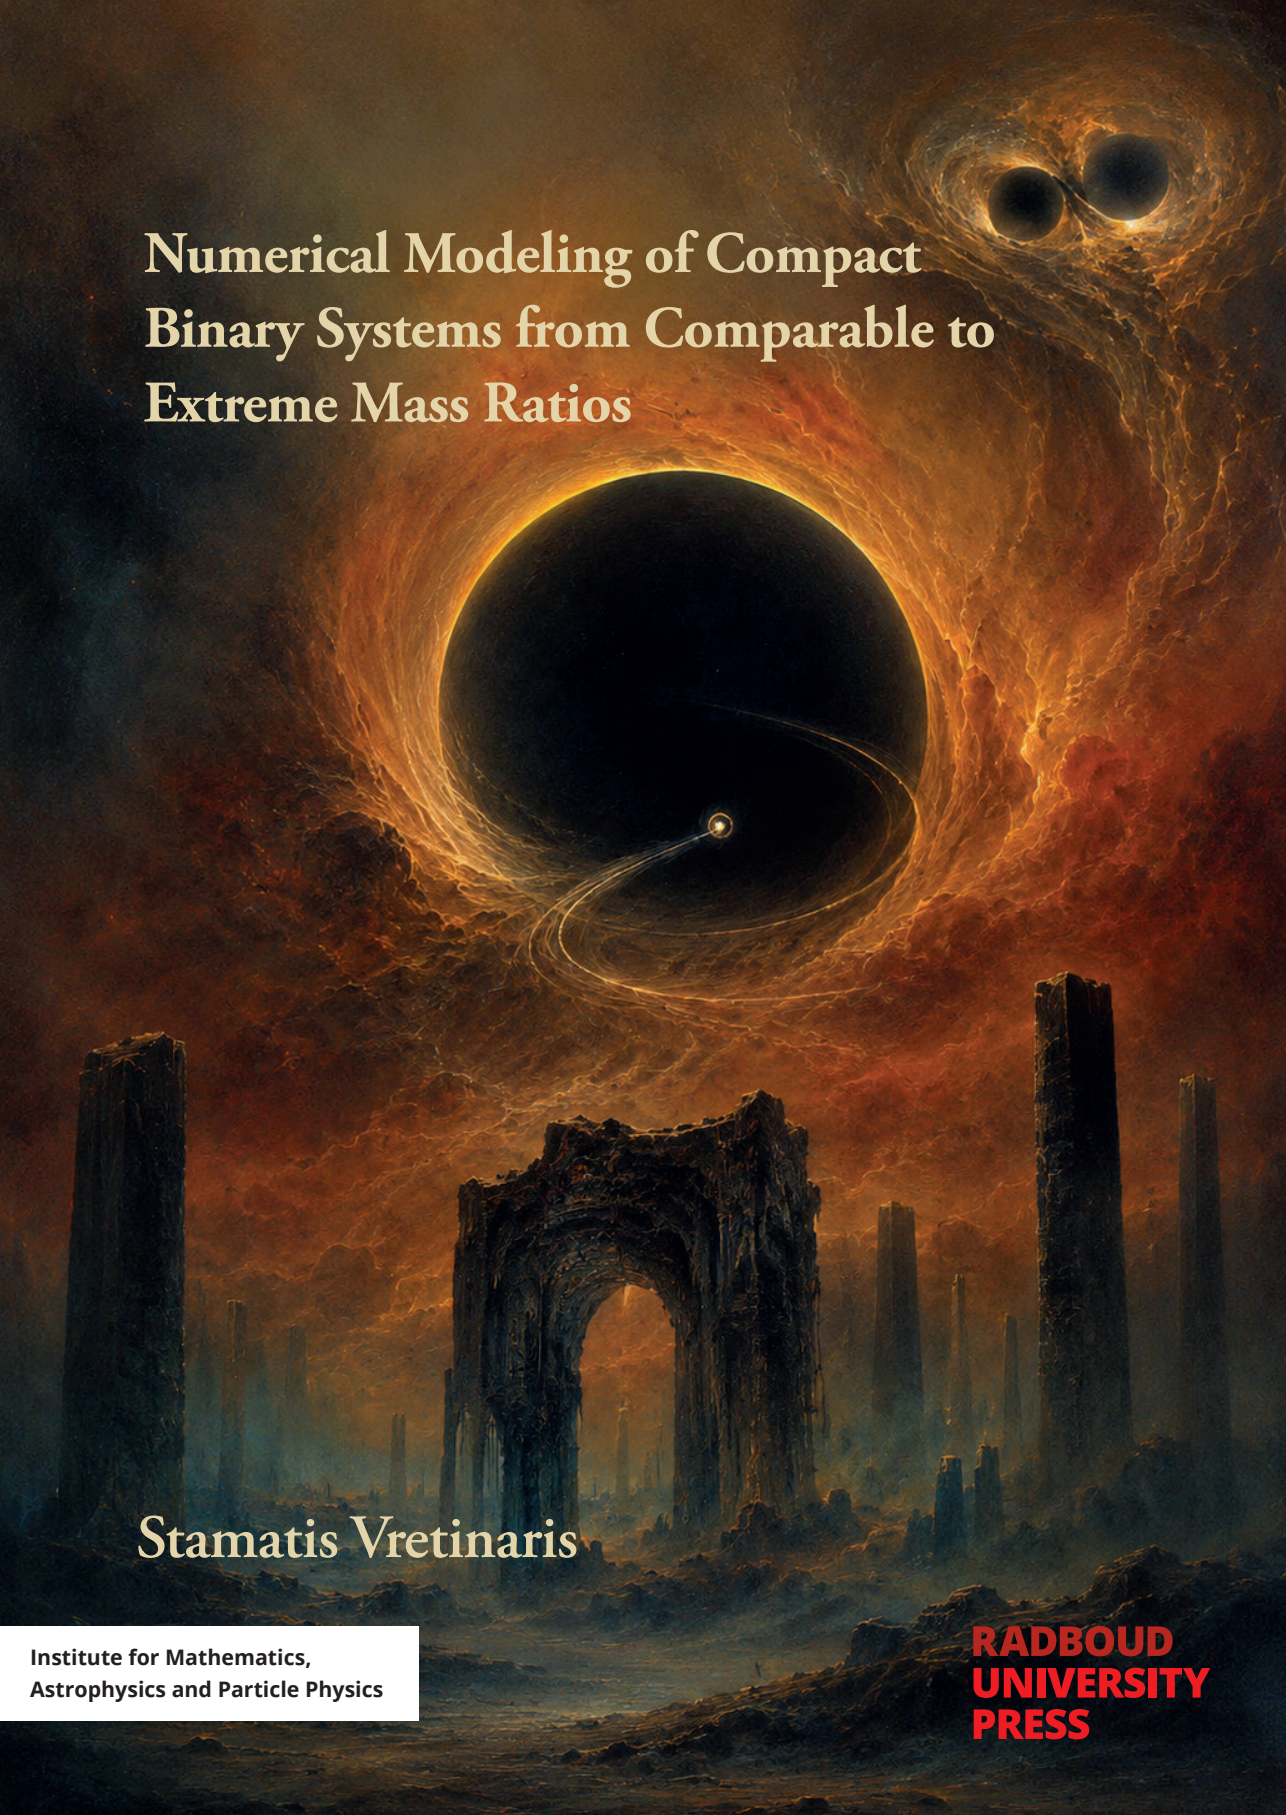
Numerical Modeling of Compact Binary Systems from Comparable to Extreme Mass Ratios



Stamatis Vretinaris

Institute for Mathematics, Astrophysics and Particle Physics

RADBOUD UNIVERSITY PRESS

Numerical Modeling of Compact Binary Systems from Comparable to Extreme Mass Ratios

Stamatis Vretinaris

Stamatis Vretinaris
Numerical Modeling of Compact Binary Systems from Comparable to Extreme Mass Ratios

Radboud Dissertation Series
ISSN: 2950-2772 (Online); 2950-2780 (Print)

Published by RADBOUD UNIVERSITY PRESS
Postbus 9100, 6500 HA Nijmegen, The Netherlands
www.radbouduniversitypress.nl

Design: Stamatis Vretinaris
Cover image: Stamatis Vretinaris
Production coordination & prepress: ProefschriftOntwerp.nl
Printing: ProefschriftMaken/Pumbo

ISBN: 9789465153445
DOI: 10.54195/9789465153445
Free download at: https://doi.org/10.54195/9789465153445

© 2026 Stamatis Vretinaris

RADBOUD
UNIVERSITY
PRESS

This is an Open Access book published under the terms of Creative Commons Attribution-Noncommercial-NoDerivatives International license (CC BY-NC-ND 4.0). This license allows reusers to copy and distribute the material in any medium or format in unadapted form only, for noncommercial purposes only, and only so long as attribution is given to the creator, see http://creativecommons.org/licenses/by-nc-nd/4.0/.

Numerical Modeling of Compact Binary Systems from Comparable to Extreme Mass Ratios

Proefschrift

ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de rector magnificus prof. dr. J.M. Sanders,
volgens besluit van het college voor promoties
in het openbaar te verdedigen op

dinsdag 19 januari 2027

om 14:30 uur precies

door

Stamatis Vretinaris

geboren op 19 januari 1994
te Düsseldorf, Duitsland

Promotor:	Prof. dr. B. Krishnan	
Copromotor:	Dr. E. Schnetter	Perimeter Institute, Canada
Manuscriptcommissie:	Prof. dr. R. Loll	
	Prof. dr. K. Kokkotas	Eberhard Karls Universität Tübingen, Duitsland
	Dr. C.F. Fernández Sopuerta	Institute of Space Sciences, Spanje

Acronyms

BH	Black hole
ASD	Amplitude spectral density
BNS	Binary neutron star
CFL	Courant-Friedrichs-Lewy
CoRe	Computational Relativity database
CPU	Central processing unit
EMRI	Extreme-mass-ratio inspiral
EOS	Equation of state
FD	Finite difference
FDSBP	Finite-difference summation by parts
FVM	Finite volume method
GW	Gravitational wave
HLV	Hanford-Livingston-Virgo network
KAGRA	Kamioka Gravitational Wave Detector
LIGO	Laser Interferometer Gravitational-Wave Observatory
LVK	LIGO-Virgo-KAGRA
MCMC	Markov chain Monte Carlo
MC	Monte Carlo
MOTS	Marginally outer trapped surface
MUSCL	Monotonic upstream-centered schemes for conservation laws
NS	Neutron star
ODE	Ordinary differential equation

PDE	Partial differential equation
PMC	Preconditioned Monte Carlo
QLH	Quasi-local horizon
RHS	Right-hand side
SAT	Simultaneous approximation term
SBP	Summation by parts
SNR	Signal-to-noise ratio
SSPRK	Strong-stability-preserving Runge-Kutta
TVD	Total variation diminishing

Contents

Acronyms **3**

1 Introduction **11**
1.1 High-Order Stable Discretizations in Curvilinear Coordinates 12
1.2 Self-Force Modeling in the Extreme-Mass-Ratio Regime 14
1.3 Black-Hole Horizon Dynamics in Binary Black-Hole Mergers 18
1.4 Binary Neutron Star Post-Merger Inference 21

2 High-Order Summation-By-Parts Schemes for First-Order Hyperbolic Systems in Curvilinear Coordinates with Singularities **25**
2.1 Introduction 27
2.2 Summation By Parts 30
2.2.1 Integration By Parts 30
2.2.2 Summation By Parts 31
2.2.3 Spherical Coordinates and Tensor Rank 32
2.2.4 Order and Degree of Derivative Operators . . . 35
2.3 Constructing Summation-by-Parts Operators for Curvilinear Coordinates with Singularities 36
2.3.1 Discrete Norm Operators 37
2.3.2 Discrete Derivative Operators 38
2.3.3 Non-Diagonal Norm Operators 39

2.3.4 Comparison to plain staggered operators 47
2.4 Spectral Stability 50
2.5 Numerical evolutions 53
2.5.1 Energy Conservation 55
2.5.2 Accuracy and Convergence 57
2.6 Conclusion 62

3 Numerical self-force calculations for scalar particles, formulated in the lab frame 65
3.1 Introduction 66
3.2 The physical system 70
3.2.1 Overview 70
3.2.2 A general remark on point particles 71
3.3 Equations of motion 72
3.3.1 A scalar field with a given charge distribution . 72
3.3.2 A particle in a given scalar field 72
3.3.3 Scalar field for a given particle trajectory 77
3.4 Numerical methods 80
3.4.1 Finite Volume Method 81
3.4.2 Spectral method 87
3.5 Numerical results 89
3.5.1 Finite Volume Method results 90
3.5.2 Spectral method results 96
3.6 Summary 103

4 Cusp Formation in Merging Black Hole Horizons 107
4.1 Introduction 108
4.2 Basic Notions: Quasi-local horizons, their mergers, and multipole moments 113
4.2.1 Quasi-local horizons 113
4.2.2 Axisymmetric MOTSs and multipole moments . 117
4.2.3 The numerical set-up 120
4.3 The area and mass at the merger 121
4.4 The geometry of $\mathcal{S}_{\text{inner}}$ near t_{touch} 124

4.5 The mass multipole moments at the merger 128
4.5.1 The mass dipole moment at the merger 128
4.5.2 The mass quadrupole and higher moments . . . 130
4.5.3 Evolution of the mass multipole moments . . . 133
4.6 Conclusions 138
4.7 Appendix: Ricci scalar at the poles 142
4.8 Appendix: Derivatives of the Ricci scalar near the poles 145

5 Robust and Fast Parameter Estimation for Gravitational Waves from Binary Neutron Star Merger Remnants 149
5.1 Introduction 150
5.2 Methods 153
5.2.1 Numerical Waveforms and Injections 153
5.2.2 Analytic Post-merger Waveform Model 157
5.2.3 Empirical Relations and Classification 159
5.2.4 Bayesian inference 162
5.2.5 Informed Gaussian Priors 163
5.2.6 Preconditioned Monte Carlo Sampling 165
5.3 Parameter Estimation 167
5.3.1 SNR = 50 167
5.3.2 SNR = 16 175
5.3.3 SNR = 8 175
5.4 Fitting Factors and Reconstruction 176
5.5 Summary and Discussion 178
5.6 Appendix: Robustness of parameter estimation 181
5.7 Appendix: Contribution of the 4th oscillator 182

6 Conclusions 187
6.1 Stable Discretization as Infrastructure 187
6.2 Self-Force Formulated in the Lab Frame 188
6.3 Horizon Geometry as a Merger Diagnostic 189
6.4 From Post-Merger Signals to Inference 190
6.5 Taken Together 191

Summary **193**
High-order SBP operators for curvilinear singularities 193
Finite-resolution scalar self-force modeling 194
Black-hole horizon mergers and cusp formation 196
Post-merger inference for binary neutron stars 197
Synthesis and outlook 198

Samenvatting **201**
Hoog-orde SBP-operatoren voor kromlijnige singulariteiten . 202
Eindige-resolutie-modellering van scalaire self-force 203
Fusies van zwarte-gathorizonten en cuspvorming 204
Post-mergerinferentie voor binaire neutronensterren 205
Synthese en vooruitblik 207

Curriculum Vitae **209**

Levensloop **211**

Description of Research Data Management **213**
Scope of the research data 213
Processing of the research data 214
MOTS chapter 214
BNS chapter 214
Self-force and SBP chapters 215
Availability of data and code 215
Data availability 215
Code availability 216
Field-specific compliance 216

Acknowledgements **217**
Financing 217
Acknowledgements 217

Bibliography **219**


1

Chapter 1

Introduction

Compact binaries provide direct access to strong-field gravity, relativistic fluid dynamics, and the generation of gravitational radiation [276, 85, 31, 264]. They also require numerical models across markedly different physical regimes. Comparable-mass systems demand fully nonlinear calculations, while extreme-mass-ratio systems are most naturally described within perturbation theory but require very high long-term phase accuracy [224, 101, 36]. Across this range, reliable modeling depends on resolving a common numerical problem: physically natural formulations often involve singular coordinates, singular sources or singular geometric behavior that defeats naive discretizations. The four chapters of this thesis address these obstacles at different points along the mass-ratio spectrum.

