analysis indicates that possible effects due to neutrinos or turbulent viscosity in the remnant core can be detected in the GW signal only in the post-merger.

However, GW models used for matched filtering that are informed on NR simulations [70, 71] at LR would not be accurate enough to detect such effects. In particular, differences due to the simulations' finite resolution would be dominant in such GW models. At SR and HR, mismatches between waveforms of runs performed with different physics schemes are comparable to the ones due to finite resolution. GW templates constructed with these data might be able to distinguish such differences in the signal from $\varrho \gtrsim 3$, according to Eq. (2.3). Notably, this precision might be sufficient for third generation observations, since differences in the signals due to variations in the EOS at extreme matter densities are potentially observable at post-merger SNR ~ 8 [71].

Overall, these results indicate that simulations at SR or HR are necessary in order to distinguish possible differences due to neutrinos or turbulent viscosity in the remnant. In particular, the high-resolution M1 simulations do not show any evidence for significant out-of-equilibrium and bulk viscosity effects in the waveforms. This is in agreement with the findings of [194] that were obtained at LR, but it is in contrast with Refs. [192, 193]. The simulations performed for the latter works do not consider weak interactions or use a LK scheme and are performed at a maximum resolution of 400 m, which is much lower than the LR employed in this work.

2.4 Ejecta mass

The material ejected from the system on dynamical timescales and up to ~ 20 ms post-merger consists of the full dynamical ejecta component and the early portion of the spiral-wave wind component. The dynamical ejecta is composed of a tidal component originating from tidally unbound NS material and a shocked component originating from the first bounce after the collision of the two cores [96]. The tidal component is launched mostly across the equatorial plane and is characterised by a low $Y_e \approx 0.05 - 0.15$ and low entropy, $s \lesssim 5 k_B \text{ baryon}^{-1}$. The shocked component has higher entropy than the tidal component and peak temperature of tens of MeV, which produces large amount of electron-positron pairs with consequent increase of Y_e due to positron captures by neutrons. Neutrinos radiated from the remnant and absorbed by neutrons in the ejecta can further increase the Y_e of this ejecta component, especially at high latitudes where neutrino emission is more efficient. The shock-heated ejecta expand over the entire solid angle due to interaction with the tidal ejecta, hydrodynamic shocks and weak interaction, with a preferential emission at small angles from the equatorial plane.




FIGURE 2.9: Time evolution of the ejecta mass extracted at $R = 443$ km comparing the Bernoulli and geodesic criteria. Mass is reported in logarithmic scale and compared across resolution.

Other mechanisms can unbind material from the disc and they act generally on longer timescales. Spiral-wave winds originate from non-axisymmetric density waves from the NS remnant [98]. The remnant’s spiral arms transport angular momentum outwards in the disc and material then gets unbound from the disc edge. On longer timescales disc winds can develop, also powered by neutrino reabsorption [e.g. 100, 102, 183, 245, 246] but the simulations in this work are not sufficiently long to capture this component.

In the literature there are two main ways to identify the unbound material from simulations, namely the geodesic and the Bernoulli criterion (see, e.g., [247] for a recent work on this topic). The geodesic criterion assumes that the ejecta follow spacetime geodesics in a time-independent, asymptotically flat spacetime. Therefore, a particle is considered unbound if $u_t < -1$, where u_t is the time component of the particle’s 4-velocity. According to the Bernoulli criterion, and using the definitions of thermodynamics quantities in §A.3, a fluid element is considered unbound if $hu_t \leq -1$, where h is the fluid specific enthalpy. The asymptotic velocity of the unbound particle is calculated as $v_\infty \simeq \sqrt{2(h(E_\infty + 1) - 1)}$. This criterion assumes that hu_t is constant along a streamline of a steady-state flow. This assumption is correct if the metric and the flow are both stationary. Even though this is not formally true for merger outflows, this criterion is considered sufficient to account for the gain in kinetic energy of the expanding matter in the outflow due to thermal and nuclear binding energy. However, it neglects the energy lost to neutrinos during r -process nucleosynthesis, which can affect the total ejecta mass and its velocity profile [247]. The geodesic and Bernoulli criteria can be used to conventionally separate the dynamical ejecta from the wind ejecta [248].

Fig. 2.9 compares the evolution of the ejecta mass in the simulations according to both the geodesic and Bernoulli criteria. At 20 ms post-merger the ejecta masses calculated with the geodesic criterion are saturated, except for the M1 runs. As expected, the ejecta mass calculated with the Bernoulli criterion is larger than the one estimated with the geodesic criterion at comparable times. The ejecta mass in the Bernoulli case keeps increasing at later time due to the contributions of the spiral-wave winds. In the rest of this section only the Bernoulli ejecta are discussed.

The ejecta mass shows a steep increase up to ~ 10 ms post-merger in all the runs and then it tends to saturate at few tens of ms after merger. Within ~ 20 ms post-merger a mass of $\gtrsim 2 \times 10^{-3} M_{\odot}$ is typically ejected. The quantitative values for all the runs at a fixed time are reported in Tab. 2.1. For HY runs, $\gtrsim 8 \times 10^{-3} M_{\odot}$ of matter is expelled, which represents the largest matter emission among all the runs. The ejecta mass in LK runs is systematically one order of magnitude lower than that of all the other runs, consistently with [221]. This happens because the neutrino cooling reduces the enthalpy of the material and as a result the emission is largely decreased. When neutrino reabsorption is included through the M0 scheme, the effect of cooling is counteracted by the neutrino energy deposition in the shock-heated ejecta and M_{ejecta} becomes larger than the LK case, reaching values $\gtrsim 10^{-2} M_{\odot}$. The evolution of M_{ejecta} in VM0 runs follows a similar behaviour. M_{ejecta} measured in M1 and M0 runs are comparable, within a few tens of percent.

Focusing on the effects of finite resolution, M_{ejecta} shows a monotonic decrease for increasing resolution for all the runs. Since the onset of BH collapse stops the matter ejection, smaller final ejecta masses in HR simulations are found, compared to the other cases. The maximum variations in the ejected mass at a fixed time of 20 ms post-merger due to resolution are found between VM0-SR and VM0-HR and amount to $\sim 50\%$.

The most salient properties of the ejected material are summarised in the histograms of Fig. 2.10. The histograms produced using the geodesic criterion do not differ significantly from those obtained with the Bernoulli criterion and therefore only the latter are discussed. The mass-weighted averages of the relevant quantities are reported in Tab. 2.1. Most of the mass is emitted almost uniformly in the interval $0^{\circ} \leq \theta \lesssim 50^{\circ}$ (second column of Fig. 2.10). The peak at $\theta \approx 45^{\circ}$ is due to an artefact in the mass extraction and it is not physical. At larger angles, the mass emission is slightly more suppressed in LK runs with respect to the other cases [221]. The average emission angle for all runs is enclosed in $\theta \in [27^{\circ}, 37^{\circ}]$ and is systematically lower for




FIGURE 2.10: Histograms of the ejecta extracted at $R = 443$ km. Each row shows the fraction of ejecta mass in a bin normalised to the total ejecta mass in logarithmic scale for different resolutions. In each column is represented, respectively, the electron fraction, the latitudinal distribution and the asymptotic velocity of the ejecta. In each frame the ejecta properties among all runs are compared. The analysis is performed until $t - t_{\text{merg}} = 20$ ms, corresponding to the earliest t_{end} of the set of simulations (see Tab. 2.1).

LK at all resolutions.

The asymptotic velocity distribution is peaked around values in the interval $0.15 \leq v_{\infty}/c \lesssim 0.22$. The velocity distribution has fast tails reaching $\sim 0.8 c$. These tails can originate a radio-X-ray afterglow to the kilonova emission, peaking at years post-merger timescales [248–251]. The mass of the fast tail of the ejecta, i.e. with asymptotic velocity $v_{\infty}/c \geq 0.6$, is $M_{\text{ej}}^{v \geq 0.6c} \approx 10^{-6} - 10^{-5} M_{\odot}$ (see Tab. 2.1).

From the simulations' data it is possible to model the function $M_{\text{ej}}(v_{\infty}/c)$ approximately with a broken power law of the kind [252]

$$M = M_0 \begin{cases} \left(\frac{\beta\gamma}{(\beta\gamma)_0} \right)^{-s_{\text{KN}}} & 0.1 < \beta\gamma < (\beta\gamma)_0 \\ \left(\frac{\beta\gamma}{(\beta\gamma)_0} \right)^{-s_{\text{FT}}} & \beta\gamma > (\beta\gamma)_0 \end{cases} \quad (2.4)$$





FIGURE 2.11: 2D snapshot of the xz -plane showing the Y_e of the material around the remnant NS at $t - t_{\text{merg}} = 20$ ms. The thickest black curve is the isodensity contour $\rho = 10^{13}$ g cm $^{-3}$ delimiting the NS remnant, while the others moving outwards correspond to densities $\rho = 10^{12}, 10^{11}, 10^{10}, \dots$ g cm $^{-3}$. Each row corresponds to different resolutions, while each column to different micro-physics prescriptions. Profiles for HY runs are not reported because neutrinos are not simulated and the electron fraction distribution is frozen at $Y_e = 0.05$.

where $\beta = v/c$, γ is the corresponding Lorentz factor and $(\beta\gamma)_0 = \beta_0 \cdot \gamma(\beta = \beta_0)$. The values of β_0 defining the “break” in the broken power vary in the range $\beta_0 \sim 0.3 - 0.45$. Fitting parameters are $M_0 \approx (3.2 - 17) \times 10^{-5} M_\odot$, $s_{\text{kN}} \approx 0.64 - 1.6$ and the ejecta tail with $v_\infty/c \gtrsim \beta_0$ can have a rather steep dependence on the velocity, with $s_{\text{FT}} \approx 4 - 11$.

The Y_e (first column of Fig. 2.10) exhibits the most complex behaviours, different among the runs. The spatial distribution of Y_e in the ejecta is depicted in Figs. 2.11 and 2.12, where 2D slices of the Y_e profiles in the xy - and xz -plane, respectively, are reported. For HY runs Y_e is frozen at ~ 0.05 because weak interactions are not simulated and the matter composition does not change throughout the run with respect to the initial neutrino-less weak equilibrium condition. For LK cases the ejecta mass composition peaks at $Y_e \approx 0.13 - 0.17$ (compare also to Tab. 2.1). No significant

FIGURE 2.12: Same as Fig. 2.11 but for the xz -plane

fraction of ejecta has $Y_e > 0.35$. The material at low $Y_e \lesssim 0.15$ is emitted at small latitudes (left-most column of Fig. 2.12), while for increasing angles Y_e increases, reaching $Y_e \lesssim 0.35$ in the lower-density region above the remnant NS. Matter at high latitudes is shock-heated ejecta, therefore hot, and is expanding in a region where the disc is not present. Under these conditions, the expanding matter becomes transparent earlier producing electron-positron pairs. Therefore, positron captures increasing Y_e are more efficient even in the absence of neutrino absorption. For both M0 and VM0 the Y_e distribution gets broader with respect to LK, with a large fraction of matter having $Y_e \in [0.2, 0.35]$. This is the effect due to neutrinos radiated by the central object and the disc that are absorbed by neutrons in the ejecta, converting of neutrons into protons. As in the previous case, the low- Y_e material is emitted at lower latitudes and the Y_e increases for increasing latitudes. The peak at $Y_e \approx 0.3$ observed in the left column of Fig. 2.10 is reached in the high-latitudes, low-density ejecta (second and third column of Fig. 2.12). This is because neutrino fluxes are significantly larger at high latitudes, due to the presence of the disc at low latitudes. In M1 runs the trend is similar but even higher values of Y_e are reached.

The histograms in Fig. 2.10 show that the peak at $Y_e \approx 0.3$ of M0 and VM0 translates to $Y_e \gtrsim 0.425$ when switching to M1. Material with such a high Y_e is found once again at large latitudes. The comparison to M0 runs indicates that accounting for neutrino transport with a more complete neutrino scheme provides more efficient proton production in the shock-heated ejecta component. One of the causes of this is that the M0 scheme uses a spherical grid that assumes neutrinos are only moving radially. On the contrary, the M1 scheme is solved in the computational grid and the radiation is evolved according to 3D transport. Neutrinos from the disc will naturally tend to escape along the z -direction, in which the gradient of the optical thickness decreases more steeply and the neutrinos mean-free path increases faster, further irradiating the high-latitude ejecta. While the M1 scheme naturally overestimate the density of neutrinos along the polar axis due to non-linear closure relation used in the scheme which causes beam crossing, such large Y_e ejecta wind is confirmed also in simulations in which the linear Eddington closure is employed [194].

Finite resolution has a clear effect on the ejecta composition, especially visible at HR. All runs at LR and SR show a peak at $Y_e \approx 0.05$ that is due to the tidal component of the ejecta, which is emitted at early times after merger and maintains the Y_e of the two initial stars. However, for HR runs this component is strongly suppressed for all but the run with viscosity. This can be explained by two different factors. First, the tidal ejecta are expected to be less massive at HR, because the tidal deformation causing this emission at merger are better resolved. Second, the discs are less massive and geometrically thinner for HR runs, compared to the others. Therefore, it is easier for neutrinos to escape from the inner regions and interact with the ejecta, increasing its Y_e . The latter explanation is supported by the fact that in the VM0-HR run the disc is not as thin as in the other HR runs and only for this case the low- Y_e peak is not heavily suppressed. This contributes to explain why the M1-HR run exhibits such a large Y_e in both the xy - and xz - planes. On the one hand, the disc is thinner because it is a HR run. On the other hand, neutrino fluxes predicted by the M1 scheme increase the Y_e in the matter more efficiently with respect to the M0 scheme.

2.5 Nucleosynthesis and Kilonova light curves

Nucleosynthesis Nucleosynthesis abundances inside the ejecta extracted from the simulations are computed with the SkyNet code [253] according to the procedure described in Radice et al. [96]. The nucleosynthesis yields are shown in Fig. 2.13, where




FIGURE 2.13: Nucleosynthesis yields comparison among the three physical prescriptions against the solar abundancies for the three resolutions in each panel. On the x -axes the mass number A is reported. The relative abundances on the y -axis are expressed in logarithmic scale. Consistently with Fig. 2.10, the analysis is performed up to $t - t_{\text{merg}} = 20$ ms for all runs.

results obtained from LK, M0 and M1 runs are compared against the solar residual r -process abundances from Arlandini et al. [254]. Abundances are normalised to the one corresponding to the third r -process peak $A \approx 190$. Once the third peak abundances have been fixed, the abundances predicted using all neutrino schemes are roughly compatible between each other and with the solar residual pattern for $A \in [125, 140]$ (i.e., for the second r -process peak) and $A \in [170, 200]$. For $A \in [140, 170]$ and $A > 200$, the yields from all the simulations significantly differ from the solar residuals. Such discrepancies are possibly due to nuclear physics inputs, as well as to a lack of suitable physical conditions to efficiently produce actinides [see, e.g., 255, 256]. The LK runs heavily underestimate the abundances for $A < 120$. This is a direct consequence of the fact that Y_e is lower in the ejecta for these cases. In the runs with M0 and M1 the abundances for $A < 120$ are closer among them and to the solar residuals, compared to LK.

Increasing the resolution does not change the abundances in the runs with LK, which also at high resolution significantly differ from the solar residual abundances for $A < 120$. For M0 and especially for M1 runs the predictions at HR better match the solar abundances for the entire range of nuclear masses (still with the exceptions discussed above).

These results confirm the relevant role of neutrino emission and absorption in shaping the nucleosynthesis yields from the early time ejecta of BNS mergers [see, e.g., 96, 189, 257, 258]. Abundances obtained in the HR simulations employing the M0 or M1 schemes are compatible between each other and reproduce well the observed

solar residual pattern. However, models featuring neutrino cooling alone underestimate the abundances of light r -process elements, since neutrino reabsorption is required to produce the ejecta conditions suitable for the production of those elements.

Kilonova light curves Synthetic kilonova light curves are computed following the approach outlined in Wu et al. [259] and using the radiation-hydrodynamics lagrangian KiloNova Explosion Code (KNEC) [260]. Accordingly, the dynamical ejecta extracted from the simulations are further evolved with KNEC up to 15 days post-merger. The corresponding light curves are presented using the AB magnitude system

$$m_{\text{AB}} = -2.5 \log_{10} \left(\frac{\int f_{\nu}(h\nu)^{-1} e(\nu) d\nu}{\int 3631 \text{Jy} (h\nu)^{-1} e(\nu) d\nu} \right) \quad (2.5)$$

where here ν is the light frequency, f_{ν} is the observed flux density at frequency ν from a distance of 40 Mpc and $e(\nu)$ are filter functions for different Gemini bands. We refer to [259] for more details.

In Fig. 2.14 the AB magnitudes at different bands are compared to the electromagnetic transient AT2017gfo associated to the BNS merger event GW170817 [40]. As input for the KNEC code, the ejecta extracted at two different times are compared: at 20 ms post-merger (dashed lines) and at the end of the simulation (solid lines). Clearly, the different simulation lengths impact on the light curve due to the different ejecta masses, but also due to the composition. AT2017gfo is significantly brighter than any of the synthetic light curves reported in the plot. Nonetheless, the hierarchy of the colours is correct at ~ 4 days, whereas the 1 day emission has a blue peak that cannot be explained with dynamical ejecta considered here. The fact that the data are not reproduced in this analysis is expected for many reasons. First, the BNS simulations are not targeted to the event GW170817; in particular it has lower mass and symmetric mass ratio, which implies smaller ejecta masses and therefore dimmer light curves. Second, the simulations are too short and cannot capture the full evolution of the post-merger disc. Therefore, ejecta emitted at secular timescales (seconds after merger) are missing. Crude estimates of later outflows emission can be made by extrapolating in time [259], but this is not attempted here. Third, multidimensional effects and viewing angle can have a strong impact on the kilonova emission [e.g., 39, 261, 262] but are neglected here. For AT2017gfo, spherically symmetric kilonova models are ruled out with high confidence [39, 40, 44].

In the following, the differences seen for different physics schemes employed are reported. For HY runs the light curves corresponding to the K_s , H and J band are




FIGURE 2.14: Light curves calculated with KNEC. The AB magnitudes as a function of days after merger are reported. The light curves correspond to several Gemini bands and are calculated from the ejecta extracted at $R = 443$ km from the system at a common time 20 ms post-merger (dashed lines) and at the end time of each simulation (solid lines). Dots correspond to the data of the kilonova event AT2017gfo for the same bands.

only a few magnitude larger than AT2017gfo data, especially when ejecta production at ~ 109 ms post-merger is considered (solid lines in the LR and SR cases). By contrast, dynamical ejecta alone produce significantly dimmer light curves (dashed lines), in particular at late time after the peaks. Despite the usually long simulation lengths, for LK runs the ejecta mass is smaller and this produces dimmer light curves compared to HY runs, considering both the early ejecta and those at the end of the simulations. The jumps that are seen in these curves are an artefact of the KNEC code. For M0, VM0 and M1 light curves are brighter with respect to LK at all bands, as a consequence of the fact that more ejecta mass, characterised by a larger Y_e , is produced. When considering only the early ejecta (dashed lines), M0, VM0 and M1 produce very compatible light curves, due to the very similar ejecta properties, see §2.4 and Tab. 2.1. M1 light curves are slightly dimmer due to the faster and less opaque ejecta, which translate to a faster kilonova evolution after the peaks. Differences become more pronounced when light curves are computed using the ejecta at the end of the simulations, since M1 runs were evolved for shorter post-merger times and produced systematically less ejecta mass.

Finite resolution effects do not significantly impact the light curves. The analysis shows that the light curves are very sensitive both to the inclusion of neutrino reabsorption in optically thin conditions and to the cumulative time during which ejecta are measured. During this time not only the ejecta mass, but also the ejecta composition changes due to the different emission mechanisms at different timescales. The better accuracy provided by the M1 scheme with respect to the M0 one seems to have a minor impact on the kilonova light curves due to the good agreement in the ejecta properties between the two schemes, when the simulations have comparable lengths. Future simulations will extend these results by also considering the winds from the viscous post-merger phase and taking into account non-spherical geometries.

2.6 Neutrino luminosity

This section investigates the impact that the various physics schemes and finite resolution effects have on the neutrino emission in the simulations. Fig. 2.15 shows the angle integrated neutrino luminosity for the three neutrino species that are simulated, comparing M0 and M1 neutrino schemes for every resolution. Hereafter, only one representative heavy flavour neutrino species is considered, denoted as ν_μ with properties




FIGURE 2.15: Neutrino luminosity comparison between M0 and M1 simulations for all three neutrino species and for the three resolutions in each panel. Neutrinos are extracted at a radius $R_{M0} = 756$ km for the M0 case and $R_{M1} = 443$ km for the M1 case. The data is smoothed using a rolling average with width 1 ms. Time is shifted by the time of merger and by the time of flight of neutrinos to the corresponding detector. The luminosity is expressed in cgs units and in logarithmic scale.




FIGURE 2.16: Same as Fig. 2.15 but for the neutrino average energy.

calculated as averages over the four neutrino species constituting ν_x . Neutrino luminosities for every species present a peak immediately after merger at $\mathcal{L} \approx 10^{52} - 10^{53}$ erg s $^{-1}$. The hierarchy $\mathcal{L}_{\nu_\mu} < \mathcal{L}_{\nu_e} < \mathcal{L}_{\bar{\nu}_e}$ observed in the neutrino luminosity evolution is consistent with previous results [see, e.g., 180, 190, 263–265] and it is explained as follows. Electron antineutrinos are the most abundant species because the positron captures on free neutrons are favoured in the neutron rich ($Y_e \approx 0.1$) matter with temperatures of tens of MeV. Electron neutrinos are produced instead mostly due to capture of electrons on protons, which are however not favoured due to the initially low proton abundance. Heavy flavour neutrinos are produced by matter with temperature of tens of MeV emitted from the bouncing remnant. The reactions producing heavy flavour neutrinos are electron-positron annihilation and plasmon decay which

are highly dependent on temperature. As the remnant stabilises and cools down, production of heavy flavour neutrinos lowers, while electron/positron captures keep happening in the highest density region of the accretion discs producing electron neutrinos and antineutrinos.

Comparing M0 and M1 runs, the same luminosity hierarchy is maintained, but the M1 scheme predicts larger neutrino brightnesses. The largest difference is observed for heavy flavour neutrinos, where neutrinos in M1 runs are $\sim 50\%$ brighter than neutrinos in M0 runs.

When a BH forms, the emission of neutrinos decreases for each species, but it abruptly stops for heavy flavour neutrinos, because this component is mostly emitted from the remnant NS. As resolution increases, all the luminosities increase at early times, within 15 ms post-merger. This is largely explained by the fact that thinner discs are formed at this resolution, which allow neutrinos to diffuse more easily and with shorter timescales. Larger electron antineutrino luminosities at HR, for both M0 and M1 schemes, are in agreement with the fact that larger electron fractions are found in the ejecta distributions for HR.

Fig. 2.16 reports the neutrino average energies for the same runs. For both M0 and M1 schemes, the energies peak at 2 – 3 ms post-merger before reaching a quasi-steady evolution at later times, and follow a hierarchy $\langle \epsilon_{\nu_e} \rangle < \langle \epsilon_{\bar{\nu}_e} \rangle < \langle \epsilon_{\nu_\mu} \rangle$. Focusing on M1 runs, neutrino average energies for heavy lepton neutrinos peak at 40 MeV and then decrease below ~ 30 MeV within few tens of ms. Electron antineutrinos and neutrinos follow a similar behaviour also with similar timescales, reaching their maxima at ~ 25 MeV and ~ 20 MeV, and decreasing to ~ 12 MeV and ~ 10 MeV, respectively. After BH collapse, the average energy of heavy flavour neutrinos drops. Runs with the M0 scheme systematically underestimate the energies in the first $t - t_{\text{merg}} = 20$ ms by $\approx 30\%$ with respect to M1 runs. We also note that the average energy of the heavy lepton neutrinos increases with time, reaching values comparable to the ones simulated with the M1 scheme, within tens of ms post-merger.

The quantitative differences in the ν_e and $\bar{\nu}_e$ energies between the two sets of runs possibly originate from different causes. M1 simulations tend to produce more massive and inflated discs. Such discs have more extended neutrino surfaces characterised by lower decoupling temperatures. At the same time, the M1 scheme properly models the diffusion of neutrinos inside the remnant up to the emission at the neutrino surface and their thermalization [194]. In M0 schemes, instead, the diffusion rate is estimated based on local properties and thermalization effects of diffusing neutrinos are not taken into

account. For heavy flavour neutrinos the situation is opposite: neutrinos decouple from matter deep inside the remnant, further diffusing through quasi-isothermal scattering inside the disc. While an M1 scheme is able to catch this effect, retaining larger ν_μ mean energies, the M0 computes the luminosities and mean energies considering neutrinos in equilibrium with matter everywhere inside the last scattering surface, providing at the same time lower mean energies and larger luminosities. With time, the disc becomes more compact and the diffusion atmosphere reduces in size, so that ν_μ mean energies become comparable.

The average neutrino energies are not largely influenced by resolution effects. This is expected because the neutrinospheres are mostly determined by the density profile inside the disc [185], which are shown to be robust with resolution (see §2.2.2).

Finally, it is interesting to notice that the value of the high electron fraction peak in the Y_e distribution in the ejecta extracted from the M1 runs (and observed at high latitudes) is close to the equilibrium electron fraction, $Y_{e,\text{eq}}$. The latter can be estimated using eq. 77 of Qian and Woosley [266]. Assuming, according to the reported neutrino luminosities and mean energies around 10 ms post-merger, $\mathcal{L}_{\bar{\nu}_e} \approx 3/2 \mathcal{L}_{\nu_e}$, $\langle \epsilon_{\nu_e} \rangle \approx 12$ MeV and $\langle \epsilon_{\bar{\nu}_e} \rangle \approx 14$ MeV, we find $Y_{e,\text{eq}} \approx 0.46$. This means that in the region above the massive NS absorption rates in the M1 runs are high enough to approach weak equilibrium, and the differences with the M0 results are mostly due to the rates values, rather than to differences in the relative luminosities or mean energies.

Chapter 3

Efficient BBH simulations with the GR-Athena++ spacetime solver

The merger of CBs is a multi-scale problem; one of the most popular approaches to treat the issue of resolving features sensitive to differing length-scales in NR simulations, inspired by the Berger-Oliger approach [228, 229], is based on the construction of nested grids. These are made of overlapping nested boxes of increasing resolution and decreasing physical extent, where the most resolved boxes typically contain the components of the CB system. During the evolution, (some of) these boxes move tracking the position of the two bodies. Subcycling in time is commonly employed for the time integration, where finer grids take multiple iterations with local, smaller time steps with respect to coarser grids. Successful examples of NR codes adopting this technique are the Cactus-based [222] codes `Llama` [230, 267], `McLachlan` [268], `LEAN` [269], `LazEv` [270], `Maya` [271], `GRHydro` [272], `WhiskyTHC` [220], which make use of the `Carpet` thorn [231], and the non-Cactus codes `BAM`, [273–275], `AMSS-NCKU` [276] and `GRChombo` [277]. The main drawback of the Berger-Oliger method is the need of data synchronization among different boxes of the mesh grid when utilizing subcycling, which can spoil parallel scaling in the modern many-cores HPC architectures when parallel execution is performed on a large number of cores [278].

One alternative to the nested grid approach is represented by block-based AMR strategies [278]. In this case the computational grid is partitioned into non-overlapping blocks, that can be logically mapped into a tree structure to enforce data locality in memory [279]. Block-based AMR strategies show high parallel efficiency and offer a greater flexibility in refining the grid compared to the nested grid approach, as demonstrated by the `Dendro-GR` code [280]. A block-based AMR strategy is at the core of `Athena++` [281–283]. The latter is a public framework for purely non- and special- relativistic MHD, as well as GRMHD for stationary spacetimes, which adopts

many of the mature and robust numerical algorithms of Stone et al. [284] and is written in a modern C++ design. Parallelism in Athena++ is exploited in two ways: through its task-based computational model, which allows an overlap between communication and computation, and through a hybrid parallelism strategy which combines message passing interface (MPI) and threading via OpenMP (OMP).

The work presented in this *Chapter* is a collaborative effort in the development of GR-Athena++, consisting of the implementation of a spacetime solver module in Athena++. My contributions in this work are the following: I collaborated in the implementation of the Z4c equations in the Athena++ code; I implemented the puncture tracker and designed and implemented the octree box-in-box AMR strategy described in §3.2; I performed the convergence tests and cross-code comparisons with the BAM code for the two punctures calibration problem reported in §3.3.2; I designed and performed the scaling tests of §3.4.

3.1 Methods

The structure of Athena++ can be summarized as follows. A class named `Mesh` contains the details about the domain of definition Ω for a given problem. The code is designed such that evolutions can be performed in spatial 1D, 2D, 3D and for different coordinatization types (e.g. Cartesian, spherical..), whose information are stored within the `Mesh`. For definiteness, hereafter only 3D problems on Cartesian coordinates are considered. However, an analogous discussion holds for lower spatial dimensions and curvilinear coordinates.

The `Mesh` class contains the number of points for the coarsest sampling along each dimension $N_M := (n_x, n_y, n_z)$ as well as information about the physical boundary conditions on $\partial\Omega$. The domain can be decomposed in rectangular cuboids with $N_B := (n_x^{\text{MB}}, n_y^{\text{MB}}, n_z^{\text{MB}})$ points each, where each element of N_B must be a divisor of each element of N_M , component-wise. The elements of this partition are called `MeshBlocks` in the Athena++ language¹. The properties of `MeshBlocks`, as well as the local values of field variables of interest, are stored in the `MeshBlock` class.

3.1.1 Tree structure of the Mesh

After the initial `Mesh` is divided into `MeshBlocks`, based on the choice of N_M and N_B , each individual `MeshBlock` can be subdivided further, as explained in detail below.

¹The superscript MB in the definition of the tuple N_B stands for `MeshBlock`.




FIGURE 3.1: Example of Mesh partitioned uniformly by MeshBlock objects (left plot) and with refinement (right plot). In both plots it is shown how MeshBlocks are indexed via Z-order and traced in red through a Z-shaped curve linking their geometric centroids. Notice that physical level p and logical level l are in general distinct.

All MeshBlocks have the same size in terms of number of points, but they can have different physical extents. However, neighbor MeshBlocks are constrained to only differ by one sub-division at most. Grids in which all MeshBlocks have the same physical extent as in the left plot of Fig. 3.1, are called *uniform grids*. Grids in which not all MeshBlocks have the same physical size, such as the one depicted in the right plot of Fig. 3.1, are called *refined grids*.

The crucial ingredient for such MeshBlock-based grid structure is an efficient inter-MeshBlock communication, which is achieved in Athena++ by means of a tree data structure of the Mesh. For 3D evolutions, the logical relationship between MeshBlock objects is stored within an oct-tree data structure. This is constructed by first selecting the minimum N such that 2^N exceeds the largest number of MeshBlocks along any dimension. The root of the tree is assigned a logical level l of 0 and the tree terminates at level N assigning every MeshBlock to an appropriate leaf. In this process, some leaves and nodes of the tree may remain empty. MeshBlocks are sorted according to Z-order [285], where multi-dimensional coordinates are recast into a linear index parametrizing a Z-shaped, space filling curve. As a result, small changes in the parameter imply spatial coordinates that are close, enforcing data locality [286]. An example of such organization for a simple uniform grid with $(2, 2, 2)$ MeshBlocks

in each direction is depicted in the left plot of Fig. 3.1.

Consider now a refined grid. Function data at a fixed physical level is transferred one level finer through use of a prolongation operator \mathcal{P} ; dually, function data may be coarsened by one physical level through restriction \mathcal{R} . The number of physical refinement levels added to a domain-decomposed Ω is controlled in Athena++ by the parameter N_L . By convention N_L starts at zero. Subject to satisfaction of problem-dependent refinement criteria, there may exist physical levels at $0, \dots, N_L$. When a given MeshBlock is refined (coarsened), 2^d MeshBlock objects are constructed (destroyed). As mentioned above, such a procedure is constrained to satisfy a 2 : 1 refinement ratio where nearest-neighbor MeshBlock objects can differ by at most one physical level. The right plot of Fig. 3.1 shows refined grid described by an initial Mesh sub-divided into (2, 2, 2) MeshBlock and subjected to a refinement introduced at the corner $(x_{\max}, y_{\min}, z_{\max})$, with $N_L = 3$.

3.1.2 Vertex-centered discretization

Natively, Athena++ implements cell-centered (CC) and face-centered (FC) description of variables, together with calculation of line-averages on cell edges [283]. GR-Athena++ extends support to allow for vertex-centered (VC) representation of variables. The modifications required to achieve this are extensive to complement all the existing functionalities of the code. The main motivations for spacetime evolution based on VC discretization is a desire to ensure each stage of the evolution maintains consistent (high) order while simultaneously maintaining efficiency of \mathcal{R} and \mathcal{P} operator choice and implementation. Moreover, VC approaches reduce reflections at the boundaries of the domain, which can introduce spurious features in CC-based codes. Details on the communication among MeshBlocks for CC and FC variables in Athena++ are present in [283]. In the following, the treatment for VC variables in GR-Athena++ is discussed.

Communication of VC variables: uniform grids In the following, N_M and N_B are considered uniform, i.e. $n_x = n_y = n_z := n^M$ and $n_x^{\text{MB}} = n_y^{\text{MB}} = n_z^{\text{MB}} := n^{\text{MB}}$. In 1D, $x \in [a, b]$ is said to be vertex-centered when discretized as $x_I = a + i\delta x$ where $\delta x := (b - a)/n^{\text{MB}}$ and $i := 0, \dots, n^{\text{MB}}$ for a total of $n^{\text{MB}} + 1$ physical nodes. Along with physical nodes, also ghost nodes need to be considered for the imposition of boundary conditions (BC). Therefore, the total number of nodes is augmented by $2\mathcal{N}_g$, i.e. \mathcal{N}_g nodes on each end of $[a, b]$.




FIGURE 3.2: Schematic of (communicated) nodes on a two-dimensional **MeshBlock** Ω_i . The ghost-layer is shaded in gray with alternating shading demarcating differing neighbor **MeshBlock** objects. Nodes marked with “ \square ” are interior to Ω_i and are unaffected as neighbor data are received. Ghost-layer multiplicities for dark-green “ \blacksquare ” are $\mu = 1$, whereas nodes “ \blacklozenge ” in light-green have $\mu = 2$. Interface nodes along edges are marked with “ \bullet ” in light-green and correspond to $\mu = 2$ whereas corner nodes marked with “ \star ” correspond to $\mu = 4$.

In 3D analogous definitions can be given considering Cartesian products of 1D intervals along the three dimensions. A **MeshBlock** then consists of a bulk of physical nodes and an outer shell of ghost points. The latter are used for inter-**MeshBlock** communication among neighboring **MeshBlocks** and to impose the physical BC at the edges of the computational domain.

Consider a domain decomposed into multiple **MeshBlocks**. Inter-**MeshBlock** communication with respect to the CC case is complicated due to the presence of physical shared nodes at the neighboring **MeshBlock** interfaces. For uniform, 3D grids, nodes that lay on the faces of **MeshBlocks** are typically shared between two blocks and have multiplicity $\mu = 2$; nodes that lay on the edges typically have $\mu = 4$; nodes at the corner can have up to $\mu = 8$. This is illustrated with the 2D example of Fig. 3.2. In this case, physical nodes can be shared at the edges of **MeshBlocks** ($\mu = 2$) and at the corners ($\mu = 4$). In the figure, $n^{\text{MB}} = 6$, $\mathcal{N}_g = 2$ and a **MeshBlock** Ω_i that is not on the physical boundary is considered.

Communication of ghost points is implemented as follows. Independent communication requests and buffers are posted for each neighbor, with no preferential order. All received data from all neighbors of Ω_i are summed and stored in an auxiliary array, based on the location of the relevant neighbors. After all data points have been received, they are averaged by dividing them by their corresponding multiplicity μ .




FIGURE 3.3: Schematic of two-dimensional $\text{MeshBlock } {}^p\Omega_A$ used to populate ghost-nodes of finer $\text{MeshBlock } {}^{p+1}\Omega_B$. Local view of the Mesh depicts nearest-neighbor MeshBlock connectivity and physical levels. Nodes over ${}^p\Omega_A$, and ${}^{p+1}\Omega_B$ together with coarse analogues are shown. Data that are to be sent are marked by “●” in dark green; these data are received and directly populate the ghost-nodes marked by “■” in dark green. Once the remaining data for the coarse representation of ${}^{p+1}\Omega_B$ – marked by “□” – are filled, and any multiplicity conditions (here suppressed) are accounted for, prolongation can be performed to fill the values at the ghost-nodes of ${}^{p+1}\Omega_B$ marked by “■” in purple. For this procedure data at nodes on the neighbor interface remain unchanged.

Communication of VC variables: refined grids Consider now a Mesh with refinement; as mentioned above, in this case neighboring MeshBlock objects may differ by (at most) a single physical level (see also §3.1.1). In the following, a MeshBlock at physical level p is denoted by ${}^p\Omega_j$. For communication purposes, when refinement is present, a coarse representation of data (hereafter coarse data) is stored for each MeshBlock . The latter consists of $(\lfloor n^{\text{MB}}/2 \rfloor + 1)^3$ samples further extended by a coarse ghost-layer comprised of \mathcal{N}_{cg} nodes. In contrast to CC and FC as implemented in Athena++, in GR-Athena++ the implementation of VC allows for \mathcal{N}_{g} and \mathcal{N}_{cg} to take odd values and be independently varied. For simplicity of discussion, hereafter $\mathcal{N}_{\text{g}} = \mathcal{N}_{\text{cg}}$ is selected.

When a Mesh involves multiple physical levels, prior to any communication of data, VC variables are initially restricted so as to have a fundamental and coarse description on each MeshBlock excluding the ghost-layers. For Cartesian grids in particular, this turns out to be an exact operation (see below). For neighboring MeshBlocks at the same physical level, data and coarse data are communicated as in the case of uniform grids.

To describe how communication is handled when neighboring physical levels differ,

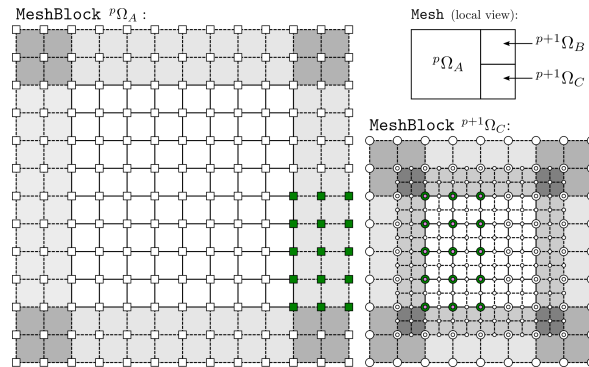


FIGURE 3.4: Schematic of two-dimensional $\text{MeshBlock } p+1\Omega_C$ used to populate ghost-layer of coarser $\text{MeshBlock } p\Omega_A$. Locally the Mesh has the same structure as in Fig. 3.3. Nodes over $p\Omega_A$, and $p+1\Omega_C$ together with coarse analogues are shown. Before the communication process starts, data of $p+1\Omega_C$ at nodes marked by “●” in purple are restricted to populate coarse data of $p+1\Omega_C$ at nodes marked by “●” in dark green. The latter are then sent and populate the nodes of $p\Omega_A$ marked by “■” in dark green. During this procedure data at nodes on the neighbor interface are (additively) updated (cf. Fig. 3.3) and multiplicity conditions (here suppressed) dynamically updated in an auxiliary array.

consider a two-dimensional Mesh where n^{MB} and $\mathcal{N}_g = 2$, in a region away from the physical boundary. Suppose $p\Omega_A$ is neighbored by $p+1\Omega_B$ and $p+1\Omega_C$ to the east and the latter two MeshBlock objects share a common edge. Data need to be communicated from $p\Omega_A$ to $p+1\Omega_B$ and $p+1\Omega_C$ (coarse-to-fine) and from $p+1\Omega_B$ and $p+1\Omega_C$ to $p\Omega_A$ (fine-to-coarse).

Fig. 3.3 shows how the ghost-layer nodes of the finer $p+1\Omega_B$ are populated. First, data are communicated from $p\Omega_A$ to the ghost layer of the coarse data representation of $p+1\Omega_B$. A blocking condition arises here, because **all** coarse ghost-layer of $p+1\Omega_B$ need to be populated before prolongation can be performed on the target MeshBlock . Nodes that are at the interface of MeshBlocks are maintained at the value of the finer level.

Fig. 3.4 depicts the dual operation, in which ghost-layer nodes of the coarser $p\Omega_A$ from $p+1\Omega_C$ are populated. Here, coarse data of $p+1\Omega_C$ are communicated and populate the ghost layers of $p\Omega_A$, without any blocking. However, non-trivial multiplicity conditions arise on the common neighbor interface, complicated by the fact that data from $p+1\Omega_B$ are also communicated. Furthermore, the coarse data of $p\Omega_A$ need to be updated and this requires an additional restriction in the coarse data of $p+1\Omega_C$; the restricted data are then communicated as before.

Restriction and Prolongation The communication process for the case of a refined `Mesh` requires restriction \mathcal{R} (fine-to-coarse) and prolongation \mathcal{P} (coarse-to-fine) of data. In `GR-Athena++`, for VC variables, such operations are implemented as products of univariate Lagrange polynomial interpolation along the three spatial dimension, centered about a target-point of interest.

The \mathcal{R} operation for variables sampled on VC is trivial, because the ghost-nodes of the coarser `MeshBlock` coincide with a subset of the physical nodes of the finer `MeshBlock` for two neighboring `MeshBlocks`, as shown in Fig. 3.4. Therefore, to populate the ghost-layers of the coarser `MeshBlock`, data can be simply copied. A potential issue here arises from the fact that a finer `MeshBlock` must communicate also its coarse data point to populate the ghosts of the coarse data of the coarser neighbor. This involves a double coarsening of data, which induces the constraint

$$n^{\text{MB}} \geq \max(4, 4\mathcal{N}_g - 2). \quad (3.1)$$

Regarding the \mathcal{P} operation, in this case function data are communicated to a finer, uniformly sampled grid of half the spacing and consequently interspersed nodes coincide, as depicted by Fig. 3.3. So, as in the previous case, for these nodes data points are simply copied. Lagrange interpolation is utilized for the non-coincident nodes. Due to the uniform structure of the source and target grids, interpolation at non-coincident nodes may be efficiently implemented through a weighted sum where weight factors can be precomputed [287]. The width of the interpolation stencil used in the code is $N = \lfloor \mathcal{N}_{\text{cg}}/2 \rfloor + 1$.

Final remarks of this section are that local operations (in the absence of distributed, MPI communications) are performed locally in memory and that the base CC and FC variables of `Athena++` continue to simultaneously function as described in Stone et al. [283].

3.1.3 Spacetime evolution

`GR-Athena++` implements the Z4c system Eqs. (A.13) (see also §A.2 in App. A) following the free evolution approach, in which the dynamical constraints of Eqs. (A.17) are not enforced, but their violation is used to assess the overall error during the evolution. On the other hand, the algebraic constraints Eqs. (A.18) are continuously enforced at each time-integrator substep (see §A.2.1), since in the absence of algebraic constraint the Z4c system is only weakly hyperbolic [288].

The numerical evolution is based on the method of lines approach. VC or CC representation of field variables may be chosen at compile time and time evolution can be performed with several time integrators described in Stone et al. [283]. Typically, the fourth order in time, four stage, low-storage RK4()4[2S] method of Ketcheson [289] is used in this work. Unlike other spacetime solvers, `GR-Athena++` features a global time step throughout the whole grid, even in presence of refined grids. In the latter case, the time step is determined by the spatial resolution on the most refined level and the choice of the CFL factor, which must satisfy the CFL condition during the evolution.

Spatial field derivatives in the bulk of the domain (away from $\partial\Omega$) are discretized with high-order, centered, finite difference (FD) stencils whereas shift advection terms use stencils lop-sided by one grid point [270, 273, 290, 291]. The implementation of the operations in the right-hand-side (RHS) of the equations is based on Alfieri et al. [292] and utilizes C++ templates to offer flexibility in problem-specific accuracy demands without loss of performance. The \mathcal{R} and \mathcal{P} operators discussed in §3.1 are similarly implemented. In the calculations involving FD and \mathcal{R} and \mathcal{P} operations, arithmetical operations are carried out in such a way to enforce the associativity property of addition, reducing the accumulation of small floating-point differences. A consistent, overall, formal order throughout the bulk of the computational domain is maintained during calculations by compile-time specification of the ghost-layer through choice of \mathcal{N}_g together with \mathcal{N}_{cg} . Throughout this work $\mathcal{N}_g = \mathcal{N}_{cg}$ is chosen, though within `GR-Athena++`, for a VC discretization, this is not a requirement. Finally, in order to suppress high-frequency numerical artifacts generated at `MeshBlock` boundaries and not present in the physical solution, the high-order Kreiss-Oliger dissipation [293, 294] of uniform factor σ over all levels is implemented. In particular, given a system of time-evolution equations for a vector of variables \mathbf{u} , the replacement $\partial_t[\mathbf{u}] \leftarrow \partial_t[\mathbf{u}] + \sigma\mathcal{D}[\mathbf{u}]$ is made, where $\mathcal{D}[\cdot]$ is proportional to a spatial derivative of order $2\mathcal{N}_g - 2$. In the case of Z4c and VC the overall order for spatial discretization in the bulk can be therefore calculated as $2(\mathcal{N}_g - 1)$.

Physical boundary conditions Artificial introduction of a boundary of the domain $\partial\Omega$ at finite distance complicates the analysis of numerical stability significantly [295]. `GR-Athena++` extends `Athena++` by providing the Sommerfeld BC [296]. Practically, BC are implemented as follows. Within every time-integrator substep an initial Lagrange extrapolation is performed so as to populate the ghost-layer at $\partial\Omega$. The order of this operation is again controlled at compile-time and typically $\mathcal{N}_g + 1$ points

are used for the extrapolation. The fields $\{\chi, \tilde{\gamma}_{ij}, \alpha, \beta^i\}$ on nodes of $\partial\Omega$ are populated using Eqs. (A.13) and the gauge conditions, whereas for $\{\hat{K}, \tilde{\Gamma}, \Theta, \tilde{A}_{ij}\}$ Sommerfeld BC are imposed as in Hilditch et al. [296]. Here, first order spatial derivatives are approximated through second order accurate, centered FD; this has proved to be crucial for numerical stability. Though this procedure is not optimal, as it is not constraint-preserving, no issues are found in this work on account of this.

All the implementation details discussed so far have been tested with the ‘‘Apples with Apples’’ test-bed suite [288, 297–299], a discussion of which is provided in the appendix of Daszuta et al. [300].

Wave extraction In GR-Athena++ wave extraction is performed by means of geodesic spheres, which are described in App. C. The Weyl scalar Ψ_4 is first calculated at all grid points of each `MeshBlock` according to Eq. (B.1) and then interpolated on the vertices of a set of geodesic spheres Q_R at given extraction radii R_Q . The multipoles $\psi_{\ell m}$ are computed according to Eq. (B.4) by performing the integrals on the geodesic spheres as described in App. C. Finally, the GW strain is obtained from the projected Weyl scalar as described in App. B, according to Eqs. (B.9), (B.10), (B.11).

3.2 Mesh refinement for punctures

GR-Athena++ makes use of the puncture formalism to model BHs, as done in BAM [273]. In NR, BHs can be treated by an initial geometry induced by the Brill-Lindquist wormhole topology of N BHs with $N + 1$ asymptotically flat ends [301]. The latter are compactified and identified with coordinate singularities on \mathbb{R}^3 , called *punctures*. Such a procedure allows one to produce initial data for an arbitrary number of BHs with associated masses, momenta and spins. The main application of this formalism is BBH evolution.

3.2.1 Punctures’ initial data

The generation of initial data for BBH in GR-Athena++ follows Brüggmann et al. [273]. Initial data (γ_{ij}, K_{ij}) on a spatial hypersurface Σ are constructed from the conformal, transverse-traceless decomposition of the initial-value equations [302]. An initial conformally flat background $\tilde{\gamma}_{ij} = \delta_{ij}$ is set in Eq. (A.11) and the maximal slicing condition is adopted for the curvature, i.e. initially $K = 0$. With this choice, the momentum constraints read $\partial_j \left(\psi^6 \tilde{A}_{ij} \right) = 0$ and admit Bowen-York solutions [303]

for an arbitrary number of BHs. The Hamiltonian constraint reduces to an elliptic equation for ψ , with generic solution for N BHs centered at r_i

$$\psi_0 = 1 + \sum_{i=1}^N \frac{m_i}{2r_i} + u. \quad (3.2)$$

ψ_0 is the initial value of ψ , which is then evolved according to Eq. (A.13a). The function u in Eq. (3.2) can be determined by solving an elliptic equation on \mathbb{R}^3 and is C^2 at the punctures and C^∞ elsewhere. The parameters m_i denote the so-called *bare mass* of the BHs, which coincide with their gravitational mass only for Schwarzschild BHs. The total ADM mass of each BH measured at the puncture is given by:

$$M_i = m_i \left(1 + u_i + \sum_{i \neq j} \frac{m_j}{2d_{ij}} \right), \quad (3.3)$$

where u_i is the value of u at each puncture and d_{ij} is the coordinate distance between a given pair of punctures. Ultimately, $M := \sum_i^N M_i$ denotes the total mass of the system and represents the physical scale of the problem.

BBH initial data for `GR-Athena++` are produced with an external `C` library based on the pseudo-spectral approach of Ansorg, Brügmann, and Tichy [304] which is also used in the `TwoPunctures`² thorn of `Cactus`.

3.2.2 Puncture tracker

The punctures' positions are tracked by solving an additional set of ordinary differential equations. Eq. (A.13a), together with $\chi|_{\mathbf{x}_p} = 0$, imply [55]

$$\dot{\mathbf{x}}_p(t) = -\beta|_{\mathbf{x}_p}(t). \quad (3.4)$$

This vectorial equation is solved after every integration time step using an explicit Euler solver. Though `BAM` implements the higher order Crank–Nicolson integration to solve Eq. (3.4), comparisons of BBH evolutions performed with the two codes show a good agreement in the trajectories of the two punctures in BBH evolutions, as can be seen in the left panel of Fig. 3.8.

²The publicly available code was cast into a stand-alone library which may be found at the URL <https://bitbucket.org/bernuzzi/twopuncturesc/>.

3.2.3 Oct-tree box-in-box




FIGURE 3.5: 2D slice at $z = 0$ of a mesh grid produced with GR-Athena++ for a BBH run. Left panel: BHs at initial position $\mathbf{x}_{p,\pm} = (\pm 3.257, 0, 0) M$. Right panel: snapshot at $50 M$ in which the trajectories of the punctures are shown in red. The colormap refers to the conformal factor χ . The black squares represent the edges of the MeshBlocks, showing the box-in-box structure of the grid. The numbers represent the refinement levels to which selected MeshBlocks belong. In the right panel such numbers are coloured in blue if the corresponding MeshBlocks are the result of refinement or in red if they result from de-refinement with respect to the left panel. For clarity, in the figure only a subset of the total slice covering the two punctures is shown, therefore only the highest refinement levels are visible.

In Athena++, the AMR is carried out by evaluating a certain refinement criterion on each MeshBlock. Depending on the result, the given MeshBlock is refined, derefined or untouched. The AMR criterion for puncture evolution implemented in GR-Athena++ mimics the box-in-box refinement (used e.g. in BAM and Cactus), in which the grid is partitioned with nested boxes that are more refined in the proximity of each puncture. The positions of the punctures in the domain Ω induce a schematic structure of nested boxes, which can be mapped into refinement levels, such that the smallest of these boxes contain the punctures and so define the highest refinement level. The boxes containing these correspond to one less refinement level and so on up to the 0th level, which coincides with the initial Mesh. The refinement criterion adopted is based on the positions of the punctures, determined by solving Eqs. (3.4) during the spacetime evolution.

Unless otherwise stated, hereafter Ω is selected as the Cartesian coordinatized cube $[-x_M, x_M]^3$. It is possible to assess to which refinement level a given `MeshBlock` belongs based on its distance from the punctures. The latter is calculated as $d_{\text{MB}} := \min_i \|\mathbf{x}_p^i - \mathbf{x}_{\text{MB}}\|_\infty$, where \mathbf{x}_p^i is the position of the i^{th} puncture with respect to the center of the grid. \mathbf{x}_{MB} denotes the position of the target `MeshBlock`, which for implementation reasons is defined as follows. Consider a rectangular cuboid concentric to the target `MeshBlock` but with edges $1/4$ the ones of the original `MeshBlock`. \mathbf{x}_{MB} are the coordinates of the corner of such cuboid closest to the closest puncture. The *theoretical* refinement level of the `MeshBlock` can be calculated as $\lfloor \log_2 \left(\frac{x_M}{d_{\text{MB}}} \right) \rfloor^3$. If the current level of the `MeshBlock` does not coincide with the theoretical one, then the `MeshBlock` is either refined or de-refined.

Practically, to define a grid in `GR-Athena++` one needs to specify the number of points constituting the initial `Mesh` and the `MeshBlocks`, i.e. the tuples N_M and N_B , respectively, and the maximum number of refinement levels N_L . Following the procedure described above, `GR-Athena++` produces a `MeshBlock`-based *oct-tree box-in-box* grid structure. Fig. 3.5 shows a visualization of such a grid with $n^M = 64$, $n^{\text{MB}} = 16$ and $N_L = 11$ for a BBH run, comparing two snapshots at different times. The initial mesh (level 0) is divided into 4^3 `MeshBlocks`. Following the procedure above, the initial `MeshBlock` configuration is refined up to refinement level 10. The final grid is composed of some of the initial `MeshBlocks`, farthest from the punctures and thus untouched, and increasingly smaller blocks (in terms of physical extent) getting closer to the two punctures. In the left panel of Fig. 3.5 are visible `MeshBlocks` belonging to levels 7, 8, 9, 10. In the right panel it is shown how, due to regriding happening during the evolution, some `MeshBlocks` initially belonging to refinement levels 7, 8, 9 have been subdivided, and their children belong to level 8, 9, 10, respectively. Analogously `MeshBlocks` initially in refinement level 10 (left panel) are de-refined and their children belong to refinement level 9 (right panel).

3.2.4 Grid configurations

In order to *accurately* and *efficiently* perform BBH merger simulations it is crucial to optimize the grid configuration for a given problem. To attain accuracy at reduced computational cost, a balance must be struck such that the strong-field dynamics are well-resolved, their effects propagate cleanly into the wave-zone for GW extraction and

³The \log_2 in the formula is a direct consequence of the requirement that neighboring `MeshBlock` differ at most by one refinement level

the maximum spatial extent x_M of the computational domain Ω is sufficiently large so as to mitigate any potential spurious effects due to the imposition of approximate physical BC. To adequately resolve the strong-field dynamics, a target resolution at the puncture δx_p can be selected with the appropriate choice of the parameter N_L for a given x_M . For the wave-zone, the resolution δx_w of the refinement level that contains the extraction sphere must be sufficiently high.

The resolution on the most refined level in the vicinity of a puncture for the grid described in §3.2.3 can be calculated as

$$\delta x_p = \frac{2x_M}{n^{M2^{N_L-1}}}. \quad (3.5)$$

The latter expression may be used to compute the wave-zone resolution δx_w with $N_L = N_L^w := \lceil \log_2 \left(\frac{2x_M}{R} \right) \rceil$, where R is the extraction radius.

3.3 Puncture tests

In this section GR-Athena++ is validated through a set of puncture evolutions tests. The results are compared against benchmark results from the BAM code and TEOBResumS. Additionally, the convergence properties of the code are demonstrated.

For the calculations presented in this work, the MeshBlock size is typically chosen as $n^{\text{MB}} = 16$. According to the constraint of Eq. (3.1), this choice allows for up to 6th order accuracy for approximants to operators pertaining to quantities appearing during spatial discretization. The gauge is set as explained in §A.2.2 of App.A.

Unless otherwise stated, throughout this section tortoise coordinates are employed, in which the evolution time t is mapped as $t \rightarrow u := t - r^*$, where $r^* = r + 2M \log \left| \frac{r}{2M} - 1 \right|$ and M , r are the total mass of the system and the Schwarzschild coordinate, respectively. In waveform plots for BBH with component masses M_1 , M_2 , quantities are suitably rescaled by M and by the symmetric mass ratio $\nu := \frac{M_1 M_2}{(M_1 + M_2)^2}$. The *time of merger* is defined as the time corresponding to the peak of the (2, 2)-mode of the strain amplitude $A_{\ell m}$, defined in Eq. (B.11).

3.3.1 Single Spinning Puncture

A first series of tests to verify the correctness of the spacetime evolution and of the GW extraction consists of a single spinning BH puncture evolution; the results are compared against the BAM code. Initial data are generated with the TwoPunctures




FIGURE 3.6: $\Re(Rh_{20}/M)$ for a single spinning puncture in GR-Athena++ and BAM with difference shown in black. Wave extracted at $R = 50M$.

library for both codes. A single spinning puncture representing a Kerr BH with mass $M = 1M$ and dimensionless spin parameter $a = 0.5$ is obtained by initializing two BHs with a small initial separation, $10^{-5}M$, masses $M_1 = 1M$ and $M_2 = 10^{-12}M$, dimensionless spins $a_1 = 0.5$, $a_2 = 0$. The two BHs merge soon after the simulation begins, producing the target Kerr BH.

A refined grid centered around the puncture is used, with a resolution at the puncture $\delta x_p = 0.0833M$ and at the wave-zone $\delta x_w = 0.667M$ respectively, matching the ones used in the corresponding BAM evolutions. The two GW signals are compared by calculating the dominant $(2, 0)$ mode of the strain, h_{20} . The double integration required to obtain the strain may add an arbitrary quadratic polynomial in time onto the strain as constants of integration [305, 306]; in the results presented here, a fit for this quadratic is obtained and then subtracted. Note that this reconstruction has been shown to introduce errors in the waveform ringdown [307].

Fig. 3.6 shows the match between the two calculated signals for $\Re(h_{20})$. These show consistency in the phasing of the signal, with slight discrepancies in the strain amplitude.

To demonstrate the convergence properties of the waveforms in GR-Athena++, the same simulation is performed at a coarse, medium and fine resolution, with finest grid spacings respectively $\delta x_c = 0.0250M$, $\delta x_m = 0.0208M$, $\delta x_f = 0.0156M$. The scaling factor Q_n for n^{th} order convergence can be calculated as

$$Q_n = \frac{\delta x_c^n - \delta x_m^n}{\delta x_m^n - \delta x_f^n}. \quad (3.6)$$




FIGURE 3.7: Left panel: convergence plot for GR-Athena++, comparing the difference between waveforms at coarse and medium resolution ($\delta x_c - \delta x_m$) and the difference between waveforms at medium and fine resolution ($\delta x_m - \delta x_f$) rescaled by the appropriate factor $Q_4 = 1.57052$. Right panel: the difference between the Richardson extrapolated waveform for BAM and GR-Athena++ is consistent with the combined error arising from the finite resolution data. In both plots, waves are extracted at $R = 50M$

The left panel of Fig. 3.7 shows the convergence properties for the case of 4th order accurate FD operators. Assuming 4th order convergence, the rescaled waveform (in orange) at early times lies below the unrescaled waveform (in blue). This means that a larger factor Q_n (corresponding to a higher order of convergence) would provide a more comparable rescaling. After $u/M \approx -15$, the rescaled curve is consistent with the unrescaled curve, suggesting 4th order convergence, though this is not a point-wise agreement. Similar tests performed with the BAM code in Hilditch et al. [296], with initial data constructed in the same manner, show analogous convergence properties for the waveform of the spinning puncture problem.

The error associated to finite resolution effects can be estimated by performing a Richardson extrapolation for both BAM and GR-Athena++ waveform data. Assuming 4th order convergence, the Richardson-extrapolated waveforms $h_{20,\text{Rich.}}$, and associated error $e = |\Re(Rh_{20,\text{Rich.}}) - \Re(Rh_{20,m})|$ are obtained comparing to the strain of the medium resolution run $\Re(Rh_{20,m})$. The Richardson extrapolation is performed by comparing data with puncture resolution $\delta x_l = 0.08333M$ to a higher resolution run with $\delta x_h = 0.08000M$. The left panel of Fig. 3.7 shows that the combined error arising from the finite resolution data is consistent with the discrepancy between the BAM and GR-Athena++ waveforms.




FIGURE 3.8: Left panel: comparison of the puncture’s trajectories between BAM and GR-Athena++; right panel: comparison of the multipole ψ_{22} obtained from the two codes at resolution $n^M = 384$. Extraction is performed at $R = 120 M$. The maximum discrepancy between BAM and GR-Athena++’s strain amplitudes is $\lesssim 1\%$.

3.3.2 Calibration evolution of two punctures

Next, a set of BBH evolutions are compared between GR-Athena++ and BAM, and convergence properties of GR-Athena++ therein are assessed. For these tests, two non-spinning BHs with initial momenta directed along the y -axis $\mathbf{p}_\pm = (0, \mp 0.133, 0) M$ and the same bare-mass $m_\pm = 0.483 M$ are initially placed on the x -axis at $x_\pm = \pm 3.257 M$. This setup results in a ~ 2.5 orbits evolution of the two BHs before merger, as can be seen in the right panel of Fig. 3.8. For the GR-Athena++ vs. BAM comparison and for convergence tests, several runs at different resolutions have been performed. The grid configurations for both codes are described in detail below.

GR-Athena++ vs. BAM comparison GR-Athena++ and BAM are compared with sets of simulations performed at 4th order FD. Athena++ and BAM implement completely different grid structures, thus for meaningful comparisons, grids as similar as possible are generated. The size of the computational domain is chosen as $x_M = 3072M$, whereas the resolution at the puncture for both sets of simulations is $\delta x_p \approx \{1.56, 1.17, 0.781, 0.586, 0.391\} \times 10^{-2}M$ and the resolution at the wave-zone δx_w matches for the two cases.




FIGURE 3.9: Comparison of the Weyl scalar's multipole ψ_{22} between BAM and GR-Athena++ for two different resolutions. The discrepancy between the amplitudes decreases with increasing resolution. Extraction is performed at $R = 120 M$.

For GR-Athena++, such a configuration is obtained with the grid parameters set as $n^{\text{MB}} = 16$, $N_L = 11$; as a consequence of Eq. (3.5), the resolution at the puncture is controlled by n^{M} , which for these tests assume the values $n^{\text{M}} = \{96, 128, 192, 256, 384\}$.

For BAM similar grid structures are achieved in each corresponding simulation by considering 6 nested outer boxes with the same value of n^{M} chosen in GR-Athena++ case and 5 smaller boxes centered around each puncture (for a total of 10 smaller boxes) with $n^{\text{M}}/2$ points along each side; the desired values of δx_p are induced by setting maximum spacings in the outermost grids $\Delta x = \{96, 64, 48, 32, 24\} M$.

Fig. 3.8 shows very good agreement among the trajectories of the two BHs evolved with the two codes (left panel). Analogously, a very good match is also obtained looking at the phase and amplitude of the Weyl scalar's multipole ψ_{22} . A maximum discrepancy $\approx 1\%$ is found between the two codes at the highest resolution considered. However, Fig. 3.9 shows that such difference converges away with resolution.

Convergence tests for GR-Athena++ Convergence tests for GR-Athena++ are performed on the set of runs described above and on an additional set of simulations which employ 6th order FD; in the latter set, a grid setup similar to the 4th order FD runs is used, but $x_M = 1536M$. This effectively doubles the values of δx_p .

The two top panels of Fig. 3.10 compare the GW strain extracted at $R = 120 M$ at different resolutions, for 4th (left) and 6th (right) order FD. In the bottom panels of




FIGURE 3.10: Convergence plots for calibration BBH evolution. Left and right plots correspond respectively to 4th order and 6th order FD. In both case waves are extracted at $R = 120M$. In bottom panels phase differences between resolutions are rescaled according to Eq. (3.6) by the appropriate factor, calculated with respect to the blue line (corresponding to lowest resolutions) and the actual numerical values can be read off in the legend.

Fig. 3.10 the GW phase differences are reported, which are used for convergence tests. Inverting Eq. (B.11) allows one to define the phase differences between two waveforms a, b as:

$$|\Delta\phi(a, b)| := |\phi[h_{22}]|_a - \phi[h_{22}]|_b|. \quad (3.7)$$

In the left-bottom panel of Fig. 3.10 the red and green lines match, demonstrating that 4th order convergence is achieved for the highest two resolutions. In the right-bottom panel of Fig. 3.10 the waveforms (and corresponding frequencies) appear very similar to each other and this translates into smaller phase differences comparing with the bottom panel of the left plot. Even though the red line, corresponding to the phase difference between the two highest resolutions, is quite noisy, 6th order convergence is




FIGURE 3.11: Self-convergence test for the calibration BBH evolution. On the y -axis the phase differences at merger with respect to the highest resolution run are reported. Diamonds correspond to 4th order series, while dots refer to 6th order series. Purple and cyan lines represent the theoretical convergence for both cases, respectively.

achieved for the three highest resolutions as in the previous case. For all extraction radii considered (not shown in these plots), the phase differences are convergent, i.e., the differences between each pair of lines decreases with increasing resolution.