The thesis is organized in two thematic blocks. The first two chapters are methodological. They focus on numerical infrastructure and evolution frameworks: first, stable high-order discretization in singular curvilinear coordinates; second, self-force evolution in the extreme-mass-ratio regime. The final two chapters turn to astrophysical phenomenology. They examine, respectively, the geometry of black-hole merger horizons in the strongly nonlinear comparable-mass regime and the extraction of dense-matter information from binary neutron star post-merger signals. The purpose of this introduction is not to

repeat the technical development of the individual publications, but to state the scientific setting of each chapter, the specific bottleneck it addresses, and the reason that bottleneck matters for compact-binary modeling.

1.1 High-Order Stable Discretizations in Curvilinear Coordinates

Many problems in computational astrophysics and numerical relativity are most naturally formulated in curvilinear coordinates. Spherical or polar descriptions are well matched to systems with approximate radial or axial structure, and they arise naturally in wave propagation, collapse problems, perturbation theory, and compact-object simulations [89, 147, 213, 190, 188]. In these settings, coordinates adapted to the geometry can reduce computational cost and represent the underlying physics more directly than Cartesian grids. This is particularly valuable when one seeks high resolution over long evolutions, where inefficiency or poor error control can quickly become prohibitive.

The motivation for the framework studied in Chapter 2 is therefore not merely geometric convenience. Reliable compact-binary modeling depends on discretizations that are both accurate and demonstrably stable. For hyperbolic problems, a central continuum tool is the existence of an energy estimate guaranteeing well-posedness. Summation-by-parts operators are important because they provide a discrete analogue of integration by parts, allowing one to construct finite-difference schemes that inherit corresponding stability properties at the semi-discrete level [178, 179, 282]. As emphasized in the broader SBP literature, SBP methods have become a standard route to provably stable higher-order discretizations for partial differential equations [99, 148].

The attraction of such a framework is especially clear in relativistic wave problems. Many formulations used in numerical relativity can be written as first-order symmetric hyperbolic systems for which contin-

uum energy estimates are available. When this structure is preserved discretely, numerical stability does not depend on taking an arbitrarily fine grid; rather, the semi-discrete system inherits a controlled energy behavior already at finite resolution. This is a stronger requirement than mere consistency. In long-time evolutions, especially those involving wave propagation over many crossing times, failure to preserve the relevant energy structure can overwhelm any formal gain in truncation order.

That general framework, however, does not by itself resolve the main difficulty addressed in this chapter. Curvilinear coordinate systems introduce singular points that are not physical features of the underlying spacetime or field configuration, but pathologies of the coordinate description in the discrete. Near the origin or symmetry axis, standard finite-difference constructions can lose accuracy or fail to preserve the discrete energy structure that makes SBP methods attractive in the first place. A variety of strategies have been used in the broader numerical literature to manage such singularities, including staggered grids, geometric transformations, special centerline conditions, and related constructions that avoid direct evaluation at the singular point [239, 216, 277, 207, 145, 92]. These methods can be effective, but they often do so by sacrificing collocation at the origin or relying on problem-specific treatments.

This difficulty is not merely a nuisance of implementation. After separation into spherical or tensor-harmonic multipoles, smooth continuum fields satisfy specific regularity conditions at the origin, and the corresponding radial equations contain terms that appear singular, schematically of the form p/r. In the continuum these singularities cancel once the regularity conditions are enforced, but a naive discretization does not guarantee the same cancellation pointwise. The numerical problem is therefore to construct discrete gradient, divergence, and norm operators that encode the correct regularity and make the cancellation happen automatically, rather than requiring ad hoc special handling at each application.

The specific problem taken up in Chapter 2 is how to retain the central advantages of SBP discretization in curvilinear coordinates with singularities while being able to have a collocated grid. In the formulation developed in the chapter, the question is not only whether one can remain stable, but whether one can do so while preserving high-order accuracy all the way to the singular point and while constructing operators systematically rather than case by case. This is the bottleneck that separates a useful local workaround from a general numerical framework.

Solving this problem matters beyond the model equations used in the chapter. If stable high-order SBP operators can be constructed directly in singular curvilinear coordinates, then one gains a more reliable numerical foundation for a broad class of simulations in which symmetry-adapted coordinates are the natural language of the problem. The resulting benefit is twofold: better numerical efficiency in applications where curvilinear coordinates are physically appropriate, and stronger control over stability and error in evolutions whose scientific value depends on long-time fidelity. In particular, the long-term goal is to apply these methods to problems such as critical collapse and hyperboloidal evolutions, where one must handle a coordinate singularity not only at the origin but, in the hyperboloidal setting, also at null infinity.

In one sentence, the research question we answer in this chapter is: **Can we construct stable high-order SBP operators in curvilinear coordinates with singularities for collocated and staggered grids, and if yes, then how?**

1.2 Self-Force Modeling in the Extreme-Mass-Ratio Regime

At one end of the mass-ratio spectrum lie extreme-mass-ratio inspirals, in which a stellar-mass compact object spirals into a much more

massive black hole. These systems are among the principal targets of future space-based gravitational-wave detectors, because they probe the geometry of massive black holes through long signals produced deep in the strong-field region [224, 101, 36]. Their astrophysical value is correspondingly high: they offer access to black-hole populations in galactic nuclei, and they provide a setting in which general relativity can be tested through detailed measurements of relativistic orbital motion and radiation [224, 36]. They are also relevant to the search for intermediate-mass black holes, whose existence has not yet been established conclusively. Detecting gravitational-wave signals from systems in this regime would provide direct evidence for such objects, and the intermediate-mass-ratio regime is precisely one in which self-force methods can fill the gap between full numerical relativity and weak-field approximation schemes.

The modeling requirements in this regime are unusually severe. An extreme-mass-ratio inspiral spends a very long time in band and can accumulate a very large number of gravitational-wave cycles before plunge. Waveform accuracy therefore depends on controlling small dynamical effects over long durations [224, 101, 36]. Direct numerical-relativity simulations are not viable in this regime because of the disparity of scales between the large black hole and the compact secondary [276, 85]. To resolve the field near the small body, the computational grid must capture very short length scales around it while still covering the much larger length scale of the primary and the full radiation zone. With explicit time evolution, the Courant-Friedrichs-Lewy condition then ties the time step to the smallest resolved scale, so resolving the compact secondary forces an extremely small time step for the entire simulation. The result is that the computational cost becomes prohibitively high long before one reaches the mass ratios relevant for these sources. Weak-field approximation schemes are likewise inadequate for the late stages of inspiral, where the dynamics unfolds in the deep gravitational well of the primary [224, 36]. The self-force framework is needed precisely because it is a perturbative

expansion in the small mass ratio that still retains the strong-field geometry of the primary: at leading order the smaller body follows a background geodesic, and at the next orders its own field drives the secular corrections that determine the inspiral [204, 242, 103, 144, 235].

The underlying physical picture is subtle. A point particle generates a retarded field that diverges on its own worldline, yet only a certain regular part of that field can consistently influence the motion. In flat spacetime this logic already appears in the classical radiation-reaction problem studied by Dirac: one removes a singular, time-symmetric contribution and identifies the remaining regular field with the part responsible for radiation reaction. In curved spacetime the Detweiler–Whiting construction provides the corresponding conceptual framework, separating the field into a singular piece that does not contribute to the force and a regular piece that does [103, 224]. Self-force theory is therefore not just a perturbation expansion; it is a prescription for extracting finite dynamics from a formally singular coupled particle-field system.

The main conceptual obstacle is that the field of the small body is singular on its worldline, even though the physically relevant motion depends only on an appropriate regularized part of that field. Much of the self-force programme has therefore been concerned with turning the formal singular-regular split into practical computational schemes [224, 101, 36]. As summarized in the self-force reviews and in Wardell's overview of computational strategies [306], the main approaches include dissipative or flux-based methods, which capture radiation-reaction information without local regularization but omit conservative effects [203, 162, 114, 285]; mode-sum regularization, which computes finite multipolar contributions and subtracts analytically derived singular pieces [39, 38]; worldline-convolution methods, which reconstruct the regular field from an integral over the particle's past history [11, 84, 307]; and effective-source or puncture methods, which regularize at the level of the field equation and evolve a finite residual field directly [37, 297, 107]. These developments have made self-force calculations

possible in a range of settings and have established self-force theory as a practical tool rather than a purely formal construction.

The specific problem addressed in Chapter 3 arises within local puncture and effective-source formulations of the self-force problem. In these methods one constructs an approximate local expansion of the singular field near the particle, subtracts it from the retarded field, and evolves the resulting finite residual field. For the scalar problem, the first self-consistent evolutions of this type were carried out by Diener et al. [107]. In that implementation, however, the local singular-field approximation was truncated at the level of velocity-dependent terms in order to maintain a numerically workable evolution system.

The next natural step is to improve these local singular-field approximations by including acceleration-dependent terms, as in later puncture and worldtube formulations [157, 311, 310]. But this is precisely where a structural difficulty appears. Once the puncture field depends on the particle's acceleration, and the acceleration is itself determined from the regularized field produced by the puncture scheme, the coupled system becomes self-referential. In the scalar case, incorporating acceleration-dependent terms into local singular-field approximations can then introduce a dependence on higher time derivatives of the particle's own acceleration [157, 105, 311, 310]. This is problematic from the perspective of an initial-value formulation: the future evolution should be determined by physically appropriate initial data for the field and the worldline, not by implicitly requiring additional higher-derivative information. Such dependence also creates a route to numerical instability in coupled particle-field evolutions [107, 157, 105].

The chapter addresses this bottleneck by developing a self-consistent scalar self-force framework that includes acceleration effects while preserving the structure of a finite initial-value problem. The aim is not to replace the broader self-force programme, but to remove a specific obstruction that appears when one tries to evolve the coupled system directly in time. This matters because phase-accurate extreme-mass-ratio modeling ultimately requires formulations that are not only

formally correct, but also numerically stable, explicit about their data requirements, and suitable for long-duration evolutions.

This issue is also broader than the scalar toy model used in the chapter. Scalar self-force calculations are valuable precisely because they isolate the structural difficulty in a simpler setting before one confronts the full tensorial complexity of gravitation. If one can formulate the particle-field evolution so that acceleration effects are incorporated without hidden higher-derivative data or iterative closure assumptions, then one gains a template for more robust effective-source formulations in electromagnetic and gravitational problems as well. In that sense, the chapter addresses a foundational numerical question whose resolution is relevant well beyond the specific scalar model studied here.

In one sentence, the research question we answer in this chapter is: **Can acceleration-dependent self-force effects be incorporated into a self-consistent scalar evolution scheme without destroying the initial-value structure of the problem, and if yes, then how?**

1.3 Black-Hole Horizon Dynamics in Binary Black-Hole Mergers

For isolated black holes, the Kerr family provides an accurate description of the spacetime and fixes all multipolar structure once the mass and spin are specified. A binary black-hole merger is the regime in which this simplicity breaks down most sharply. During coalescence, two initially distinct black holes interact nonlinearly and settle to a single remnant whose mass, angular momentum, and higher multipolar properties encode the cumulative effects of the merger [276, 85]. Chapter 4 focuses on this strongly nonlinear stage in the comparable-mass regime, where perturbative approaches are no longer sufficient and the dynamics must be extracted from full numerical-relativity simulations.

The astrophysical and gravitational-wave motivation is immediate.

The mapping from progenitor black-hole parameters to the remnant state influences population-level predictions for stellar-mass and supermassive black holes, informs interpretations of gravitational-wave observations, and determines the ringdown spectrum used in tests of general relativity [202, 256, 58, 115]. In the broader programme, the goal is to determine the final mass and angular momentum of the remnant black hole from the initial binary configuration, in particular from the individual masses and spins of the two black holes. A detailed understanding of how the remnant approaches equilibrium is also relevant to black-hole spectroscopy, where one seeks to connect observed quasi-normal-mode content to the geometry of the final black hole [212, 176].

The natural global notion of a black-hole boundary in general relativity is the event horizon, defined as the boundary of the region from which no signal can reach future null infinity. This definition is geometrically clean, but it is also teleological: locating the event horizon requires knowledge of the spacetime arbitrarily far into the future. That makes it poorly suited to genuinely dynamical situations such as gravitational collapse, binary black-hole mergers, or numerical simulations, where one would like to identify and track the black hole during the evolution rather than only after the full spacetime is known.

For this reason, the chapter adopts the quasi-local viewpoint based on marginally outer trapped surfaces (MOTSs). A MOTS is a closed two-surface whose outgoing null expansion vanishes, placing it at the threshold between escaping and trapped light rays. Three-dimensional hypersurfaces foliated by such marginally trapped surfaces provide a practical notion of quasi-local horizons, which can be located on individual time slices and followed throughout a simulation. This framework is especially useful in merger problems because it separates the local, dynamical behavior of the horizons from the global causal definition of the event horizon and provides direct access to quantities such as horizon area, deformations, and multipolar structure as matter and gravitational radiation cross the horizon.

The central problem addressed in this chapter is that the merger itself is the least understood part of the horizon dynamics. Horizon-based measures of black-hole parameters are well established for the initial individual holes and for the final remnant once it has approximately settled [116], but the transition between these regimes is geometrically singular and numerically delicate. Near merger, the relevant marginally trapped surfaces become strongly distorted and develop cusp-like features [229, 231, 233, 234, 69]. This makes it difficult to track quasi-local horizons through the instant at which the common remnant forms, and therefore difficult to determine how mass, angular momentum, and multipolar structure are transferred across the merger.

Solving this problem matters because it opens a complementary route to understanding binary black-hole coalescence. There is already a substantial literature that models the progenitor-to-remnant map using waveform information, inspiral-based information, or phenomenological fits to large sets of simulations [317, 156, 195, 159, 290, 42, 252, 253, 296, 72, 73, 82, 94, 95]. What distinguishes the present work is that it seeks to address this same remnant-prediction problem through the quasi-local horizons themselves. Even though the chapter does not yet carry out the full angular-momentum analysis, it is part of the same general programme: to infer the final mass and angular momentum directly from horizon dynamics through the merger, rather than only from asymptotic radiation or inspiral-to-ringdown fits. If one can characterize the remnant directly through the behavior of the horizons, including the discontinuous changes associated with merger, then one gains access to the nonlinear transition from a two-black-hole system to a single Kerr black hole in a form that is closely tied to the strong-field region itself. This is valuable both for interpreting numerical simulations and for clarifying how local horizon geometry is connected to the observable gravitational-wave signal [254, 170, 212, 176, 70].

In one sentence, the research question we answer in this chapter is: **Can quasi-local horizons be tracked accurately through black-hole merger so that the transfer of mass, angular momentum,**

and multipolar structure across the cusp is determined, and if yes, then how?

1.4 Binary Neutron Star Post-Merger Inference

Binary neutron star mergers occupy the comparable-mass regime, but unlike binary black holes they involve matter at supranuclear densities, shock heating, differential rotation, and the possible formation of a long-lived neutron-star remnant or a prompt collapse to a black hole [31, 264, 130]. Their gravitational-wave signal therefore contains information not only about gravity in the dynamical strong-field regime but also about the equation of state of dense matter. Chapter 5 is concerned specifically with the post-merger signal, where the surviving remnant can exhibit characteristic oscillation frequencies linked to its internal structure [45, 46].

The scientific motivation for studying this regime is twofold. First, binary neutron star mergers are central to multimessenger astrophysics and to efforts to constrain the microphysics of neutron-star matter [4, 264]. Second, the post-merger signal probes densities beyond those cleanly accessible during the inspiral, making it a particularly promising channel for distinguishing among competing equations of state and for searching for signatures of new high-density physics [46, 168, 255, 217]. The inspiral and post-merger phases are complementary in a precise sense: inspiral measurements mainly constrain the cold, beta-equilibrated equation of state through tidal effects, while the post-merger remnant samples hotter, denser, and strongly differentially rotating matter. Agreement between the two phases would therefore strengthen confidence in a common EOS interpretation, whereas tension could point to missing remnant physics or to genuinely new high-density effects such as phase transitions or composition-dependent microphysics [46, 130, 168]. These goals are tightly connected to detector develop-

ment: present instruments are only marginally sensitive to the relevant high-frequency band, but future observatories are expected to improve the prospects for post-merger detection substantially [219, 316].

This observational opportunity is tied to the characteristic structure of the post-merger spectrum. Numerical-relativity simulations show that the signal is not described only by a single dominant oscillation, even though the main peak at $f_{\rm peak}$ is often the most robust feature. The spectrum can also contain secondary peaks associated with remnant oscillation modes, mode couplings, transient spiral structures, damping, and the time evolution of the differentially rotating remnant [281, 45, 287, 302]. Empirical and quasi-universal relations involving $f_{\rm peak}$ have therefore been extremely valuable, because they expose robust correlations with bulk neutron-star properties such as radius, compactness, or tidal deformability. At the same time, reducing the post-merger signal to a small number of summary frequencies leaves part of the available physical information unused.

The central modelling problem is therefore an inverse problem: given a measured post-merger spectrum in the presence of detector noise, infer the source and EOS-sensitive parameters that produced it. Existing approaches span two broad strategies. One emphasizes empirical relations between a few spectral features and stellar properties, while the other attempts to use more of the signal through phenomenological waveform models and Bayesian inference in either the time or frequency domain [119, 120, 221]. The main bottleneck is not the absence of physical information in the signal, but the difficulty of extracting it robustly. Post-merger waveforms are short, spectrally rich, and observationally challenging. Their interpretation depends on numerical-relativity input and on empirical relations that connect the pre-merger binary properties to the remnant [45, 302]. Standard Bayesian parameter estimation becomes inefficient when the parameter space is weakly constrained and the available prior information is not used effectively. In this setting, even a physically informative model can be difficult to deploy for inference at the speed and robustness

demanded by realistic data analysis.

Addressing this bottleneck is essential if future post-merger observations are to become scientifically useful. A detection alone does not yield constraints on dense matter unless the associated inference problem can be solved reliably and efficiently. Physically motivated models that account for the dominant spectral features while remaining consistent with known EOS trends and inspiral-derived information are especially important, because they offer a way to use more of the post-merger morphology without abandoning the robust correlations already established by quasi-universal relations [119, 302, 304]. In that sense, Chapter 5 addresses a translation problem between theory and data analysis: how to make the post-merger spectrum usable for astrophysical inference rather than merely identifiable in principle.

In one sentence, the research question we answer in this chapter is: **Can prior information from the inspiral and numerical-relativity-informed spectral structure be used to make binary neutron star post-merger parameter estimation robust and practical, and if yes, then how?**


2

Chapter 2

High-Order Summation-By-Parts Schemes for First-Order Hyperbolic Systems in Curvilinear Coordinates with Singularities

Publication status

This chapter is based on the following paper:

> **S. Vretinaris**, E. Schnetter, "High-Order Summation-By-Parts Schemes for First-Order Hyperbolic Systems in Curvilinear Coordinates with Singularities", arXiv: 2606.05155, Submitted to Physical Review D.

Contributions

I contributed to the formulation of the problem, development of the method, implementation, simulations, production of figures, interpre-

tation of results, and manuscript writing.

Abstract

Formulating stable numerical methods for hyperbolic systems in curvilinear coordinate with singularities, e.g. spherical coordinates, is complicated by the presence of these singularities. We present a method for constructing high-order accurate, energy-stable finite difference operators satisfying the Summation-by-Parts (SBP) property on spherical domains, extending ideas presented by Gundlach, Martín-García, and Garfinkle [148]. We define discrete gradient and divergence operators that mirror the continuous integration-by-parts principle, even though there is a $1/r^p$ coordinate singularity present at the origin. We explicitly construct such operators up to order six. Our operators place a grid point directly on the origin. We also review how to construct stable SBP operators that straddle the origin. We analyse the accuracy and spectral radii of these operators, and we show example evolutions of the scalar wave equation to demonstrate the advantages of such operators.

2.1 Introduction

One often-used method to prove well posedness for a hyperbolic PDE is by constructing a conserved energy (or a non-increasing energy that is bounded from below). The analogue of the well posedness in the continuum is numerical stability for a discretized PDE. Kreiss and Scherer [178, 179] proposed a powerful way of constructing linearly stable schemes for solving evolution partial differential equations which admit an energy estimate at the continuum level, through the use of difference operators satisfying *summation by parts* (SBP). These operators allow deriving estimates analogous to the continuum ones, up to boundary terms.

While the former guarantees well posedness, the discrete counterparts guarantee numerical stability. Numerical stability prevents unbounded growth, i.e. prevents exponential growth with a resolution-

dependent growth rate where higher resolutions have faster growing modes.

Summation by parts is the discrete equivalent of integration by parts. A set of finite differencing operators may or may not satisfy the SBP property; if they do, they allow certain kinds of proofs and estimate to "go through" for the discretized equations in the same way as for the continuum equations. For example, if a hyperbolic partial differential equation (PDE) has a conserved energy, then discretizing this system using SBP operators often leads to discrete system where the discrete energy is also conserved. This then holds *independent of the resolution*, i.e. no matter how large the discretization error – energy conservation is then not a property obtained in the limit of infinite resolution, but it can be tested at any resolution, up to floating-point error.

There is a large body of literature for SBP operators that we briefly review further below. In this chapter we will follow ideas first presented by Gundlach, Martín-García, and Garfinkle [148] and construct SBP finite differencing operators (stencils) for curvilinear coordinate systems with singularities, for example for polar or spherical coordinates. We introduce a set of *accuracy conditions* that our operators need to satisfy near the origin, and then explicitly construct such operators for fourth and sixth order accuracy. Our operators retain their full accuracy at the origin; in fact, we are even able to demand a somewhat higher accuracy near the origin to select from a family of operators.

It goes without saying that spherical symmetry, or multipole expansion in spherical coordinates, or in general curvilinear coordinate systems with singularities are immensely useful in describing many physical systems. One famous example in general relativity is Choptuik's study of gravitational critical collapse [89, 147, 146]. Many systems in general relativity exhibit symmetries that can be well captured by a curvilinear coordinate system, such as supernova core collapse simulations [213], black hole mimickers like boson stars [190], black strings [188], etc.

The method we describe in this chapter has the potential to improve accuracy and guarantee discrete stability in all cases where high-order finite differences are applicable.

Issues related to coordinate singularities have been studied outside numerical relativity. Prochnow et al. proposed a high-order SBP treatment for the polar singularity in axisymmetric wave propagation using staggered SBP operators; their method avoids evolving the unknown at r=0 and uses weighted SBP norms so that the origin contributes no spurious boundary work [239]. O'Reilly and Petersson developed energy-conservative SBP discretizations for acoustic wave propagation in covariant form on staggered curvilinear grids [216]. Mimetic finite-difference methods provide another staggered-grid framework with discrete divergence/gradient operators and compatible quadratures [277]. Other approaches include geometric transformations or pole conditions that avoid direct evaluation at the polar point [207], centerline equations obtained with l'Hospital's rule [145], series-expansion treatments of cylindrical-axis singularities [92], and specialized finite-element basis modifications at polar grid centers [63].

Since the singularity at the origin is not physical, but a mere pathology of the coordinate system, it is possible to construct appropriate SBP operators that treat this pathology. This was first shown in [148], where the authors constructed stable second and fourth order SBP operators for spherical decompositions of the wave equation.

In this chapter we extend the work of [148] on spherical/curvilinear SBP operators and develop a general method to construct SBP operators of arbitrary order that do not lose accuracy at coordinate singularities. We formulate our operator construction by demanding *accuracy conditions*, which are linear conditions on SBP the operators, generalizing the recurrence relations introduced in [148]. We rely on existing Cartesian SBP operators, and then transform these operators to be applicable in curvilinear coordinate systems, finding appropriate origin closures for the norm operators such that the resulting operators are stable and accurate.

We give a general overview of the summation-by-parts idea and its relevance to discrete stability in section 2.2. We also review some basic relevant properties of curvilinear coordinates and tensor multipole moments there. In section 2.3 we describe our construction and show two examples for fourth- and sixth-order accurate SBP operators for spherical symmetry. Finally, we show example evolutions of the scalar wave equation with these operators in section 2.5, demonstrating the expected convergence and accuracy, in particular at and near the coordinate singularity at the origin.

2.2 Summation By Parts

2.2.1 Integration By Parts

Given a scalar function Π and a vector-valued function Ψ on some domain Ω in a suitable function space, the well-known *integration by parts* theorem states that

$$\int_\Omega (\operatorname{grad} \Pi) \cdot \Psi \, dV + \int_\Omega \Pi \, (\operatorname{div} \Psi) \, dV \;=\; \oint_{\partial\Omega} \Pi \, \Psi \cdot ds \qquad (2.1)$$

Here, the gradient is the partial derivative of a scalar field, whereas the divergence is the metric-dependent covariant divergence of a vector field. This distinction is essential in curvilinear coordinates.

Integration by parts often plays a useful role in proving the well-posedness of wave-type hyperbolic partial differential equations (PDEs). For example, the wave equation

$$\partial_t \Pi = \operatorname{div} \Psi \qquad (2.2)$$

$$\partial_t \Psi = \operatorname{grad} \Pi \qquad (2.3)$$

has the conserved energy

$$E = \frac{1}{2} \int_\Omega \Pi^2 + \Psi^2 \, dV \qquad (2.4)$$

A sufficient condition for this PDE to be well-posed is the existence of a conserved (or non-increasing) energy [132, 135, 133] (which is bounded from below). We show that this is the case here:

$$\begin{aligned} \partial_t E &= \int_\Omega \Pi \left(\partial_t \Pi\right) + \Psi \cdot \left(\partial_t \Psi\right) dV && (2.5) \\ &= \int_\Omega \Pi \left(\operatorname{div} \Psi\right) + \Psi \cdot \left(\operatorname{grad} \Pi\right) dV && (2.6) \end{aligned}$$

and use integration by parts to find

$$E = \oint_{\partial\Omega} \Pi \, \Psi \cdot ds \,. \tag{2.7}$$

Given suitable boundary conditions, this energy is conserved (in particular for either $\Pi = 0$ or $\Psi = 0$) or non-increasing.

We can also write the integration-by-parts theorem in operator notation as

$$\langle \operatorname{grad} \Pi | H_V | \Psi \rangle + \langle \Pi | H_S | \operatorname{div} \Psi \rangle = \langle \Pi | B | \Psi \rangle \tag{2.8}$$

where we allow for a non-trivial volume element (integration weight, norm) H_S and incorporate the dot product into H_V. bnd is an operator describing the integral over the boundary. This can also be written concisely as operator equation

$$\operatorname{grad}^T H_V + H_S \operatorname{div} = \operatorname{bnd} . \tag{2.9}$$

It is now evident that integration by parts is a condition on the derivative operators (gradient and divergence), norms (for scalar and vector-valued functions), and boundary operator.

2.2.2 Summation By Parts

Summation by parts is the discrete equivalent of integration by parts, i.e. it is applicable to discretized functions, and is similarly useful there.

A sufficient condition for stability of a spatial discretization of a PDE is the existence of a discrete energy estimate, namely a conserved or non-increasing discrete energy. This approach is commonly referred to as the energy method [150, 282].

We discretize (2.2) in space but leave time continuous. We use an asterisk $\Box^*$ to indicate discrete quantities and discrete operators.

This system is then

$$\partial_t \Pi^* = \text{grad}^* \Psi^* \tag{2.10}$$

$$\partial_t \Psi^* = \text{div}^* \Pi^* \tag{2.11}$$

with the discrete energy

$$E^* = \frac{1}{2} \left(\Pi^* H_S^* \Pi^* + \Psi^* H_V^* \Psi^* \right) . \tag{2.12}$$

The expressions for the energy corresponds to the integral above, written in operator notation. The operators H_S^* and H_V^* are discrete quadrature rules [98].

If we now demand that our discrete operators satisfy the summation-by-parts property

$$(\text{grad}^*)^T H_V^* + H_S^* \text{div}^* = \text{bnd}^* \tag{2.13}$$

then this discrete energy is also conserved up to boundary terms:

$$\partial_t E^* = \Pi^* \text{bnd}^* \Psi^* . \tag{2.14}$$

The proof is exactly equivalent to the continuum proof above. This is the main feature of summation-by-parts operators: They allow translating a continuum proof to a proof for a discrete system, and this simplifies finding a stable discretization for a well-posed hyperbolic PDE.

2.2.3 Spherical Coordinates and Tensor Rank

In spherical coordinates $x^i = (r, \theta^A)$ in d dimensions, the volume element is $dV = r^{d-1} \sqrt{\det \Omega}\, dr\, d^{d-1}\theta$, where Ω_{AB} is the metric on the

unit sphere S^{d-1}. For a physical tensor field $\mathbf{\Phi}$, the continuous L^2 energy norm is defined by:

$$\|\mathbf{\Phi}\|^2 = \int_\Sigma |\mathbf{\Phi}|^2 \, r^{d-1} \, dr \, d\Omega \tag{2.15}$$

When the angular dependence of the field is projected onto a basis of tensor (hyper)spherical harmonics of degree l, the radial profile of a rank-s tensor multipole scales with a leading power r^k as $r \to 0$. (This is necessary so that its Cartesian tensor components remain smooth).

For a generic tensor field, k takes the form

$$k = l + \delta_s \tag{2.16}$$

where l is the multipole degree, and the integer shift $\delta_s \in \{-s, -s+1, \ldots, s\}$ is determined by the tensor rank s and the specific parity of the tensor harmonic component (e.g., electric vs. magnetic parity). For a scalar field ($s = 0$), the shift is $\delta_0 = 0$, yielding $k = l$.

To construct a numerical scheme that automatically enforces this required analyticity, we define a regularized tensor variable $\tilde{\mathbf{\Phi}}$ by factoring out the geometric scaling:

$$\mathbf{\Phi}(r) = r^k \tilde{\mathbf{\Phi}}(r) \tag{2.17}$$

We substitute the regularized variable back into the continuous L^2 energy norm and find

$$\begin{aligned} \|\mathbf{\Phi}\|^2 &= \int (r^k \tilde{\mathbf{\Phi}})^2 \, r^{d-1} \, dr \, d\Omega \\ &= \int \tilde{\mathbf{\Phi}}^2 \, r^{2k+d-1} \, dr \, d\Omega \end{aligned} \tag{2.18}$$

Equation (2.18) implies that factoring out the analytic regularity condition maps the dynamics of the smooth variable $\tilde{\mathbf{\Phi}}$ into a new effective function space with the modified volume element $\tilde{\gamma}$:

$$\tilde{\gamma} \, dr = r^{2k+d-1} \, dr = r^p \, dr \tag{2.19}$$

where we define $p = 2k + d - 1$.

In our first-order symmetric hyperbolic system, the evolution is governed by the gradient and divergence operators. We define here a new set of operators that act on the regularized variables. These operators are rescaled versions of their Cartesian counterparts, and are defined as follows:

$$\tilde{G} = \text{grad} \tag{2.20}$$
$$\tilde{S} = r^p H_S \tag{2.21}$$
$$\tilde{V} = r^p H_V \tag{2.22}$$
$$\tilde{B} = r^p \text{bnd} \tag{2.23}$$

We demand that the rescaled operators should satisfy integration by parts

$$\tilde{S}\tilde{D} + \tilde{G}^T\tilde{V} = \tilde{B}\,, \tag{2.24}$$

and the rescaled divergence operator can thus be defined via

$$\tilde{D} = \tilde{S}^{-1}\left(\tilde{B} - \tilde{G}^T\tilde{V}\right)\,. \tag{2.25}$$

where we used the fact that $\tilde{H}_S$ is a norm and thus invertible.