Similarly to §3.3.1, the accuracy of the convergence tests performed here is checked by evaluating the error on the phase differences, estimated for each line as the difference between the phase given by the Richardson extrapolation formula and the phase corresponding to the highest resolution used to calculate the Richardson extrapolated phase [see also 308]. For each corresponding phase difference line, this error is always $\lesssim 50\% \Delta\phi$.

Finally, Fig. 3.11 shows a convergence plot in which phase differences are calculated at merger. This figure further confirms the clean 4th order convergence of GR-Athena++ for the highest resolution cases. For the 6th order case, comparing with the other, phase differences are smaller and the convergence is faster. However, the noise in the phase differences with respect to the highest resolution discussed above makes the convergence assessment less clean.

3.3.3 Two punctures evolution of ten orbits

The inspiral of BBH in most physical scenarios is well described by quasi-circular trajectories of the two BHs [1, 2]. This is consistent with the detections made so far by LIGO and Virgo detectors [7, 8, 46, 178, 309]. Consequently it is of considerable interest to also test the performance of GR-Athena++ in this scenario.

TABLE 3.1: A distinct label $\rho_{(\cdot)}$ is assigned to each run with corresponding maximum number of refinement levels N_L and fixed physical extent x_M of the `Mesh` (see §3.2.4). Resultant puncture resolutions δx_p and total number of `MeshBlock` objects initially partitioning a `Mesh` are provided.

$\rho_{(\cdot)}$	N_L	x_M	$\delta x_p \times 10^{-2}[M]$	#MB
vv1	10	768	4.6875	1072
vl	11	1152	3.515625	1352
l	11	768	2.34375	1184
ml	12	1152	1.757812	1464
m	12	768	1.171875	1296
mh	13	1152	0.878906	1576
h	13	768	0.585938	1744

To this end non-spinning, equal-mass, low-eccentricity initial data based on Hannam et al. [310] are evolved; bare-mass parameters are $m_{\pm} = 0.488479 M$ and the punctures are initially on-axis at $x_{\pm} = \pm 6.10679 M$ with instantaneous momenta $\mathbf{p}_{\pm} = (\mp 5.10846 \times 10^{-4}, \pm 8.41746 \times 10^{-2}, 0) M$. This choice of parameters results in ~ 10 orbits prior to merger at $t \approx 2145 M$. In comparison to the calibration evolutions of §3.3.2, this evolution is of significantly longer duration and therefore it is of interest to investigate how waveform accuracy is affected for a selection of `Mesh` parameters that reduce computational resource requirements. In order to provide another comparison that is independent of `BAM`, here the quality of waveforms computed with `GR-Athena++` is assessed against the NR-informed, EOB model `TEOBResumS` [146].

Setup The convergence studies performed for the calibration BBH merger problem provide a guide as to how to choose resolution at the puncture δx_p . Here the `MeshBlock` sampling is $n^{\text{MB}} = 16$ and 6th order FD is used. Regarding the `Mesh` sampling, $n^{\text{M}} = 64$ and a sequence of grid configurations is constructed in such a way that each value of δx_p is reduced by a factor of 3/2 compared to the previous, as reported in Tab. 3.1.

That the choice of parameters in Tab. 3.1 reduce overall computational resource requirements can be understood as follows. Consider the choice of parameters made in ρ_{h} and suppose n^{M} and N_L are varied while maintaining δx_p fixed. With this, the number of `MeshBlock` objects required to partition the initial `Mesh` changes. For example, taking $n^{\text{M}} = 128$ and $N_L = 12$ results in an initial number of `MeshBlock` objects of $\# \text{MB} = 8352$, whereas selecting $N_M = 256$ and $N_L = 11$ leads to $\# \text{MB} = 58752$ initially. Generally, it is found that $\# \text{MB}$ scales approximately cubically with




FIGURE 3.12: Coordinate trajectories of both punctures ($\mathbf{x}_+(t)$ in blue and $\mathbf{x}_-(t)$ in orange) for parameter choice ρ_h .

n^M , which is related to the dimensionality of the problem.

In this section an extraction radius of $R = 100 M$ is used. The CFL condition is 0.25 and a Kreiss-Oliger dissipation of $\sigma = 0.5$ is chosen. Constraint damping parameters are selected as $\kappa_1 = 0.02$ and $\kappa_2 = 0$.

The coordinate trajectories of the punctures for a calculation utilizing grid parameters ρ_h of Tab. 3.1 can be seen to satisfy ten orbits in Fig. 3.12. This provides an initial verification of expected qualitative properties [310] of the BBH inspiral and merger. The quality of the evolution is evaluated by assessing the violation of the constraints of Eq. (A.7). In particular, the collective constraint is used here, defined from the expressions of the constraints of §A.2.1, as

$$\mathcal{C} := \sqrt{\mathcal{H}^2 + \gamma_{ij} \mathcal{M}^i \mathcal{M}^j + \Theta^2 + 4\gamma_{ij} Z^i Z^j}. \quad (3.8)$$

Values of \mathcal{C} in the orbital plane ($z = 0$) at fixed times $t = 500 M$ and $t = 2100 M$ are shown in Fig. 3.13. In all cases it is found that away from the punctures, values of \mathcal{C} decrease on average as the physical boundary $\partial\Omega$ is approached. In particular, for the calculation involving parameters ρ_h and during $500 M \leq t \leq 2200 M$ as $\varrho := \sqrt{x^2 + y^2} \rightarrow 100 M$ it is found that $\mathcal{C} \sim 10^{-8}$ which continues to decrease as $\varrho \rightarrow 300 M$ to $\mathcal{C} \sim 10^{-10}$ thereafter plateauing at $\mathcal{C} \sim 10^{-11}$ towards the boundaries. The general properties of \mathcal{C} discussed here are shared between other simulations utilizing parameters from Tab. 3.1. Crucially, this means that increasing refinement in




FIGURE 3.13: Values of the (normalised) collective constraint $\widehat{\mathcal{C}}(x, y, z) := \mathcal{C}(x, y, 0) / \max_{x, y} \mathcal{C}(x, y, 0)$ over the orbital plane $z = 0$ for a simulation with ρ_h of Tab. 3.1. Left panel: evolution time is $t = 500 M$ and $\max_{x, y} \mathcal{C}(x, y, 0) \simeq 111.3$. Right panel: evolution time is $t = 2100 M$ and $\max_{x, y} \mathcal{C}(x, y, 0) \simeq 3.3$. As can be seen in both cases, constraint violation is greatest in the vicinity of the punctures.

the vicinity of the puncture does not contaminate the rest of the physical domain. Qualitatively similar behaviour is found for the Hamiltonian constraint \mathcal{H} .

Next, the GW strain is investigated for all choices of Tab. 3.1. To this end Eq. (B.10) for h_{22} is solved in the frequency domain making use of the fixed frequency integration method of Reisswig and Pollney [311]. A cut-off frequency of $f_0 = 1/300 \times 3/4$ is chosen, which is physically motivated by inspecting the early time puncture trajectories of the inspiral. The resulting h_{22} are displayed in Fig. 3.14.

For a grid parameter choice of ρ_h , merger occurs at $u = 2037.5 M$. The maximum deviation from this value for the parameters investigated in Fig. 3.14 is obtained for ρ_m , resulting in $\Delta u = 10.3 M$. In order to directly quantify how the choice of δx_p affects the phase error in the strain waveform as merger time is approached, $\Delta\phi$ is computed according to Eq. (3.7) and the results are shown in Fig. 3.15.

In Fig. 3.15 (upper panel) clean 6th order convergence in $\Delta\phi$ is not found for all u upon rescaling with the appropriate factors determined through Eq. (3.6). An additional issue that complicates the discussion here is that the parameters of Tab. 3.1 also vary the spatial extent of the computational domain potentially introducing a source of systematic error. For example, at merger $\Delta\phi(\rho_m, \rho_h) \simeq 4 \times 10^{-1}$ and $\Delta\phi(\rho_m, \rho_h) \simeq 6 \times 10^{-3}$, though $\delta x_p(\rho_m) > \delta x_p(\rho_m)$. In order to compensate for this effect, also phase differences at fixed x_M are considered. In particular, the lower panel of Fig. 3.15 displays phase differences at merger, where ρ_h is taken as the reference value for all comparisons. While displayed $\Delta\phi$ are compatible with a 6th order trend,




FIGURE 3.14: The (2,2) multipole of the GW strain normalised to the symmetric mass ratio $\nu = 1/4$ and M , computed for simulations based on parameters of Tab. 3.1. Peak amplitude for a choice of ρ_h occurs at $u/M = 2037.5$ which indicates the end of the inspiral [312]. Dephasing as merger time is approached reduces rapidly with increased resolution (see also Fig. 3.15 though note that the legend differs there). Note: horizontal axis scale changes at $u/M = 1500$.

the ρ_{vvl} choice of parameters is likely of too low resolution to make a robust claim on convergence order with respect to varying δx_p . However, as in these tests the global resolution over the computational domain was not modified, it is not clear what sort of convergence should be expected. Furthermore, the extent to which a time-integrator order below the order of the spatial discretization affects GW waveform quality can also be somewhat delicate (see e.g. the super-convergence discussion of Reisswig and Pollney [311]). Nonetheless, this analysis suggests that it is possible to reduce the GW phase error by accurately selecting the refinement level and hence resolution at the punctures δx_p .

EOB comparison As a last test, the GWs from the 10 orbit simulation are compared to the state-of-the-art EOB model `TEOBResumS` [146]. The latter is a waveform model informed by several existing NR datasets and faithfully models the two-body




FIGURE 3.15: Phase differences $\Delta\phi$ between simulations involving parameters of Tab. 3.1. Upper panel: A trend of $\Delta\phi$ accumulating with time is present. Merger time corresponding to ρ_h is indicated with a vertical black line at $u/M = 2037.5$. A decrease in $\Delta\phi$ occurs as δx_p pairs of decreasing values are compared. In order to mitigate a systematic effect of varied spatial extent in the computational domain, also phase differences at fixed x_M such as $\Delta\phi(\rho_l, \rho_m)$ and $\Delta\phi(\rho_m, \rho_h)$ are computed. These two differences are also shown rescaled with Q_6 and Q'_6 , computed according to Eq. (3.6) under the assumption of 6th order spatial discretization. Lower panel: Phase differences at merger computed with reference data taken from the ρ_h run are depicted as a function of puncture resolution. Data on the black reference curve would obey a 6th order convergence trend.

dynamics and radiation of spin-aligned BBH multipolar waveforms for a wide variety of mass ratio and spin magnitudes. For this comparison, the GR-Athena++ ψ_{22} output mode is first extrapolated to null infinity using the asymptotic extrapolation formula [78, 313]:

$$\lim_{r \rightarrow \infty} r\psi_{22} \sim A \left(r\psi_{22} - \frac{(l-1)(l+2)}{2r} \int r\psi_{22} dt \right), \quad (3.9)$$

where $A(r) := 1 - 2M/r$ and $r := R(1 + M/(2R))^2$, with R being the (finite) GW extraction radius of an NR simulation. The extrapolated result is then integrated as in the previous section to obtain the GW strain. The waveform comparison is




FIGURE 3.16: Comparison between the GR-Athena++ BBH $q = 1$ waveform and the semi-analytical EOB model TEOBResumS. The plot shows the $(2, 2)$ multipole of the GW strain normalised to the symmetric mass ratio $\nu = 1/4$ and M and the instantaneous GW frequency. The time is shifted to the mode amplitude peak that approximately defines the merger time. The GR-Athena++ waveform corresponds to the highest resolution simulation (ρ_h of Tab. 3.1), extracted at coordinate radius $R = 100 M$ and extrapolated to null infinity using Eq. (3.9). Note: horizontal axis scale changes at $(t - t_{\text{mrg}})/M = -400$.

performed by suitably aligning the two waveforms; the time and phase relative shifts are determined by minimizing the L^2 norm of the phase differences [234].

Fig. 3.16 shows that the two waveforms are compatible within the NR errors. The accumulated EOB-NR phase differences are of order ~ 0.1 rad to merger and ~ 0.4 rad to the ringdown for the highest resolution GR-Athena++ simulation. The larger inaccuracy of the ringdown part is a resolution effect related to the higher frequency of the wave in this regime. The latter issue could be potentially improved by adding a refinement level so as to better resolve the BH remnant. The maximum relative amplitude difference is of order ~ 0.01 . The same comparison using the lowest resolution gives ~ 0.4 rad at merger (~ 1 rad during the ringdown) and maximum relative amplitude differences of ~ 0.01 . These results comparable favourably with those reported in Hannam et al. [314]. Overall, this analysis demonstrates that GR-Athena++ can produce high-quality data for waveform modeling.

3.4 Scaling tests

The scalability properties of GR-Athena++ are assessed with sets of weak and strong scaling tests conducted for the physical problem setup of §3.3.2. In particular, BBH evolutions of 20 RK time steps are considered, with full AMR and full production grids choosing grid parameters $n^{\text{MB}} = 16$, $N_L = 11$, $x_M = 1536 M$. Hybrid MPI and




FIGURE 3.17: Strong scaling tests for **GR-Athena++** with 4th order FD for different regimes of number of cores. Left plot: speed-up (top panel) and parallel efficiency (bottom panel) calculated with respect to the first point of the series. In each series, the second point corresponds to a theoretical speed-up by a factor of 2, the third point a factor of 4 and so on. Right plot: parallel efficiency as a function of the ratio $\frac{\#\text{MeshBlocks}}{\#\text{cores}}$.

OMP parallelization is used for these tests. These tests are performed on the cluster SuperMUC-NG at LRZ, on nodes equipped with 48 Intel Skylake Xeon Platinum 8174 cores each. The runs are performed by filling up the nodes launching 8 MPI tasks with 6 OMP threads per node. Very good scaling properties are demonstrated using up to 2048 nodes, for a total of $\sim 10^5$ cores. The results obtained compare favourably to the ones presented using the Dendro-GR code [280].

3.4.1 Scaling tests for 4th order FD

In the first set of tests **GR-Athena++** is configured with $\mathcal{N}_g = 3$, therefore utilizing 4th order FD stencils.

Strong scaling tests Strong scaling tests are conducted for the 6 resolutions determined by $n^M = \{64, 96, 128, 192, 256, 384\}$. For each resolution, a series of runs is performed in which, for each run, the number of cores is increased with respect to the previous one.

There are two constraints that induce upper and lower limits in the choice of number of nodes and cores that must be taken into account. On one hand, a minimum number of nodes must be employed for a given problem in order to have enough memory to store all the variables for the given grid configuration. On the other hand,




FIGURE 3.18: Weak scaling tests for GR-Athena++ with 4th order FD, performed compiling the code with GNU compiler and Intel compiler. The top panel reports the total CPU time measured by MPI rank 0 directly in the code through the C++ function `clock_t`. CPU time in the bottom panel is measured instead using the profiling tool `gprof`; here, times smaller than 2% of the total CPU time are neglected. Light-coloured (left) bars represent results obtained with GNU compiler, while dark-coloured (right) ones are for Intel compiler.

at most one OMP thread can be assigned to each `MeshBlock`, as a consequence of the coarse-grained parallelization strategy of `Athena++`.

The left panel of Fig. 3.17 shows that excellent strong scaling performances are obtained up to $\sim 1.5 \times 10^4$ cores, with efficiencies above 90%. For a given resolution, a degradation in performances is observed when the number of cores increases above a certain value. The right panel of Fig. 3.17 shows that the efficiency heavily depends on the ratio between `MeshBlock` number (hereafter `#MeshBlock`) and number of cores (hereafter `#cores`). For every resolution, efficiencies above 90% are obtained when $\#MeshBlocks/\#cores \gtrsim 10$. This is a consequence of the parallelization strategy of `Athena++`, in which one or more `MeshBlocks` are assigned to each OMP thread and it turns out to be more efficient when many `MeshBlocks` are handled by each thread.

Comparing these results to the strong scaling plot of Fernando et al. [280], in the




FIGURE 3.19: Strong scaling tests for GR-Athena++ with 6th order FD for different regimes of number of cores. Left plot: speed-up (top panel) and parallel efficiency (bottom panel) calculated with respect to the first point of the series. In each series, the second point corresponds to a theoretical speed-up by a factor of 2, the third point a factor of 4 and so on. Right plot: parallel efficiency as a function of the ratio $\#MeshBlocks/\#cores$.

latter case the parallel efficiency decreases faster than the highest-resolution case of Fig. 3.17, although the two results are obtained for slightly different $\#$ cores.

Weak scaling tests In order to ensure that the utilized resources increase at the same rate as the number of overall computations, weak scaling tests are performed by considering asymmetric grids in terms of N_M . The weak scaling series of runs is conducted as follows. The first run is performed on 1 node with a grid described by the tuple $N_M = (128, 64, 64)$. The second run is performed on 2 nodes, with n_y doubled ($N_M = (128, 128, 64)$). For the run on 4 nodes n_z is also doubled ($N_M = (128, 128, 128)$) and for 8 nodes n_x is doubled again ($N_M = (256, 128, 128)$). This procedure is continued up to 2048 nodes, and it allows to maintain an almost constant ratio of ~ 33 MeshBlocks/core.

Two independent series of tests are performed, once with the code compiled with the Intel compiler and once with the code compiled with the GNU compiler.

The results are displayed in Fig. 3.18. The top panel of Fig. 3.18 shows that the total CPU time per MPI task remains constant up to $\sim 10^5$ cores, demonstrating excellent weak scalability in an unprecedented regime of number of cores for a NR code. The bottom panel shows the execution time of the main computational kernels. As expected, most of the computational time is spent in the calculation of the RHS, which indicates the absence of race conditions or other bottlenecks elsewhere in the




FIGURE 3.20: Weak scaling tests for GR-Athena++ with 6th order FD, performed on SuperMUC-NG at LRZ compiling the code with the GNU compiler. The top panel reports the total CPU time measured by MPI rank 0 directly in the code through the C++ function `clock_t`. CPU time in the bottom panel is measured instead using the profiling tool `gprof`; here, times smaller than 2% of the total CPU time are neglected.

code. The discrepancy of the total CPU time between the top and bottom panels gives an estimate of the communication time among the MPI processes; the comparison between the two plots suggests that also the communication time presents a good scaling behaviour.

3.4.2 Scaling tests for 6th order FD

Additional scaling tests are performed analogously to §3.4 but for the code configured with a 6th order FD scheme ($\mathcal{N}_g = 4$). There are a few differences with respect to the previous series of tests, due to the larger overall memory involved in this case. In particular, the ratio `#MeshBlocks/#cores` needs to be reduced.

Strong scaling tests In the case of strong scaling tests, increasing the number of communicated ghosts causes a translation of the curves of Figs. 3.17, as shown in

Fig. 3.19. In the latter case it is possible to perform tests up to 1024 nodes, for a total of $\approx 5 \times 10^4$ cores. The parallel efficiency (right panel) is above 90% as long as the ratio $\#MeshBlocks/\#core \gtrsim 6$, analogously to the 4th order case.

Weak scaling tests For weak scaling tests a slightly different construction of the grid is required. The initial grid for the run on 1 node is $N_M = (64, 64, 64)$ and then the same procedure described in the case of 4th order FD is followed to conduct the test. This allows one to decrease the load to ~ 25 MeshBlocks/core. Similar weak scaling results to the 4th order case are depicted in Fig. 3.20. The larger number of communicated ghosts does not affect sensitively the percentage of time spent in ghost communications and \mathcal{R}/\mathcal{P} operations with respect to the fraction of time spent in RHS calculations, which still largely dominates. Also, the fraction of communication time among MPI processes does not get worse. Overall the CPU time per MPI task is approximately constant in both top and bottom panels, which indicates excellent parallel performances also in this case.

Chapter 4

GR-AthenaK: first performance-portable dynamical spacetime solver

The latest advances in terms of heterogeneous HPC platforms based on GPUs push towards the creation of NR codes capable of efficiently using them. Two recent examples of such facilities are the Leonardo HPC System (CINECA, Italy) ¹, and the Perlmutter supercomputer (NERSC, USA) ². The former provides 3456 computing nodes, each equipped with four NVIDIA Ampere GPUs that are expected to provide a total computational performance over 240 Pflops; the latter provides 1789 computing nodes mounting four NVIDIA Ampere GPUs for a total of ~ 60 Pflops.

The desirable property of new codebases is performance portability, i.e. the capability of achieving comparable performance across different platforms with the least amount of architecture-dependent code. In particular, this includes the possibility of executing a single codebase on both GPU accelerators and CPUs without substantial performance loss. Currently, the Kokkos [315–317] programming model is arguably the best option to achieve this goal. Kokkos is a C++ library based on templated programming and provides an easy way to write portable parallel codes through abstractions of parallel instructions and memory hierarchy.

The excellent scalability properties of Athena++, together with its intrinsic modular structure, made it the perfect candidate to be refactored through Kokkos, resulting in the performance-portable MHD solver K-Athena [318]. More recently, the AMR infrastructure of Athena++ has been extended to the standalone and general-purpose Parthenon, based on Kokkos [319].

The work presented in this *Chapter* is an effort towards the realization of GR-AthenaK, the first performance-portable dynamical spacetime solver. GR-AthenaK implements

¹<https://leonardo-supercomputer.cineca.eu/hpc-system/>

²<https://docs.nersc.gov/systems/perlmutter/architecture/>

the main features of GR-Athena++ in AthenaK, which is another Kokkos version of Athena++ still under development.

This *Chapter* is organized as follows. §4.1 introduces the Kokkos framework; in §4.2 AthenaK and its main differences with respect to Athena++ are described; §4.3 similarly details the refactorization of specific features of GR-Athena++ into GR-AthenaK; in §4.4 GR-AthenaK is tested with cross-code comparisons against GR-Athena++, and its convergence and performance properties are assessed; eventually, §4.5 discusses the challenges and feasibility of BBH evolutions with GR-AthenaK.

4.1 Kokkos framework

Focusing on the two main target execution devices, i.e. CPUs and GPUs, the two main differences in the code execution reside in the parallel execution and memory management. The Kokkos library offers three abstractions for each of these instances, allowing one to specify them with ease at compile time.

Parallel execution The first abstraction for parallel execution is the *execution space*, i.e. on which *device* the parallel regions of the code are going to be executed. Examples of possible device specifications are `serial`, OMP for CPUs, CUDA for NVIDIA GPUs. The different code structures that can be executed in parallel constitute the second abstraction, called *execution pattern*. These correspond for instance to parallel `for` loops, parallel reductions and similar constructs, that are commonly referred to as *kernels*. The third abstraction is the *execution policy*, which specifies how parallel work is organized and distributed. There are two main types of execution policies: the `RangePolicy/MDRangePolicy` and the `TeamPolicy`. The former policies are the analogue of OMP parallelization of `for` loops of the types `# omp pragma/# omp pragma collapse (N)`. The latter policy distributes the parallel work among `members` of `teams` with shared memory, allowing to exploit hierarchical parallelism and more optimal memory management for specific purposes.

Memory management Regarding memory management, the first abstraction is the *memory space*, which specifies where data are located during the execution. Ideally, execution space and *memory space* should coincide for best performances. In practice GPUs are commonly mounted on CPU nodes, which represent the *host* in Kokkos language. If memory space and execution space do not match, expensive deep copies between the device and the host typically heavily penalize the performances. Different

devices have different optimal mappings between multidimensional array indices and memory. This can be controlled through the *memory layout*, specifying whether data should be mapped by row-major order (optimal for CPUs) or column-major order (optimal for GPUs). Finally, the *memory traits* abstraction allows one to specify how data are accessed.

Despite the great flexibility offered by the six abstractions above, the programmer is not typically required to explicitly specify each one of them. While writing the code, the developer needs to appropriately choose the execution patterns and the execution policies that better suit the problem at hand. However, at compilation time the other four abstractions will be automatically specified by default as the most optimal ones, if they have not been explicitly determined. For instance, when compiling for CUDA device, this will be automatically selected as the execution space and memory space and the memory layout will be chosen as column-major order.

Finally, Kokkos provides an additional abstraction to manage the data through multi-dimensional arrays called **Views**. These constitute one of the the main building blocks of a Kokkos code and allow one to easily allocate, deallocate, copy and access data on different devices.

4.2 Kokkos version of Athena++: AthenaK

This work is based on a previous refactoring of Athena++ with Kokkos performed by the main authors of Athena++, which resulted in the new performance-portable code AthenaK. In this section the main features of AthenaK are discussed.

To write a code with Kokkos, all arrays must be defined as Views and all parallel loops have to be written as Kokkos kernels. Additionally, all support functions used in kernels must be defined as inline functions using the Kokkos macro `KOKKOS_INLINE_FUNCTION`, which allows one to automatically compile such functions for CUDA with the required `__device__` attribute.

The MeshBlock-based structure of Athena++ as described in *Chapter 3* is maintained in AthenaK, but with one fundamental difference between the two codes. The coarse-grained MeshBlock-based OMP parallelization of Athena++ is abandoned in AthenaK, because in this case it would require a large number of expensive kernel launches, slowing down the execution of the code. In AthenaK MeshBlocks are packed together in the new MeshBlockPack class and the main data structure of the code

is a five-dimensional `View` containing all `MeshBlocks` (first index) and all the evolved variables (second index) along the three spatial dimensions (last 3 indices).

`AthenaK` implements parallel kernels in two ways. The first is based on `RangePolicy` parallelizing over the indices of the 5D `Views`. The other is based on `TeamPolicy` to exploit hierarchical parallelism: the first four indices of a 5D `View` are mapped to teams. The members of each team work simultaneously on 1D buffers of memory corresponding to the last spatial index. Loops in this latter index are vectorized using the `Kokkos::TeamVectorRange` policy.

Lastly, `Athena++` makes use of persistent ghost buffers to allow asynchronous and fast communication among different MPI processes. In `AthenaK` such buffers are defined as 2D `Views` and stored in the memory space, in order to minimize deep copies between device and host. Moreover, ghost communication is handled within appropriate parallel kernels.

4.3 GR-AthenaK

The main goal of this work is to refactor `GR-Athena++` into a `Kokkos` version, by porting the Z4c system Eqs. (A.13) into `AthenaK` whilst maintaining the main features of `GR-Athena++`.

4.3.1 Memory requirements

The biggest challenge when running on CUDA devices is represented by the limited RAM capacity of most GPUs, which can hinder the execution of memory-intense problems. The most extensive GPUs available to date in terms of memory are the Ampere GPUs, with up to 80 GB of RAM.

The implementation of `GR-Athena++` for the spacetime evolution makes use of a total of 22 Z4c evolved variables, 13 ADM variables, used for initial data and to impose the algebraic constraints and 7 `Con` variables that are used to monitor the constraint violation during the evolution. For the evolved variables, 3 storages are required for the RK time integration. When using refined grids (or multigrid), extra storage is required for ghost communication of evolved variables. This storage holds half of the points per side for each `MeshBlock` (see §3.1.2). In the following, cubic `MeshBlocks` with size n^{MB} are considered. The memory per `MeshBlock` required to perform a spacetime evolution can be calculated as

$$\mathcal{M} = B \cdot \left[(n^{\text{MB}} + 2\mathcal{N}_g)^3 \cdot (3 \cdot 22 + 13 + 7) + 22\sigma \left(\frac{n^{\text{MB}}}{2} + 2\mathcal{N}_g \right)^3 \right], \quad (4.1)$$

where B is the number of bytes of the floating point representation chosen and $\sigma = 0$ for unigrid case and $\sigma = 1$ for the multigrid case. The term proportional to σ takes into account the additional coarse representation storage for multigrid case.

The design of **AthenaK** for communication between **MeshBlocks** inherited from **Athena++** requires the allocation of extra memory buffers, as mentioned at the end of §4.2. For uniform grids, the total memory associated to such buffers simply corresponds to the shell of ghost points surrounding the physical points of **MeshBlocks**. For refined grids, a given **MeshBlock** can have more neighbors than the unigrid case (see Figs. 3.3, 3.4), belonging to finer refinement levels. In this case additional buffers are stored for each one of the possible neighbors. The buffer memory per **MeshBlock** amounts to

$$\mathcal{M}_{\text{Buffer}} = 2 \cdot 22 \cdot B \left\{ (1 + \sigma) \left[(n^{\text{MB}} + 2\mathcal{N}_g)^3 - (n^{\text{MB}})^3 \right] - 8\sigma (\mathcal{N}_g)^3 \right\} \quad (4.2)$$

where the factor 2 in front accounts for both receiving and sending buffer points, which are stored separately and only the 22 Z4c evolved variables have associated buffers. The difference between the two cubic terms calculates the shell of ghost points constituting the buffers in unigrid ($\sigma = 0$). In the multigrid case ($\sigma = 1$) this quantity is almost³ doubled due to the particular implementation discussed above. The total amount of memory required for a spacetime evolution is thus

$$\mathcal{M}_{\text{tot}} = (\mathcal{M} + \mathcal{M}_{\text{Buffer}}) \cdot N_{\text{MB}}(n^{\text{MB}}), \quad (4.3)$$

where $N_{\text{MB}}(n^{\text{MB}})$ is the total number of **MeshBlocks**, which for a given mesh is a (decreasing) function of their size n^{MB} . For instance, a spacetime evolution with double precision ($B = 8$), 6th order FD ($\mathcal{N}_g = 4$), 1024 cubic **MeshBlocks** ($N_{\text{MB}} = 1024$) with 32^3 points ($n_x^{\text{MB}} = n_y^{\text{MB}} = n_z^{\text{MB}} = 32$) for a refined grid ($\sigma = 1$) requires $\mathcal{M}_{\text{tot}} \approx 65.1$ GB of storage.

Eq. (4.3) implies that to fit the memory constraints of GPUs it is important to increase n^{MB} for a given configuration, to decrease N_{MB} and especially the memory used to store buffer ghosts $\mathcal{M}_{\text{Buffer}}$. Differently from **Athena++**, decreasing N_{MB} does

³Note the last term subtracted in the multigrid case.

not lead to a performance hit, due to the different parallelization strategy, and has the further advantage of reducing the overall communication. The limit in decreasing the number of `MeshBlock` is dictated by the use of the AMR, which requires a sufficiently fine partition of the grid in order to adequately resolve the desired features.

4.3.2 AthenaTensor views

In the `Z4c` system, numerous quantities are mathematically defined as tensors with specific symmetries, requiring common tensor operations in the RHS such as contractions and index raising. Tensor symmetries can be used in the code to make calculations more efficient and to save memory space. To achieve this in `GR-AthenaK`, new data structures need to be defined.

The first step is the definition of the `AthenaTensor` class for tensor variables that are stored in the whole spatial grid. Taking as an example the `Z4c` variables, these are stored in a 5D `View` whose second index corresponds to the variable index. The latter is conveniently sliced using the `Kokkos::subview` command. The resulting `Views` are then remapped into tensor `View` structures with appropriate dimensions and symmetries as follows. Scalar quantities such as the lapse function α result in 4D slices of the original 5D `View`, and constitute 4D `AthenaTensors`. Vector quantities such as the shift function β are 5D slices of the original 5D `View`, and constitute 5D `AthenaTensors` with 1 tensor index running from 0 to 2. 2D tensors in the `Z4c` system are all symmetric and thus have only 6 independent components; hence, they result in 5D slices of the original 5D `View`, with variable index running from 0 to 5. The latter is then remapped into 6D `AthenaTensors` with 2 tensor indices running from 0 to 2 but with the appropriate symmetry.

The second step is the definition of the `AthenaScratchTensor` class for temporary variables used in the calculations. The latter are defined as 2D arrays in the `Kokkos Scratch Memory Space`. The first index is the tensor index, which is remapped similarly to `AthenaTensors`. In this case, however, tensors can have up to 4 indices and different kinds of symmetries. The second index is the spatial index, spanning the x direction. Depending on the dimensionality of the tensor, `AthenaScratchTensors` can be up to 5D `Views`.