Expanded and simplified, this leads to

$$\tilde{G} = \partial_r \tag{2.26}$$
$$\tilde{D} = \partial_r + \frac{p}{r}\,. \tag{2.27}$$

The term $\frac{p}{r}$ is the expected Christoffel symbol, which is singular at the origin. The main point of designing summation-by-parts operators is to ensure that the formal singularities in the divergence operator are guaranteed to cancel not only in the continuum but also in the discrete, so that a discrete divergence operator can be applied without worrying about the origin. As the equation above suggests, this means

that our set of discrete operators will contain a (covariant) divergence operator, and not separate partial derivatives and Christoffel symbols.

The parameter p combines all of the spatial dimension d, the multipole degree l, and the tensor rank s into a single integer. For example, applying this to the 3D scalar wave equation ($d = 3, \delta_0 = 0$) leads to the expected $p = 2l + 2$ [148]. In the next section we develop a method to construct SBP operators that handle the singular p/r term in the divergence operator and lead to a stable discretization, applicable to any tensorial hyperbolic system, in any spatial dimension, for any multipole expansion.

2.2.4 Order and Degree of Derivative Operators

For classical finite-difference summation-by-parts (FD-SBP) operators, the *degree* of the operator is defined as the maximum degree of the monomials that are differentiated exactly [100], i.e. without any discretization error. The discrete derivative approximation introduces a leading truncation error at each node that scales with some power of the grid spacing h. The *order* of the operator is defined by the minimum exponent of h that characterizes these local truncation errors. For an operator approximating the m-th derivative, the relationship between order and degree is

$$\text{order} = \text{degree} - m + 1. \tag{2.28}$$

Because we are here primarily concerned with first derivatives ($m = 1$), the order of accuracy and the degree of exactness are equivalent for classical FD-SBP operators in Cartesian coordinates.

However, this direct equivalence does not persist when transforming to curvilinear coordinates. While the gradient operator retains this property, the exactness degree of the divergence operator changes. The relationship becomes

$$\text{degree} = \text{order} - p \tag{2.29}$$

where p is the geometric parameter introduced above. Although both the gradient and divergence operators share the same formal order of accuracy, their polynomial degrees of exactness differ. This discrepancy arises from the metric term r^p in the continuous (covariant) divergence operator (e.g., $p = 1$ for cylindrical coordinates and $p = 2$ for spherical coordinates):

$$\mathcal{D}f(r) = \frac{1}{r^p}\frac{\partial}{\partial r}\left(r^p f(r)\right). \tag{2.30}$$

When applying $\mathcal{D}$ to a monomial $f(r) = r^k$, the derivative ∂_r acts on r^{k+p}. Consequently, for the divergence operator to be exact for a polynomial of degree k, the underlying first-derivative FDSBP operator must be exact for polynomials up to degree $k + p$, thereby lowering the effective exactness degree of the operator by p. Assuming an FD stencil of order q, the resulting divergence operator remains q-th order accurate, while its polynomial exactness is reduced according to Eq. (2.29).

2.3 Constructing Summation-by-Parts Operators for Curvilinear Coordinates with Singularities

As described in the previous section, proving the energy stability of a hyperbolic PDE relies on the fact that the spatial derivative operators, i.e. the gradient $\tilde{G} = \partial_r$ and the weighted divergence $\tilde{D} = \partial_r + p/r$ satisfy the integration-by-parts property with respect to the volume element r^p.

Summation-by-Parts (SBP) operators achieve discrete stability by ensuring that the discrete matrix operators satisfy an algebraic analog.

For simplicity, we omit the tilde and asterisk for rescaled and discrete operators and quantities.

2.3.1 Discrete Norm Operators

We use finite differences to construct the discrete operators. We map the radial domain $r \in [0, R]$ onto a grid with N nodes, $r_i = (i-1)\,h$. This places the first grid point $i = 1$ exactly at the origin.

For simplicity and elegance we will present our method for non-staggered grids, i.e. for grids which place one grid point exactly at the origin. Our method can also be applied to staggered grids.

The origin ($r = 0$) becomes a boundary point in our discrete domain. This is not a physical boundary. The boundary conditions there are given by symmetry conditions. Different variables satisfy different symmetry conditions there:

- **Scalar fields:** Regularized scalars are even functions ($\Pi(-r) = \Pi(r)$), and their value at the origin is non-zero in general ($\Pi(0) \neq 0$). Their gradient is odd and vanishes at the origin.
- **Vector fields:** Regularized vectors are odd functions ($\Psi(-r) = -\Psi(r)$), and they vanish at the origin. Their gradient is even and does in general not vanish at the origin.

Above, we already used a two different norm operators (mass matrices) H_S and H_V for the scalar and vector components in our state vector. Here we also define two different discrete SPD norm operators for scalar and vector fields, respectively [148]:

- **The Scalar Norm Operator (S):** Defines the discrete inner product for the even parity fields.
- **The Vector Norm Operator (V):** Defines the discrete inner product for odd parity fields. Because regularized vectors vanish at the origin, their geometric energy contribution at the origin is zero. This decouples the weight at the origin V_{11} from the standard geometric quadrature rules. In principle one could remove the value at the origin from the state vector for odd functions. We

do not do that because this would increase the complexity of our setup for very little practical gain.

In the one-dimensional Cartesian case there is only a single norm H because there is no distinction between scalars and vectors and between gradient and divergence.

Since we include a grid point at $r = 0$, it is tempting to choose the value of S at the origin $S_{11} = 0$. This would make S singular; it would violate the positivity condition S needs to satisfy to be a norm. It is therefore necessary to regularize the operators. To resolve this while preserving discrete stability, we construct a modified scalar norm matrix where the weight at the origin is positive, $S_{11} = \mathcal{O}(h^{p+1}) > 0$. This regularized weight guarantees that S is invertible (which will be necessary to construct the divergence operator). It also maintains the formal order of accuracy of volume integration.

2.3.2 Discrete Derivative Operators

With these parity-specific norms, the continuous derivatives grad and div are approximated by discrete operators G and D. The generalized Summation-by-Parts property then demands

$$SD + G^T V \quad = \quad B \tag{2.31}$$

where $B = \mathrm{diag}(0, \ldots, 0, R^p)$ is a diagonal matrix that projects onto the outer boundary point.

We choose to construct the discrete divergence operator D from the other operators using the generalized curvilinear SBP relation (2.31):

$$D \quad = \quad S^{-1}(B - G^T V)\,. \tag{2.32}$$

That is, we first choose the gradient G, the norms S and V, and the boundary operator B, and then construct the divergence D from these, ensuring that the SBP property holds. Since S is a norm, it is by definition symmetric and positive and thus invertible.

To construct the discrete derivative operators, we start by defining the s-th order accurate cartesian SBP operator G. We do so by starting with a given Cartesian SBP operator. Near the origin, where the stencil coefficients would naively reach "across the origin" into negative radii, we use the symmetry of scalar functions $\Pi(r)$ to map the operator back into the domain. It is important to notice that we do *not* treat the origin as a boundary. The behavior of G near the origin is defined only in terms of the symmetry of scalar functions. We also note in passing that the gradient is an operator that maps scalar functions to vector-valued functions while the divergence maps vector-valued functions back to scalar functions.

We also apply this symmetry condition to the Cartesian norm H, which in practice means that the matrix element H_{11} picks up a factor of $1/2$ compared to its Cartesian value. H_{11} is not directly applicable to our domain because it does not take the volume element into account. It is only a starting point for defining the norm operators S and V.

Finally we define our boundary operator B to be the Cartesian boundary operator, scaled by R^p.

2.3.3 Non-Diagonal Norm Operators

It appears to be impossible to find a consistent set of higher-order operators G, D, S, V, and B that satisfy summation-by-parts while keeping the norm operators diagonal, if one also insists that the operators retain their degree of exactness and convergence order near the origin. (We tried and failed – we have not proved that this is impossible.) Since the origin is "difficult" to handle anyway, it seems prudent to us to insist that our scheme should be as accurate as possible there, and we do not want to loosen this requirement. Instead we thus choose to allow the vector norm operator to have a few off-diagonal elements near the origin, leading to *non-diagonal norm operators*.

In practice this means that there are cross-terms when integrating over the domain. For example, a term $V_{2,3}$ means that the norm $\Psi V \Psi$

contains not only terms such as $V_2\Psi_2^2 + V_3\Psi_3^2$, but that there is also an additional term $\Psi_2 V_{2,3}\Psi_3$ contributing to the norm.

These off-diagonal elements have very little effect in practice. A typical evolution scheme will only use gradient, divergence, and boundary operators, and these retain the usual structure. The norm operators are mostly relevant if one wants to integrate over the domain, e.g. to calculate the total energy. This is usually not part of the evolution scheme, but is only used as analysis method, and the small additional complication posed by a few non-diagonal terms do not pose a large programming burden.

Since we compute the divergence from the SBP relation, which contains the inverse of the scalar norm, we keep the scalar norm S simple and diagonal.

$$S \quad = \quad h \operatorname{diag}(s_1, \ldots, s_N) \tag{2.33}$$

The values of s_i will be determined below.

We construct a non-diagonal vector norm V as described in the following subsections. Given that V is defined for vector-valued functions, which are odd functions and are thus zero at the origin, the elements of the first row and column of V do not matter since they are always applied to a value of $u_1 = u(0) = 0$. We choose to set $\mathrm{V}_{1,1} = 1$ for simplicity.

Below, we construct the operators while setting $h = 1$ for simplicity. (It is straightforward to rescale them by the respective factors h^p.)

Simple and convenient second order operators were constructed in [148]. We do not describe second order operators here and start by constructing fourth-order operators.

4th order accurate non-diagonal norm operators

For a fourth order accurate vector norm V we use the ansatz

$$\mathrm{V} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & v_2 & v_{2,3} & 0 & 0 & 0 & \cdots & 0 \\ 0 & v_{2,3} & v_3 & v_{3,4} & 0 & 0 & \cdots & 0 \\ 0 & 0 & v_{3,4} & v_4 & v_{4,5} & 0 & \cdots & 0 \\ 0 & 0 & 0 & v_{4,5} & v_5 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & s_6 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \cdots & s_n \end{pmatrix} \tag{2.34}$$

We add a small region with one additional symmetric non-zero band near the origin. Further away from the origin we set $\mathrm{V} = \mathrm{S}$. A similar off-diagonal structure is used in [148] for the 4th order operators presented there.

To determine the values of s_i and v_{ij}, we first formally calculate the divergence D from the SBP relation, and we then impose the following *accuracy constraints*:

1. Everywhere, except close to the outer boundary, we demand that the divergence is exact when applied to the linear function $f(r) = r$, i.e. $\mathcal{D}\, r = p + 1$. We cannot demand this near the outer boundary because the underlying Cartesian SBP gradient loses accuracy near the outer boundary.

2. Near the origin we also demand that the divergence is exact for $f(r) = r^3$. We do this to increase the number of conditions to match the number of unknowns (see below).

3. Finally we demand that the volume of the domain is exact, i.e. that $\int f(r)\, r^p\, dr$ is exact for $f(r) = 1$.

We do so because the first and third conditions are special cases, and it seems advantageous to handle them exactly.

We find a sufficient number of off-diagonal non-zero elements by trial and error. The ansatz here leads to a fourth-order operator with good properties. It is certainly possible to allow for more off-diagonal elements, and we do so for defining a sixth-order operator in the next section.

In detail, the accuracy conditions above lead to the equations

$$(\mathrm{D}r)_i = p+1 \qquad i = 1 \ldots N - N_c \tag{2.35}$$

$$(\mathrm{D}r^3)_i = (p+3)\, r^2 \quad i = 1 \ldots 5 \tag{2.36}$$

$$\int_0^R \mathbf{1}\, r^p\, dr = \sum_i s_i = \frac{R^{p+1}}{p+1} \tag{2.37}$$

where N_c is the width of boundary closure (i.e. the number of "special" points near the boundary) of the underlying Cartesian operators.

The operator D above is defined as $\mathrm{D} = \mathrm{S}^{-1}(\mathrm{B} - \mathrm{G}^T\mathrm{V})$, where G and B are known from their respective Cartesian operators. We insert this definition into the conditions above and then multiply both sides by S, leading to

$$(\mathrm{B}^i{}_j - \mathrm{G}_k{}^i \mathrm{V}^k{}_j)\, r^j = (p+1)\, \mathrm{S}^i{}_j\, 1^j \qquad i = 1 \ldots N - N_c \tag{2.38}$$

$$(\mathrm{B}^i{}_j - \mathrm{G}_k{}^i \mathrm{V}^k{}_j)\, (r^3)^j = (p+3)\, \mathrm{S}^i{}_j\, (r^2)^j \quad i = 1 \ldots 5 \tag{2.39}$$

$$1_i\, S^i{}_j\, 1^j = \frac{R^{p+1}}{p+1} \tag{2.40}$$

where 1^i, r^i, $(r^2)^i$, $(r^3)^i$ are vectors of grid point values for the functions $f(r) = 1$, $f(r) = r$, etc.

The only unknowns in these equations are the values s_i and v_{ij}. This system of equations is linear and can be solved in a straightforward manner. We solve this system using rational numbers (not floating-point numbers) to ensure we find an exact solution.

One needs to be careful when counting the number of linearly independent equations because r^i and $(r^3)^i$ are zero at the origin,

reducing the number of linearly independent equations. For this fourth-order operator we demand third-order exactness at 5 points near the origin to ensure that the number of linearly independent equations matches the number of unknowns in S and V. If one wanted to modify our ansatz, then this number 5 might need to be modified accordingly.

Perhaps surprisingly, the solution for the scalar norm S is $s_i = r_i^p H_i$. These are exactly the values that one obtains by scaling the Cartesian norm H by the volume element. The scalar norm has non-trivial entries near the origin (as expected) and near the outer boundary (where the Cartesian operators have a reduced accuracy).

As example, we show the resulting operators for a domain size of $R = 30$ and with a grid spacing $h = 1$ (for simplicity), in spherical symmetry (i.e. $p = 2$), using the Mattsson et. al. operators [198]: The Cartesian norm is

$$\mathrm{H} = \mathrm{diag}\left(\frac{1}{2}, 1, \ldots, 1, \frac{49}{48}, \frac{43}{48}, \frac{59}{48}, \frac{17}{48}\right). \tag{2.41}$$

Our scalar norm:

$$\begin{aligned}\mathrm{S} = \mathrm{diag}\big(&\frac{3714185}{6311896}, \frac{876978}{788987}, \frac{2673132}{788987}, \frac{7012892}{788987}, \frac{101138679}{6311896},\\ &5^2, 6^2, \ldots, 26^2,\\ &\frac{11907}{16}, \frac{2107}{3}, \frac{49619}{48}, \frac{1275}{4}\big).\end{aligned} \tag{2.42}$$

Our vector norm:

$$\mathrm{V} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \frac{2689591}{788987} & \frac{8002}{46411} & 0 & 0 \\ 0 & \frac{8002}{46411} & \frac{3706327}{928220} & \frac{904416}{3944935} & 0 \\ 0 & 0 & \frac{904416}{3944935} & \frac{251857497}{27614545} & -\frac{1132080}{5522909} \\ 0 & 0 & 0 & -\frac{1132080}{5522909} & \frac{89215604}{5522909} \end{pmatrix} \oplus \operatorname{diag}\big(5^2, 6^2, \ldots, 26^2, \frac{11907}{16}, \frac{2107}{3}, \frac{49619}{48}, \frac{1275}{4}\big). \tag{2.43}$$

6th order accurate non-diagonal norm operators

To construct a sixth-order operator, we start with a similar ansatz for the vector norm V, with more non-zero off-diagonal elements:

$$\mathrm{V} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & v_{2,2} & v_{2,3} & v_{2,4} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & v_{2,3} & v_{3,3} & v_{3,4} & v_{3,5} & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & v_{2,4} & v_{3,4} & v_{4,4} & v_{4,5} & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & v_{3,5} & v_{4,5} & v_{5,5} & v_{5,6} & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & v_{5,6} & v_{6,6} & v_{6,7} & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & v_{6,7} & v_{7,7} & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & v_{8,8} & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & v_{9,9} & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & v_{10,10} & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & s_{11} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & s_n \end{pmatrix} \tag{2.44}$$

We impose an additional set of accuracy conditions, namely that the divergence of $f(r) = r^3$ also be exact everywhere except near the outer boundary.