With the above data structures, calculations can be performed hierarchically as in `GR-Athena++`, following Alfieri et al. [292]. In `GR-AthenaK` the main computational kernels are organized as follows:

1. Outer loops over MeshBlock index m and spatial indices k (z -direction) and j (y -direction)
2. Tensor loops over tensor indices $a, b, c \dots$
3. Inner vectorized loops over the spatial index i (x -direction)

Such hierarchical computations can be performed in Kokkos only by selecting the TeamPolicy execution policy. An example of such calculations in GR-Athena++ is showed in the code snippet of Listing 4.1.

```

1 par_for_outer("z4c rhs loop", DevExeSpace(), scr_size, scr_level, 0, nmb
   -1, ks, ke, js, je, KOKKOS_LAMBDA(TeamMember_t member, const int m,
   const int k, const int j) {
2 ...
3 for(int a = 0; a < 3; ++a)
4   for(int b = a; b < 3; ++b) {
5     par_for_inner(member, is, ie, [&](const int i) {
6       rhs.g_dd(m, a, b, k, j, i) = - 2. * z4c.alpha(m, k, j, i)
7                                     * z4c.A_dd(m, a, b, k, j, i)
8                                     + Lg_dd(a, b, i);
9
10    });
11  }
12 }
13 ...
14 });

```

LISTING 4.1: RHS construction of the Z4c metric variable, showing the hierarchical structure of the main kernels. The conformal metric `g_dd` is a 6D AthenaTensor, where `m` is the MeshBlock index, `a, b` are the tensor indices and `k, j, i` are the spatial indices. `par_for_outer` is the outer loop carried out with TeamPolicy. Tensor loops over `a` and `b` contain the internal `par_for_inner` vectorized loops along the x -direction.

`rhs.g_dd`, `z4c.A_dd` are 6D AthenaTensors, `z4c.alpha` is a 4D AthenaTensor and `Lg_dd` is a 3D AthenaScratchTensor with two tensor indices. In the code, AthenaScratchTensors are defined and stored in the `par_for_outer` loop and recycled after usage. The main advantages of this implementation are the very efficient usage of memory and the high readability of the code.

4.3.3 Finite differences

In `GR-AthenaK`, derivatives are approximated using FD, but their implementation differs from that in `GR-Athena++`. The latter takes advantage of the very efficient strided memory access. However, this is not possible within `GR-AthenaK`, because the memory structure is architecture-dependent and additionally, the usage of `Kokkos::subview` for slicing `Views` complicates the data location pattern in memory. For these reasons, `GR-AthenaK` data are accessed through `View` indices. Similarly to `GR-Athena++`, FD in `GR-AthenaK` are implemented in a way that preserves the associative property of addition, thereby minimizing the truncation error associated with this operation.

4.4 GR-AthenaK tests

In this section three different sets of tests with BH evolutions are performed on `GR-AthenaK`. 3D initial data are constructed by setting a static puncture at the origin of the grid which is subsequently evolved in time according to the Z4c evolution equations. The correctness of the evolutions obtained with `GR-AthenaK` is assessed by comparing sets of `GR-AthenaK` runs against benchmark `GR-Athena++` runs and with one series of self-convergence tests. Lastly, the performances of `GR-AthenaK` running on CPUs and GPUs are compared to `GR-Athena++`.

For these tests `GR-Athena++` and `GR-AthenaK` have been set up in the same way. In particular, RK3 is used as time integrator, with CFL factor 0.25; CC and double precision representation of data is used; $\mathcal{N}_g = 2$ is selected, corresponding to an overall 2nd order accuracy. The Kreiss-Oliger dissipation parameter is selected as $\sigma_{KO} = 0.1$. Simple outflow BC are employed.

For runs on CPUs, `GR-Athena++` and `GR-AthenaK` have been compiled with the GNU compiler; most of the simulations are performed on Broadwell (BDW) cores, if not otherwise stated. For runs on GPUs, `GR-AthenaK` has been compiled with the `nvcc` wrapper compiler, and CUDA module version V11.4.120 is used. The GPU used for these tests is an Ampere100 GPU with 80 GB RAM.

4.4.1 Comparison tests

The solutions obtained with the `GR-Athena++` and `GR-AthenaK` codes are compared with sets of unigrid runs in which a cubic mesh with $n^M = 80$ and physical size $x_M = 80$ M is partitioned in three different ways, i.e. using cubic `MeshBlocks` with




FIGURE 4.1: Comparison tests between GR-Athena++ and GR-AthenaK on unigrid meshes partitioned with MeshBlocks of different size. Left plot: results for GR-AthenaK run on CPUs. Right plot: results for GR-AthenaK run on one GPU. In both plots, top panel: direct comparison of x -slices of χ obtained with the two codes; bottom panel: point-wise differences between the two solutions. Results are reported at $t = 100M$. Analogous results are found for all directions and for all the variables.

sizes $n^{\text{MB}} = \{20, 40, 80\}$. This results in a total of $\{64, 8, 1\}$ MeshBlocks, respectively. The one puncture initial data are evolved up to $t = 100M$ and GR-AthenaK simulations are run on both CPUs and GPUs. This test checks that the solution is correct and does not depend neither on how the mesh is partitioned, nor on the parallel device employed for the evolution.

Fig. 4.1 shows the comparison of x -slices of the field variable χ evolved with GR-Athena++ and GR-AthenaK after several time steps. The bottom panels of both plots show that for evolutions of GR-AthenaK on CPUs and GPUs, differences in the solutions given by the two codes are of the order of 10^{-15} at $t = 100M$ for all grid configurations. Notably, the difference curves in the bottom panel of both plots are symmetric with respect to the origin of the axis. Analogous results are found for y - and z - directions and for all grid variables. This indicates that the implementation of the equations and FD are correct and the spacetime evolution is robust. To obtain such results, it was necessary to implement explicit barriers in the computational kernels, after each `par_for_inner` execution, causing a potential performance hit discussed below.




FIGURE 4.2: Convergence plot showing 2nd-3rd order convergence for the one puncture problem. Top panel: x -slice of χ in the innermost refinement level for the three different resolutions considered here, where medium and fine solutions are interpolated with cubic polynomials on the coarse grid; bottom panel: absolute difference between medium and coarse, and between fine and medium solutions and corresponding rescaling demonstrating the convergence order. The simulations to produce these data are run on GPUs. Similar results are obtained considering y - and z - slices and any other field variable.

4.4.2 Convergence tests

Convergence tests on GR-AthenaK are performed following Brüggmann et al. [273]. These tests require one to have a large resolution in the proximity of the puncture, which is achieved by using refined grids. In particular, three different simulations are considered, with refined grids constructed starting from a cubic mesh with $x_M = 192 M$ and three resolutions, indicated by $n^M = \{64, 96, 128\}$ respectively. The initial mesh is divided into cubic MeshBlocks with $n^{\text{MB}} = 32$ and refined 7 times around the origin of the axes, for a total of 8 refinement levels. This procedure results in maximum resolutions at the puncture of $\delta x = \{6M, 4M, 3M\} \times 2^{-7}$, respectively. The simulations corresponding to these grid spacings are referred to as coarse, medium, fine, respectively. The results presented here are obtained with runs on GPUs.

Considering the setup described at the beginning of this section, 2nd order convergence is expected. Communication among MeshBlocks for static mesh refinement (SMR) is handled with 2nd order \mathcal{R}/\mathcal{P} operators.

Convergence is assessed by considering the χ variable in the innermost (and most

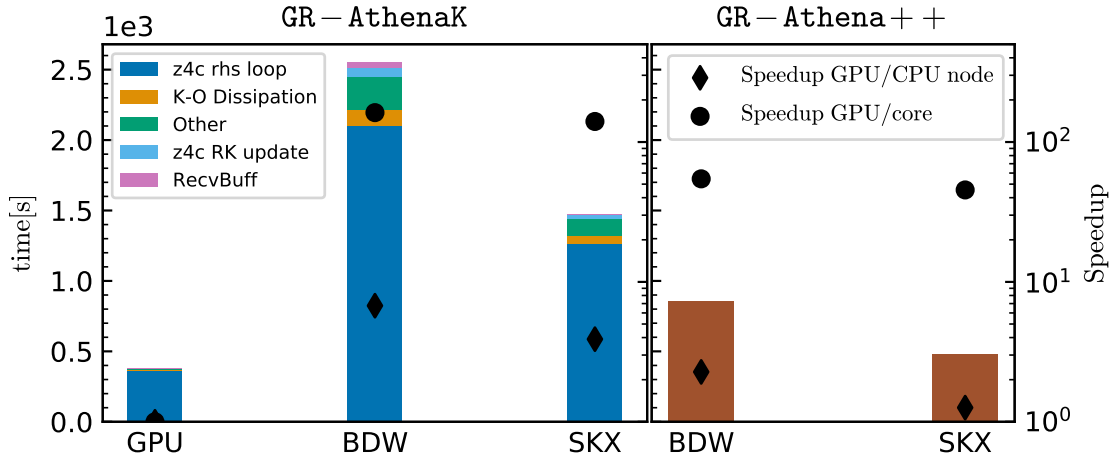


FIGURE 4.3: Performances of GR-AthenaK comparing the execution time on a GPU to the execution time of both GR-Athena++ and GR-AthenaK on full BDW and SKX CPU nodes. Left panel: GR-AthenaK runs; the total execution time is measured with `Kokkos::Timer`, while the time spent in each kernel is measured with `simple-kernel-timer` of `kokkos-tools`. Right panel reports data for GR-Athena++ runs; the total execution time is measured with the C++ function `omp_get_wtime()`. The black markers in both panels represent the speedup obtained by executing the code on one GPU.

resolved) refinement level. Here, the following procedure is applied. A common time of $t = 48M$ is considered, as this is a multiple of the integration time steps of all the three simulations. The x -slice closest to the puncture⁴ is selected in the coarse simulation; in the medium and coarse simulations, the corresponding samples of χ are interpolated on the same physical points selected in the coarse simulation. Interpolation is performed with cubic polynomials, which has a higher order accuracy than the expected convergence order.

Fig. 4.2 shows that a convergence order between 2 and 3 is achieved, when rescaling the fine-medium curve by the appropriate Q_n calculated as in Eq. (3.6). The convergence order is almost 3 at the puncture and decreases to about 2 close to the MeshBlock boundaries, where communication happens.

4.4.3 Performance tests

The performances of GR-AthenaK are evaluated by comparing the execution of the same problem on a single GPU to the execution of both GR-AthenaK and GR-Athena++ on CPU nodes using OMP parallelization. Runs on CPU nodes are performed on one BDW node (24 cores) and one SKX node (36 cores), utilizing as many OMP threads as available cores in one node.

The grid configuration selected for this test is chosen to achieve saturation of the available memory of the GPU. To this aim, a refined grid is considered. According to Eq. 4.3, using cubic MeshBlocks of size $n^{\text{MB}} = 32$ with SMR and $\mathcal{N}_g = 2$, a total of MeshBlocks 1985 fit into the GPU memory for a spacetime evolution. The grid is constructed by considering a cubic mesh with $n^{\text{M}} = 128$ and $x_M = 1536M$ and cubic MeshBlocks of size $n^{\text{MB}} = 32$. Two refinement centers are considered, at positions $r_{\pm} = (\pm 3.75, 0, 0)$ with 14 refinement levels for both centers. This results in a total of 1744 MeshBlocks and $\sim 5.71 \times 10^7$ physical cells, or 1.25×10^9 physical data points considering all the Z4c variables. The runs are conducted for 20 time steps.

Fig. 4.3 shows the results of the performance tests. The left panel demonstrates that most of the execution time of GR-AthenaK is spent in the kernel calculating the Z4c RHS, as expected. The execution times reported on the plot for BDW and SKX correspond to runs performed with full nodes. To compare the performances of runs on one GPU, the speedup with respect to CPU runs is calculated as the ratio between the total execution times on CPU and the corresponding ones of GR-AthenaK on the GPU. Speedups of ~ 6 and ~ 4 are found comparing GR-AthenaK on one GPU to GR-AthenaK on one BDW and SKX full node, respectively (black diamonds, left panel). Comparing the execution time on a single GPU to the one on a single core, instead, speedups $\sim 140 - 150$ are found (black dot, left panel). Focusing now on the right panel, comparing GR-AthenaK run on one GPU to GR-Athena++ on full BDW and SKX nodes gives speedup GPU/core of $\sim 45 - 55$. Overall, the number of cells updated per second and per core is $\mathcal{O}(10^4)$ for both GR-Athena++ and GR-AthenaK, while up to 2.4×10^6 cells are updated per second on a GPU.

These results show that running the same problem on a single GPU with the new GR-AthenaK code is slightly faster (by a factor 1.4 – 2) than running it on a full CPU node with the GR-Athena++ code, while a speedup ~ 50 is obtained comparing the execution of one core to one GPU. However, GR-AthenaK on CPUs is a factor ~ 3

⁴The puncture is defined analytically at the origin, but the CC grid used here does not include the point $\mathcal{O} = (0, 0, 0)$.

slower than `GR-Athena++`. This can be due to several factors. First, the coarse-grained parallelism approach used in `GR-Athena++` is particularly efficient for executions on CPUs. Second, the explicit synchronization of team members by means of barriers, required for a correct evolution on GPUs (see end of §4.4.1), have the side effect of slowing down the computations. Third, FD are particularly efficient in `GR-Athena++`, due to strided access to data through pointers; in `GR-AthenaK` on the other hand, data are accessed less efficiently via `View` indices; additionally, in `GR-AthenaK` \mathcal{N}_g is regulated via a runtime parameter and `if` conditions are used to select the correct FD stencil to be used.

4.5 Towards BBH evolutions

In the current version of `GR-AthenaK`, BBH initial data are included through the external `C` library based on the pseudo-spectral approach of Ansorg, Brüggmann, and Tichy [304], as in `GR-Athena++` (see §3.2.1). In particular, initial data are calculated on CPUs, stored in the host and then deep copied in the device. This procedure does not affect the overall performances as it is done only once when the problem is initialized.

However, AMR is required for BBH evolutions, in which the highest-refined regions follow the two punctures. At the time of writing this *Dissertation*, AMR is being implemented and tested for the specific problem of BBH evolution. In particular, the box-in-box oct-tree AMR criterion as described in §3.2.3 is being included.

The missing features required for BBH simulations in `GR-AthenaK` are the implementation of the Sommerfeld BC coupled with the extrapolated outflows BC, high-order \mathcal{R}/\mathcal{P} operators and wave extraction. However, for the latter all the necessary infrastructure (lagrange interpolation, geodesic grid) are already in place and tested.

The remainder of this section is a discussion of the feasibility of BBH runs on GPUs in terms of memory, taking the calibration runs of §3.3.2 as benchmark. For all the tests presented in §3.3 of *Chapter 3*, cubic `MeshBlocks` of size $n^{\text{MB}} = 16$ have been chosen. For the calibration runs in particular, 11 refinement levels have been used. The resulting number of `MeshBlocks` for the different resolutions of such runs are reported in the second column of Tab. 4.1.

Due to the coarse-grained parallelism approach of `GR-Athena++`, each `MeshBlock` can be handled at most by one OMP thread, and this limits the number of threads that can be launched for a given grid setup (see discussion in strong scaling tests of §3.4). Hence, in `GR-Athena++` it is more efficient to partition the mesh with many

TABLE 4.1: Recap of MeshBlock number for the calibration runs of §3.3.2 and corresponding number of GPUs required for such simulations according to Eq. 4.3. Different MeshBlock sizes n^{MB} are considered. Dashes correspond to impossible configurations, while entries that have not been considered are marked with "?". Highlighted in bold are the entries that are more relevant and discussed in the main text.

n^{MB}	16		24		32	
n^{M}	#MB	#GPUs	#MB	#GPUs	#MB	#GPUs
64	1184	1	-	-	?	?
96	5144	2	1184	1	1224	1
128	6336	2	-	-	1184	1
192	22392	5	6336	3	?	?
256	50688	10	-	-	6336	6
384	175104	35	50688	22	?	?

small MeshBlocks to be able to employ many threads, rather than having fewer large MeshBlocks. In this case, a fixed choice of n^{MB} regardless of the value of n^{M} is good, because this implies that for increasing resolution the number of MeshBlocks increases and the number of threads can also be increased accordingly.

In GR-AthenaK, parallelism is based on the multi-index loops and the number of threads that can be launched is limited only by the number of grid points to be updated. Here, having many small MeshBlocks implies that a large number of ghost points are stored, increasing the required memory, that rapidly becomes prohibitive for runs on GPUs. In this case, adapting the MeshBlock size to the resolution is beneficial for better performances and for the feasibility of higher resolution runs.

For instance, in the calibration run $n^{\text{M}} = 64$, $n^{\text{MB}} = 16$ the initial mesh in GR-Athena++ is partitioned into $\left(\frac{64}{16}\right)^3 = 64$ MeshBlocks, which results in a total of 1184 MeshBlocks (see Tab. 4.1). According to Eq. 4.3, it is feasible to perform this simulation on 1 GPU (assuming Ampere GPU 80GB). Analogously, for $n^{\text{M}} = 96$, 2 GPUs would be required to handle the 5144 MeshBlocks generated in this case. However, by setting $n^{\text{MB}} = 24$ the total number of MeshBlocks is reduced to 1184, and the simulation can be performed on a single GPU. Similarly, the higher resolution run with $n^{\text{M}} = 128$ can be performed on one GPU, provided that $n^{\text{MB}} = 32$, (see columns 5-6 of Tab. 4.1). Finally, it is possible to run the simulation $n^{\text{M}} = 192$ on 3 GPUs if $n^{\text{MB}} = 48$, $n^{\text{M}} = 256$ on 5 GPUs with $n^{\text{MB}} = 64$ and $n^{\text{M}} = 384$ on 14 GPUs with $n^{\text{MB}} = 96$, compared to 5 GPUs, 10 GPUs, 35 GPUs respectively that would be required by selecting $n^{\text{MB}} = 16$.

Conclusion

In the first part of this *Dissertation*, results from NR simulations are used to produce models to characterize the remnant BH of BHNS mergers. These are demonstrated to be crucial for GW modeling and observations. In the context of BNS mergers, the systematic study of the evolution of a binary in different setups highlights the importance of high resolution simulations and advanced input physics to make reliable astrophysical predictions based on NR. On the other hand, the feasibility of high-resolution simulations that include advanced physics treatments is limited by the consequent increase of computational cost and execution time. To perform such simulations new NR codes, capable of efficiently making use of the latest many-core hardware infrastructures, are needed. In the second part of this *Dissertation* the new oct-tree based spacetime solver `GR-Athena++` is presented, demonstrating its correctness with cross-code comparisons and convergence tests. `GR-Athena++` represents a promising tool to overcome the limitations mentioned above, showing excellent scalability properties up to $\sim \mathcal{O}(10^5)$ cores. Lastly, the development of the first performance-portable spacetime solver `GR-AthenaK` is discussed; the latter is a version of `GR-Athena++` based on the `Kokkos` framework, and can be compiled virtually for any hardware architecture, including heterogeneous platforms. Full 3D spacetime simulations run on a single GPU demonstrate speedup ~ 50 compared to the same simulations run with `GR-Athena++` on a single CPU core.

Remnant BH of BHNS mergers The application of the remnant model on synthetic populations of BHNS indicates that light and moderately spinning BH remnants surrounded by low-mass accretion discs are the most likely outcome for BHNS if $\Lambda \lesssim 1000$ and the BH has aligned spin $a_{\text{BH}} \lesssim 0.75$. The observation of GW170817 rules out NSs with $\Lambda \gtrsim 1800$ ($\gtrsim 2600$) for the low-(high-)spin prior cases [17]. Similarly, large aligned BH initial spins are disfavoured by current GW observations [178]. BHNS mergers are not expected to be accompanied by bright EM counterparts, since remnant discs large enough to power a SGRB are found to be rare from this analysis, unless BHNSs are characterized by large and aligned BH initial spins, very stiff EOS

and/or compact objects with mass $2 - 5 M_{\odot}$ (i.e. within the mass gap suggested by X-ray binaries). The BH remnant model constructed in this work is a key building block of GW models for BHNSs [72, 130, 168]; notably, the extension of `TEOBResumS` in [72] is the first GW model for BHNS mergers with higher modes and spin precession effect. GW models based on the remnant model of this work [130, 168] have been used to characterize the sources of the signals GW200105 and GW200115 [45], under the hypothesis of BHNS mergers. For GW200105, the remnant mass and spin resulted as $M_{\bullet} = 10.4_{-2.0}^{+2.7}$ and $a_{\bullet} = 0.43_{-0.03}^{+0.04}$, while for GW200115, $M_{\bullet} = 7.8_{-1.6}^{+1.4}$ and $a_{\bullet} = 0.38_{-0.02}^{+0.04}$.

Impact of input physics and resolution on BNS merger simulations This work represents the first systematic study of the impact of different neutrino treatments (namely LK, M0 and M1 schemes), turbulent viscosity and finite grid resolution on the computation of multi-messenger observables from BNS merger simulations. The overall dynamics and thermodynamics of the system are robust, implying that the maximum remnant density can be inferred from GW observations. However, the time of BH collapse varies significantly with the grid resolution. Viscosity tends to stabilize the remnant NS, in agreement with previous findings. Differently from recent work, possible signatures of out-of-thermodynamical equilibrium effects are identified in the GW only at the highest resolutions considered; new high-resolution simulations will be thus required to build accurate GW templates to observe these effects. Different neutrino schemes impact significantly the mass, geometry and composition of the remnant’s disc and ejecta. M1 simulations show systematically larger proton fractions, reaching maximum values $\gtrsim 0.4$. r -process nucleosynthesis yields reflect the different ejecta compositions, and are robust only if both neutrino emission and absorption are simulated. Nuclear abundances obtained adopting M0 or M1 neutrino schemes are in agreement with the solar residual pattern, unlike the ones obtained with LK scheme, which are smaller for elements with $A \lesssim 130$. Kilonova light curves computed with spherically-symmetric radiation-hydrodynamics evolutions up to 15 days post-merger are mostly sensitive to the ejecta mass and electron fraction. Accounting for multiple ejecta components appears necessary for reliable predictions.

This work indicates that advanced neutrino transport schemes are absolutely necessary in future long-term disc evolutions, and LK schemes should be abandoned. At the same time, post-merger simulations at maximum mesh resolutions above 200 m are insufficient to deliver quantitative results for astrophysical predictions. Advanced

neutrino schemes, like the M1, and sub-grid-MHD effects are demonstrated to be necessary ingredients for accurate predictions of the winds from long-term post-merger simulations and, ultimately, for more complete understanding of the EM counterparts to BNS mergers.

GR-Athena++ In this work **GR-Athena++** is presented, a new high-performance space-time solver with block-based AMR, built on top of **Athena++**. The correctness of the code is assessed through cross-comparisons of full 3D BH evolutions against analogous results obtained with the benchmark code **BAM**; these tests include a single spinning puncture evolution and the BBH inspiral calibration problem of Brüggmann et al. [273]. Additionally, 4th and 6th order convergence is demonstrated with series of BBH mergers with full AMR. As an independent test, a quasi-circular ten orbit BBH inspiral with the parameters of Hannam et al. [310] is performed. The resulting generated waveforms are compared against waveforms computed with the GW model **TEOBResumS** [146], obtaining accumulated phase differences between the two waveforms of the order of ~ 0.1 rad to merger and ~ 0.4 rad to the ringdown phase for the highest resolution **GR-Athena++** run. These tests prove that **GR-Athena++** is accurate and robust for BBH inspiral calculations and capable of providing high-quality data for waveform modeling. The impressive scalability properties of **Athena++** are maintained in **GR-Athena++**, which demonstrates strong scaling efficiencies above 95% for up to $\sim 1.2 \times 10^4$ cores, whereas in weak scaling tests almost perfect scaling is achieved up to $\sim 10^5$ cores. For the high resolutions and consequent resources required for a potential calculation describing an intermediate mass ratio BBH inspiral, **GR-Athena++** scalability properties compare favorably with the code-base of **Dendro-GR** [280].

GR-AthenaK **GR-AthenaK** is the first performance-portable spacetime solver, obtained by refactoring **GR-Athena++** with **Kokkos**. The comparison of full 3D BH evolutions against the benchmark code **GR-Athena++** demonstrates the correctness of **GR-AthenaK** at floating point level. Expected 2nd order convergence is shown in full SMR runs executed on GPUs. Comparing performances of a 3D SMR BH run on a single GPUs to the analogous run executed with **GR-Athena++** on a single core, a speedup ~ 50 is obtained. However, a speedup of ~ 150 is achieved comparing **GR-AthenaK** on 1 GPU vs. **GR-AthenaK** on 1 CPU core. This translates into 2.28×10^6 cell-updates/s for **GR-AthenaK**, against 4.19×10^4 cell-updates/s of **GR-Athena++** for the same problem. Finally, by conveniently tuning the partition of the mesh grid, it is possible to execute




FIGURE 4: Example of an AMR spherical grid produced with **Athena++**, with geometric spacing along the radial coordinate. Such grids are well-suited to track the ejecta of BNS and BHNS mergers.

BBH simulations on $\mathcal{O}(1)$ GPUs for medium-high resolution and $\mathcal{O}(10)$ GPUs for very high resolution runs.

Outlook and future work The beginning of the new GW observing run O4 is planned in May 2023 and will be carried out collaboratively by LIGO’s detectors, Virgo’s detector in Italy and KAGRA’s detector in Japan. LIGO and Virgo project sensitivity goals for BNS mergers of 160-190 Mpc and 80-115 Mpc, respectively. Improved GW models for CB mergers are necessary for this incoming campaign.

TEOBResumS is one of the most advanced GW models for BHNS mergers and will be crucial for off-line parameter estimation of such events. The accuracy of this model is currently limited by the lack of available simulations. The model would benefit from NR data covering the BHNS parameter space in the regions $\nu \gtrsim 0.83$, $a_{\text{BH}} \lesssim 0$ and $\Lambda \gtrsim 4000$ [72]. One possibility to perform such simulations is offered by interfacing the recently available BHNS initial data solver Elliptica [320] to **GR-Athena++**. This would allow on the one hand to produce high-quality waveforms and on the other hand to study remnant discs and matter ejection in BHNS merger simulations taking advantage of the flexibility offered by the **GR-Athena++** AMR, as explained below.

The development of the matter sector of **GR-Athena++** is currently ongoing, together with efforts towards the inclusion of magnetic fields [321] and advanced neutrino transport using the M1 scheme of Radice et al. [194]. The goal is to achieve high-resolution, multi-physics, long-term NR simulations of BNS mergers that can

capture their full complexity. The AMR of **GR-Athena++** offers greater flexibility in conducting these simulations, allowing to better resolve disc and to track the ejecta by locally increasing the grid resolution with e.g. a density-based AMR criterion. Additionally, **Athena++** provides the option for using spherical grids with geometric spacing which are well suited for following the long-term ejecta evolution. A visual example of such grids is given in Fig. 4. In the near future, it is expected that these advancements will enable detailed simulations of BNS mergers within **GR-Athena++**, allowing accurate predictions of EM counterparts and r -process nucleosynthesis.

Currently, the AMR infrastructure for **AthenaK** is being developed by the authors of the code. The other missing aspects for stable BBH evolutions concern the treatment of BC and the implementation of high-order restrict/prolong operators. One example of an immediate physical application of **GR-AthenaK** will be the simulation of intermediate mass ratio BBH coalescence up to a mass ratio of 1:1000 [322]. The latter is a very computationally demanding problem, but it is of increasing relevance because such events are expected to fall in the sensitivity band of LISA [323].

Appendix A

Numerical Relativity

This Appendix recaps the key concepts of NR, which is the fundamental tool employed in this *Dissertation*. Throughout this Appendix geometric units $G = c = 1$ are used; indices denoted with greek letters run from 0 to 3, while indices denoted with latin letters from 1 to 3; Einstein summation convention is used for the sum on repeated indices.

A.1 3+1 decomposition of spacetime

A key feature of the EFE (Eqs. (1)) is that they are fully covariant, and there is no distinction between space and time. However, to be able to evolve spacetime in order to make predictions it is necessary to refactorize such equations as a Cauchy problem, in which a set of initial data and appropriate boundary conditions for the gravitational field can be specified initially and then evolved in time. This requires making a clear separation between space and time. One of the most popular approaches to this problem is the *3+1 formalism*, [302, 324, 325]. In the following, only spacetimes with metric $g_{\mu\nu}$ that are globally hyperbolic, i.e. that admits a Cauchy surface are considered. Such spacetimes can be completely foliated into three-dimensional spacelike hypersurfaces usually labeled as Σ_t , where t is a parameter that can be considered a *universal time function*. Considering two close hypersurfaces Σ_t and Σ_{t+dt} , the spacetime is locally characterized by three quantities. First, the three-dimensional metric γ_{ij} , that defines proper distances on a given hypersurface $dl := \gamma_{ij}x^ix^j$; second, the *lapse function* $\alpha(t, x^i)$, which gives the proper time $d\tau := \alpha(t, x^i)dt$ measured by *Eulerian observers*, i.e. observers moving in the normal direction to hypersurfaces; third, the *shift vector* $\beta^i(t, x^i)$, which is the velocity between Eulerian observers and the lines of constant spatial coordinates $x^i_{t+dt} := x^i_t - \beta^i(t, x^i)dt$. Because α and β^i contain the information about the choice of coordinates, they play the role of *gauge functions*, and

can be freely specified. This encodes the fact that EFE are invariant under a change of coordinates. The expression of the spacetime metric in terms of α , β^i and γ_{ij} reads

$$g_{\mu\nu} = \begin{pmatrix} -\alpha^2 + \beta_k \beta^k & \beta_i \\ \beta_j & \gamma_{ij} \end{pmatrix} \quad (\text{A.1})$$

while the unit normal vector to Σ_t can be expressed as

$$n^\mu = (1/\alpha, -\beta^i/\alpha) \quad (\text{A.2})$$

A.1.1 Spacetime evolution and constraint equations

In order to find the evolution equations, the concept of extrinsic curvature is needed. This is a measurement of the change of the normal vector under parallel transport within an hypersurface. Defining the projector operator onto hypersurfaces

$$P_\beta^\alpha := \delta_\beta^\alpha + n^\alpha n_\beta, \quad (\text{A.3})$$

the *extrinsic curvature tensor* is defined as

$$K_{\mu\nu} := -P_\mu^\alpha \nabla_\alpha n_\nu. \quad (\text{A.4})$$

The extrinsic curvature tensor is purely spatial by definition and also symmetric, i.e. $K_{\mu\nu} = K_{\nu\mu}$. It can be shown that $K_{\mu\nu}$ corresponds to a sort of velocity of the spatial metric as seen by an Eulerian observer, i.e.

$$K_{\mu\nu} = -\frac{1}{2} \mathcal{L}_n \gamma_{\mu\nu} \quad (\text{A.5})$$

where \mathcal{L}_n denotes the Lie derivative with respect to the unit normal vector. From here the evolution equations for the spatial metric in the coordinates adapted to the foliation can be derived, which read

$$\partial_t \gamma_{ij} = -2\alpha K_{ij} + D_i \beta_j + D_j \beta_i, \quad (\text{A.6})$$

where D_i denotes the covariant derivative associated to γ_{ij} . The latter equations are purely kinematic and were obtained purely from geometrical considerations. The actual spacetime dynamics are encoded in the EFE, which give the evolution equations for K_{ij} .

To rewrite the EFE in the 3+1 formalism, the first step is to project the intrinsic curvature described by the Riemann tensor $\mathcal{R}_{\alpha\beta\gamma\delta}$ onto the hypersurfaces Σ_t . The projections $P_\lambda^\alpha P_\rho^\beta P_\sigma^\gamma P_\eta^\delta \mathcal{R}_{\alpha\beta\gamma\delta}$ and $P_\lambda^\alpha P_\rho^\beta P_\sigma^\gamma n^\delta \mathcal{R}_{\alpha\beta\gamma\delta}$ lead respectively to the Gauss-Codazzi and Codazzi-Mainardi equations. Combining such equations with the EFE gives respectively the following equations, written here in the coordinates adapted to the foliation:

$${}^{(3)}\mathcal{R} + K^2 - K_{ij}K^{ij} = 16\pi\rho_{\text{ADM}} \quad (\text{A.7a})$$

$$D_j (K^{ij} - \gamma^{ij}K) = 8\pi j^i \quad (\text{A.7b})$$

where ${}^{(3)}\mathcal{R}$ is the Ricci scalar ${}^{(3)}\mathcal{R}^{ij}{}^{(3)}\mathcal{R}_{ij}$ calculated with γ_{ij} , K is the trace of the extrinsic curvature and $\rho_{\text{ADM}} := n^\mu n^\nu T_{\mu\nu}$, $j^i := -P^{i\mu} n^\nu T_{\mu\nu}$ are respectively the local energy and momentum measured by Eulerian observers. The key feature of the 4 equations Eqs. (A.7) is that they do not contain the lapse and the shift function. Thus, they do not give information about the evolution of the system but they must be satisfied at all times during evolution and are known as *constraint equations*. In particular, the first row and the second row are respectively the *Hamiltonian* and *momentum* constraints. The 6 evolution equations for the extrinsic curvature are obtained by combining the projection $P_\lambda^\alpha P_\rho^\beta n^\gamma n^\delta \mathcal{R}_{\alpha\beta\gamma\delta}$ together with the Gauss-Codazzi equations and the EFE. In the coordinates adapted to the foliation the result is

$$\begin{aligned} \partial_t K_{ij} &= \beta^k \partial_k K_{ij} + K_{ki} \partial_j \beta^k + K_{kj} \partial_i \beta^k - D_i D_j \alpha \\ &+ \alpha [{}^{(3)}\mathcal{R}_{ij} + K K_{ij} - 2K_{ik} K_j^k] + 4\pi\alpha [\gamma_{ij} (S_{\text{ADM}} - \rho_{\text{ADM}}) - 2S_{ij}], \end{aligned} \quad (\text{A.8})$$

where $S_{\text{ADM}\alpha\beta} := P_\alpha^\mu P_\beta^\nu T_{\mu\nu}$ and $S_{\text{ADM}} := S_{\text{ADM}\mu}^\mu$. Eqs. (A.8), together with Eqs. (A.6), allow one to rewrite the EFE as a Cauchy problem. The result is a system of partial differential equations that is first order in time and second order in space, and is therefore suitable for numerical evolution. Note that there are no evolution equations for α and β^i because these are gauge functions that can be freely specified. The evolution equations Eqs. (A.6) and (A.8) are known as *standard* ADM equations and arise from a non trivial rewriting of the original ADM equations developed in [326] due to York [302].