The detailed accuracy conditions are

$$(\mathrm{D}\, r)_i = p + 1 \qquad i = 1 \ldots N, \tag{2.45}$$

$$(\mathrm{D}\, r^3)_i = (p+3)\, r^2 \quad i = 1 \ldots N - N_c, \tag{2.46}$$

$$(\mathrm{D}\, r^5)_i = (p+5)\, r^4 \quad i = 1 \ldots 3, \tag{2.47}$$

$$\int \mathbf{1}\, dx = \sum_i s_i = \frac{R^{p+1}}{p+1}, \tag{2.48}$$

where N_c is again the boundary closure width. As above, the number 3 in the third condition is chosen to ensure that the number of linearly

independent equations matches the number of unknowns. We construct and solve a linear system for the values s_i and v_{ij} in the same way as in the previous subsection. We again find that $s_i = r_i^p h_i$ outside of the block near the origin.

Using the Cartesian operators from Mattsson et. al. [198] (Note one can use various cartesian operators to start from, such as the ones in [106, 199, 197, 196]), we start with the norm H

$$\mathrm{H} = \mathrm{diag}\left(\frac{1}{2}, 1, \ldots, 1, \frac{43801}{43200}, \frac{7877}{8640}, \frac{5359}{4320}, \frac{2711}{4320}, \frac{12013}{8640}, \frac{13649}{43200}\right), \tag{2.49}$$

the scalar norm S is

$$\begin{aligned}\mathrm{S} = \mathrm{diag}\Bigg(&\frac{1430827147971}{4664636476456}, \frac{42904697232509}{51311001241016}, \frac{1539566352701}{424057861496}, \frac{474102956851443}{51311001241016},\\ &\frac{37409168459741}{2332318238228}, \frac{1279079894309069}{51311001241016}, \frac{168033579790977}{4664636476456},\\ &\frac{1257018716149729}{25655500620508}, \frac{298538403436579}{4664636476456}, \frac{4156190224075161}{51311001241016},\\ &10^2, 11^2, \ldots, 24^2,\\ &\frac{1095025}{1728}, \frac{1331213}{2160}, \frac{144693}{160}, \frac{132839}{270}, \frac{10102933}{8640}, \frac{13649}{48}\Bigg).\end{aligned} \tag{2.50}$$

The diagonal entries of the vector norm V (again for $p = 2$) are then:

$$\begin{aligned}\mathrm{V}_{\mathrm{diag}} = \mathrm{diag}\Bigg(&1, \frac{69612410558321}{51311001241016}, \frac{3923059854503}{1509147095324}, \frac{3014861355653139}{359177008687112},\\ &\frac{11554098239601139}{718354017374224}, \frac{29105204758165}{1166159119114}, \frac{461864744704269}{12827750310254},\\ &7^2, 8^2, \ldots, 23^2, 24^2, \frac{1095025}{1728}, \frac{1331213}{2160}, \frac{144693}{160},\\ &\frac{132839}{270}, \frac{10102933}{8640}, \frac{13649}{48}\Bigg),\end{aligned} \tag{2.51}$$

and the off diagonal entries of V are:

$$\mathrm{V}_{2,3} = \frac{4996740529431}{6413875155127} \tag{2.52}$$

$$\mathrm{V}_{2,4} = -\frac{13306004610507}{51311001241016} \tag{2.53}$$

$$\mathrm{V}_{3,4} = \frac{949724456067}{1166159119114} \tag{2.54}$$

$$\mathrm{V}_{3,5} = -\frac{11638692514107}{51311001241016} \tag{2.55}$$

$$\mathrm{V}_{4,5} = \frac{11797110150741}{359177008687112} \tag{2.56}$$

$$\mathrm{V}_{5,6} = \frac{239755863585}{4664636476456} \tag{2.57}$$

$$\mathrm{V}_{6,7} = -\frac{70992217935}{12827750310254} \tag{2.58}$$

2.3.4 Comparison to plain staggered operators

Instead of constructing derivative operators that include a grid point on the origin, one can also construct staggered operators. We consider here the "plain" case, i.e. defining such operators without imposing any additional accuracy constraints. We show that such operators, if they satisfy the SBP property, are stable, but have significantly larger errors than the operators we defined above.

The points of a staggered grid lie at $r_i = (i - \frac{1}{2})\, h$. Assuming a Cartesian norm operator H, we can define scalar and vectors norms via

$$S = V \;\; = \;\; r_i^p \, H \tag{2.59}$$

i.e. by scaling the Cartesian norm in a straightforward manner. Since all $r_i > 0$ this defines a norm.

The divergence is then defined as above via

$$D = S^{-1} \left(B - G^T V \right) \tag{2.60}$$

This leads to a set of SBP operators.

One might be tempted to define the divergence directly via $D_{\text{bad}} = G + p/r$. While well-defined, such operators will in general not satisfy the SBP property, and they lead in general to unstable (growing) modes. The problem is that the identity $r^{-p}\,\partial_r\, r^p = p/r$ holds in the continuum, but not for the discrete derivative operator. This is not a suitable way to discretize the divergence operator.

However, it *is* possible to split the divergence into a partial derivative plus a discrete Christoffel symbol while retaining the SBP property, leading to a stable setup. For a Cartesian SBP operator we have

$$HG + (HG)^T \;=\; B\,. \tag{2.61}$$

We introduce $R = \mathrm{diag}(r_i^p)$ as an operator that scales with the effective volume element. We multiply by R from the right, and insert RR^{-1} into the first term. We obtain

$$HRR^{-1}GR + (RHG)^T \;=\; BR\,. \tag{2.62}$$

We now define $S := HR$, $V := RH$, $D := R^{-1}GR$, and rescale B. We thus find a new set of SBP operators that have been rescaled by R:

$$SD + (VG)^T \;=\; B\,. \tag{2.63}$$

That is, $D = R^{-1}GR$ is a consistent rescaling that defines a divergence that satisfies SBP. We can trivially rewrite this as

$$D \;=\; G + \Gamma \tag{2.64}$$

$$\Gamma \;=\; R^{-1}GR - G\,. \tag{2.65}$$

Of course, this is just a rewritten version (2.60) with exactly the same properties, apart from floating-point accuracy. This ansatz defines the Christoffel symbol Γ in such a way that it exactly complements the discrete partial derivative G. Γ is a linear operator that discretizes the operation "multiply by (p/r)", which is indeed linear operation.

To compare this plain (without imposing additional accuracy conditions) staggered divergence to the operators we constructed in Equation 2.60, we apply these operators to the functions $f(r) = r^3$, $f(r) = r^5$, and $f(r) = r^7$, and compare the discretization errors. See Figures 2.1 and 2.2. Recall that a discrete divergence loses, a priori, p degrees of exactness compared to the discrete gradient since the divergence implicitly divides by r^p, as discussed in Section 2.2.4. We also compare the spectral radius of these operators, which affects the CFL factor, in the next section.

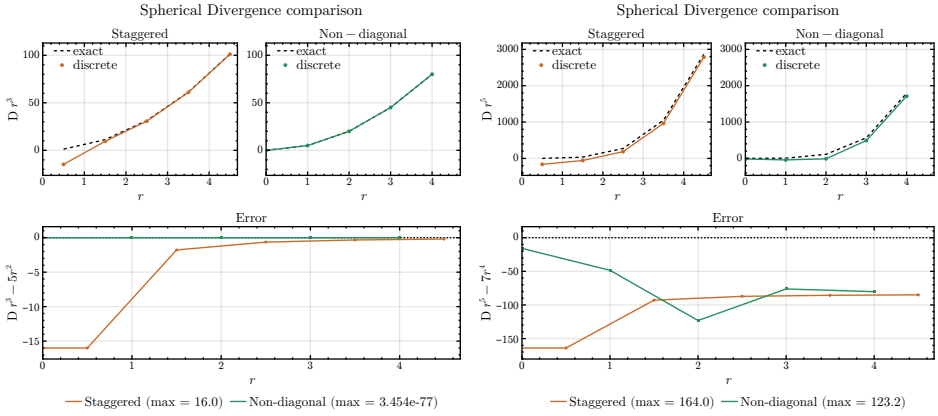


Figure 2.1: Discretization error when applying 4th-order discrete divergence operator to $f(r) = r^3$ and $f(r) = r^5$, for spherical symmetry ($p = 2$). We compare a plain staggered operator to our non-diagonal operator. Our operator is exact for r^3 near the origin, and has a much smaller error near the origin for r^5.

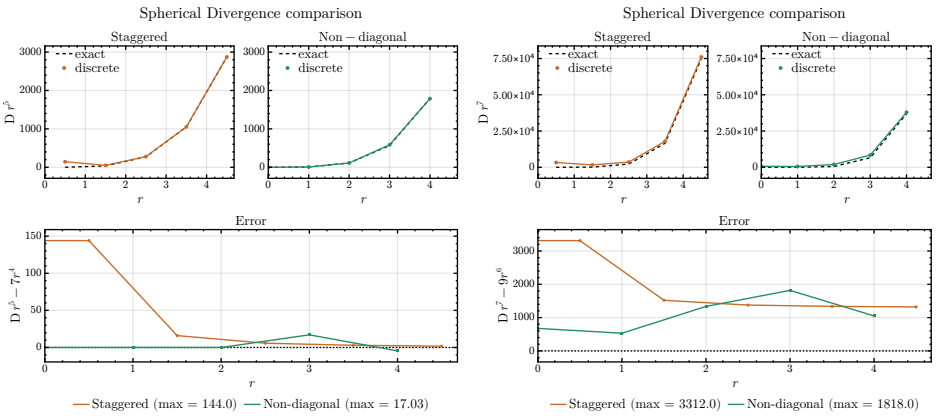


Figure 2.2: Discretization error when applying 6th-order discrete divergence operator to $f(r) = r^5$ and $f(r) = r^7$, for spherical symmetry ($p = 2$). We compare a plain staggered operator to our non-diagonal operator. Our operator is exact for r^5 near the origin, and has a significantly smaller error near the origin for r^7.

2.4 Spectral Stability

In the section above we discussed the stability of the semi-discrete system, i.e. we assumed a discretization in space while keeping the time direction continuous. In this section we comment on the stability of a fully discrete system, assuming that one uses an ODE integrator for time stepping.

The stability of a numerical algorithm solving the wave equation on the spectral properties of its spatial discretization. Its eigenvalues scaled by the timestep, $\lambda \cdot \Delta t$, need to lie within the stability region of the time integrator to avoid growing modes [152]. This imposes constraints on the eigenvalues of the spatial operators. (The eigenvalues will in general be complex numbers.) For example, the eigenvalues cannot have a positive real part, since this corresponds to a growing mode. Moreover, eigenvalues with a large magnitude need to be scaled by

a small enough timestep to ensure they fall into the stability region of the time integrator. Eigenvalues with a large magnitude make an operator *stiff*, i.e. they impose a restriction on the allowed time step size.

The continuous Laplace operator is symmetric and negative semi-definite, guaranteeing purely real, non-positive eigenvalues. In the discrete setting, constructing the Laplacian by multiplying two first-derivative SBP operators ($\mathrm{L}_{\mathrm{lap}} = \mathrm{D} \cdot \mathrm{G}$) places certain demands on the operators that we discuss below.

Transforming the second-order wave equation into a system of first-order equations changes the spatial discretization into a skew-symmetric block structure:

$$\mathrm{L}_{\mathrm{block}} = \begin{bmatrix} 0 & \mathrm{G} \\ \mathrm{D} & 0 \end{bmatrix} . \tag{2.66}$$

Energy stability relies on the SBP property of the derivative operators. The eigenvalues of this block matrix are purely imaginary, representing undamped physical oscillation and propagation of numerical waves.

Applying boundary conditions modifies the real components of the block operator's spectrum.

- **Reflective (energy-conserving) boundaries:** A properly formulated reflective boundary condition preserves the discrete skew-symmetry of the block matrix. The real parts of the eigenvalues remain bounded at machine zero ($|\Re(\lambda)| \approx 10^{-14}$) and prevent unphysical growth or decay. The imaginary eigenvalue pairs with the largest magnitude define the highest resolvable frequencies and lead to the temporal stability bounds for explicit time integrators.

- **Radiative (energy-dissipating) boundaries:** To model an open domain, the radiative boundaries (e.g. applied as Simultaneous Approximation Terms (SAT)) break the global skew-symmetry by introducing a negative-definite penalty. This shifts

the spectrum to the left half of the complex plane. While $\max \Re(\lambda)$ remains zero and guarantees global stability, $\min \Re(\lambda)$ becomes negative, depending on the rate of energy dissipation as the wave exits the computational domain.

To avoid numerical blow-up, the time step size Δt must be constrained so that the scaled eigenvalues $\lambda \cdot \Delta t$ fall within the stability region of the time integrator. For a standard explicit scheme such as the classic fourth-order Runge-Kutta (`RK4`) method [267, 151], the stability region extends along the imaginary axis up to approximately $\pm 2.82i$. Therefore, the stable time step sizes are restricted to

$$\Delta t \lesssim \frac{2.82}{\max |\lambda|} \tag{2.67}$$

The 8th-order accurate `DP8` method by Dormand & Prince [151] that we use in the next section has a larger stability region with size ~ 5.9, allowing either larger time step sizes, or allowing using operators with a larger spectral radius.

We compare the largest eigenvalues and spectral radius of several discrete derivative operators in Table 2.1, where we also examine the staggered operators described in section 2.3.4 above. As the table shows, the SBP operators are stable, and radiative boundaries dissipate energy, as expected. As described in the previous section, our new SBP operators have an increased accuracy near the origin, as prescribed in our accuracy conditions.

The staggered non-SBP operators we introduced in the previous section are unstable. This is not a CFL instability – there are one or more growing modes in the discretization, independent of the time discretization. We did not investigate whether adding artificial dissipation can stabilize the system.

Order	Setup	reflecting (Dirichlet) boundaries			radiative (SAT) boundaries		
		$\max \Re(\lambda)$	$\min \Re(\lambda)$	$\lvert\rho(\lambda)\rvert$	$\max \Re(\lambda)$	$\min \Re(\lambda)$	$\lvert\rho(\lambda)\rvert$
4th	our operators	0	0	1.94	0	−0.293	1.851
	SBP staggered	0	0	2.0	0	−0.291	2.0
	non-SBP staggered	0.025	−0.025	1.98	0.024	−0.233	1.4
6th	our operators	0	0	2.47	0	−0.331	1.981
	SBP staggered	0	0	2.47	0	−0.329	2.20
	non-SBP staggered	0.028	−0.028	2.44	0.036	−0.260	1.930

Table 2.1: Largest real eigenvalues and spectral radii ρ for several derivative operators using the cartesian operators from [198], assuming the first-order discretization in (2.66). As discussed in the main text, a larger spectral radius requires smaller time step sizes. Eigenvalues with a positive real part indicate a growing mode (an instability), a negative real part indicates damped modes (possibly from outgoing boundary conditions). All SBP operators are stable, and the non-SBP staggered operators are unstable (they have a growing mode). The radiative boundaries dissipate energy, as expected. Other cartesian operators produced the same qualitative results.