A.2 The Z4c formulation of EFE

It turns out that the *standard* ADM equations are not well posed, and their numerical evolution present violent instabilities in many cases. However, the 3+1 evolution

equations obtained above are not unique, because it is always possible to obtain new valid equations by adding multiples of the constraint equations to them. The resulting equations will be *mathematically* different from the starting equations, but they will still describe the same spacetime. This non uniqueness is key in the development of the BSSNOK formulation of the EFE [327–329], which is based on the conformal decomposition of the spacetime metric and promotion of conformal Christoffel symbols to evolved variables, and has been proven to be particularly robust for many different spacetimes, with and without the presence of matter. Another approach is the so-called Z4 formulation [330], in which an auxiliary dynamical vector field Z^α is included in the EFE, improving the mathematical properties of their 3+1 decomposition at the cost of adding a new constraint equation for this field. One of the most popular approaches used nowadays is based on the Z4c formulation [296, 331–333] and combines the strengths of BSSNOK and Z4 [288], as discussed below. The Z4c approach is used in most of the codes considered in this *Dissertation*, with the exception of *Chapter 1*. The data used in the latter were produced with the SACRA code [147], which implements the BSSNOK formulation of the EFE.

A.2.1 The Z4c equations

The constraint damped Z4 formulation [330, 334, 335] replaces the EFE by

$$\mathcal{R}_{\alpha\beta} + \nabla_\alpha Z_\beta + \nabla_\beta Z_\alpha = 8\pi \left(T_{\alpha\beta} - \frac{1}{2} g_{\alpha\beta} T \right) + \kappa_1 [t_\alpha Z_\beta + t_\beta Z_\alpha - (1 + \kappa_2) g_{\alpha\beta} t_c Z^c] \quad (\text{A.9})$$

where Z^α is an additional field, t^α is a timelike vector field and κ_1, κ_2 are *constraint damping parameters*. Solutions of Eq. (A.9) are also solutions of the original EFE if the constraint $Z_\alpha = 0$ is satisfied.

In order to obtain stable evolution equations, the following definitions are used. First of all, Z_α is decomposed as

$$\Theta := -n_\alpha Z^\alpha, \quad Z_i := P_i^\alpha Z_\alpha; \quad (\text{A.10})$$

Next, as in the BSSNOK formulation, a conformal decomposition of the 3-metric and of the extrinsic curvature is introduced, via the so-called *conformal factor* $\psi := \gamma^{1/12}$, where $\gamma := \det \gamma_{ij}$:

$$\tilde{\gamma}_{ij} := \psi^{-4} \gamma_{ij}, \quad \tilde{A}_{ij} := \psi^{-4} \left(K_{ij} - \frac{1}{3} K \gamma_{ij} \right); \quad (\text{A.11})$$

where the quantity \tilde{A}_{ij} represents the conformally-rescaled trace-free part of K_{ij} . Furthermore, the following quantities come into play

$$\begin{aligned}\chi &:= \gamma^{-1/3}, & \hat{K} &:= K - 2\Theta, \\ \tilde{\Gamma}^i &:= 2\tilde{\gamma}^{ij}Z_j + \tilde{\gamma}^{ij}\tilde{\gamma}^{kl}\partial_l[\tilde{\gamma}_{jk}], & \hat{\Gamma}^i &:= \tilde{\gamma}^{jk}\tilde{\Gamma}^i{}_{jk};\end{aligned}\tag{A.12}$$

where $\tilde{\Gamma}^i{}_{jk}$ is the Christoffel symbols of the conformal metric and the definition of χ implies that $\chi = \psi^{-4}$. Finally, using these definition, Eqs. (A.9) can be 3+1 decomposed considering the set of dynamical variables $(\chi, \tilde{\gamma}_{ij}, \hat{K}, \tilde{A}_{ij}, \Theta, \tilde{\Gamma}^i)$. The result of this procedure is the Z4c system of equations, which reads:

$$\partial_t\chi = \frac{2}{3}\chi \left[\alpha(\hat{K} + 2\Theta) - \partial_i\beta^i \right] + \beta^i\partial_i\chi,\tag{A.13a}$$

$$\begin{aligned}\partial_t\tilde{\gamma}_{ij} &= -2\alpha\tilde{A}_{ij} + \beta^k\partial_k\tilde{\gamma}_{ij} - \frac{2}{3}\tilde{\gamma}_{ij}\partial_k\beta^k \\ &\quad + 2\tilde{\gamma}_{k(i}\partial_{j)}\beta^k,\end{aligned}\tag{A.13b}$$

$$\begin{aligned}\partial_t\hat{K} &= -D^iD_i\alpha + \alpha \left[\tilde{A}_{ij}\tilde{A}^{ij} + \frac{1}{3}(\hat{K} + 2\Theta)^2 \right] \\ &\quad + \beta^i\partial_i\hat{K} + \alpha\kappa_1(1 - \kappa_2)\Theta + 4\pi\alpha[S_{\text{ADM}} + \rho_{\text{ADM}}],\end{aligned}\tag{A.13c}$$

$$\begin{aligned}\partial_t\tilde{A}_{ij} &= \chi[-D_iD_j\alpha + \alpha(\mathcal{R}_{ij} - 8\pi S_{\text{ADM}ij})]^{\text{TF}} \\ &\quad + \alpha[(\hat{K} + 2\Theta)\tilde{A}_{ij} - 2\tilde{A}^k{}_i\tilde{A}_{kj}] + \beta^k\partial_k\tilde{A}_{ij} \\ &\quad + 2\tilde{A}_{k(i}\partial_{j)}\beta^k - \frac{2}{3}\tilde{A}_{ij}\partial_k\beta^k,\end{aligned}\tag{A.13d}$$

$$\partial_t\Theta = \frac{\alpha}{2} \left[\tilde{\mathcal{H}} - 2\kappa_1(2 + \kappa_2)\Theta \right] + \beta^i\partial_i\Theta,\tag{A.13e}$$

$$\begin{aligned}\partial_t\tilde{\Gamma}^i &= -2\tilde{A}^{ij}\partial_j\alpha + 2\alpha \left[\tilde{\Gamma}^i{}_{jk}\tilde{A}^{jk} - \frac{3}{2}\tilde{A}^{ij}\partial_j[\ln(\chi)] \right. \\ &\quad \left. - \kappa_1(\tilde{\Gamma}^i - \hat{\Gamma}^i) - \frac{1}{3}\tilde{\gamma}^{ij}\partial_j[2\hat{K} + \Theta] - 8\pi\tilde{\gamma}^{ij}S_{\text{ADM}j} \right] \\ &\quad + \tilde{\gamma}^{jk}\partial_k\partial_j\beta^i + \frac{1}{3}\tilde{\gamma}^{ij}\partial_j\partial_k\beta^k \\ &\quad + \beta^j\partial_j\tilde{\Gamma}^i - \hat{\Gamma}^j\partial_j\beta^i + \frac{2}{3}\hat{\Gamma}^i\partial_j\beta^j.\end{aligned}\tag{A.13f}$$

where $S_{\text{ADM}i} := j_i$ defined below Eq. (A.7) and $[\dots]^{\text{TF}}$ means that one has to compute the trace-free part of the expression in parentheses. Moreover, parentheses enclosing indexes mean that symmetrization is performed over such indexes.

The Ricci tensor appearing in Eqs. (A.13d) is decomposed as follows:

$$\mathcal{R}_{ij} = \tilde{\mathcal{R}}^{\chi}_{ij} + \tilde{\mathcal{R}}_{ij}, \quad (\text{A.14})$$

where in terms of the covariant derivative \tilde{D}_i compatible with $\tilde{\gamma}_{jk}$ the two terms read

$$\begin{aligned} \tilde{\mathcal{R}}^{\chi}_{ij} = & \frac{1}{2\chi} \left[\tilde{D}_i \tilde{D}_j \chi + \tilde{\gamma}_{ij} \tilde{D}^l \tilde{D}_l \chi - \frac{1}{2\chi} \tilde{D}_i \chi \tilde{D}_j \chi \right] \\ & - \frac{3}{4\chi^2} \tilde{D}^l \chi \tilde{D}_l \chi \tilde{\gamma}_{ij}, \end{aligned} \quad (\text{A.15})$$

$$\begin{aligned} \tilde{\mathcal{R}}_{ij} = & -\frac{1}{2} \tilde{\gamma}^{lm} \partial_l \partial_m \tilde{\gamma}_{ij} + \tilde{\gamma}_{k(i} \partial_j) \tilde{\Gamma}^k + \hat{\Gamma}^k \tilde{\Gamma}_{(ij)k} \\ & + \tilde{\gamma}^{lm} (2\tilde{\Gamma}^k{}_{l(i} \tilde{\Gamma}_{j)km} + \tilde{\Gamma}^k{}_{im} \tilde{\Gamma}_{klj}). \end{aligned} \quad (\text{A.16})$$

As already mentioned, the 3+1 decomposition of EFE includes the constraint equations, which must be satisfied at all times during the evolution. However, a common approach is to carry out free evolutions, in which the constraints are used to assess the quality of the numerical calculations. The *dynamical constraints* in terms of transformed variables $(\tilde{\mathcal{H}}, \tilde{\mathcal{M}}_i, \Theta, Z^i)$ read:

$$\tilde{\mathcal{H}} := \mathcal{R} - \tilde{A}_{ij} \tilde{A}^{ij} + \frac{2}{3} (\hat{K} + 2\Theta)^2 - 16\pi\rho_{\text{ADM}} = 0, \quad (\text{A.17a})$$

$$\begin{aligned} \tilde{\mathcal{M}}_j := & \tilde{D}_i \tilde{A}^i{}_j - \frac{3}{2} \tilde{A}^i{}_j \partial_i [\ln(\chi)] \\ & - \frac{2}{3} \partial_j [\hat{K} + 2\Theta] - 8\pi S_{\text{ADM}j} = 0, \end{aligned} \quad (\text{A.17b})$$

$$\Theta = 0, \quad Z^i = \tilde{\Gamma}^i - \hat{\Gamma}^i = 0. \quad (\text{A.17c})$$

Additionally, the definitions of Eqs. (A.11) and Eqs. (A.12) imply that the following *algebraic constraints* must be satisfied:

$$\ln(\det \tilde{\gamma}) = 0, \quad \tilde{\gamma}^{ij} \tilde{A}_{ij} = 0. \quad (\text{A.18})$$

To ensure consistency and stable evolution, the latter need to be continuously enforced during numerical evolution [288].

A.2.2 Gauge choice and hyperbolicity

The gauge functions α and β^i can be freely specified in principle, as already mentioned in §A.1.1. However, Ref.[331] shows that the Z4c system is strongly hyperbolic, i.e. it admits a first-order strongly hyperbolic reduction [336, 337] when coupled with the puncture gauge condition. The latter consists of the Bona-Másson lapse [338] and the gamma-driver shift [339]:

$$\begin{aligned}\partial_t \alpha &= -\mu_L \alpha^2 \hat{K} + \beta^i \partial_i \alpha, \\ \partial_t \beta^i &= \mu_S \alpha^2 \tilde{\Gamma}^i - \eta \beta^i + \beta^j \partial_j \beta^i.\end{aligned}\tag{A.19}$$

A common choice is to specify the $1 + \log$ lapse $\mu_L = 2/\alpha$, together with $\mu_S = 1/\alpha^2$. Initially a “precollapsed” lapse and zero-shift is set:

$$\alpha|_{t=0} = \psi^{-2}|_{t=0}, \quad \beta^i|_{t=0} = 0;\tag{A.20}$$

where the choice is motivated by a resulting reduction in initial gauge dynamics [55]. The shift damping parameter η in Eq.(A.19) reduces long-term drifts in the metric variables [339] and effectively magnifies the spatial resolution near a massive feature, thus reducing the noise in its local motion and extracted gravitational waveforms [273]. A common choice that leads to successful time evolution of BBH of comparable masses [273] and improves stability in general [288] is $\eta = 2/M$, where M is the total ADM mass [326] of the system.

A.3 General relativistic hydrodynamics

In order to simulate the dynamics of compact objects such as NSs, a formulation for GRHD that is consistent with the Z4c system is needed. Matter is considered a perfect fluid composed of a single particle species, described by the stress-energy tensor

$$T_{\alpha\beta} = \rho h u_\alpha u_\beta + p g_{\alpha\beta}.\tag{A.21}$$

In the latter expression, ρ is the rest-mass density of the fluid, $h := 1 + \epsilon + p/\rho$ is the specific enthalpy, ϵ is the specific internal energy, p is the pressure and u_α is the 4-velocity of the fluid, normalised as $u_\alpha u^\alpha = -1$. The *total* energy density is $e = \rho(1 + \epsilon)$. The GRHD evolution equations for a perfect fluid follow from the local conservation law of $T_{\alpha\beta}$, the conservation of the baryon number and the EOS of the

fluid, which read respectively:

$$\nabla_\alpha T^{\alpha\beta} = 0, \quad (\text{A.22})$$

$$\nabla_\alpha (\rho u^\alpha) = 0, \quad (\text{A.23})$$

$$P(\rho, \epsilon) = p. \quad (\text{A.24})$$

The first step to obtain a system of evolution equations is to define the *conservative* variables $\vec{q} := \sqrt{\gamma}(D, S_k, \tau)$ where $D := W\rho$ is the rest-mass density of the fluid, $S_k := W^2\rho h v_k$ is the momentum density and $\tau := (W^2\rho h - p) - D = \rho_{\text{ADM}} - D$ is the internal energy. All these quantities are defined for Eulerian observers. In these definitions, $v^i := u^i/W + \beta^i/\alpha$ is the fluid velocity in the reference frame of the Eulerian observer and $W = 1/\sqrt{1 - \gamma_{ij}v^iv^j}$ is the Lorentz factor between the fluid frame and the Eulerian observer. Now, Eqs. (A.22) and (A.23) can be written in first-order flux-conservative form [340] as

$$\partial_t \vec{q} + \partial_i \vec{f}^i(\vec{q}) = \vec{s}(\vec{q}). \quad (\text{A.25})$$

\vec{f}^i represent the fluxes, which read

$$\vec{f}^i = \sqrt{-g} \begin{pmatrix} D \left(v^i - \frac{\beta^i}{\alpha} \right) \\ S_k \left(v^i - \frac{\beta^i}{\alpha} \right) + p \delta_k^i \\ \tau \left(v^i - \frac{\beta^i}{\alpha} \right) + p v^i \end{pmatrix} \quad (\text{A.26})$$

and \vec{s} are the source terms

$$\begin{aligned} \vec{s}^i &= \sqrt{-g} \begin{pmatrix} 0 \\ T^{\alpha\beta} (\partial_\alpha g_{\beta k} - \Gamma_{\alpha\beta}^\delta g_{\delta k}) \\ \alpha (T^{\alpha 0} \partial_\alpha \ln \alpha - T^{\alpha\beta} \Gamma_{\alpha\beta}^0) \end{pmatrix} \\ &= \sqrt{-g} \begin{pmatrix} 0 \\ T^{00} \left(\frac{1}{2} \beta^i \beta^j \partial_k \gamma_{ij} - \alpha \partial_k \alpha \right) + T^{0i} \beta^j \partial_k \gamma_{ij} + T_i^0 \partial_k \beta^i + \frac{1}{2} T^{ij} \partial_k \gamma_{ij} \\ T^{00} (\beta^i \beta^j K_{ij} - \beta^i \partial_i \alpha) + T^{0i} (2\beta^j K_{ij} - \partial_i \alpha) + T^{ij} K_{ij} \end{pmatrix} \end{aligned} \quad (\text{A.27})$$

where $g := \det g_{\alpha\beta}$. Note that both \vec{s} and \vec{f}^i depend on the so-called *primitive* variables $\vec{w} := (p, \rho, \epsilon, v^i)$. Eq. (A.25) implies the conservation of the *rest mass* (also known as baryonic mass) M_b

$$M_b := \int d^3x q^0 = \int d^3x \sqrt{\gamma} D. \quad (\text{A.28})$$

The ADM matter variables defined in the previous two sections are given by the expressions

$$\rho_{\text{ADM}} = \tau + D, \quad S_{\text{ADM}}^i = S^i, \quad S_{\text{ADM}}^{ij} = \rho h W^2 v^i v^j + \gamma^{ij} p. \quad (\text{A.29})$$

The system (A.25) is strongly hyperbolic if the EOS Eq. (A.24) is causal, i.e. the speed of sound $c_s := \left(\frac{\partial P}{\partial \rho} + \frac{P}{\rho^2} \frac{\partial P}{\partial \epsilon} \right)^{\frac{1}{h}}$ is smaller than c [340]. The EOS describes the properties of the NS fluid. The simplest example of EOS is the polytropic EOS, which is of the form $P(\rho) = K \rho^\Gamma$, where K and Γ are known as polytropic constant and adiabatic index, respectively. Assuming that for a NS in equilibrium the matter is in hydrostatic and thermodynamic equilibrium, highly degenerate and cold, it is appropriate to describe the fluid composing the NSs with polytropic EOSs. In the literature the NS matter for BNS merger simulations has been widely described with piece-wise polytropic EOSs, in which the parameters K and Γ above assume different values for different density ranges. The assumption of cold matter is not anymore appropriate when the two NSs collide, and finite-temperature EOSs ought to be used for more realistic simulations.

Appendix B

Gravitational wave extraction

In GR-Athena++, GW extraction is performed by first calculating the Weyl scalar Ψ_4 , i.e. the projection of the Weyl tensor onto an appropriately chosen null tetrad k, l, m, \bar{m} . Following the definitions in Brüggmann et al. [273]:

$$\Psi_4 = -R_{\mu\nu\rho\sigma} k^\mu \bar{m}^\nu k^\rho \bar{m}^\sigma, \quad (\text{B.1})$$

where the Riemann tensor is used instead of the Weyl tensor because GWs are extracted in vacuum. The 4D Riemann tensor is constructed from 3 + 1 split ADM variables using the Gauss-Codazzi relations as detailed in Brüggmann et al. [273]. The null tetrad is computed starting from a purely spatial coordinate basis:

$$\phi^i = (-y, x, 0), \quad r^i = (x, y, z), \quad \theta^i = \epsilon_{kl}^i \phi^k r^l; \quad (\text{B.2})$$

which is then Gram-Schmidt orthonormalised. The components of the newly formed orthonormal, spatial triad are extended to 4-vectors with 0th components set to 0. The null tetrad is finally constructed as:

$$\begin{aligned} k &= \frac{1}{\sqrt{2}}(n - \hat{r}), & l &= \frac{1}{\sqrt{2}}(n + \hat{r}); \\ m &= \frac{1}{\sqrt{2}}(\hat{\theta} + i\hat{\phi}), & \bar{m} &= \frac{1}{\sqrt{2}}(\hat{\theta} - i\hat{\phi}); \end{aligned} \quad (\text{B.3})$$

where n^μ is the normal vector defined in Eq. (A.2). The Weyl scalar is then decomposed in multipoles onto spherical harmonics of spin-weight $s = -2$ as follows¹:

$$\psi_{\ell m} = \int_0^{2\pi} \int_0^\pi -{}_2\bar{Y}_{\ell m} \Psi_4 \sin \vartheta d\vartheta d\varphi, \quad (\text{B.4})$$

$$-{}_2Y_{\ell m} = \sqrt{\frac{2\ell+1}{4\pi}} d_{\ell m}(\vartheta) e^{im\varphi}, \quad (\text{B.5})$$

$$d_{\ell m}(\vartheta) = \sum_{k=k_1}^{k_2} \frac{(-1)^k ((\ell+m)! (\ell-m)! (\ell-2)! (\ell+2)!)^{1/2}}{(\ell+m-k)! (\ell-k-2)! k! (k-m+2)!} \\ \times \left(\cos\left(\frac{\vartheta}{2}\right) \right)^{2\ell+m-2-2k} \left(\sin\left(\frac{\vartheta}{2}\right) \right)^{2k-m+2} \quad (\text{B.6})$$

$$k_1 = \max(0, m-2), \quad (\text{B.7})$$

$$k_2 = \min(\ell+m, \ell-2). \quad (\text{B.8})$$

The modes of the gravitational wave strain $h_{\ell m}$ are computed from the Weyl scalar Ψ_4 projected on coordinate spheres and decomposed in $s = -2$ spin weighted spherical harmonics, $\psi_{\ell m}$ by solving

$$\psi_{\ell m} = \ddot{h}_{\ell m}, \quad (\text{B.9})$$

using the method of [311]; the strain is then given by the mode-sum:

$$R(h_+ - ih_\times) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h_{\ell m}(t) -{}_2Y_{\ell m}(\vartheta, \varphi). \quad (\text{B.10})$$

where R is the finite radius at which GWs are extracted. Following the convention of the LIGO algorithms library [342, 343], the GW amplitude and phase are defined through

$$Rh_{\ell m} = A_{\ell m} \exp(-i\phi_{\ell m}), \quad (\text{B.11})$$

and the correspondent pulsation and frequency are, respectively $\omega_{\ell m} = d\phi_{\ell m}/dt$, $f_{\ell m} = \omega_{\ell m}/2\pi$.

¹The convention here is that of Goldberg et al. [341] up to a Condon-Shortley phase factor of $(-1)^m$.

Appendix C

Geodesic spheres

The computation of several quantities of interest such as GW extraction in spacetime evolutions, requires integration over spherical surfaces. Typically, the most common coordinatization for a 2-sphere \mathbb{S}_R^2 consists in working in polar spherical coordinates and uniformly sampling the the polar and azimuthal angles. This procedure, though, has the undesirable effects of leading to clustering of points towards the poles.

In **GR-Athena++** triangulated *geodesic spheres* are used. A geodesic sphere Q_R of radius R_Q may be viewed as the boundary of a convex polyhedron embedded in \mathbb{R}^3 with triangular faces, which is homeomorphic to \mathbb{S}_R^2 . Increasing the number of vertices of a geodesic sphere leads to more accurate approximantions to \mathbb{S}_R^2 , as depicted in Fig. C.1.

The construction of geodesic grids used in **GR-Athena++** starts from a regular icosahedron with 12 vertices and 20 plane equilateral triangular faces, embedded in a unit sphere. The initial icosahedron is then refined using the “non-recursive” approach outlined in Wang and Lee [344]. Considering one of the equilateral triangles that constitute the surface of the icosahedron, each side of the triangle is split into n_Q equal segments. The partitioning points on each side of the triangle are then connected with straight lines parallel to the sides of the initial triangle and the result is that the initial plane equilateral triangle is divided into n_Q^2 smaller equilateral triangles. The vertices of such partitioning triangles are then projected onto the unit sphere, and together with the original 12 vertices of the icosahedron they form the convex polyhedron used as a grid. The resulting polyhedron has $10n_Q^2 + 2$ vertices, in which the physical quantities to be integrated are defined. The left panel of Fig. C.1 shows a grid consisting of 92 vertices, corresponding to $n_Q = 3$, while the right panel shows the grid consisting of 9002 vertices ($n_Q = 30$) that is typically used in production runs.

Integrals on the sphere are computed with numerical quadratures. For geodesic spheres, this is done as follows. Each grid point is a shared corner among 6 (or




FIGURE C.1: Structure of the geodesic grid used by GR-Athena++. Left panel: low resolution example with 92 vertices. Right panel: a grid used for GW extraction in production simulations (9002 vertices).

5, as explained below) triangles. A solid angle is associated to each grid point by constructing cells formed by connecting the circumcenters of all the triangles that share the target vertex point. The resulting cells (called Voronoi cells) are mostly hexagons, except the ones corresponding to the 12 vertices of the original icosahedron, which have only five neighbors and are therefore pentagons; the total number of cells coincides with the total number of vertices, i.e. $10n_Q^2 + 2$. The solid angles subtended by the cells at the center of the sphere are used as weighting coefficients when computing the averages. The logical connection between the neighboring cells is implemented as described in Randall et al. [345].

Using geodesic grids ensures more even tiling of the sphere compared to the uniform latitude-longitude grid of similar resolution. This might have an impact when comparing the two methods of calculation of quantities on the sphere. However, possible discrepancies are expected to converge away with spherical grid resolution.

Appendix D

Input physics in WhiskyTHC

In WhiskyTHC the proton and neutron number densities n_p and n_n are evolved separately according to

$$\nabla_a(J_{p,n}^a) = R_{p,n} \quad (\text{D.1})$$

where $J_{p,n}^a := n_{p,n}u^a$ is the four-current associated to $n_{p,n}$ and $R_p = -R_n$ is the net lepton number deposition rate due to absorption and emission of neutrinos and antineutrinos. n_b denotes the total baryon number density, such that $n_b = n_p + n_n$ while Y_e is the electron fraction, defined as the net number density of electrons and positrons, normalised to n_b . Under the assumption of charge neutrality, $n_p = Y_e n_b$. The expressions for $R_{p,n}$ depend on the particular neutrino treatment employed, which will be discussed in this appendix.

D.1 Neutrino schemes

In the work of *Chapter 2*, weak interactions and neutrino radiation are simulated with three different schemes, namely the LK scheme, the M0 scheme (which is always coupled with the LK scheme), and the M1 transport scheme. In all schemes, three different neutrino species are explicitly modelled: ν_e , $\bar{\nu}_e$, and ν_x , where the latter represents a collective species describing heavy flavour neutrinos and antineutrinos. Moreover, all schemes are grey, i.e. the explicit dependence on the neutrino energy is integrated out for all the relevant quantities.

Neutrino LK scheme The LK scheme [221, 348] accounts for the net emission of neutrinos that are produced as a result of weak interactions happening during and after the NS collision. The reactions that are typically considered are summarised in Tab. D.1. Due to the large variety of conditions experienced by matter in BNS mergers, neutrinos that are produced in the process can be roughly divided in two components.

TABLE D.1: List of weak reactions typically considered in *WhiskyTHC*. ν denotes a generic neutrino species amongst electron neutrino ν_e , electron antineutrino $\bar{\nu}_e$ or heavy flavour neutrinos ν_x . The latter is an effective neutrino species containing muon and tau neutrinos and their antineutrinos collected together. N and A indicate respectively nucleons and generic nuclei.

Reaction	Reference
$\nu_e + n \leftrightarrow p + e^-$	[346]
$\bar{\nu}_e + p \leftrightarrow n + e^+$	[346]
$e^+ + e^- \rightarrow \nu + \bar{\nu}$	[180]
$\gamma + \gamma \rightarrow \nu + \bar{\nu}$	[180]
$\nu + N \rightarrow \nu + N$	[180]
$N + N \rightarrow \nu + \bar{\nu} + N + N$	[347]
$\nu + A \rightarrow \nu + A$	[15]

A first component gets trapped in the high-density and optically thick regions of the NS remnant, with the possibility of diffusing out on the diffusion timescale. Such component is close to thermodynamical and weak equilibrium with matter. A second component streams freely from the low-density, optically thin regions, with a small probability to further interact with the surrounding matter. The LK scheme uses a phenomenological formula to interpolate between the diffusion rate and the production rate, where the former (latter) is the relevant one in optically thick (thin) conditions. The scheme crucially relies on the evaluation of the optical depth inside the computational domain. The resulting effective rates correspond to neutrinos leaving the system, carrying away energy and lepton number. In particular, the particle emission rates correspond to the rates appearing on the right-hand side of Eq. (D.1), while the total energy emission rate, Q , is included in the simulations as a source term in the Euler equations

$$\nabla_b T^{ab} = Qu^a. \quad (\text{D.2})$$

Technical details on the numerical schemes employed for the discretization of Eq. (D.1) and Eq. (D.2) can be found in Radice et al. [96]. Such a LK scheme catches the essential cooling effect in NS matter provided by the emission of neutrinos. Moreover, it also affects the matter composition by allowing the conversion of neutrons into protons, and viceversa. However, neutrinos are not explicitly transported and the possible interaction of streaming neutrinos with matter in optically thin condition is neglected. Additionally, no neutrino trapped component is explicitly modelled in it (i.e., neutrino radiation is not included in the stress-energy tensor), since the density of particles

and energy of equilibrated neutrinos are used only to compute the diffusion rates. The non-inclusion of a neutrino trapped component in the remnant NS excludes the correct modelling of out-of-equilibrium effects that might manifest due to the transition from a neutrino-less beta equilibrium to a new equilibrium state with the presence of neutrinos. Moreover, the formation and presence of a trapped neutrino gas might change the pressure in the remnant and therefore potentially have an impact on its stability [82].

Neutrino reabsorption: LK+M0 and M1 schemes The interaction of the free-streaming neutrino component with matter in optically thin conditions can be simulated in `WhiskyTHC` using the M0 scheme, as described in Radice et al. [221]. The M0 scheme accounts for possible re-absorption of the emitted neutrinos, as computed by the LK scheme, and the consequent change in matter's composition (i.e., Y_e) and temperature. In the simulations of *Chapter 2*, the M0 scheme is implemented on a spherical grid centred at the centre of the computational grid, with outer radius ~ 756 km.

A more appropriate way to include neutrinos in the simulations is the M1 scheme, which is an approximated approach to neutrino transport that applies to neutrino radiation in all relevant regimes. The Boltzmann equations describing neutrino transport are first cast into a system of 3+1 equations, similar to the hydrodynamics equations, using a moment-based approach [349, 350]. The resulting equations are then evolved consistently to the hydrodynamics and spacetime equations. In the M1 scheme, the terms that describe neutrino interactions with matter are included directly in the stress-energy tensor of EFE. `WhiskyTHC` implements the module `THC_M1`, presented in Radice et al. [194]. This scheme requires a closure, i.e. an expression for the pressure in terms of the energy and the flux. Here the approximate analytic *Minerbo closure* is adopted. The latter is exact in the optically thick limit (matter and radiation in thermodynamic equilibrium) and in the optically thin limit (radiation streaming at the speed of light in the direction of the radiation flux) if the system has some symmetries (slab, spherical). The two limits are then connected by means of the *Eddington* factor as described in Radice et al. [194]. The weak interactions considered in `THC_M1` are the same ones included in the LK scheme, listed in Tab. D.1.

D.2 Turbulent viscosity

An effective treatment to simulate turbulent viscosity is implemented with the GRLES method [208]. In particular, this is a sub-grid model which is used to include the effect of magnetic-induced viscosity and consequent angular momentum transport in the NS matter. The viscosity parameter is related to a density-dependent mixing parameter, which is estimated from high-resolution MHD simulations in full GR from Kiuchi et al. [197], as described in detail in Radice [208].

Bibliography

- [1] P. C. Peters and J. Mathews. “Gravitational radiation from point masses in a Keplerian orbit”. In: *Phys. Rev.* 131 (1963), pp. 435–439. DOI: 10.1103/PhysRev.131.435.
- [2] P. C. Peters. “Gravitational Radiation and the Motion of Two Point Masses”. In: *Phys. Rev.* 136 (1964), B1224–B1232. DOI: 10.1103/PhysRev.136.B1224.
- [3] C. V. Vishveshwara. “Scattering of Gravitational Radiation by a Schwarzschild Black-hole”. In: *Nature* 227 (1970), pp. 936–938.
- [4] E. Berti, V. Cardoso, and A. O. Starinets. “Quasinormal modes of black holes and black branes”. In: *Class. Quant. Grav.* 26 (2009), p. 163001. DOI: 10.1088/0264-9381/26/16/163001. arXiv: 0905.2975 [gr-qc].
- [5] R. Hulse and J. Taylor. “Discovery of a pulsar in a binary system”. In: *Astrophys.J.* 195 (1975), pp. L51–L53. DOI: 10.1086/181708.
- [6] J. Weisberg, D. Nice, and J. Taylor. “Timing Measurements of the Relativistic Binary Pulsar PSR B1913+16”. In: *Astrophys.J.* 722 (2010), pp. 1030–1034. DOI: 10.1088/0004-637X/722/2/1030. arXiv: 1011.0718 [astro-ph.GA].
- [7] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Phys. Rev. Lett.* 116.6 (2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102. arXiv: 1602.03837 [gr-qc].
- [8] B. P. Abbott et al. “Properties of the Binary Black Hole Merger GW150914”. In: *Phys. Rev. Lett.* 116.24 (2016), p. 241102. DOI: 10.1103/PhysRevLett.116.241102. arXiv: 1602.03840 [gr-qc].
- [9] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Phys. Rev. Lett.* 119.16 (2017), p. 161101. DOI: 10.1103/PhysRevLett.119.161101. arXiv: 1710.05832 [gr-qc].

- [10] A. Goldstein et al. “An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi-GBM Detection of GRB 170817A”. In: *Astrophys. J.* 848.2 (2017), p. L14. DOI: 10.3847/2041-8213/aa8f41. arXiv: 1710.05446 [astro-ph.HE].
- [11] V. Savchenko et al. “INTEGRAL Detection of the First Prompt Gamma-Ray Signal Coincident with the Gravitational-wave Event GW170817”. In: *Astrophys. J.* 848.2 (2017), p. L15. DOI: 10.3847/2041-8213/aa8f94. arXiv: 1710.05449 [astro-ph.HE].
- [12] B. P. Abbott et al. “Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A”. In: *Astrophys. J. Lett.* 848.2 (2017), p. L13. DOI: 10.3847/2041-8213/aa920c. arXiv: 1710.05834 [astro-ph.HE].
- [13] J. Antoniadis et al. “A Massive Pulsar in a Compact Relativistic Binary”. In: *Science* 340 (2013), p. 6131. DOI: 10.1126/science.1233232. arXiv: 1304.6875 [astro-ph.HE].
- [14] T. M. Tauris et al. “Formation of Double Neutron Star Systems”. In: *Astrophys. J.* 846.2 (2017), p. 170. DOI: 10.3847/1538-4357/aa7e89. arXiv: 1706.09438 [astro-ph.HE].
- [15] S. L. Shapiro and S. A. Teukolsky. *Black holes, white dwarfs, and neutron stars: The physics of compact objects*. New York, USA: Wiley, 1983. ISBN: 978-0-471-87316-7.
- [16] B. P. Abbott et al. “Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A”. In: *Astrophys. J.* 848.2 (2017), p. L13. DOI: 10.3847/2041-8213/aa920c. arXiv: 1710.05834 [astro-ph.HE].
- [17] B. P. Abbott et al. “Properties of the binary neutron star merger GW170817”. In: *Phys. Rev. X* 9.1 (2019), p. 011001. DOI: 10.1103/PhysRevX.9.011001. arXiv: 1805.11579 [gr-qc].
- [18] B. P. Abbott et al. “Multi-messenger Observations of a Binary Neutron Star Merger”. In: *Astrophys. J.* 848.2 (2017), p. L12. DOI: 10.3847/2041-8213/aa91c9. arXiv: 1710.05833 [astro-ph.HE].
- [19] B. P. Abbott et al. “Estimating the Contribution of Dynamical Ejecta in the Kilonova Associated with GW170817”. In: *Astrophys. J.* 850.2 (2017), p. L39. DOI: 10.3847/2041-8213/aa9478. arXiv: 1710.05836 [astro-ph.HE].

- [20] I. Arcavi et al. “Optical emission from a kilonova following a gravitational-wave-detected neutron-star merger”. In: *Nature* 551 (2017), p. 64. DOI: 10.1038/nature24291. arXiv: 1710.05843 [astro-ph.HE].
- [21] D. A. Coulter et al. “Swope Supernova Survey 2017a (SSS17a), the Optical Counterpart to a Gravitational Wave Source”. In: *Science* (2017). [Science358,1556(2017)]. DOI: 10.1126/science.aap9811. arXiv: 1710.05452 [astro-ph.HE].
- [22] M. R. Drout et al. “Light Curves of the Neutron Star Merger GW170817/SSS17a: Implications for R-Process Nucleosynthesis”. In: *Science* 358 (2017), pp. 1570–1574. DOI: 10.1126/science.aaq0049. arXiv: 1710.05443 [astro-ph.HE].
- [23] P. A. Evans et al. “Swift and NuSTAR observations of GW170817: detection of a blue kilonova”. In: *Science* 358 (2017), p. 1565. DOI: 10.1126/science.aap9580. arXiv: 1710.05437 [astro-ph.HE].
- [24] G. Hallinan et al. “A Radio Counterpart to a Neutron Star Merger”. In: *Science* 358 (2017), p. 1579. DOI: 10.1126/science.aap9855. arXiv: 1710.05435 [astro-ph.HE].
- [25] M. M. Kasliwal et al. “Illuminating Gravitational Waves: A Concordant Picture of Photons from a Neutron Star Merger”. In: *Science* 358 (2017), p. 1559. DOI: 10.1126/science.aap9455. arXiv: 1710.05436 [astro-ph.HE].
- [26] M. Nicholl et al. “The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. III. Optical and UV Spectra of a Blue Kilonova From Fast Polar Ejecta”. In: *Astrophys. J.* 848 (2017), p. L18. DOI: 10.3847/2041-8213/aa9029. arXiv: 1710.05456 [astro-ph.HE].
- [27] S. J. Smartt et al. “A kilonova as the electromagnetic counterpart to a gravitational-wave source”. In: *Nature* (2017). DOI: 10.1038/nature24303. arXiv: 1710.05841 [astro-ph.HE].
- [28] M. Soares-Santos et al. “The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. I. Dark Energy Camera Discovery of the Optical Counterpart”. In: *Astrophys. J.* 848.2 (2017), p. L16. DOI: 10.3847/2041-8213/aa9059. arXiv: 1710.05459 [astro-ph.HE].
- [29] N. R. Tanvir et al. “The Emergence of a Lanthanide-Rich Kilonova Following the Merger of Two Neutron Stars”. In: *Astrophys. J.* 848 (2017), p. L27. DOI: 10.3847/2041-8213/aa90b6. arXiv: 1710.05455 [astro-ph.HE].

- [30] E. Troja et al. “The X-ray counterpart to the gravitational wave event GW 170817”. In: *Nature* (2017). DOI: 10.1038/nature24290. arXiv: 1710.05433 [astro-ph.HE].
- [31] K. P. Mooley et al. “Superluminal motion of a relativistic jet in the neutron-star merger GW170817”. In: *Nature* 561.7723 (2018), pp. 355–359. DOI: 10.1038/s41586-018-0486-3. arXiv: 1806.09693 [astro-ph.HE].
- [32] G. Ghirlanda et al. “Compact radio emission indicates a structured jet was produced by a binary neutron star merger”. In: *Science* 363 (2019), p. 968. DOI: 10.1126/science.aau8815. arXiv: 1808.00469 [astro-ph.HE].
- [33] J. J. Ruan et al. “Brightening X-Ray Emission from GW170817/GRB 170817A: Further Evidence for an Outflow”. In: *Astrophys. J.* 853.1 (2018), p. L4. DOI: 10.3847/2041-8213/aaa4f3. arXiv: 1712.02809 [astro-ph.HE].
- [34] J. D. Lyman et al. “The optical afterglow of the short gamma-ray burst associated with GW170817”. In: *Nat. Astron.* 2.9 (2018), pp. 751–754. DOI: 10.1038/s41550-018-0511-3. arXiv: 1801.02669 [astro-ph.HE].
- [35] R. Chornock et al. “The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. IV. Detection of Near-infrared Signatures of r-process Nucleosynthesis with Gemini-South”. In: *Astrophys. J.* 848.2 (2017), p. L19. DOI: 10.3847/2041-8213/aa905c. arXiv: 1710.05454 [astro-ph.HE].
- [36] P. S. Cowperthwaite et al. “The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. II. UV, Optical, and Near-infrared Light Curves and Comparison to Kilonova Models”. In: *Astrophys. J.* 848.2 (2017), p. L17. DOI: 10.3847/2041-8213/aa8fc7. arXiv: 1710.05840 [astro-ph.HE].
- [37] M. Tanaka et al. “Kilonova from post-merger ejecta as an optical and near-infrared counterpart of GW170817”. In: *Publ. Astron. Soc. Jap.* (2017). DOI: 10.1093/pasj/psx121. arXiv: 1710.05850 [astro-ph.HE].
- [38] Y. Utsumi et al. “J-GEM observations of an electromagnetic counterpart to the neutron star merger GW170817”. In: *Publ. Astron. Soc. Jap.* 69.6 (2017), p. 101. DOI: 10.1093/pasj/psx118. arXiv: 1710.05848 [astro-ph.HE].

- [39] A. Perego, D. Radice, and S. Bernuzzi. “AT2017gfo: An Anisotropic and Three-component Kilonova Counterpart of GW170817”. In: *Astrophys. J.* 850.2 (2017), p. L37. DOI: 10.3847/2041-8213/aa9ab9. arXiv: 1711.03982 [astro-ph.HE].
- [40] 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: Unified Data Set, Analytic Models, and Physical Implications”. In: *Astrophys. J.* 851.1 (2017), p. L21. DOI: 10.3847/2041-8213/aa9c84. arXiv: 1710.11576 [astro-ph.HE].
- [41] E. Waxman et al. “Constraints on the ejecta of the GW170817 neutron-star merger from its electromagnetic emission”. In: *Mon. Not. Roy. Astron. Soc.* 481.3 (2018), pp. 3423–3441. DOI: 10.1093/mnras/sty2441. arXiv: 1711.09638 [astro-ph.HE].
- [42] B. D. Metzger, T. A. Thompson, and E. Quataert. “A magnetar origin for the kilonova ejecta in GW170817”. In: *Astrophys. J.* 856.2 (2018), p. 101. DOI: 10.3847/1538-4357/aab095. arXiv: 1801.04286 [astro-ph.HE].
- [43] K. Kawaguchi, M. Shibata, and M. Tanaka. “Radiative transfer simulation for the optical and near-infrared electromagnetic counterparts to GW170817”. In: *Astrophys. J.* 865.2 (2018), p. L21. DOI: 10.3847/2041-8213/aade02. arXiv: 1806.04088 [astro-ph.HE].
- [44] M. Breschi et al. “AT2017gfo: Bayesian inference and model selection of multicomponent kilonovae and constraints on the neutron star equation of state”. In: *Mon. Not. Roy. Astron. Soc.* 505.2 (2021), pp. 1661–1677. DOI: 10.1093/mnras/stab1287. arXiv: 2101.01201 [astro-ph.HE].
- [45] R. Abbott et al. “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences”. In: *Astrophys. J. Lett.* 915.1 (2021), p. L5. DOI: 10.3847/2041-8213/ac082e. arXiv: 2106.15163 [astro-ph.HE].
- [46] R. Abbott et al. “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run”. In: (Nov. 2021). arXiv: 2111.03606 [gr-qc].
- [47] B. P. Abbott et al. “GW150914: First results from the search for binary black hole coalescence with Advanced LIGO”. In: *Phys. Rev. D* 93.12 (2016), p. 122003. DOI: 10.1103/PhysRevD.93.122003. arXiv: 1602.03839 [gr-qc].

- [48] J. Veitch and A. Vecchio. “Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network”. In: *Phys.Rev.* D81 (2010), p. 062003. DOI: 10.1103/PhysRevD.81.062003. arXiv: 0911.3820 [astro-ph.CO].
- [49] J. Veitch et al. “Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library”. In: *Phys. Rev.* D91.4 (2015), p. 042003. DOI: 10.1103/PhysRevD.91.042003. arXiv: 1409.7215 [gr-qc].
- [50] B. P. Abbott et al. “A guide to LIGO-Virgo detector noise and extraction of transient gravitational-wave signals”. In: (2019). arXiv: 1908.11170 [gr-qc].
- [51] E. Poisson and C. M. Will. “Gravitational waves from inspiraling compact binaries: Parameter estimation using second postNewtonian wave forms”. In: *Phys.Rev.* D52 (1995), pp. 848–855. DOI: 10.1103/PhysRevD.52.848. arXiv: gr-qc/9502040 [gr-qc].
- [52] L. Blanchet. “Second postNewtonian generation of gravitational radiation”. In: *Phys.Rev.* D51 (1995), pp. 2559–2583. DOI: 10.1103/PhysRevD.51.2559. arXiv: gr-qc/9501030 [gr-qc].
- [53] F. Pretorius. “Evolution of binary black hole spacetimes”. In: *Phys. Rev. Lett.* 95 (2005), p. 121101. DOI: 10.1103/PhysRevLett.95.121101. arXiv: gr-qc/0507014.
- [54] J. G. Baker et al. “Consistency of post-Newtonian waveforms with numerical relativity”. In: *Phys.Rev.Lett.* 99 (2007), p. 181101. DOI: 10.1103/PhysRevLett.99.181101. arXiv: gr-qc/0612024 [gr-qc].
- [55] M. Campanelli et al. “Accurate Evolutions of Orbiting Black-Hole Binaries Without Excision”. In: *Phys. Rev. Lett.* 96 (2006), p. 111101. DOI: 10.1103/PhysRevLett.96.111101. arXiv: gr-qc/0511048.
- [56] M. Shibata and K. Uryu. “Simulation of merging binary neutron stars in full general relativity: $\Gamma = 2$ case”. In: *Phys. Rev.* D61 (2000), p. 064001. DOI: 10.1103/PhysRevD.61.064001. arXiv: gr-qc/9911058.
- [57] P. Ajith et al. “Phenomenological template family for black-hole coalescence waveforms”. In: *Class.Quant.Grav.* 24 (2007), S689–S700. DOI: 10.1088/0264-9381/24/19/S31. arXiv: 0704.3764 [gr-qc].

-
- [58] L. Santamaria et al. “Matching post-Newtonian and numerical relativity waveforms: systematic errors and a new phenomenological model for non-precessing black hole binaries”. In: *Phys.Rev.* D82 (2010), p. 064016. DOI: 10.1103/PhysRevD.82.064016. arXiv: 1005.3306 [gr-qc].
- [59] M. Hannam et al. “Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms”. In: *Phys. Rev. Lett.* 113.15 (2014), p. 151101. DOI: 10.1103/PhysRevLett.113.151101. arXiv: 1308.3271 [gr-qc].
- [60] S. Khan et al. “Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects”. In: *Phys. Rev.* D100.2 (2019), p. 024059. DOI: 10.1103/PhysRevD.100.024059. arXiv: 1809.10113 [gr-qc].
- [61] G. Pratten et al. “Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes”. In: *Phys. Rev. D* 103.10 (2021), p. 104056. DOI: 10.1103/PhysRevD.103.104056. arXiv: 2004.06503 [gr-qc].
- [62] G. Pratten et al. “Setting the cornerstone for a family of models for gravitational waves from compact binaries: The dominant harmonic for nonprecessing quasicircular black holes”. In: *Phys. Rev. D* 102.6 (2020), p. 064001. DOI: 10.1103/PhysRevD.102.064001. arXiv: 2001.11412 [gr-qc].
- [63] C. García-Quirós et al. “Multimode frequency-domain model for the gravitational wave signal from nonprecessing black-hole binaries”. In: *Phys. Rev. D* 102.6 (2020), p. 064002. DOI: 10.1103/PhysRevD.102.064002. arXiv: 2001.10914 [gr-qc].
- [64] S. E. Field et al. “Fast prediction and evaluation of gravitational waveforms using surrogate models”. In: *Phys.Rev.* X4.3 (2014), p. 031006. DOI: 10.1103/PhysRevX.4.031006. arXiv: 1308.3565 [gr-qc].
- [65] J. Blackman et al. “A Surrogate Model of Gravitational Waveforms from Numerical Relativity Simulations of Precessing Binary Black Hole Mergers”. In: *Phys. Rev.* D95.10 (2017), p. 104023. DOI: 10.1103/PhysRevD.95.104023. arXiv: 1701.00550 [gr-qc].
- [66] V. Varma et al. “Surrogate models for precessing binary black hole simulations with unequal masses”. In: *Phys. Rev. Research.* 1 (2019), p. 033015. DOI: 10.1103/PhysRevResearch.1.033015. arXiv: 1905.09300 [gr-qc].

- [67] A. Buonanno and T. Damour. “Transition from inspiral to plunge in binary black hole coalescences”. In: *Phys. Rev. D* 62 (2000), p. 064015. DOI: 10.1103/PhysRevD.62.064015. arXiv: gr-qc/0001013.
- [68] T. Damour, P. Jaranowski, and G. Schaefer. “On the determination of the last stable orbit for circular general relativistic binaries at the third postNewtonian approximation”. In: *Phys. Rev. D* 62 (2000), p. 084011. DOI: 10.1103/PhysRevD.62.084011. arXiv: gr-qc/0005034 [gr-qc].
- [69] T. Damour. “Coalescence of two spinning black holes: An effective one-body approach”. In: *Phys. Rev. D* 64 (2001), p. 124013. DOI: 10.1103/PhysRevD.64.124013. arXiv: gr-qc/0103018.
- [70] M. Breschi et al. “kiloHertz gravitational waves from binary neutron star remnants: time-domain model and constraints on extreme matter”. In: *Phys. Rev. D* 100.10 (2019), p. 104029. DOI: 10.1103/PhysRevD.100.104029. arXiv: 1908.11418 [gr-qc].
- [71] M. Breschi et al. “Kilohertz Gravitational Waves From Binary Neutron Star Mergers: Numerical-relativity Informed Postmerger Model”. In: (May 2022). arXiv: 2205.09112 [gr-qc].
- [72] A. Gonzalez et al. “Numerical-Relativity-Informed Effective-One-Body model for Black-Hole-Neutron-Star Mergers with Higher Modes and Spin Precession”. In: (Dec. 2022). arXiv: 2212.03909 [gr-qc].
- [73] B. P. Abbott et al. “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA”. In: *Living Reviews in Relativity* 23.1 (2020). ISSN: 1433-8351. DOI: 10.1007/s41114-020-00026-9. URL: <http://dx.doi.org/10.1007/s41114-020-00026-9>.
- [74] T. Akutsu et al. “Overview of KAGRA: Calibration, Detector Characterization, Physical Environmental Monitors, and the Geophysics Interferometer”. In: *arXiv:2009.09305 [astro-ph, physics:gr-qc]* (Sept. 2020). arXiv: 2009.09305 [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2009.09305> (visited on 01/07/2021).
- [75] Amaro-Seoane et al. “Laser Interferometer Space Antenna”. In: *arXiv:1702.00786 [astro-ph]* (Feb. 2017). arXiv: 1702.00786 [astro-ph]. URL: <http://arxiv.org/abs/1702.00786> (visited on 01/08/2021).

- [76] M. Punturo et al. “The Einstein Telescope: A third-generation gravitational wave observatory”. In: *Class.Quant.Grav.* 27 (2010), p. 194002. DOI: 10.1088/0264-9381/27/19/194002.
- [77] B. P. Abbott et al. “Exploring the Sensitivity of Next Generation Gravitational Wave Detectors”. In: *Class. Quant. Grav.* 34.4 (2017), p. 044001. DOI: 10.1088/1361-6382/aa51f4. arXiv: 1607.08697 [astro-ph.IM].
- [78] C. O. Lousto et al. “Intermediate-mass-ratio black hole binaries: Intertwining numerical and perturbative techniques”. In: *Phys.Rev.* D82 (2010), p. 104057. DOI: 10.1103/PhysRevD.82.104057. arXiv: 1008.4360 [gr-qc].
- [79] H. Nakano et al. “Intermediate-mass-ratio black hole binaries II: Modeling Trajectories and Gravitational Waveforms”. In: *Phys.Rev.* D84 (2011), p. 124006. DOI: 10.1103/PhysRevD.84.124006. arXiv: 1108.4421 [gr-qc].
- [80] M. Shibata and K. Taniguchi. “Coalescence of Black Hole-Neutron Star Binaries”. In: *Living Rev. Rel.* 14 (2011), p. 6. DOI: 10.12942/lrr-2011-6.
- [81] D. Radice et al. “Long-lived Remnants from Binary Neutron Star Mergers”. In: *Mon. Not. Roy. Astron. Soc.* 481.3 (2018), pp. 3670–3682. DOI: 10.1093/mnras/sty2531. arXiv: 1803.10865 [astro-ph.HE].
- [82] A. Perego, S. Bernuzzi, and D. Radice. “Thermodynamics conditions of matter in neutron star mergers”. In: *Eur. Phys. J.* A55.8 (2019), p. 124. DOI: 10.1140/epja/i2019-12810-7. arXiv: 1903.07898 [gr-qc].
- [83] D. Radice, S. Bernuzzi, and A. Perego. “The Dynamics of Binary Neutron Star Mergers and of GW170817”. In: *Ann. Rev. Nucl. Part. Sci.* 70 (2020). DOI: 10.1146/annurev-nucl-013120-114541. arXiv: 2002.03863 [astro-ph.HE].
- [84] M. Shibata, K. Taniguchi, and K. Uryu. “Merger of binary neutron stars with realistic equations of state in full general relativity”. In: *Phys. Rev.* D71 (2005), p. 084021. DOI: 10.1103/PhysRevD.71.084021. arXiv: gr-qc/0503119.
- [85] M. Shibata and K. Taniguchi. “Merger of binary neutron stars to a black hole: disk mass, short gamma-ray bursts, and quasinormal mode ringing”. In: *Phys.Rev.* D73 (2006), p. 064027. DOI: 10.1103/PhysRevD.73.064027. arXiv: astro-ph/0603145 [astro-ph].
- [86] K. Hotokezaka et al. “Binary Neutron Star Mergers: Dependence on the Nuclear Equation of State”. In: *Phys.Rev.* D83 (2011), p. 124008. DOI: 10.1103/PhysRevD.83.124008. arXiv: 1105.4370 [astro-ph.HE].

- [87] A. Bauswein, T. Baumgarte, and H. T. Janka. “Prompt merger collapse and the maximum mass of neutron stars”. In: *Phys. Rev. Lett.* 111.13 (2013), p. 131101. DOI: 10.1103/PhysRevLett.111.131101. arXiv: 1307.5191 [astro-ph.SR].
- [88] F. Zappa et al. “Gravitational-wave luminosity of binary neutron stars mergers”. In: *Phys. Rev. Lett.* 120.11 (2018), p. 111101. DOI: 10.1103/PhysRevLett.120.111101. arXiv: 1712.04267 [gr-qc].
- [89] M. Agathos et al. “Inferring Prompt Black-Hole Formation in Neutron Star Mergers from Gravitational-Wave Data”. In: *Phys. Rev. D* 101.4 (2020), p. 044006. DOI: 10.1103/PhysRevD.101.044006. arXiv: 1908.05442 [gr-qc].
- [90] S. Bernuzzi et al. “Accretion-induced prompt black hole formation in asymmetric neutron star mergers, dynamical ejecta and kilonova signals”. In: *Mon. Not. Roy. Astron. Soc.* 497.2 (2020), pp. 1488–1507. DOI: 10.1093/mnras/staa1860. arXiv: 2003.06015 [astro-ph.HE].
- [91] M. Kölsch et al. “Investigating the mass-ratio dependence of the prompt-collapse threshold with numerical-relativity simulations”. In: *Phys. Rev. D* 106.4 (2022), p. 044026. DOI: 10.1103/PhysRevD.106.044026. arXiv: 2112.11851 [gr-qc].
- [92] R. Kashyap et al. “Numerical relativity simulations of prompt collapse mergers: Threshold mass and phenomenological constraints on neutron star properties after GW170817”. In: *Phys. Rev. D* 105.10 (2022), p. 103022. DOI: 10.1103/PhysRevD.105.103022. arXiv: 2111.05183 [astro-ph.HE].
- [93] A. Perego et al. “Probing the Incompressibility of Nuclear Matter at Ultrahigh Density through the Prompt Collapse of Asymmetric Neutron Star Binaries”. In: *Phys. Rev. Lett.* 129.3 (2022), p. 032701. DOI: 10.1103/PhysRevLett.129.032701. arXiv: 2112.05864 [astro-ph.HE].
- [94] V. Nedora et al. “Mapping dynamical ejecta and disk masses from numerical relativity simulations of neutron star mergers”. In: *Class. Quant. Grav.* 39.1 (2022), p. 015008. DOI: 10.1088/1361-6382/ac35a8. arXiv: 2011.11110 [astro-ph.HE].
- [95] K. Kiuchi et al. “Self-consistent picture of the mass ejection from a one second-long binary neutron star merger leaving a short-lived remnant in general-relativistic neutrino-radiation magnetohydrodynamic simulation”. In: (Nov. 2022). arXiv: 2211.07637 [astro-ph.HE].

- [96] D. Radice et al. “Binary Neutron Star Mergers: Mass Ejection, Electromagnetic Counterparts and Nucleosynthesis”. In: *Astrophys. J.* 869.2 (2018), p. 130. DOI: 10.3847/1538-4357/aaf054. arXiv: 1809.11161 [astro-ph.HE].
- [97] A. Perego et al. “Production of Very Light Elements and Strontium in the Early Ejecta of Neutron Star Mergers”. In: *Astrophys. J.* 925.1 (2022), p. 22. DOI: 10.3847/1538-4357/ac3751. arXiv: 2009.08988 [astro-ph.HE].
- [98] V. Nedora et al. “Numerical Relativity Simulations of the Neutron Star Merger GW170817: Long-Term Remnant Evolutions, Winds, Remnant Disks, and Nucleosynthesis”. In: *Astrophys. J.* 906.2 (2021), p. 98. DOI: 10.3847/1538-4357/abc9be. arXiv: 2008.04333 [astro-ph.HE].
- [99] A. Perego, F. K. Thielemann, and G. Cescutti. “r-Process Nucleosynthesis from Compact Binary Mergers”. In: *Handbook of Gravitational Wave Astronomy*. 2021, 1, p. 1. DOI: 10.1007/978-981-15-4702-7\13-1.
- [100] S. Rosswog and O. Korobkin. “Heavy elements and electromagnetic transients from neutron star mergers”. In: (Aug. 2022). arXiv: 2208.14026 [astro-ph.HE].
- [101] R. Fernández et al. “The interplay of disc wind and dynamical ejecta in the aftermath of neutron star–black hole mergers”. In: *Mon. Not. Roy. Astron. Soc.* 449.1 (2015), pp. 390–402. DOI: 10.1093/mnras/stv238. arXiv: 1412.5588 [astro-ph.HE].
- [102] O. Just et al. “Comprehensive nucleosynthesis analysis for ejecta of compact binary mergers”. In: *Mon. Not. Roy. Astron. Soc.* 448.1 (2015), pp. 541–567. DOI: 10.1093/mnras/stv009. arXiv: 1406.2687 [astro-ph.SR].
- [103] D. M. Siegel and B. D. Metzger. “Three-Dimensional General-Relativistic Magnetohydrodynamic Simulations of Remnant Accretion Disks from Neutron Star Mergers: Outflows and r -Process Nucleosynthesis”. In: *Phys. Rev. Lett.* 119.23 (2017), p. 231102. DOI: 10.1103/PhysRevLett.119.231102. arXiv: 1705.05473 [astro-ph.HE].
- [104] S. Fujibayashi et al. “Mass Ejection from the Remnant of a Binary Neutron Star Merger: Viscous-Radiation Hydrodynamics Study”. In: *Astrophys. J.* 860.1 (2018), p. 64. DOI: 10.3847/1538-4357/aabafd. arXiv: 1711.02093 [astro-ph.HE].

- [105] R. Fernández et al. “Long-term GRMHD simulations of neutron star merger accretion discs: implications for electromagnetic counterparts”. In: *Mon. Not. Roy. Astron. Soc.* 482.3 (2019), pp. 3373–3393. DOI: 10.1093/mnras/sty2932. arXiv: 1808.00461 [astro-ph.HE].
- [106] A. Janiuk. “The r-process nucleosynthesis in the outflows from short GRB accretion disks”. In: (July 2019). DOI: 10.3847/1538-4357/ab3349. arXiv: 1907.00809 [astro-ph.HE].
- [107] J. M. Miller et al. “Full Transport Model of GW170817-Like Disk Produces a Blue Kilonova”. In: *Phys. Rev. D* 100.2 (2019), p. 023008. DOI: 10.1103/PhysRevD.100.023008. arXiv: 1905.07477 [astro-ph.HE].
- [108] S. Fujibayashi et al. “Mass ejection from disks surrounding a low-mass black hole: Viscous neutrino-radiation hydrodynamics simulation in full general relativity”. In: *Phys. Rev. D* 101.8 (2020), p. 083029. DOI: 10.1103/PhysRevD.101.083029. arXiv: 2001.04467 [astro-ph.HE].
- [109] O. Just et al. “Neutrino absorption and other physics dependencies in neutrino-cooled black hole accretion discs”. In: *Mon. Not. Roy. Astron. Soc.* 509.1 (2021), pp. 1377–1412. DOI: 10.1093/mnras/stab2861. arXiv: 2102.08387 [astro-ph.HE].
- [110] E. A. Huerta et al. “Supporting High-Performance and High-Throughput Computing for Experimental Science”. en. In: *Computing and Software for Big Science* 3.1 (Feb. 2019), p. 5. ISSN: 2510-2044. DOI: 10.1007/s41781-019-0022-7. URL: <https://doi.org/10.1007/s41781-019-0022-7> (visited on 01/07/2021).
- [111] S. Rosswog. “Mergers of neutron star black hole binaries with small mass ratios: Nucleosynthesis, gamma-ray bursts and electromagnetic transients”. In: *Astrophys. J.* 634 (2005), p. 1202. DOI: 10.1086/497062. arXiv: astro-ph/0508138 [astro-ph].
- [112] K. Kyutoku et al. “Gravitational waves from spinning black hole-neutron star binaries: dependence on black hole spins and on neutron star equations of state”. In: *Phys. Rev. D* 84 (2011), p. 064018. DOI: 10.1103/PhysRevD.84.064018. arXiv: 1108.1189 [astro-ph.HE].
- [113] F. Foucart et al. “Black hole-neutron star mergers at realistic mass ratios: Equation of state and spin orientation effects”. In: *Phys. Rev. D* 87 (2013), p. 084006. DOI: 10.1103/PhysRevD.87.084006. arXiv: 1212.4810 [gr-qc].

-
- [114] F. Foucart et al. “Neutron star-black hole mergers with a nuclear equation of state and neutrino cooling: Dependence in the binary parameters”. In: *Phys. Rev. D* 90 (2014), p. 024026. DOI: 10.1103/PhysRevD.90.024026. arXiv: 1405.1121 [astro-ph.HE].
- [115] F. Foucart et al. “Low mass binary neutron star mergers : gravitational waves and neutrino emission”. In: *Phys. Rev. D* 93.4 (2016), p. 044019. DOI: 10.1103/PhysRevD.93.044019. arXiv: 1510.06398 [astro-ph.HE].
- [116] V. Paschalidis. “General relativistic simulations of compact binary mergers as engines for short gamma-ray bursts”. In: *Class. Quant. Grav.* 34.8 (2017), p. 084002. DOI: 10.1088/1361-6382/aa61ce. arXiv: 1611.01519 [astro-ph.HE].
- [117] M. Bhattacharya, P. Kumar, and G. Smoot. “Mergers of black hole–neutron star binaries and rates of associated electromagnetic counterparts”. In: (2018). arXiv: 1809.00006 [astro-ph.HE].
- [118] K. Kyutoku, M. Shibata, and K. Taniguchi. “Gravitational waves from non-spinning black hole-neutron star binaries: dependence on equations of state”. In: *Phys. Rev. D* 82 (2010). [Erratum: *Phys. Rev. D* 84,049902(2011)], p. 044049. DOI: 10.1103/PhysRevD.82.044049, 10.1103/PhysRevD.84.049902. arXiv: 1008.1460 [astro-ph.HE].
- [119] K. Kyutoku, K. Ioka, and M. Shibata. “Anisotropic mass ejection from black hole-neutron star binaries: Diversity of electromagnetic counterparts”. In: *Phys. Rev. D* 88.4 (2013), p. 041503. DOI: 10.1103/PhysRevD.88.041503. arXiv: 1305.6309 [astro-ph.HE].
- [120] K. Kyutoku et al. “Dynamical mass ejection from black hole-neutron star binaries”. In: *Phys. Rev. D* 92 (2015), p. 044028. DOI: 10.1103/PhysRevD.92.044028. arXiv: 1502.05402 [astro-ph.HE].
- [121] F. Foucart et al. “Black hole-neutron star mergers: effects of the orientation of the black hole spin”. In: *Phys. Rev. D* 83 (2011), p. 024005. DOI: 10.1103/PhysRevD.83.024005. arXiv: 1007.4203 [astro-ph.HE].
- [122] K. Kawaguchi et al. “Black hole-neutron star binary merger: Dependence on black hole spin orientation and equation of state”. In: *Phys. Rev. D* 92.2 (2015), p. 024014. DOI: 10.1103/PhysRevD.92.024014. arXiv: 1506.05473 [astro-ph.HE].