2.5 Numerical evolutions

We demonstrate the higher-order stencils by solving the scalar wave equation in spherical symmetry. This corresponds to setting $p = 2$. We set up a Gaussian wave packet that starts away from the origin, moves inwards, "reflects" at the origin (or moves through the origin, depending on point of view), and then expands outwards again.

The general solution for this PDE is given by

$$\Phi(t, r) = \frac{g(t + r) - g(t - r)}{r} \tag{2.68}$$

for an arbitrary differentiable function g which depends only on a single argument. This ansatz satisfies the scalar wave equation in spherical symmetry, and by inspection it is clear that Φ does not have any

singularities at the origin. Numerically, l'Hôpital's rule needs to be used at the origin.

We choose g to be a Gaussian

$$g(u) \quad = \quad A \exp\left(-\frac{(u - r_c)^2}{d^2} \right) . \tag{2.69}$$

with amplitude $A = 1$, width $d = 2$, and initial position $r_c = 10$ for a domain $[0, 40]$. We use reflecting boundary conditions at the outer boundary $r = R$ (although we are not interested in the behavior near the outer boundary here – we want to demonstrate convergence and accuracy near the origin).

Time integration uses Hairer's 8/5/3 Dormand-Prince method `DP8` [151]. It is implemented in `OrdinaryDiffEq.jl` [243], with fixed time step Δt and adaptive stepping disabled. The CFL factor is $\mathrm{cfl} = 0.5$.

We show in Figure 2.3 a spacetime plot which gives an overview over the evolution of this system. The wave packet starts at $r_c = 10$, and then moves inwards as it amplitude increases like $1/r$. It interacts with the origin and then moves outwards again, its amplitude falling off as $1/r$. The overall evolution looks clean, although this is not a very strong statement given the coarse scale of the color map.

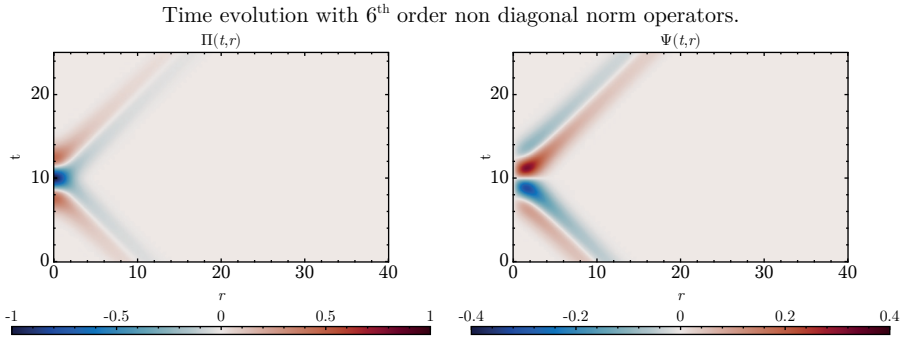


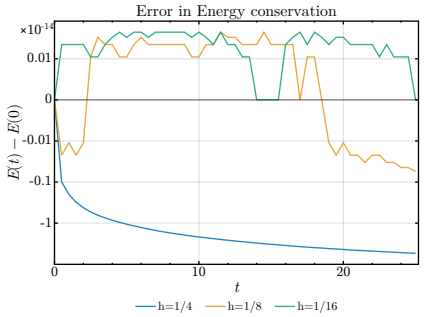
Figure 2.3: Time evolution for the 6th-order non-diagonal SBP operator in spherical symmetry ($p = 2$), shown at resolution $h = 1/16$ for $0 \leq t \leq 25$. The corresponding time evolutions for the other operators are visually indistinguishable.

2.5.1 Energy Conservation

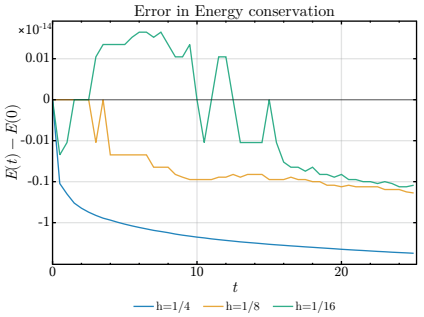
The discrete energy of the system is

$$E(t) = \frac{1}{2} \left(\Pi S \Pi + \Psi V \Psi \right). \tag{2.70}$$

We use reflecting boundary conditions, and thus energy must be conserved. We show in figure 2.4 the error in energy conservation $E(t) - E(0)$ up to $t = 25$. At this time the wave packet has reflected off the origin. The graphs show that energy is indeed conserved up to floating-point round-off errors, demonstrating the stability of our discretization.

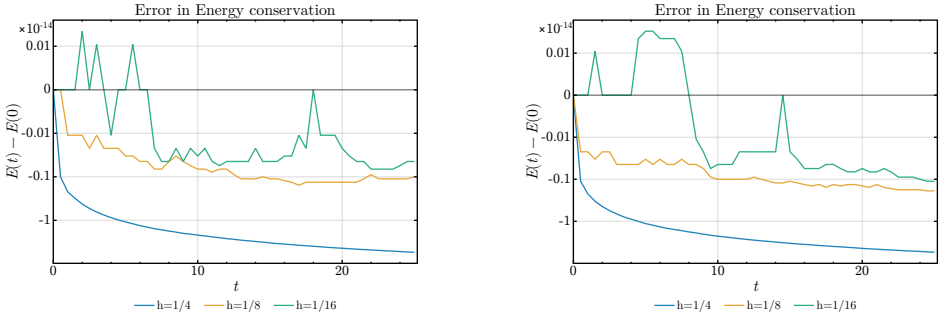


(a) 4th order non-diagonal operator



(b) 6th order non-diagonal operator

Figure 2.4: Error in energy conservation for the non-diagonal-norm SBP wave evolutions. The left and right subfigures show the 4th- and 6th-order operators using $p = 2$ and $R = 40$. Three resolutions are shown: $h = (1/4, 1/8, 1/16)$ over the time interval $0 \leq t \leq 25$, using a CFL factor of $1/2$. Note the graphs are scaled by 10^{-14}, i.e. they are showing only floating-point round-off errors. The lowest resolution $h = 1/4$ shows a slight decrease in the total energy over time. We confirmed (not shown here) that this can also cured by reducing the CFL factor.



(a) 4th order staggered operator

(b) 6th order staggered operator

Figure 2.5: Error in energy conservation for the staggered SBP wave evolutions. The left and right subfigures show the 4th- and 6th-order operators using $p = 2$ and $R = 39.5$. Three resolutions are shown: $h = (1/4, 1/8, 1/16)$ over the time interval $0 \leq t \leq 25$, using a CFL factor of $1/2$. Note the graphs are scaled by 10^{-14}, i.e. they are showing only floating-point round-off errors. The lowest resolution $h = 1/4$ shows a slight decrease in the total energy over time. We confirmed (not shown here) that this can also cured by reducing the CFL factor.

2.5.2 Accuracy and Convergence

We investigate the accuracy and convergence of the operators by showing the scaled differences to the exact solution for different resolutions at $t = 10$ and $t = 25$. At $t = 10$, the pulse is interacting with the origin. At $t = 25$, the pulse has reflected off the origin and is moving outwards. Figures 2.6 and 2.8 shows the scaled errors at these times. The errors have been scaled according to the expected convergence order (4 and 6, respectively), and good convergence is indicated by scaled errors that lie on top of each other. At $t = 10$, near the origin, the convergence is not quite clean, and a small amount of noise is visible. At late times ($t = 25$), convergence is clean again, and there is a small amount of noise visible that propagates faster than the pulse. Please consider that these figures show the error, not the solution, and

that the high-resolution errors have been scaled by $16^6 \approx 1.6 \cdot 10^{-7}$, i.e. these errors are quite small. The figures also show that the noise converges away with increasing resolution.

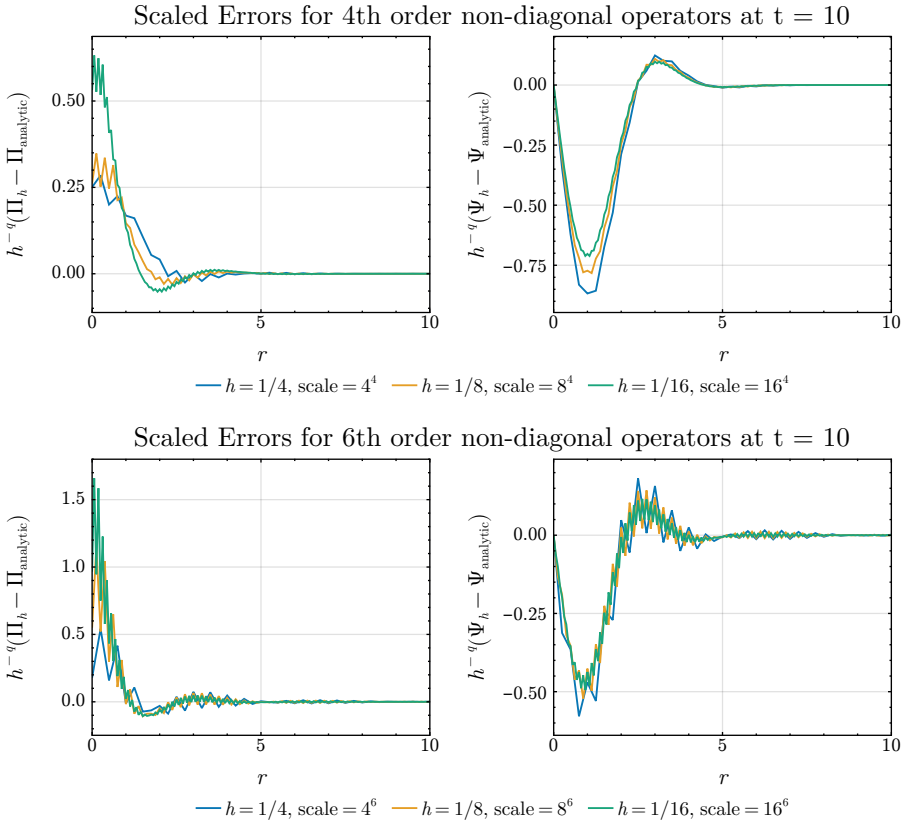


Figure 2.6: Convergence diagnostics for the 4th-order and 6th-order non-diagonal-norm SBP operators at $t = 10$ for the resolutions $h = (1/4, 1/8, 1/16)$. Shows scaled errors (not the solution!), scaled by h^{-4}. This is the configuration $p = 2$ on $r = [0, 40]$. The noise amplitude for the highest resolution $h = 1/16$ is at the level of 10^{-8} and converges away with increase resolution.

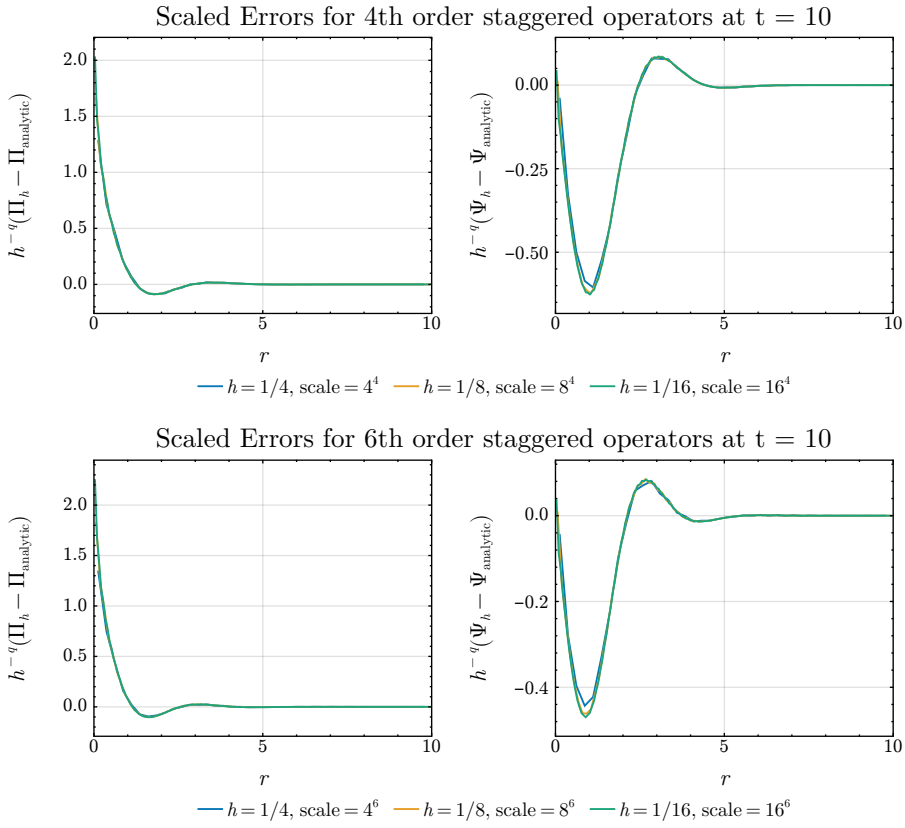


Figure 2.7: Convergence diagnostics for the 4th-order and 6th-order staggered SBP operators at $t = 10$ for the resolutions $h = (1/4, 1/8, 1/16)$. Shows scaled errors (not the solution!), scaled by h^{-4}. This is the configuration $p = 2$ on $r = [0, 39.5]$. The noise amplitude for the highest resolution $h = 1/16$ is at the level of 10^{-8} and converges away with increase resolution.

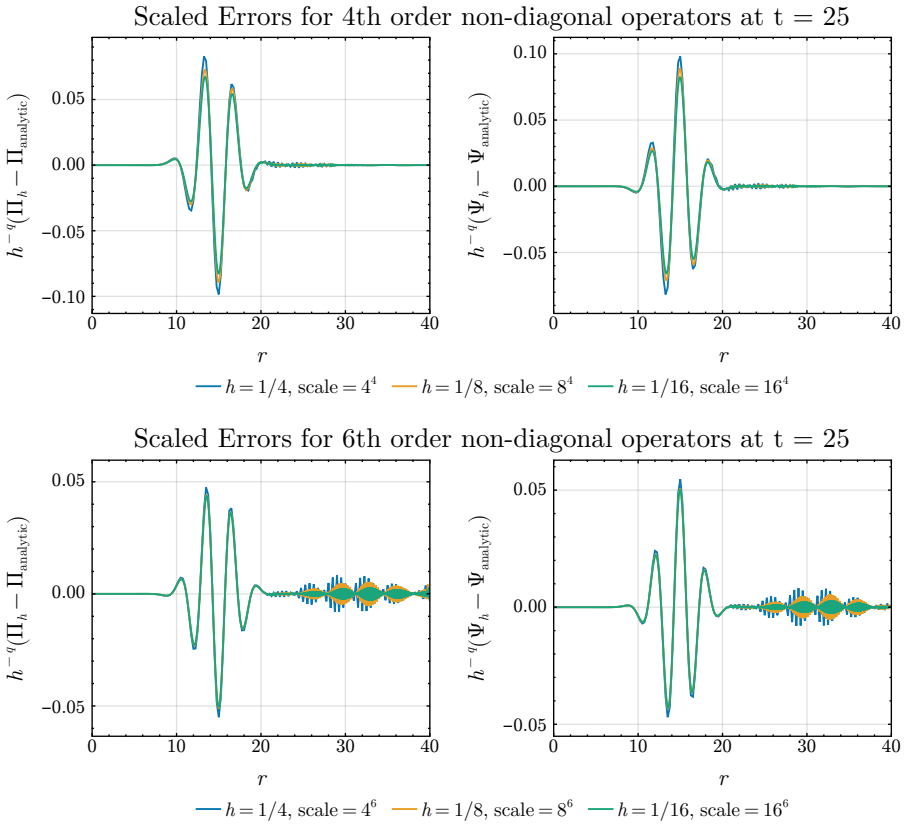


Figure 2.8: Convergence diagnostics for the 4th-order and 6th-order non-diagonal-norm SBP operator at $t = 25$ for the resolutions $h = (1/4, 1/8, 1/16)$. Shows scaled errors (not the solution!), scaled by h^{-4}. This is the configuration $p = 2$ on $r = [0, 40]$. The noise amplitude for the highest resolution $h = 1/16$ is at the level of 10^{-9} and converges away with increased resolution.