- [123] Z. B. Etienne et al. “General relativistic simulations of black-hole-neutron- star mergers: Effects of black-hole spin”. In: *Phys. Rev. D* 79 (2009), p. 044024. DOI: 10.1103/PhysRevD.79.044024. arXiv: 0812.2245 [astro-ph].
- [124] Z. B. Etienne et al. “Fully General Relativistic Simulations of Black Hole- Neutron Star Mergers”. In: *Phys. Rev. D* 77 (2008), p. 084002. DOI: 10.1103/PhysRevD.77.084002. arXiv: 0712.2460 [astro-ph].
- [125] Z. B. Etienne et al. “General relativistic simulations of black hole-neutron star mergers: Effects of magnetic fields”. In: *Phys. Rev. D* 85 (2012), p. 064029. DOI: 10.1103/PhysRevD.85.064029. arXiv: 1112.0568 [astro-ph.HE].
- [126] Z. B. Etienne, V. Paschalidis, and S. L. Shapiro. “General relativistic simulations of black hole-neutron star mergers: Effects of tilted magnetic fields”. In: *Phys. Rev. D* 86 (2012), p. 084026. DOI: 10.1103/PhysRevD.86.084026. arXiv: 1209.1632 [astro-ph.HE].
- [127] V. Paschalidis, Z. B. Etienne, and S. L. Shapiro. “General relativistic simulations of binary black hole-neutron stars: Precursor electromagnetic signals”. In: *Phys. Rev. D* 88.2 (2013), p. 021504. DOI: 10.1103/PhysRevD.88.021504. arXiv: 1304.1805 [astro-ph.HE].
- [128] M. Bejger et al. “Impact of the nuclear equation of state on the last orbits of binary neutron stars”. In: *Astron. Astrophys.* 431 (2005), p. 297. DOI: 10.1051/0004-6361:20041441. arXiv: astro-ph/0406234.
- [129] T. Dietrich et al. “Gravitational waves and mass ejecta from binary neutron star mergers: Effect of the mass-ratio”. In: *Phys. Rev. D* 95.2 (2017), p. 024029. DOI: 10.1103/PhysRevD.95.024029. arXiv: 1607.06636 [gr-qc].
- [130] A. Matas et al. “Aligned-spin neutron-star-black-hole waveform model based on the effective-one-body approach and numerical-relativity simulations”. In: *Phys. Rev. D* 102.4 (2020), p. 043023. DOI: 10.1103/PhysRevD.102.043023. arXiv: 2004.10001 [gr-qc].
- [131] F. Pannarale. “Black hole remnant of black hole-neutron star coalescing binaries”. In: *Phys. Rev. D* 88.10 (2013), p. 104025. DOI: 10.1103/PhysRevD.88.104025. arXiv: 1208.5869 [gr-qc].
- [132] F. Pannarale. “Black hole remnant of black hole-neutron star coalescing binaries with arbitrary black hole spin”. In: *Phys. Rev. D* 89.4 (2014), p. 044045. DOI: 10.1103/PhysRevD.89.044045. arXiv: 1311.5931 [gr-qc].

- [133] A. Buonanno, L. E. Kidder, and L. Lehner. “Estimating the final spin of a binary black hole coalescence”. In: *Phys. Rev. D* 77 (2008), p. 026004. DOI: 10.1103/PhysRevD.77.026004. arXiv: 0709.3839 [astro-ph].
- [134] F. Foucart. “Black Hole-Neutron Star Mergers: Disk Mass Predictions”. In: *Phys. Rev. D* 86 (2012), p. 124007. DOI: 10.1103/PhysRevD.86.124007. arXiv: 1207.6304 [astro-ph.HE].
- [135] F. Foucart, T. Hinderer, and S. Nissanke. “Remnant baryon mass in neutron star-black hole mergers: Predictions for binary neutron star mimickers and rapidly spinning black holes”. In: *Phys. Rev. D* 98.8 (2018), p. 081501. DOI: 10.1103/PhysRevD.98.081501. arXiv: 1807.00011 [astro-ph.HE].
- [136] J. M. Bardeen, W. H. Press, and S. A. Teukolsky. “Rotating black holes: Locally nonrotating frames, energy extraction, and scalar synchrotron radiation”. In: *Astrophys. J.* 178 (1972), p. 347. DOI: 10.1086/151796.
- [137] T. Hinderer et al. “Tidal deformability of neutron stars with realistic equations of state and their gravitational wave signatures in binary inspiral”. In: *Phys. Rev. D* 81 (2010), p. 123016. DOI: 10.1103/PhysRevD.81.123016. arXiv: 0911.3535 [astro-ph.HE].
- [138] T. Damour and A. Nagar. “Effective One Body description of tidal effects in inspiralling compact binaries”. In: *Phys. Rev. D* 81 (2010), p. 084016. DOI: 10.1103/PhysRevD.81.084016. arXiv: 0911.5041 [gr-qc].
- [139] B. Paczynski. “Gamma-ray bursters at cosmological distances”. In: *Astrophys. J.* 308 (1986), pp. L43–L46.
- [140] D. Eichler et al. “Nucleosynthesis, Neutrino Bursts and Gamma-Rays from Coalescing Neutron Stars”. In: *Nature* 340 (1989), pp. 126–128. DOI: 10.1038/340126a0.
- [141] L.-X. Li and B. Paczynski. “Transient events from neutron star mergers”. In: *Astrophys. J.* 507 (1998), p. L59. DOI: 10.1086/311680. arXiv: astro-ph/9807272 [astro-ph].
- [142] B. D. Lackey et al. “Extracting equation of state parameters from black hole-neutron star mergers. I. Nonspinning black holes”. In: *Phys. Rev. D* 85 (2012), p. 044061. DOI: 10.1103/PhysRevD.85.044061. arXiv: 1109.3402 [astro-ph.HE].

- [143] F. Pannarale et al. “Nonspinning black hole-neutron star mergers: a model for the amplitude of gravitational waveforms”. In: *Phys. Rev. D* 88.8 (2013), p. 084011. DOI: 10.1103/PhysRevD.88.084011. arXiv: 1307.5111 [gr-qc].
- [144] F. Pannarale et al. “Aligned spin neutron star-black hole mergers: a gravitational waveform amplitude model”. In: *Phys. Rev. D* 92.8 (2015), p. 084050. DOI: 10.1103/PhysRevD.92.084050. arXiv: 1509.00512 [gr-qc].
- [145] T. Hinderer et al. “Effects of neutron-star dynamic tides on gravitational waveforms within the effective-one-body approach”. In: *Phys. Rev. Lett.* 116.18 (2016), p. 181101. DOI: 10.1103/PhysRevLett.116.181101. arXiv: 1602.00599 [gr-qc].
- [146] A. Nagar et al. “Time-domain effective-one-body gravitational waveforms for coalescing compact binaries with nonprecessing spins, tides and self-spin effects”. In: *Phys. Rev. D* 98.10 (2018), p. 104052. DOI: 10.1103/PhysRevD.98.104052. arXiv: 1806.01772 [gr-qc].
- [147] T. Yamamoto, M. Shibata, and K. Taniguchi. “Simulating coalescing compact binaries by a new code SACRA”. In: *Phys. Rev. D* 78 (2008), p. 064054. DOI: 10.1103/PhysRevD.78.064054. arXiv: 0806.4007 [gr-qc].
- [148] F. Zappa et al. “Black-hole remnants from black-hole–neutron-star mergers”. In: *Phys. Rev. Lett.* 123.4 (2019), p. 041102. DOI: 10.1103/PhysRevLett.123.041102. arXiv: 1903.11622 [gr-qc].
- [149] L. Bildsten and C. Cutler. “Tidal interactions of inspiraling compact binaries”. In: *Astrophys. J.* 400 (Nov. 1992), pp. 175–180. DOI: 10.1086/171983.
- [150] C. S. Kochanek. “Coalescing binary neutron stars”. In: *Astrophys. J.* 398 (1992), p. 234. DOI: 10.1086/171851.
- [151] X. Jiménez-Forteza et al. “Hierarchical data-driven approach to fitting numerical relativity data for nonprecessing binary black holes with an application to final spin and radiated energy”. In: *Phys. Rev. D* 95.6 (2017), p. 064024. DOI: 10.1103/PhysRevD.95.064024. arXiv: 1611.00332 [gr-qc].
- [152] K. Hayashi et al. “Properties of the remnant disk and the dynamical ejecta produced in low-mass black hole-neutron star mergers”. In: *Phys. Rev. D* 103.4 (2021), p. 043007. DOI: 10.1103/PhysRevD.103.043007. arXiv: 2010.02563 [astro-ph.HE].

- [153] E. Barausse and L. Rezzolla. “Predicting the direction of the final spin from the coalescence of two black holes”. In: *Astrophys. J.* 704 (2009), pp. L40–L44. DOI: 10.1088/0004-637X/704/1/L40. arXiv: 0904.2577 [gr-qc].
- [154] F. Hofmann, E. Barausse, and L. Rezzolla. “The final spin from binary black holes in quasi-circular orbits”. In: *Astrophys. J.* 825.2 (2016), p. L19. DOI: 10.3847/2041-8205/825/2/L19. arXiv: 1605.01938 [gr-qc].
- [155] N. Giacobbo, M. Mapelli, and M. Spera. “Merging black hole binaries: the effects of progenitor’s metallicity, mass-loss rate and Eddington factor”. In: *Mon. Not. Roy. Astron. Soc.* 474.3 (2018), pp. 2959–2974. DOI: 10.1093/mnras/stx2933. arXiv: 1711.03556 [astro-ph.SR].
- [156] N. Giacobbo and M. Mapelli. “The progenitors of compact-object binaries: impact of metallicity, common envelope and natal kicks”. In: *Mon. Not. Roy. Astron. Soc.* 480.2 (2018), pp. 2011–2030. DOI: 10.1093/mnras/sty1999. arXiv: 1806.00001 [astro-ph.HE].
- [157] N. Giacobbo and M. Mapelli. “The impact of electron-capture supernovae on merging double neutron stars”. In: *Mon. Not. Roy. Astron. Soc.* 482.2 (2019), pp. 2234–2243. DOI: 10.1093/mnras/sty2848. arXiv: 1805.11100 [astro-ph.SR].
- [158] M. Vogelsberger et al. “Properties of galaxies reproduced by a hydrodynamic simulation”. In: *Nature* 509 (2014), pp. 177–182. DOI: 10.1038/nature13316. arXiv: 1405.1418 [astro-ph.CO].
- [159] M. Vogelsberger et al. “Introducing the Illustris Project: Simulating the coevolution of dark and visible matter in the Universe”. In: *Mon. Not. Roy. Astron. Soc.* 444.2 (2014), pp. 1518–1547. DOI: 10.1093/mnras/stu1536. arXiv: 1405.2921 [astro-ph.CO].
- [160] D. Nelson et al. “The Illustris Simulation: Public Data Release”. In: *Astron. Comput.* 13 (2015), pp. 12–37. DOI: 10.1016/j.ascom.2015.09.003. arXiv: 1504.00362 [astro-ph.CO].
- [161] M. Mapelli et al. “The cosmic merger rate of stellar black hole binaries from the Illustris simulation”. In: *Mon. Not. Roy. Astron. Soc.* 472.2 (2017), pp. 2422–2435. DOI: 10.1093/mnras/stx2123. arXiv: 1708.05722 [astro-ph.GA].
- [162] M. Mapelli and N. Giacobbo. “The cosmic merger rate of neutron stars and black holes”. In: *Mon. Not. Roy. Astron. Soc.* 479.4 (2018), pp. 4391–4398. DOI: 10.1093/mnras/sty1613. arXiv: 1806.04866 [astro-ph.HE].

- [163] M. Mapelli et al. “The properties of merging black holes and neutron stars across cosmic time”. In: *Mon. Not. Roy. Astron. Soc.* 487.1 (2019), pp. 2–13. DOI: 10.1093/mnras/stz1150. arXiv: 1902.01419 [astro-ph.HE].
- [164] F. Ozel et al. “The Black Hole Mass Distribution in the Galaxy”. In: *Astrophys. J.* 725 (2010), pp. 1918–1927. DOI: 10.1088/0004-637X/725/2/1918. arXiv: 1006.2834 [astro-ph.GA].
- [165] W. M. Farr et al. “The Mass Distribution of Stellar-Mass Black Holes”. In: *Astrophys. J.* 741 (2011), p. 103. DOI: 10.1088/0004-637X/741/2/103. arXiv: 1011.1459 [astro-ph.GA].
- [166] N. Stone, A. Loeb, and E. Berger. “Pulsations in Short GRBs from Black Hole-Neutron Star Mergers”. In: *Phys. Rev. D* 87 (2013), p. 084053. DOI: 10.1103/PhysRevD.87.084053. arXiv: 1209.4097 [astro-ph.HE].
- [167] F. Pannarale and F. Ohme. “Prospects for joint gravitational-wave and electromagnetic observations of neutron-star–black-hole coalescing binaries”. In: *Astrophys. J.* 791 (2014), p. L7. DOI: 10.1088/2041-8205/791/1/L7. arXiv: 1406.6057 [gr-qc].
- [168] J. E. Thompson et al. “Modeling the gravitational wave signature of neutron star black hole coalescences: PhenomNSBH”. In: *Phys. Rev. D* 101.12 (2020), p. 124059. DOI: 10.1103/PhysRevD.101.124059. arXiv: 2002.08383 [gr-qc].
- [169] A. Buonanno and T. Damour. “Effective one-body approach to general relativistic two-body dynamics”. In: *Phys. Rev. D* 59 (1999), p. 084006. DOI: 10.1103/PhysRevD.59.084006. arXiv: gr-qc/9811091.
- [170] T. Damour and A. Nagar. “New effective-one-body description of coalescing nonprecessing spinning black-hole binaries”. In: *Phys. Rev. D* 90.4 (2014), p. 044018. DOI: 10.1103/PhysRevD.90.044018. arXiv: 1406.6913 [gr-qc].
- [171] T. Damour, P. Jaranowski, and G. Schäfer. “Fourth post-Newtonian effective one-body dynamics”. In: *Phys. Rev. D* 91.8 (2015), p. 084024. DOI: 10.1103/PhysRevD.91.084024. arXiv: 1502.07245 [gr-qc].
- [172] A. Nagar et al. “Energetics and phasing of nonprecessing spinning coalescing black hole binaries”. In: *Phys. Rev. D* 93.4 (2016), p. 044046. DOI: 10.1103/PhysRevD.93.044046. arXiv: 1506.08457 [gr-qc].

- [173] S. Akcay et al. “Effective-one-body multipolar waveform for tidally interacting binary neutron stars up to merger”. In: *Phys. Rev. D* 99.4 (2019), p. 044051. DOI: 10.1103/PhysRevD.99.044051. arXiv: 1812.02744 [gr-qc].
- [174] A. Nagar et al. “A Multipolar Effective One Body Model for Non-Spinning Black Hole Binaries”. In: (2019). arXiv: 1904.09550 [gr-qc].
- [175] A. Nagar et al. “Multipolar effective one body waveform model for spin-aligned black hole binaries”. In: *Phys. Rev. D* 102.2 (2020), p. 024077. DOI: 10.1103/PhysRevD.102.024077. arXiv: 2001.09082 [gr-qc].
- [176] G. Riemenschneider et al. “Assessment of consistent next-to-quasicircular corrections and postadiabatic approximation in effective-one-body multipolar waveforms for binary black hole coalescences”. In: *Phys. Rev. D* 104.10 (2021), p. 104045. DOI: 10.1103/PhysRevD.104.104045. arXiv: 2104.07533 [gr-qc].
- [177] D. Keitel et al. “The most powerful astrophysical events: Gravitational-wave peak luminosity of binary black holes as predicted by numerical relativity”. In: *Phys. Rev. D* 96.2 (2017), p. 024006. DOI: 10.1103/PhysRevD.96.024006. arXiv: 1612.09566 [gr-qc].
- [178] B. P. Abbott et al. “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs”. In: *Phys. Rev. X* 9.3 (2019), p. 031040. DOI: 10.1103/PhysRevX.9.031040. arXiv: 1811.12907 [astro-ph.HE].
- [179] S. Bernuzzi. “Neutron Star Merger Remnants”. In: *Gen. Rel. Grav.* 52.11 (2020), p. 108. DOI: 10.1007/s10714-020-02752-5. arXiv: 2004.06419 [astro-ph.HE].
- [180] M. Ruffert et al. “Coalescing neutron stars: A Step towards physical models. 2. Neutrino emission, neutron tori, and gamma-ray bursts”. In: *Astron. Astrophys.* 319 (1997), pp. 122–153. arXiv: astro-ph/9606181 [astro-ph].
- [181] S. Rosswog and M. Liebendoerfer. “High resolution calculations of merging neutron stars. 2: Neutrino emission”. In: *Mon. Not. Roy. Astron. Soc.* 342 (2003), p. 673. DOI: 10.1046/j.1365-8711.2003.06579.x. arXiv: astro-ph/0302301 [astro-ph].
- [182] Y. Sekiguchi et al. “Gravitational waves and neutrino emission from the merger of binary neutron stars”. In: *Phys. Rev. Lett.* 107 (2011), p. 051102. DOI: 10.1103/PhysRevLett.107.051102. arXiv: 1105.2125 [gr-qc].

- [183] A. Perego et al. “Neutrino-driven winds from neutron star merger remnants”. In: *Mon. Not. Roy. Astron. Soc.* 443 (2014), p. 3134. DOI: 10.1093/mnras/stu1352. arXiv: 1405.6730 [astro-ph.HE].
- [184] C. Palenzuela et al. “Effects of the microphysical Equation of State in the mergers of magnetized Neutron Stars With Neutrino Cooling”. In: *Phys. Rev. D* 92.4 (2015), p. 044045. DOI: 10.1103/PhysRevD.92.044045. arXiv: 1505.01607 [gr-qc].
- [185] A. Endrizzi et al. “Thermodynamics conditions of matter in the neutrino decoupling region during neutron star mergers”. In: *Eur. Phys. J. A* 56.1 (2020), p. 15. DOI: 10.1140/epja/s10050-019-00018-6. arXiv: 1908.04952 [astro-ph.HE].
- [186] B. D. Metzger and R. Fernández. “Red or blue? A potential kilonova imprint of the delay until black hole formation following a neutron star merger”. In: *Mon. Not. Roy. Astron. Soc.* 441 (2014), p. 3444. DOI: 10.1093/mnras/stu802. arXiv: 1402.4803 [astro-ph.HE].
- [187] D. Martin et al. “Neutrino-driven winds in the aftermath of a neutron star merger: nucleosynthesis and electromagnetic transients”. In: *Astrophys. J.* 813.1 (2015), p. 2. DOI: 10.1088/0004-637X/813/1/2. arXiv: 1506.05048 [astro-ph.SR].
- [188] J. Lippuner et al. “Signatures of hypermassive neutron star lifetimes on r-process nucleosynthesis in the disc ejecta from neutron star mergers”. In: *Mon. Not. Roy. Astron. Soc.* 472.1 (2017), pp. 904–918. DOI: 10.1093/mnras/stx1987. arXiv: 1703.06216 [astro-ph.HE].
- [189] S. Wanajo et al. “Production of all the r -process nuclides in the dynamical ejecta of neutron star mergers”. In: *Astrophys. J.* 789 (2014), p. L39. DOI: 10.1088/2041-8205/789/2/L39. arXiv: 1402.7317 [astro-ph.SR].
- [190] F. Foucart et al. “Impact of an improved neutrino energy estimate on outflows in neutron star merger simulations”. In: *Phys. Rev. D* 94.12 (2016), p. 123016. DOI: 10.1103/PhysRevD.94.123016. arXiv: 1607.07450 [astro-ph.HE].
- [191] M. G. Alford et al. “Viscous Dissipation and Heat Conduction in Binary Neutron-Star Mergers”. In: *Phys. Rev. Lett.* 120.4 (2018), p. 041101. DOI: 10.1103/PhysRevLett.120.041101. arXiv: 1707.09475 [gr-qc].

-
- [192] E. R. Most et al. “Emergence of microphysical viscosity in binary neutron star post-merger dynamics”. In: (July 2022). arXiv: 2207.00442 [astro-ph.HE].
- [193] P. Hammond, I. Hawke, and N. Andersson. “Detecting the impact of nuclear reactions on neutron star mergers through gravitational waves”. In: (May 2022). arXiv: 2205.11377 [astro-ph.HE].
- [194] D. Radice et al. “A new moment-based general-relativistic neutrino-radiation transport code: Methods and first applications to neutron star mergers”. In: *Mon. Not. Roy. Astron. Soc.* 512.1 (2022), pp. 1499–1521. DOI: 10.1093/mnras/stac589. arXiv: 2111.14858 [astro-ph.HE].
- [195] K. Kiuchi et al. “High resolution numerical-relativity simulations for the merger of binary magnetized neutron stars”. In: *Phys.Rev.* D90.4 (2014), p. 041502. DOI: 10.1103/PhysRevD.90.041502. arXiv: 1407.2660 [astro-ph.HE].
- [196] K. Kiuchi et al. “Efficient magnetic-field amplification due to the Kelvin-Helmholtz instability in binary neutron star mergers”. In: *Phys. Rev. D* 92.12 (2015), p. 124034. DOI: 10.1103/PhysRevD.92.124034. arXiv: 1509.09205 [astro-ph.HE].
- [197] K. Kiuchi et al. “Global simulations of strongly magnetized remnant massive neutron stars formed in binary neutron star mergers”. In: *Phys. Rev.* D97.12 (2018), p. 124039. DOI: 10.1103/PhysRevD.97.124039. arXiv: 1710.01311 [astro-ph.HE].
- [198] S. de Haas et al. “Magnetic field effects on nucleosynthesis and kilonovae from neutron star merger remnants”. In: (Aug. 2022). arXiv: 2208.05330 [astro-ph.HE].
- [199] L. Combi and D. Siegel. “GRMHD simulations of neutron-star mergers with weak interactions: r-process nucleosynthesis and electromagnetic signatures of dynamical ejecta”. In: (June 2022). arXiv: 2206.03618 [astro-ph.HE].
- [200] M. D. Duez et al. “General relativistic hydrodynamics with viscosity: Contraction, catastrophic collapse, and disk formation in hypermassive neutron stars”. In: *Phys.Rev.* D69 (2004), p. 104030. DOI: 10.1103/PhysRevD.69.104030. arXiv: astro-ph/0402502 [astro-ph].
- [201] M. D. Duez et al. “Evolving black hole-neutron star binaries in general relativity using pseudospectral and finite difference methods”. In: *Phys. Rev.* D78 (2008), p. 104015. DOI: 10.1103/PhysRevD.78.104015. arXiv: 0809.0002 [gr-qc].

- [202] K. Hotokezaka et al. “Remnant massive neutron stars of binary neutron star mergers: Evolution process and gravitational waveform”. In: *Phys.Rev.* D88.4 (2013), p. 044026. DOI: 10.1103/PhysRevD.88.044026. arXiv: 1307.5888 [astro-ph.HE].
- [203] R. Ciolfi et al. “First 100 ms of a long-lived magnetized neutron star formed in a binary neutron star merger”. In: *Phys. Rev.* D100.2 (2019), p. 023005. DOI: 10.1103/PhysRevD.100.023005. arXiv: 1904.10222 [astro-ph.HE].
- [204] D. M. Siegel and B. D. Metzger. “Three-dimensional GRMHD simulations of neutrino-cooled accretion disks from neutron star mergers”. In: *Astrophys. J.* 858.1 (2018), p. 52. DOI: 10.3847/1538-4357/aabaec. arXiv: 1711.00868 [astro-ph.HE].
- [205] D. Radice. “General-Relativistic Large-Eddy Simulations of Binary Neutron Star Mergers”. In: *Astrophys. J.* 838.1 (2017), p. L2. DOI: 10.3847/2041-8213/aa6483. arXiv: 1703.02046 [astro-ph.HE].
- [206] M. Shibata and K. Kiuchi. “Gravitational waves from remnant massive neutron stars of binary neutron star merger: Viscous hydrodynamics effects”. In: *Phys. Rev.* D95.12 (2017), p. 123003. DOI: 10.1103/PhysRevD.95.123003. arXiv: 1705.06142 [astro-ph.HE].
- [207] R. Aguilera-Miret et al. “Turbulent magnetic-field amplification in the first 10 milliseconds after a binary neutron star merger: Comparing high-resolution and large-eddy simulations”. In: *Phys. Rev. D* 102.10 (2020), p. 103006. DOI: 10.1103/PhysRevD.102.103006. arXiv: 2009.06669 [gr-qc].
- [208] D. Radice. “Binary Neutron Star Merger Simulations with a Calibrated Turbulence Model”. In: *Symmetry* 12.8 (2020), p. 1249. DOI: 10.3390/sym12081249. arXiv: 2005.09002 [astro-ph.HE].
- [209] F. Douchin and P. Haensel. “A unified equation of state of dense matter and neutron star structure”. In: *Astron. Astrophys.* 380 (2001), pp. 151–167. eprint: astro-ph/0111092.
- [210] A. S. Schneider, L. F. Roberts, and C. D. Ott. “Open-source nuclear equation of state framework based on the liquid-drop model with Skyrme interaction”. In: *Phys. Rev.* C96.6 (2017), p. 065802. DOI: 10.1103/PhysRevC.96.065802. arXiv: 1707.01527 [astro-ph.HE].

- [211] H. T. Cromartie et al. “Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar”. In: *Nat. Astron.* 4.1 (2019), pp. 72–76. DOI: 10.1038/s41550-019-0880-2. arXiv: 1904.06759 [astro-ph.HE].
- [212] E. Fonseca et al. “Refined Mass and Geometric Measurements of the High-mass PSR J0740+6620”. In: *Astrophys. J. Lett.* 915.1 (2021), p. L12. DOI: 10.3847/2041-8213/ac03b8. arXiv: 2104.00880 [astro-ph.HE].
- [213] M. C. Miller et al. “PSR J0030+0451 Mass and Radius from *NICER* Data and Implications for the Properties of Neutron Star Matter”. In: *Astrophys. J.* 887.1 (2019), p. L24. DOI: 10.3847/2041-8213/ab50c5. arXiv: 1912.05705 [astro-ph.HE].
- [214] T. E. Riley et al. “A *NICER* View of PSR J0030+0451: Millisecond Pulsar Parameter Estimation”. In: *Astrophys. J.* 887.1 (2019), p. L21. DOI: 10.3847/2041-8213/ab481c. arXiv: 1912.05702 [astro-ph.HE].
- [215] R. Abbott et al. “Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo”. In: (Dec. 2019). arXiv: 1912.11716 [gr-qc].
- [216] E. Gourgoulhon et al. In: (2016). Paris Observatory, Meudon section - LUTH laboratory.
- [217] D. Radice and L. Rezzolla. “THC: a new high-order finite-difference high-resolution shock-capturing code for special-relativistic hydrodynamics”. In: *Astron. Astrophys.* 547 (2012), A26. DOI: 10.1051/0004-6361/201219735. arXiv: 1206.6502 [astro-ph.IM].
- [218] D. Radice, L. Rezzolla, and F. Galeazzi. “Beyond second-order convergence in simulations of binary neutron stars in full general-relativity”. In: *Mon.Not.Roy.Astron.Soc.* 437 (2014), pp. L46–L50. DOI: 10.1093/mnrasl/s1t137. arXiv: 1306.6052 [gr-qc].
- [219] D. Radice, L. Rezzolla, and F. Galeazzi. “High-Order Numerical-Relativity Simulations of Binary Neutron Stars”. In: *ASP Conf. Ser.* 498 (2015), pp. 121–126. arXiv: 1502.00551 [gr-qc].
- [220] D. Radice, L. Rezzolla, and F. Galeazzi. “High-Order Fully General-Relativistic Hydrodynamics: new Approaches and Tests”. In: *Class.Quant.Grav.* 31 (2014), p. 075012. DOI: 10.1088/0264-9381/31/7/075012. arXiv: 1312.5004 [gr-qc].

- [221] D. Radice et al. “Dynamical Mass Ejection from Binary Neutron Star Mergers”. In: *Mon. Not. Roy. Astron. Soc.* 460.3 (2016), pp. 3255–3271. DOI: 10.1093/mnras/stw1227. arXiv: 1601.02426 [astro-ph.HE].
- [222] T. Goodale et al. “The Cactus Framework and Toolkit: Design and Applications”. In: *Vector and Parallel Processing – VECPAR’2002, 5th International Conference, Lecture Notes in Computer Science*. Berlin: Springer, 2003.
- [223] E. Schnetter et al. “Cactus Framework: Black Holes to Gamma Ray Bursts”. In: (2007). arXiv: 0707.1607 [cs.DC].
- [224] C. Reisswig et al. “Formation and Coalescence of Cosmological Supermassive Black Hole Binaries in Supermassive Star Collapse”. In: *Phys.Rev.Lett.* 111 (2013), p. 151101. DOI: 10.1103/PhysRevLett.111.151101. arXiv: 1304.7787 [astro-ph.CO].
- [225] F. Loffler et al. “The Einstein Toolkit: A Community Computational Infrastructure for Relativistic Astrophysics”. In: *Class. Quant. Grav.* 29 (2012), p. 115001. DOI: 10.1088/0264-9381/29/11/115001. arXiv: 1111.3344 [gr-qc].
- [226] <http://www.einsteintoolkit.org>. EinsteinToolkit: A Community Toolkit for Numerical Relativity.
- [227] S. Gottlieb and C.-W. Ketcheson David I.and Shu. “High Order Strong Stability Preserving Time Discretizations”. In: *Journal of Scientific Computing* 38.3 (2009), pp. 251–289. ISSN: 1573-7691. DOI: 10.1007/s10915-008-9239-z. URL: <https://doi.org/10.1007/s10915-008-9239-z>.
- [228] M. J. Berger and J. Olinger. “Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations”. In: *J.Comput.Phys.* 53 (1984), p. 484.
- [229] M. J. Berger and P. Colella. “Local adaptive mesh refinement for shock hydrodynamics”. In: *Journal of Computational Physics* 82 (May 1989), pp. 64–84. DOI: 10.1016/0021-9991(89)90035-1.
- [230] C. Reisswig et al. “Three-Dimensional General-Relativistic Hydrodynamic Simulations of Binary Neutron Star Coalescence and Stellar Collapse with Multi-patch Grids”. In: *Phys. Rev.* D87.6 (2013), p. 064023. DOI: 10.1103/PhysRevD.87.064023. arXiv: 1212.1191 [astro-ph.HE].
- [231] E. Schnetter, S. H. Hawley, and I. Hawke. “Evolutions in 3-D numerical relativity using fixed mesh refinement”. In: *Class.Quant.Grav.* 21 (2004), pp. 1465–1488. DOI: 10.1088/0264-9381/21/6/014. arXiv: gr-qc/0310042 [gr-qc].

- [232] T. Damour et al. “Energy versus Angular Momentum in Black Hole Binaries”. In: *Phys.Rev.Lett.* 108 (2012), p. 131101. DOI: 10.1103/PhysRevLett.108.131101. arXiv: 1110.2938 [gr-qc].
- [233] S. Bernuzzi et al. “Tidal effects in binary neutron star coalescence”. In: *Phys.Rev. D* 86 (2012), p. 044030. DOI: 10.1103/PhysRevD.86.044030. arXiv: 1205.3403 [gr-qc].
- [234] S. Bernuzzi, M. Thierfelder, and B. Brügmann. “Accuracy of numerical relativity waveforms from binary neutron star mergers and their comparison with post-Newtonian waveforms”. In: *Phys.Rev. D* 85 (2012), p. 104030. DOI: 10.1103/PhysRevD.85.104030. arXiv: 1109.3611 [gr-qc].
- [235] S. Bernuzzi et al. *How loud are neutron star mergers?* July 2016. DOI: 10.5281/zenodo.57844. URL: <https://doi.org/10.5281/zenodo.57844>.
- [236] M. Shibata, K. Kiuchi, and Y.-i. Sekiguchi. “General relativistic viscous hydrodynamics of differentially rotating neutron stars”. In: *Phys. Rev. D* 95 (2017), p. 083005. DOI: 10.1103/PhysRevD.95.083005. arXiv: 1703.10303 [astro-ph.HE].
- [237] D. Radice et al. “Probing Extreme-Density Matter with Gravitational Wave Observations of Binary Neutron Star Merger Remnants”. In: *Astrophys. J.* 842.2 (2017), p. L10. DOI: 10.3847/2041-8213/aa775f. arXiv: 1612.06429 [astro-ph.HE].
- [238] W. Kastaun, R. Ciolfi, and B. Giacomazzo. “Structure of Stable Binary Neutron Star Merger Remnants: a Case Study”. In: *Phys. Rev. D* 94.4 (2016), p. 044060. DOI: 10.1103/PhysRevD.94.044060. arXiv: 1607.02186 [astro-ph.HE].
- [239] M. Hanauske et al. “Rotational properties of hypermassive neutron stars from binary mergers”. In: *Phys. Rev. D* 96.4 (2017), p. 043004. DOI: 10.1103/PhysRevD.96.043004. arXiv: 1611.07152 [gr-qc].
- [240] M. Shibata et al. “Modeling GW170817 based on numerical relativity and its implications”. In: *Phys. Rev. D* 96.12 (2017), p. 123012. DOI: 10.1103/PhysRevD.96.123012. arXiv: 1710.07579 [astro-ph.HE].
- [241] A. Camilletti et al. “Numerical relativity simulations of the neutron star merger GW190425: microphysics and mass ratio effects”. In: (Apr. 2022). DOI: 10.1093/mnras/stac2333. arXiv: 2204.05336 [astro-ph.HE].

- [242] P. Mösta et al. “A magnetar engine for short GRBs and kilonovae”. In: *Astrophys. J. Lett.* 901 (2020), p. L37. DOI: 10.3847/2041-8213/abb6ef. arXiv: 2003.06043 [astro-ph.HE].
- [243] T. Damour, A. Nagar, and M. Trias. “Accuracy and effectualness of closed-form, frequency-domain waveforms for non-spinning black hole binaries”. In: *Phys. Rev. D* 83 (2011), p. 024006. DOI: 10.1103/PhysRevD.83.024006. arXiv: 1009.5998 [gr-qc].
- [244] K. Chatziioannou et al. “Constructing Gravitational Waves from Generic Spin-Precessing Compact Binary Inspirals”. In: *Phys. Rev. D* 95.10 (2017), p. 104004. DOI: 10.1103/PhysRevD.95.104004. arXiv: 1703.03967 [gr-qc].
- [245] L. Dessart et al. “Neutrino signatures and the neutrino-driven wind in Binary Neutron Star Mergers”. In: *Astrophys. J.* 690 (2009), p. 1681. DOI: 10.1088/0004-637X/690/2/1681. arXiv: 0806.4380 [astro-ph].
- [246] S. Fujibayashi et al. “Properties of Neutrino-driven Ejecta from the Remnant of a Binary Neutron Star Merger: Pure Radiation Hydrodynamics Case”. In: *Astrophys. J.* 846.2 (2017), p. 114. DOI: 10.3847/1538-4357/aa8039. arXiv: 1703.10191 [astro-ph.HE].
- [247] F. Foucart et al. “Estimating outflow masses and velocities in merger simulations: impact of r-process heating and neutrino cooling”. In: (Sept. 2021). arXiv: 2109.00565 [astro-ph.HE].
- [248] V. Nedora et al. “Dynamical ejecta synchrotron emission as possible contributor to the rebrightening of GRB170817A”. In: *preprint (ArXiv:2104.04537)* (Apr. 2021). arXiv: 2104.04537 [astro-ph.HE].
- [249] E. Nakar and T. Piran. “Implications of the radio and X-ray emission that followed GW170817”. In: *Mon. Not. Roy. Astron. Soc.* 478.1 (2018), pp. 407–415. DOI: 10.1093/mnras/sty952. arXiv: 1801.09712 [astro-ph.HE].
- [250] K. Hotokezaka et al. “Synchrotron radiation from the fast tail of dynamical ejecta of neutron star mergers”. In: *Astrophys. J.* 867.2 (2018), p. 95. DOI: 10.3847/1538-4357/aadf92. arXiv: 1803.00599 [astro-ph.HE].
- [251] A. Hajela et al. “Evidence for X-Ray Emission in Excess to the Jet-afterglow Decay 3.5 yr after the Binary Neutron Star Merger GW 170817: A New Emission Component”. In: *Astrophys. J. Lett.* 927.1 (2022), p. L17. DOI: 10.3847/2041-8213/ac504a. arXiv: 2104.02070 [astro-ph.HE].

- [252] G. Sadeh, O. Guttman, and E. Waxman. “Non-thermal emission from mildly relativistic dynamical ejecta of neutron star mergers”. In: (July 2022). arXiv: 2207.05746 [astro-ph.HE].
- [253] J. Lippuner and L. F. Roberts. “SkyNet: A modular nuclear reaction network library”. In: *Astrophys. J. Suppl.* 233.2 (2017), p. 18. DOI: 10.3847/1538-4365/aa94cb. arXiv: 1706.06198 [astro-ph.HE].
- [254] C. Arlandini et al. “Neutron capture in low mass asymptotic giant branch stars: cross-sections and abundance signatures”. In: *Astrophys. J.* 525 (1999), pp. 886–900. DOI: 10.1086/307938. arXiv: astro-ph/9906266 [astro-ph].
- [255] M. R. Mumpower et al. “Reverse engineering nuclear properties from rare earth abundances in the r process”. In: *J. Phys. G* 44.3 (2017), p. 034003. DOI: 10.1088/1361-6471/44/3/034003. arXiv: 1609.09858 [nucl-th].
- [256] M.-R. Wu and P. Banerjee. “The production of actinides in neutron star mergers”. In: *AAPPS Bull.* 32.1 (2022), p. 19. DOI: 10.1007/s43673-022-00048-1. arXiv: 2205.11806 [astro-ph.HE].
- [257] S. Goriely et al. “Impact of weak interactions of free nucleons on the r -process in dynamical ejecta from neutron-star mergers”. In: *Mon. Not. Roy. Astron. Soc.* 452.4 (2015), pp. 3894–3904. DOI: 10.1093/mnras/stv1526. arXiv: 1504.04377 [astro-ph.SR].
- [258] D. Martin et al. “The role of weak interactions in dynamic ejecta from binary neutron star mergers”. In: *Class. Quant. Grav.* 35.3 (2018), p. 034001. DOI: 10.1088/1361-6382/aa9f5a. arXiv: 1710.04900 [astro-ph.HE].
- [259] Z. Wu et al. “Radiation hydrodynamics modelling of kilonovae with SNEC”. In: *Mon. Not. Roy. Astron. Soc.* 512.1 (2022), pp. 328–347. DOI: 10.1093/mnras/stac399. arXiv: 2111.06870 [astro-ph.HE].
- [260] V. Morozova et al. “Light Curves of Core-Collapse Supernovae with Substantial Mass Loss using the New Open-Source SuperNova Explosion Code (SNEC)”. In: *Astrophys. J.* 814.1 (2015), p. 63. DOI: 10.1088/0004-637X/814/1/63. arXiv: 1505.06746 [astro-ph.HE].
- [261] K. Kawaguchi, M. Shibata, and M. Tanaka. “Diversity of Kilonova Light Curves”. In: *Astrophys. J.* 889.2, 171 (Feb. 2020), p. 171. DOI: 10.3847/1538-4357/ab61f6. arXiv: 1908.05815 [astro-ph.HE].

- [262] O. Korobkin et al. “Axisymmetric Radiative Transfer Models of Kilonovae”. In: (Mar. 2020). arXiv: 2004.00102 [astro-ph.HE].
- [263] S. Rosswog, E. Ramirez-Ruiz, and M. B. Davies. “High Resolution Calculations of Merging Neutron Stars. 3. Gamma-Ray Bursts”. In: *Mon. Not. Roy. Astron. Soc.* 345 (2003), p. 1077. DOI: 10.1046/j.1365-2966.2003.07032.x. arXiv: astro-ph/0306418 [astro-ph].
- [264] Y. Sekiguchi et al. “Dynamical mass ejection from binary neutron star mergers: Radiation-hydrodynamics study in general relativity”. In: *Phys.Rev.* D91.6 (2015), p. 064059. DOI: 10.1103/PhysRevD.91.064059. arXiv: 1502.06660 [astro-ph.HE].
- [265] M. Cusinato et al. “Neutrino emission from binary neutron star mergers: characterizing light curves and mean energies”. In: (Nov. 2021). DOI: 10.1140/epja/s10050-022-00743-5. arXiv: 2111.13005 [astro-ph.HE].
- [266] Y. Qian and S. Woosley. “Nucleosynthesis in neutrino driven winds: 1. The Physical conditions”. In: *Astrophys. J.* 471 (1996), pp. 331–351. DOI: 10.1086/177973. arXiv: astro-ph/9611094.
- [267] D. Pollney et al. “High accuracy binary black hole simulations with an extended wave zone”. In: *Phys. Rev.* D83 (2011), p. 044045. DOI: 10.1103/PhysRevD.83.044045. arXiv: 0910.3803 [gr-qc].
- [268] D. Brown et al. “Turduckening Black Holes: An Analytical and Computational Study”. In: *Physical Review D* 79.4 (Feb. 2009), p. 044023. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.79.044023. arXiv: 0809.3533. URL: <http://arxiv.org/abs/0809.3533> (visited on 07/06/2020).
- [269] U. Sperhake. “Binary Black-Hole Evolutions of Excision and Puncture Data”. In: *Physical Review D* 76.10 (Nov. 2007), p. 104015. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.76.104015. arXiv: gr-qc/0606079. URL: <http://arxiv.org/abs/gr-qc/0606079> (visited on 07/03/2020).
- [270] Y. Zlochower et al. “Accurate Black Hole Evolutions by Fourth-Order Numerical Relativity”. In: *Physical Review D* 72.2 (July 2005), p. 024021. DOI: 10.1103/PhysRevD.72.024021. URL: <https://link.aps.org/doi/10.1103/PhysRevD.72.024021> (visited on 12/18/2020).

- [271] F. Herrmann et al. “Unequal Mass Binary Black Hole Plunges and Gravitational Recoil”. In: *Classical and Quantum Gravity* 24.12 (June 2007), S33–S42. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/24/12/S04. arXiv: gr-qc/0601026. URL: <http://arxiv.org/abs/gr-qc/0601026> (visited on 01/15/2021).
- [272] P. Mösta et al. “GRHydro: A new open source general-relativistic magnetohydrodynamics code for the Einstein Toolkit”. In: *Class.Quant.Grav.* 31 (2014), p. 015005. DOI: 10.1088/0264-9381/31/1/015005. arXiv: 1304.5544 [gr-qc].
- [273] B. Brügmann et al. “Calibration of Moving Puncture Simulations”. In: *Phys.Rev.D* 77 (2008), p. 024027. DOI: 10.1103/PhysRevD.77.024027. arXiv: gr-qc/0610128 [gr-qc].
- [274] P. Galaviz, B. Bruegmann, and Z. Cao. “Numerical Evolution of Multiple Black Holes with Accurate Initial Data”. In: *Physical Review D* 82.2 (July 2010), p. 024005. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.82.024005. arXiv: 1004.1353. URL: <http://arxiv.org/abs/1004.1353> (visited on 01/07/2021).
- [275] M. Thierfelder, S. Bernuzzi, and B. Brügmann. “Numerical relativity simulations of binary neutron stars”. In: *Phys.Rev.D* 84 (2011), p. 044012. DOI: 10.1103/PhysRevD.84.044012. arXiv: 1104.4751 [gr-qc].
- [276] Z. Cao, H.-J. Yo, and J.-P. Yu. “A Reinvestigation of Moving Punctured Black Holes with a New Code”. In: *Physical Review D* 78.12 (Dec. 2008), p. 124011. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.78.124011. arXiv: 0812.0641. URL: <http://arxiv.org/abs/0812.0641> (visited on 01/07/2021).
- [277] K. Clough et al. “GRChombo : Numerical Relativity with Adaptive Mesh Refinement”. In: (2015). arXiv: 1503.03436 [gr-qc].
- [278] Q. F. Stout et al. “Adaptive Blocks: A High Performance Data Structure”. In: *Proceedings of the 1997 ACM/IEEE Conference on Supercomputing. SC '97*. New York, NY, USA: Association for Computing Machinery, Nov. 1997, pp. 1–10. ISBN: 978-0-89791-985-2. DOI: 10.1145/509593.509650. URL: <https://doi.org/10.1145/509593.509650> (visited on 01/07/2021).
- [279] C. Burstedde, L. C. Wilcox, and O. Ghattas. “p4est: Scalable Algorithms for Parallel Adaptive Mesh Refinement on Forests of Octrees”. In: *SIAM Journal on Scientific Computing* 33.3 (2011), pp. 1103–1133. DOI: 10.1137/100791634.

- [280] M. Fernando et al. “Massively Parallel Simulations of Binary Black Hole Intermediate-Mass-Ratio Inspirals”. In: (July 2018). DOI: 10.1137/18M1196972. arXiv: 1807.06128 [gr-qc].
- [281] C. J. White, J. M. Stone, and C. F. Gammie. “An Extension of the Athena++ Code Framework for GRMHD Based on Advanced Riemann Solvers and Staggered-Mesh Constrained Transport”. In: *The Astrophysical Journal Supplement Series* 225.2 (Aug. 2016), p. 22. ISSN: 1538-4365. DOI: 10.3847/0067-0049/225/2/22. arXiv: 1511.00943. URL: <http://arxiv.org/abs/1511.00943> (visited on 07/10/2020).
- [282] K. G. Felker and J. M. Stone. “A Fourth-Order Accurate Finite Volume Method for Ideal MHD via Upwind Constrained Transport”. In: *Journal of Computational Physics* 375 (Dec. 2018), pp. 1365–1400. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2018.08.025. URL: <http://adsabs.harvard.edu/abs/2018JCoPh.375.1365F> (visited on 07/10/2020).
- [283] J. M. Stone et al. “The Athena \mathbb{P} Adaptive Mesh Refinement Framework: Design and Magnetohydrodynamic Solvers”. en. In: *The Astrophysical Journal Supplement Series* 249.1 (June 2020), p. 4. ISSN: 0067-0049. DOI: 10.3847/1538-4365/ab929b. URL: <https://doi.org/10.3847/1538-4365/ab929b> (visited on 07/03/2020).
- [284] J. M. Stone et al. “Athena: A New Code for Astrophysical MHD”. en. In: *The Astrophysical Journal Supplement Series* 178.1 (Sept. 2008), p. 137. ISSN: 0067-0049. DOI: 10.1086/588755. URL: <https://iopscience.iop.org/article/10.1086/588755/meta> (visited on 07/10/2020).
- [285] G. M. Morton. *A computer oriented geodetic data base and a new technique in file sequencing*. Tech. rep. 1966.
- [286] C. Burstedde, J. Holke, and T. Isaac. “On the Number of Face-Connected Components of Morton-Type Space-Filling Curves”. en. In: *Foundations of Computational Mathematics* 19.4 (Aug. 2019), pp. 843–868. ISSN: 1615-3383. DOI: 10.1007/s10208-018-9400-5. URL: <https://doi.org/10.1007/s10208-018-9400-5> (visited on 12/10/2020).
- [287] J.-P. Berrut and L. N. Trefethen. “Barycentric Lagrange Interpolation”. In: *SIAM Review* 46.3 (Jan. 2004), pp. 501–517. ISSN: 0036-1445, 1095-7200. DOI: 10.1137/S0036144502417715. URL: <http://epubs.siam.org/doi/abs/10.1137/S0036144502417715> (visited on 02/11/2013).

- [288] Z. Cao and D. Hilditch. “Numerical stability of the Z4c formulation of general relativity”. In: *Phys.Rev. D* 85 (2012), p. 124032. DOI: 10.1103/PhysRevD.85.124032. arXiv: 1111.2177 [gr-qc].
- [289] D. I. Ketcheson. “Runge–Kutta Methods with Minimum Storage Implementations”. en. In: *Journal of Computational Physics* 229.5 (Mar. 2010), pp. 1763–1773. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2009.11.006. URL: <http://www.sciencedirect.com/science/article/pii/S0021999109006251> (visited on 11/03/2020).
- [290] S. Husa et al. “Reducing phase error in long numerical binary black hole evolutions with sixth order finite differencing”. In: *Class. Quant. Grav.* 25 (2008), p. 105006. DOI: 10.1088/0264-9381/25/10/105006. arXiv: 0706.0740 [gr-qc].
- [291] M. Chirvasa and S. Husa. “Finite Difference Methods for Second Order in Space, First Order in Time Hyperbolic Systems and the Linear Shifted Wave Equation as a Model Problem in Numerical Relativity”. en. In: *Journal of Computational Physics* 229.7 (Apr. 2010), pp. 2675–2696. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2009.12.016. URL: <http://www.sciencedirect.com/science/article/pii/S0021999109006950> (visited on 12/18/2020).
- [292] R. Alfieri et al. “Optimization of Finite-Differencing Kernels for Numerical Relativity Applications”. In: *Journal of Low Power Electronics and Applications*. Special Issue "Energy Aware Scientific Computing on Low Power and Heterogeneous Architectures" 8.2 (2018). ISSN: 2079-9268. DOI: 10.3390/jlpea8020015. URL: <http://www.mdpi.com/2079-9268/8/2/15>.
- [293] H. O. Kreiss and J. Olinger. *Methods for the approximate solution of time dependent problems*. Geneva: International Council of Scientific Unions, World Meteorological Organization, 1973.
- [294] B. Gustafsson, H.-O. Kreiss, and J. Olinger. *Time-dependent problems and difference methods; 2nd ed.* Pure and applied mathematics a wiley series of texts, monographs and tracts. Somerset: Wiley, 2013. URL: <https://cds.cern.ch/record/2122877>.
- [295] D. Hilditch and M. Ruiz. “The initial boundary value problem for free-evolution formulations of General Relativity”. In: *Class. Quant. Grav.* 35.1 (2018), p. 015006. DOI: 10.1088/1361-6382/aa96c6. arXiv: 1609.06925 [gr-qc].

- [296] D. Hilditch et al. “Compact binary evolutions with the Z4c formulation”. In: *Phys. Rev. D* 88 (2013), p. 084057. DOI: 10.1103/PhysRevD.88.084057. arXiv: 1212.2901 [gr-qc].
- [297] M. Alcubierre et al. “Toward standard testbeds for numerical relativity”. In: *Class. Quant. Grav.* 21.2 (2004), pp. 589–613. DOI: 10.1088/0264-9381/21/2/019. arXiv: gr-qc/0305023.
- [298] M. Babiuc et al. “Implementation of standard testbeds for numerical relativity”. In: *Class. Quant. Grav.* 25 (2008), p. 125012. DOI: 10.1088/0264-9381/25/12/125012. arXiv: 0709.3559 [gr-qc].
- [299] D. Daverio, Y. Dirian, and E. Mitsou. “Apples with Apples comparison of 3+1 conformal numerical relativity schemes”. In: (Oct. 2018). arXiv: 1810.12346 [gr-qc].
- [300] B. Daszuta et al. “GR-Athena++: Puncture Evolutions on Vertex-centered Oct-tree Adaptive Mesh Refinement”. In: *Astrophys. J. Supp.* 257.2 (2021), p. 25. DOI: 10.3847/1538-4365/ac157b. arXiv: 2101.08289 [gr-qc].
- [301] D. R. Brill and R. W. Lindquist. “Interaction Energy in Geometrostatics”. In: *Phys. Rev.* 131 (1 1963), pp. 471–476. DOI: 10.1103/PhysRev.131.471. URL: <https://link.aps.org/doi/10.1103/PhysRev.131.471>.
- [302] J. W. York. “Kinematics and Dynamics of General Relativity”. In: *Sources of gravitational radiation*. Ed. by L. L. Smarr. Cambridge, UK: Cambridge University Press, 1979, pp. 83–126. ISBN: 0-521-22778-X.
- [303] J. M. Bowen and J. W. York Jr. “Time asymmetric initial data for black holes and black hole collisions”. In: *Phys. Rev. D* 21 (1980), pp. 2047–2056. DOI: 10.1103/PhysRevD.21.2047.
- [304] M. Ansorg, B. Brügmann, and W. Tichy. “A single-domain spectral method for black hole puncture data”. In: *Phys. Rev. D* 70 (2004), p. 064011. DOI: 10.1103/PhysRevD.70.064011. arXiv: gr-qc/0404056.
- [305] T. Damour et al. “Accurate Effective-One-Body waveforms of inspiralling and coalescing black-hole binaries”. In: *Phys. Rev. D* 78 (2008), p. 044039. DOI: 10.1103/PhysRevD.78.044039. arXiv: 0803.3162 [gr-qc].
- [306] L. Baiotti et al. “Gravitational-Wave Extraction from Neutron Stars Oscillations: comparing linear and nonlinear techniques”. In: *Phys. Rev. D* 79 (2009), p. 024002. DOI: 10.1103/PhysRevD.79.024002. arXiv: 0808.4002 [gr-qc].

- [307] T. Dietrich and S. Bernuzzi. “Simulations of rotating neutron star collapse with the puncture gauge: end state and gravitational waveforms”. In: *Phys. Rev. D* 91.4 (2015), p. 044039. DOI: 10.1103/PhysRevD.91.044039. arXiv: 1412.5499 [gr-qc].
- [308] S. Bernuzzi and T. Dietrich. “Gravitational waveforms from binary neutron star mergers with high-order weighted-essentially-nonscillatory schemes in numerical relativity”. In: *Phys. Rev. D* 94.6 (2016), p. 064062. DOI: 10.1103/PhysRevD.94.064062. arXiv: 1604.07999 [gr-qc].
- [309] R. Abbott et al. “GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run”. In: *Phys. Rev. X* 11 (2021), p. 021053. DOI: 10.1103/PhysRevX.11.021053. arXiv: 2010.14527 [gr-qc].
- [310] M. Hannam et al. “Simulations of black-hole binaries with unequal masses or nonprecessing spins: Accuracy, physical properties, and comparison with post-Newtonian results”. In: *Phys. Rev. D* 82 (2010), p. 124008. DOI: 10.1103/PhysRevD.82.124008. arXiv: 1007.4789 [gr-qc].
- [311] C. Reisswig and D. Pollney. “Notes on the integration of numerical relativity waveforms”. In: *Class. Quant. Grav.* 28 (2011), p. 195015. DOI: 10.1088/0264-9381/28/19/195015. arXiv: 1006.1632 [gr-qc].
- [312] S. Bernuzzi et al. “Quasi-universal properties of neutron star mergers”. In: *Phys. Rev. Lett.* 112 (2014), p. 201101. DOI: 10.1103/PhysRevLett.112.201101. arXiv: 1402.6244 [gr-qc].
- [313] H. Nakano. “A note on gravitational wave extraction from binary simulations”. In: *Classical and Quantum Gravity* 32.17 (2015), p. 177002. ISSN: 1361-6382. DOI: 10.1088/0264-9381/32/17/177002. URL: <http://dx.doi.org/10.1088/0264-9381/32/17/177002>.
- [314] M. Hannam et al. “The Samurai Project: verifying the consistency of black-hole-binary waveforms for gravitational-wave detection”. In: *Phys. Rev. D* 79 (2009), p. 084025. DOI: 10.1103/PhysRevD.79.084025. arXiv: 0901.2437 [gr-qc].
- [315] H. C. Edwards, C. R. Trott, and D. Sunderland. “Kokkos: Enabling many-core performance portability through polymorphic memory access patterns”. In: *Journal of Parallel and Distributed Computing* 74.12 (2014). Domain-Specific Languages and High-Level Frameworks for High-Performance Computing, pp. 3202

- 3216. ISSN: 0743-7315. DOI: <https://doi.org/10.1016/j.jpdc.2014.07.003>. URL: <http://www.sciencedirect.com/science/article/pii/S0743731514001257>.
- [316] C. R. Trott et al. “Kokkos 3: Programming Model Extensions for the Exascale Era”. In: *IEEE Transactions on Parallel and Distributed Systems* 33.4 (2022), pp. 805–817. DOI: 10.1109/TPDS.2021.3097283.
- [317] C. Trott et al. “The Kokkos EcoSystem: Comprehensive Performance Portability for High Performance Computing”. In: *Computing in Science Engineering* 23.5 (2021), pp. 10–18. DOI: 10.1109/MCSE.2021.3098509.
- [318] P. Grete, F. W. Glines, and B. W. O’Shea. “K-Athena: A Performance Portable Structured Grid Finite Volume Magnetohydrodynamics Code”. In: *arXiv:1905.04341 [astro-ph, physics:physics]* (May 2019). arXiv: 1905.04341 [astro-ph, physics:physics]. URL: <http://arxiv.org/abs/1905.04341> (visited on 07/10/2020).
- [319] P. Grete et al. “Parthenon—a performance portable block-structured adaptive mesh refinement framework”. In: *The International Journal of High Performance Computing Applications* (2022), p. 109434202211437. DOI: 10.1177/10943420221143775. URL: <https://doi.org/10.1177%2F10943420221143775>.
- [320] A. Rashti et al. “New pseudospectral code for the construction of initial data”. In: *Phys. Rev. D* 105.10 (2022), p. 104027. DOI: 10.1103/PhysRevD.105.104027. arXiv: 2109.14511 [gr-qc].
- [321] A. Sur et al. “Long-term general relativistic magnetohydrodynamics simulations of magnetic field in isolated neutron stars”. In: *Mon. Not. Roy. Astron. Soc.* 511.3 (2022), pp. 3983–3993. DOI: 10.1093/mnras/stac353. arXiv: 2108.11858 [astro-ph.HE].
- [322] C. O. Lousto and J. Healy. “Study of the Intermediate Mass Ratio Black Hole Binary Merger up to 1000:1 with Numerical Relativity”. In: (Mar. 2022). arXiv: 2203.08831 [gr-qc].
- [323] P. Amaro-Seoane et al. “Astrophysics with the Laser Interferometer Space Antenna”. In: *Living Rev. Rel.* 26.1 (2023), p. 2. DOI: 10.1007/s41114-022-00041-y. arXiv: 2203.06016 [gr-qc].
- [324] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. San Francisco: W. H. Freeman, 1973. ISBN: 9780716703440, 9780691177793.

- [325] M. Alcubierre. *Introduction to 3+1 Numerical Relativity*. Ed. by M. Alcubierre. Oxford University Press, 2008.
- [326] R. L. Arnowitt, S. Deser, and C. W. Misner. “The Dynamics of general relativity”. In: *Gen. Rel. Grav.* 40 (2008), pp. 1997–2027. DOI: 10.1007/s10714-008-0661-1. arXiv: gr-qc/0405109 [gr-qc].
- [327] T. Nakamura, K. Oohara, and Y. Kojima. “General Relativistic Collapse to Black Holes and Gravitational Waves from Black Holes”. In: *Prog. Theor. Phys. Suppl.* 90 (1987), pp. 1–218.
- [328] M. Shibata and T. Nakamura. “Evolution of three-dimensional gravitational waves: Harmonic slicing case”. In: *Phys. Rev. D* 52 (1995), pp. 5428–5444. DOI: 10.1103/PhysRevD.52.5428.
- [329] T. W. Baumgarte and S. L. Shapiro. “On the numerical integration of Einstein’s field equations”. In: *Phys. Rev. D* 59 (1999), p. 024007. DOI: 10.1103/PhysRevD.59.024007. arXiv: gr-qc/9810065.
- [330] C. Bona et al. “General-covariant evolution formalism for Numerical Relativity”. In: *Phys. Rev. D* 67 (2003), p. 104005. DOI: 10.1103/PhysRevD.67.104005. arXiv: gr-qc/0302083.
- [331] S. Bernuzzi and D. Hilditch. “Constraint violation in free evolution schemes: comparing BSSNOK with a conformal decomposition of Z4”. In: *Phys. Rev. D* 81 (2010), p. 084003. DOI: 10.1103/PhysRevD.81.084003. arXiv: 0912.2920 [gr-qc].
- [332] M. Ruiz, D. Hilditch, and S. Bernuzzi. “Constraint preserving boundary conditions for the Z4c formulation of general relativity”. In: *Phys. Rev. D* 83 (2011), p. 024025. DOI: 10.1103/PhysRevD.83.024025. arXiv: 1010.0523 [gr-qc].
- [333] A. Weyhausen, S. Bernuzzi, and D. Hilditch. “Constraint damping for the Z4c formulation of general relativity”. In: *Phys. Rev. D* 85 (2012), p. 024038. DOI: 10.1103/PhysRevD.85.024038. arXiv: 1107.5539 [gr-qc].
- [334] C. Bona and C. Palenzuela. “Dynamical shift conditions for the Z4 and BSSN hyperbolic formalisms”. In: *Phys. Rev. D* 69 (2004), p. 104003. DOI: 10.1103/PhysRevD.69.104003. arXiv: gr-qc/0401019.
- [335] C. Gundlach et al. “Constraint damping in the Z4 formulation and harmonic gauge”. In: *Class. Quant. Grav.* 22 (2005), pp. 3767–3774. DOI: 10.1088/0264-9381/22/17/025. arXiv: gr-qc/0504114.

- [336] C. Gundlach and J. M. Martin-Garcia. “Symmetric hyperbolic form of systems of second-order evolution equations subject to constraints”. In: *Phys. Rev.* D70 (2004), p. 044031. DOI: 10.1103/PhysRevD.70.044031. arXiv: gr-qc/0402079.
- [337] C. Gundlach and J. M. Martin-Garcia. “Hyperbolicity of second-order in space systems of evolution equations”. In: *Class. Quant. Grav.* 23 (2006), S387–S404. DOI: 10.1088/0264-9381/23/16/S06. arXiv: gr-qc/0506037.
- [338] C. Bona et al. “New Formalism for Numerical Relativity”. In: *Phys. Rev. Lett.* 75 (1995), pp. 600–603. eprint: gr-qc/9412071.
- [339] M. Alcubierre et al. “Gauge conditions for long term numerical black hole evolutions without excision”. In: *Phys.Rev.* D67 (2003), p. 084023. DOI: 10.1103/PhysRevD.67.084023. arXiv: gr-qc/0206072 [gr-qc].
- [340] F. Banyuls et al. “Numerical 3+1 General Relativistic Hydrodynamics: A Local Characteristic Approach”. In: *Astrophys. J.* 476 (1997), p. 221.
- [341] J. N. Goldberg et al. “Spin s spherical harmonics and edth”. In: *J. Math. Phys.* 8 (1967), p. 2155.
- [342] LIGO Scientific Collaboration. *LIGO Algorithm Library - LALSuite*. free software (GPL). 2018. DOI: 10.7935/GT1W-FZ16.
- [343] C. Cutler and E. E. Flanagan. “Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral wave form?” In: *Phys.Rev.* D49 (1994), pp. 2658–2697. DOI: 10.1103/PhysRevD.49.2658. arXiv: gr-qc/9402014 [gr-qc].
- [344] N. Wang and J.-L. Lee. “Geometric properties of the icosahedral-hexagonal grid on the two-sphere”. In: *SIAM Journal of Scientific Computing* 33.5 (2011), pp. 2536–2559.
- [345] D. A. Randall et al. “Climate modeling with spherical geodesic grids”. In: *Comput. Sci. Eng.* 4 (2002), pp. 32–41.
- [346] S. W. Bruenn. “Stellar core collapse: Numerical model and infall epoch”. In: *Astrophys. J. Suppl.* 58 (1985), pp. 771–841. DOI: 10.1086/191056.
- [347] A. Burrows, S. Reddy, and T. A. Thompson. “Neutrino opacities in nuclear matter”. In: *Nucl. Phys.* A777 (2006), pp. 356–394. DOI: 10.1016/j.nuclphysa.2004.06.012. arXiv: astro-ph/0404432 [astro-ph].

Bibliography

- [348] F. Galeazzi et al. “Implementation of a simplified approach to radiative transfer in general relativity”. In: *Phys.Rev.* D88 (2013), p. 064009. DOI: 10.1103/PhysRevD.88.064009. arXiv: 1306.4953 [gr-qc].
- [349] K. S. Thorne. “Relativistic radiative transfer - Moment formalisms”. In: *Mon. Not. Roy. Astron. Soc.* 194 (Feb. 1981), pp. 439–473.
- [350] M. Shibata et al. “Truncated Moment Formalism for Radiation Hydrodynamics in Numerical Relativity”. In: *Prog.Theor.Phys.* 125 (2011), pp. 1255–1287. DOI: 10.1143/PTP.125.1255. arXiv: 1104.3937 [astro-ph.HE].