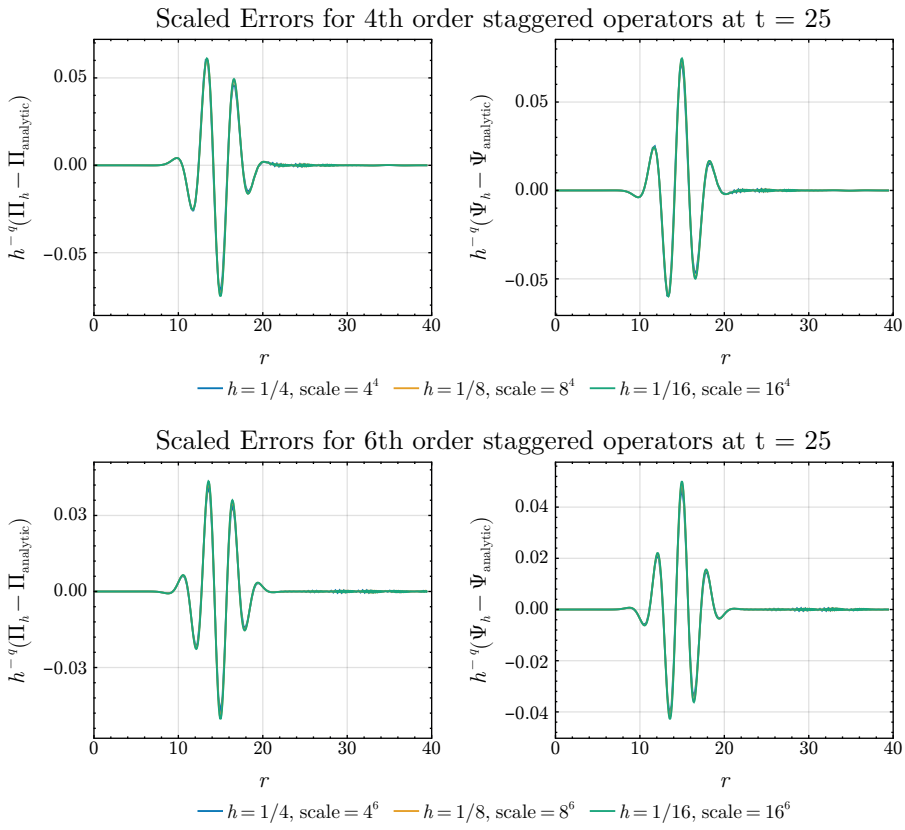


Figure 2.9: Convergence diagnostics for the 4th-order and 6th-order staggered SBP operator at $t = 25$ for the resolutions $h = (1/4, 1/8, 1/16)$. Shows scaled errors (not the solution!), scaled by h^{-4}. This is the configuration $p = 2$ on $r = [0, 39.5]$. The noise amplitude for the highest resolution $h = 1/16$ is at the level of 10^{-9} and converges away with increased resolution.

2.6 Conclusion

We describe a mechanism to construct summation-by-parts operators for finite differences in curvilinear coordinate system with singularities. Our ansatz generalizes ideas presented in [148]. Starting from existing Cartesian SBP operators, we distinguish between a scalar and a vector-valued norm, allow the vector-valued norm to have non-diagonal elements, and choose a set of accuracy conditions near the origin that determines the elements of these operators. We then derive the covariant divergence operator from the SBP property. We show examples for sets of operators in three-dimensional spherical symmetry, demonstrating that these operators lead to excellent (accurate, smooth, higher-order convergent) solutions at and near the origin, even though we place a grid point directly onto the origin. We also briefly reviewed how to construct stable SBP operators for staggered grids and compare to these operators.

Our ansatz is applicable to operators for any tensor rank and any spacetime dimension.

We make the operators available at `SphericalSBPOperators.jl` repository [301].

Plots were generated using `Makie.jl` [96], time evolutions were performed with `OrdinaryDiffEq.jl` [243]. We used extended precision types from `MultiFloats.jl` [314] and eigenvalues computed with `GenericSchur.jl` [272]. The cartesian SBP operators were computed with `SummationByPartsOperators.jl` [251].


3

Chapter 3

Numerical self-force calculations for scalar particles, formulated in the lab frame

Publication status

This chapter is based on the following paper:

> **S. Vretinaris**, E. Schnetter, B. Krishnan, "Numerical self-force calculations for scalar particles, formulated in the lab frame", arXiv: 2606.06487, Submitted to Physical Review D.

Contributions

I contributed to the formulation of the problem, development of the method, implementation, code writing, simulations, production of figures, interpretation of results, and manuscript writing.

Abstract

We derive equations of motion for scalar particles self-consistently interacting with a scalar field, including the radiation produced by the particles' acceleration. Our approach differs in three key aspects from current methods: (1) we assume a small but finite discretization length scale h, which allows us to treat the particle as a small but finite object, (2) we choose the state vector for the system before deriving equations of motion, and (3) we formulate the equations explicitly in the lab frame and not in a manifestly covariant manner. This approach, which is self-consistent, happens to greatly simplify the resulting equations and their derivation, and is directly suitable for numerical calculations. The result is an effective source method which generalizes to electrodynamics or general relativity in a straightforward manner (although we do not consider this here). We then provide two possible discretizations of these equations, based on finite volumes and spectral methods, and show results of one-dimensional calculations. These calculations show excellent agreement with analytic solutions.

3.1 Introduction

Finding solutions for physical systems which couple point particles and fields is difficult because point particles act as formally singular sources for the field. Doing so numerically is even more complicated. Yet such systems are physically of great interest because point particles are excellent approximations for small extended bodies. For example, in relativistic astrophysics, Extreme Mass Ratio Inspirals (EMRIs) are such system: There is a large black hole (with size M) and a secondary, much smaller compact object of size $m \ll M$. If the internal degrees of freedom of the small object can be neglected, then it is well approximated by a point particle.

While it is theoretically possible to discretize such a system using a

length scale $h \ll m$ to resolve both bodies, the computational cost is prohibitive. For a standard discretization with a cutoff at wavelength h, the cost scales as $O(1/h^4)$ in four-dimensional spacetime. Numerical calculations for such a system are thus intractable as the mass ratio m/M decreases. Current state-of-the-art simulations find a mass ratio of 10^{-1} computationally feasible, and while 10^{-2} has been attempted, more extreme mass ratios are unfeasible with standard methods.

The gravitational self-force approach describes the motion of a small, compact body of mass $m \ll M$ as a perturbation on a fixed background spacetime. At leading order, the small body follows a geodesic of the background geometry; at first order, its own metric perturbation modifies that motion. The central difficulty is that the retarded field of a point mass diverges on its worldline, so any physically meaningful equation of motion requires a regularization prescription. Building on earlier radiation-reaction analyses in flat and curved spacetime [112, 104], Mino, Sasaki and Tanaka and Quinn and Wald derived the first-order gravitational self-force equations now known as the MiSaTaQuWa equations [204, 242]. The Detweiler–Whiting decomposition then clarified the split of the retarded field into a singular piece, which exerts no force, and a regular piece that leads to acceleration [103]. Subsequent matched-asymptotic and self-consistent derivations established the gravitational equation of motion [144, 235], while extensions to second perturbative order developed the nonlinear gravitational self-force framework needed for phase-accurate EMRI waveforms [258, 259, 236, 143, 102]. For a comprehensive review of the self-force problem, see [224, 101, 36, 66].

To make self-force calculations computationally feasible, one must define a finite regular field, or the force derived from it, without evaluating divergent quantities on the worldline. Dissipative and flux-based schemes avoid local regularization by using radiative fields or wave fluxes at infinity and the horizon, providing efficient adiabatic inspirals while omitting conservative effects [203, 162, 114, 285]. For local self-force calculations, three regularization strategies are commonly

used: Mode-sum regularization decomposes the retarded field into finite spherical-harmonic modes and subtracts analytic regularization parameters before summing over modes [39, 38]. Worldline-convolution methods compute the regular field from a Green function integral over the particle's past history, using quasilocal expansions together with frequency or time domain Green function calculations to make the approach practical [225, 11, 84, 307, 313]. Effective-source or puncture methods instead subtract an approximate singular field at the level of the differential equation, yielding a finite residual field suitable for time or frequency domain evolutions in $1+1$, $2+1$, or $3+1$ dimensions [37, 297, 113, 108, 305].

For the scalar self-force problem, self-consistent evolutions were first attempted by Diener et al. [107] using an effective-source method. In that implementation, however, numerical stability required omitting parts of the local singular-field approximation beyond velocity-dependent terms. Later, Heffernan et al. [157] extended the singular field to include acceleration-dependent contributions, but these terms introduce a self-referential dependence on higher time derivatives of the particle's self-acceleration. As discussed in Diener [105], retaining such higher-derivative terms can lead to instabilities in the coupled particle-field evolution.

The first fully self-consistent evolutions of a scalar-charged particle on a Schwarzschild background were reported by Wittek et al. [311, 310] using a worldtube excision framework. In their formulation, the acceleration-dependent part of the local analytical field introduces an implicit coupling: the puncture field depends on the particle's self-acceleration and its derivatives, while the self-acceleration is determined from the regularized field obtained through the matching procedure. At the order considered, evaluating the acceleration-dependent puncture and its time derivative requires quantities such as derivatives of the self-force, the first and second time derivatives of the four-velocity, and higher time derivatives of the regular field on the worldline. Rather than treating these higher derivatives as independent evolution variables

or additional initial data, Wittek et al. close the system through an iterative scheme, initialized with the geodesic acceleration, in which the puncture, regular-field coefficients, and particle acceleration are updated successively. They further note that the acceleration terms can lead to instabilities in parts of parameter space, possibly reflecting the fact that the coupled point-particle system becomes effectively higher than second order in time; their iteration can therefore be viewed as analogous to a reduction-of-order treatment.

From the perspective of an initial-value formulation, the appearance of self-referential higher derivatives of the particle's acceleration is problematic. A well-posed evolution scheme should determine the future state of the coupled particle–field system from appropriate initial data for the field, the particle position, and the particle velocity. If the local singular-field approximation entering the effective source depends on time derivatives of the self-acceleration, then the force law is no longer a standard second-order equation of motion coupled to a hyperbolic field equation. Rather, the acceleration required to advance the worldline depends implicitly on higher derivatives of that same acceleration.

In this work, we develop a new framework for self-consistent scalar self-force evolutions that incorporates acceleration effects while preserving the structure of an initial-value formulation. The state of the coupled particle-field system is specified by a finite set of initial data, and its subsequent evolution is determined by equations that do not require higher time derivatives of the particle's acceleration. In contrast to formulations in which acceleration-dependent singular-field terms generate self-referential higher-derivative contributions, our approach avoids introducing additional derivative data beyond those required by the physical initial-value problem.

The framework is closely related in spirit to the effective-source method, but it does not rely on an expansion in powers of the particle's motion to construct the regularized evolution system. As a result, acceleration effects are incorporated naturally into the dynamics of the

particle, rather than entering through higher-order local time-derivative corrections. This yields a self-consistent evolution scheme capable of accounting for accelerated motion while remaining compatible with a stable initial-value formulation.

Below, after giving details of the physical system we consider in section 3.2, we derive the equations of motion in stages: we first consider a scalar field with a given charge distribution, then a particle moving in a given background field, and then we couple the two systems in section 3.3. We pay particular attention to explicitly factoring out a *particle field* in section 3.3.3, as is common in self-force approaches. (Since we consider a finite-sized particle, this particle field is non-singular.) Finally, we describe our numerical methods in section 3.4 and present results in section 3.5.

3.2 The physical system

3.2.1 Overview

We adopt a framework based on the hierarchy of scales $M \gg h \gg m$. Because we are not concerned with physics at scales smaller than h, short-wavelength and high-frequency phenomena – such as the quasi-normal modes of the secondary – are neglected. We approximate the small black hole by a particle of effective size $\ell \lesssim h$. Within this approximation, ℓ is a free parameter independent of m, provided it remains significantly smaller than the grid scale h, so that the internal structure of the secondary can be neglected.

To simplify the present exposition, we forgo the complexities of general relativity and consider a scalar particle coupled to a scalar field in flat spacetime.

3.2.2 A general remark on point particles

Before describing our ansatz, we comment on the notion of a *point particle*. Point particles were probably first introduced in a self-consistent manner by Newton, who showed that the gravitational field of a spherically symmetric mass distribution is equal to that of a point mass. This statement also holds in electrodynamics and in general relativity. Point particles have been a hugely successful concept in physics.

However, none of the currently known particles have been shown to be a point particle by observational evidence, nor will this ever be possible. Any observation and any experiment have a finite accuracy, and we will thus only ever be able to put an upper limit on a particle's size, and will never be able to show that a particle is actually a point particle. This means that we use the concept of a point particle only because it is mathematically convenient, not because it is necessary for consistency with physics. For example, the current upper limit for the electron radius is about 1 attometer, and thus any calculation that assumes a finite electron radius smaller than that will produce results that are consistent with all known observations.

For simplicity, we will forgo the notion of a point particle here, and instead consider spherically symmetric particles with a small but finite radius ℓ.

Finite-sized rigid particles are in principle incompatible with special relativity. A force acting on one end of the particle would need to accelerate that side of the particle before information about the force has reached the other side, and a particle with internal degrees of freedom (e.g. compressibility) seems the only option. However, since we only study the spacetime with a finite resolution h, we cannot see the details of such an interaction which happens at a length scale $\ell \lesssim h$. Such approximations are ubiquitous in numerical methods: Consider e.g. a finite volume method where the manifold is split into cells of size h, and where the fluid in each cell is described by a single, averaged velocity v. Such discretizations are consistent, and in the same way a rigid particle is consistent with relativity in our case. We discuss

convergence in the limit $h \to 0$ in section 3.4 below.

3.3 Equations of motion

3.3.1 A scalar field with a given charge distribution

To set up notation we briefly describe the scalar field we consider without particles present. The scalar field Φ is described by the action

$$S = \frac{1}{2}\int \eta^{ab}(\partial_a\Phi)(\partial_b\Phi) + \rho_{\mathrm{e}}\Phi \tag{3.1}$$

with an external (given) scalar charge distribution ρ_{e}. We assume that ρ_{e} is given as a function of space and time. After a $3+1$ split, the equations of motion are then

$$\partial_t\Phi = \Pi \tag{3.2}$$
$$\partial_t\Pi = \delta^{ij}\partial_i\partial_j\Phi - \rho_{\mathrm{e}}\,. \tag{3.3}$$

Here we introduced a new variable $\Pi := \partial_t\Phi$ to express the equations of motion as a system that is first order in time. The state vector of the system, i.e. the instantaneous conditions that define the future evolution of the system, consists of instantaneous values for Φ and Π.

3.3.2 A particle in a given scalar field

Next we examine the motion of a scalar particle in a fixed background field Φ_b which is given as a function of space and time. We neglect the radiation emitted by the particle in this section.

We pay particular attention to defining the finite-sized particle in a consistent manner. We define the particle's charge distribution in two steps: First we consider a particle at rest at the origin, then a particle moving with constant velocity.

Particle at rest

We define the charge distribution for a particle at rest at the origin to be $\rho_0(x^i)$. As described above, we assume that ρ_0 has a finite support of size ℓ, with $\ell \lesssim h$ where h is a given length cutoff parameter.

This charge distribution is associated with a scalar field $\Phi_0(x^i)$ through the Poisson equation

$$\delta^{ij}\partial_i\partial_j\Phi_0 \;:=\; \rho_0\,. \tag{3.4}$$

This particle has a charge Q and a mass (energy) M given by

$$Q \;:=\; \int \rho_0 \tag{3.5}$$

$$M_0 \;:=\; m_0 + \frac{1}{2}\int \delta^{ij}(\partial_i\Phi_0)(\partial_j\Phi_0) + \int \rho_0\Phi_0 \tag{3.6}$$

$$=\; m_0 + \frac{1}{2}\int \rho_0\Phi_0\,. \tag{3.7}$$

Here, m_0 represents a possible intrinsic mass; we will set it to zero below. Since we assume that the particle has a finite radius ℓ, all quantities above are well-defined and finite. These conditions are *constraints* on the possible choices of the charge distribution function $\rho_0(x^i)$, if one wants to model a particle with a given charge Q and rest mass M_0.

Moving particle

We consider a particle located at position z^i and moving with a constant three-velocity v^i. We describe the particle at a fixed instance of time in the lab frame. This will allow us below to make the particle part of the state vector; a state vector describes a physical system at an instance of time.

The charge distribution and potential are given by boosting the quantities ρ_0 and Φ_0 defined above from the particle rest frame to the

lab frame:

$$\rho_{\text{lab}}(z^i, v^i; x^i) := \gamma\rho_0(\Lambda^i_j(x^i - z^i)) \tag{3.8}$$
$$\Phi_{\text{lab}}(z^i, v^i; x^i) := \Phi_0(\Lambda^i_j(x^i - z^i)) \tag{3.9}$$

where $\Lambda^i_j = \Lambda^i_j(v^k)$ is the spatial part of the Lorentz transformation for the velocity v^i, and $\gamma = \gamma(v^i) = 1/\sqrt{1 - v^2}$ is the respective Lorentz factor. We treat ρ_0 as a scalar density that scales with γ so that the total particle charge Q remains invariant under a Lorentz boost ($Q_{\text{lab}} = Q$). A calculation shows that $M_{\text{lab}} = \gamma M_0$ as expected.

To avoid ambiguity in the notation, ρ_{lab} and Φ_{lab} are defined to be functions of three arguments: the particle position and velocity as well as a position in space. As defined, these functions do not have an explicit time dependence. These functions define, instantaneously in the lab frame, the charge distribution and field for a respective particle that moves uniformly.

We will also be careful in our notation when using partial and total derivatives. A partial derivative is always applied to a function, and denotes the derivative with respect to one of the function's arguments. A total derivative is applied to an expression, and denotes the derivative with respect to one of the variables appearing in that expression. (This is the same distinction made e.g. in Mathematica between `Derivative` and `D`.) For example, we can write $\partial/\partial x_i\, \rho_0$ (because ρ_0 is a function that depends on an argument named x_i), or we can write $d/dt\, \sin(t^2)$ (since the variable t appears in the expression). It does *not* make sense to write $\partial/\partial t\, \rho_{\text{lab}}$ since ρ_{lab} does not have an argument named t. It *does* make sense to write $d/dt\, \Phi_{\text{lab}}(z^i + tv^i, v^i; x^i)$ because the variable t appears in that expression.

We will later also need the time derivatives of the particle field in the lab frame. We define Π_{lab} and Ξ_{lab} to be the first and second time derivatives of Φ_{lab} that would be seen in the lab frame for this

non-accelerating particle:

$$\begin{aligned}
\Pi_{\text{lab}}(z^i, v^i; x^i) &:= \frac{d}{dt}\Phi_{\text{lab}}(z^i + tv^i, v^i; x^i) && (3.10)\\
&= v^i \partial_{z^i}\Phi_{\text{lab}}(z^i, v^i; x^i) && (3.11)\\
&= -v^i \partial_i \Phi_{\text{lab}}(z^i, v^i; x^i) && (3.12)\\
\Xi_{\text{lab}}(z^i, v^i; x^i) &:= \frac{d}{dt}\Pi_{\text{lab}}(z^i + tv^i, v^i; x^i) && (3.13)\\
&= v^i \partial_{z^i}\Pi_{\text{lab}}(z^i, v^i; x^i) && (3.14)\\
&= -v^i \partial_i \Pi_{\text{lab}}(z^i, v^i; x^i) && (3.15)\\
&= v^i v^j \partial_i \partial_j \Phi_{\text{lab}}(z^i, v^i; x^i) && (3.16)
\end{aligned}$$

where we write ∂_i for ∂_{x^i}.

With these definitions, the wave equation holds for the non-accelerating particle:

$$-\Xi_{\text{lab}} + \delta^{ij}\partial_i \partial_j \Phi_{\text{lab}} = \rho_{\text{lab}}\,. \qquad (3.17)$$

All quantities defined here so far (ρ_0, Φ_0, ρ_{lab}, Π_{lab}, Ξ_{lab}) are defined without explicit time dependence. In other words, these are functions that map a given particle state (position, velocity) to a charge distribution and field *instantaneously* in the lab frame. This is an important point where this work differs from other approaches which use radar coordinates, retarded Green's functions, or similar methods that propagate information causally.

Zeno's Arrow Paradox

In the arrow paradox [187, 300], Zeno of Elea (c. 490–430 BC) argues:

> ... If everything when it occupies an equal space is at rest at that instant of time, and if that which is in locomotion is always occupying such a space at any moment, the flying arrow is therefore motionless at that instant of time and at the next

> instant of time but if both instants of time are taken as the same instant or continuous instant of time then it is in motion (Aristotle Physics, 239b5-7) [161].
>
> ... What is in motion moves neither in the place it is nor in one in which it is not (Diogenes Laertius, Lives of Famous Philosophers, ix.72) [185].

In Zeno's arrow paradox, the arrow is momentarily at rest for each time instance. The paradox arises from the attempt to reconcile the discrete nature of time (as a sequence of instants) with the continuous nature of motion. This instantaneity is also present in numerical simulations at a given time step. (Zeno's paradox is resolved for us because our state vector holds both the position and the velocity of the particle. Hence the particle is not "at rest at that instant in time" — it retains its velocity, even if only a single instance of time is considered.)

Our definition of the particle field follows a similar philosophy as Zeno's paradox. We define it in the instantaneous rest frame of the particle. This could lead to confusion: since we assume uniform motion every time we define the particle field, one might be led to think that there is no acceleration, in the same way as in Zeno's paradox there is no motion. To better grasp the idea behind our definition of the particle field, one can imagine Zeno's arrow paradox, but instead of uniform motion, an arrow that is accelerating. At each time instance the arrow would not accelerate but move with constant velocity. (This is not a paradox since the state vector does not contain the acceleration.)

Particle equations of motion

We describe the particle trajectory by the functions $z^i_{\mathrm{p}}(t)$ and $v^i_{\mathrm{p}}(t)$. The well-known equations of motion for the particle in a given external

(background) potential Φ_e are then

$$\partial_t z_p^i = v_p^i \tag{3.18}$$

$$\gamma\left(M_0 + Q\Phi_e\right)\partial_t v_p^i = -\delta^{ij}\int \rho_p \partial_j \Phi_e \tag{3.19}$$

$$\approx -Q\delta^{ij}\partial_j \Phi_e(z_p^i) \tag{3.20}$$

with the particle charge distribution

$$\rho_p(t, x^i) := \rho_{lab}(z_p^i(t), v_p^i(t); x^i)\,. \tag{3.21}$$

In this approximation, we assumed that the particle has size ℓ (i.e. that ρ_p has compact support of size ℓ) and that Φ_e only varies slowly compared to $h \gg \ell$, so that the curvature of Φ_e can be neglected (i.e. that the particle acceleration a satisfies $a \ll 1/h$).

These equations of motion are equivalent to the covariant formulation [224, 101, 36]:

$$\left(M_0 - Q\Phi_e\right)\frac{\partial}{\partial\tau}u^a = -Q\left(\eta^{ab} + u^a u^b\right)\partial_b \Phi_e\,. \tag{3.22}$$

We choose to keep the charge of the particle constant in time.

3.3.3 Scalar field for a given particle trajectory

We now discuss the evolution of a scalar field in the presence of a scalar particle with a given trajectory. The emphasis in this section lies on handling the field Φ with a given spatial resolution h in the presence of a particle with a length scale ℓ where $\ell \lesssim h$. The particle's charge distribution introduces a very small length scale ℓ that needs to be handled well to avoid a large discretization error.

The particle trajectory (position and velocity) is given as functions of time via $z_p(t)$ and $v_p(t)$, respectively.

Describing this system in a self-consistent manner is purely a matter of choosing a suitable discretization method. As mentioned earlier, if

we were to choose a discretization length scale $h \ll \ell$, then this would be straightforward. Our task here is to invent a discretization method that is consistent and accurate for the case $h \gtrsim \ell$, i.e. a method which can handle a small particle and the associated large field gradients near the particle.

To do so, we use the well-known approach to split the physical fields Φ and Π into two parts:

$$\Phi = \Phi_{\rm r} + \Phi_{\rm p} \tag{3.23}$$
$$\Pi = \Pi_{\rm r} + \Pi_{\rm p} \tag{3.24}$$

where we define the *particle fields* $\Phi_{\rm p}$ and $\Pi_{\rm p}$ in a convenient manner so that they "handle" the large field gradients. The *remainder fields* $\Phi_{\rm r}$ and $\Pi_{\rm r}$ will then not have any large gradients. Since we are introducing new degree of freedom here, we are free to choose $\Phi_{\rm p}$ and $\Pi_{\rm p}$ (in principle) arbitrarily, and this choice then defines $\Phi_{\rm r}$ and $\Pi_{\rm r}$.

It is important to note that neither $\Phi_{\rm r}$ nor $\Phi_{\rm p}$ can be measured; only the total field Φ is physical. We can thus define $\Phi_{\rm p}$ in any way we deem convenient. Neither $\Phi_{\rm r}$ nor $\Phi_{\rm p}$ must satisfy any causality conditions, only their sum Φ needs to do so. In a manner of speaking, $\Phi_{\rm p}$ and $\Pi_{\rm p}$ can be viewed as additional numerical basis functions that we introduce to keep $\Phi_{\rm r}$ and $\Pi_{\rm r}$ smooth.

This singular–regular field split is standard in self-force and effective-source formulations [103, 37, 297, 107].

Representing the particle field

We choose to define the particle field $\Phi_{\rm p}$ (at a time t) as the field of a non-accelerating particle (at that same time t), as seen in the lab frame. We already defined these quantities in section 3.3.2 above:

$$\Phi_{\rm p}(t, x^i) := \Phi_{\rm lab}(z^i_{\rm p}(t), v^i_{\rm p}(t); x^i) \tag{3.25}$$
$$\Pi_{\rm p}(t, x^i) := \Pi_{\rm lab}(z^i_{\rm p}(t), v^i_{\rm p}(t); x^i) \tag{3.26}$$
$$\rho_{\rm p}(t, x^i) := \rho_{\rm lab}(z^i_{\rm p}(t), v^i_{\rm p}(t); x^i)\,. \tag{3.27}$$

We remind the reader that $z^i_{\rm p}(t)$ and $v^i_{\rm p}(t)$ are functions of time, while z^i and v^i are scalar values.

These definitions lead to the following time derivatives for $\Phi_{\rm p}$ and $\Pi_{\rm p}$:

$$\begin{aligned}
\partial_t \Phi_{\rm p}(t,x^i) &= \frac{d}{dt}\Phi_{\rm lab}(z^i_{\rm p}(t),v^i_{\rm p}(t);x^i) && (3.28)\\
&= (\partial_t z^i_{\rm p})\,\partial_{z^i}\Phi_{\rm lab}(z^i_{\rm p}(t),v^i_{\rm p}(t);x^i) + (\partial_t v^i_{\rm p})\,\partial_{v^i}\Phi_{\rm lab}(z^i_{\rm p}(t),v^i_{\rm p}(t);x^i) && (3.29)\\
&= \Pi_{\rm p}(t,x^i) + S_\Phi && (3.30)\\
\partial_t \Pi_{\rm p}(t,x^i) &= \frac{d}{dt}\Pi_{\rm lab}(z^i_{\rm p}(t),v^i_{\rm p}(t);x^i) && (3.31)\\
&= (\partial_t z^i_{\rm p})\,\partial_{z^i}\Pi_{\rm lab}(z^i_{\rm p}(t),v^i_{\rm p}(t);x^i) + (\partial_t v^i_{\rm p})\,\partial_{v^i}\Pi_{\rm lab}(z^i_{\rm p}(t),v^i_{\rm p}(t);x^i) && (3.32)\\
&= \Xi_{\rm p}(t,x^i) + S_\Pi && (3.33)\\
&= \delta^{ij}\partial_i\partial_j\Phi_{\rm p}(t,x^i) - \rho_{\rm p}(t,x^i) + S_\Pi(t,x^i) && (3.34)
\end{aligned}$$

with the effective source terms [37, 297, 107]

$$\begin{aligned}
S_\Phi &:= (\partial_t v^i_{\rm p})\,\partial_{v^i}\Phi_{\rm lab}(z^i_{\rm p}(t),v^i_{\rm p}(t);x^i) && (3.35)\\
S_\Pi &:= (\partial_t v^i_{\rm p})\,\partial_{v^i}\Pi_{\rm lab}(z^i_{\rm p}(t),v^i_{\rm p}(t);x^i)\,. && (3.36)
\end{aligned}$$

The source terms are proportional to the particle acceleration $\partial_t v_{\rm p}$. The source terms are fully defined by the functions $\Phi_{\rm lab}$ and $\Pi_{\rm lab}$ which we introduced above. The source terms are known functions of the particle position, velocity, and acceleration. They describe how the instantaneous field of the particle ($\Phi_{\rm p}$ and $\Pi_{\rm p}$) changes when the particle's velocity changes.

We note that the source terms depend only on the change of velocity $\partial_t v^i$ and not on the change of position $\partial_t x^i$. The reason is that the latter terms cancel because we intentionally chose the particle field to satisfy the Poisson equation in a moving reference frame. In principle, other choices for the particle field would be possible, and in that case, the position-dependent source terms might not cancel.

Equations of motion

The equations of motion for the remainder fields $\Phi_{\rm r}$ and $\Pi_{\rm r}$ are now automatically also defined. They follow from the equations of motion

for the physical fields Φ and Π. They are

$$\begin{aligned}
\partial_t \Phi_{\mathrm{r}} &= \partial_t \Phi - \partial_t \Phi_{\mathrm{p}} && (3.37)\\
&= \Pi - \Pi_{\mathrm{p}} - S_\Phi && (3.38)\\
&= \Pi_{\mathrm{r}} - S_\Phi && (3.39)\\
\partial_t \Pi_{\mathrm{r}} &= \partial_t \Pi - \partial_t \Pi_{\mathrm{p}} && (3.40)\\
&= \delta^{ij}\partial_i\partial_j \Phi - \rho - \delta^{ij}\partial_i\partial_j \Phi_{\mathrm{p}} + \rho_{\mathrm{p}} - S_\Pi && (3.41)\\
&= \delta^{ij}\partial_i\partial_j \Phi_{\mathrm{r}} - S_\Pi && (3.42)
\end{aligned}$$

These are the same equations as for a scalar field without a particle, but with added source terms.

While we can assume that Φ_{r} and Π_{r} are smooth, i.e. have no important features below the length scale h, this is not the case for the source terms. The source terms contain nontrivial features at small length scales down to ℓ. It is therefore necessary to filter or smoothen them in (3.37) to avoid contamination from small wavelength / high frequency noise.

3.4 Numerical methods

For simplicity, we study the scalar particle and its fields only in one dimension, not in three dimensions. With "one dimension" we refer to a truly 1+1-dimensional spacetime, not to a spherically symmetric reduction of a 3+1-dimensional spacetime.

We consider two distinct numerical implementations for the field equations: a finite volume approximation and a spectral approximation. These methods offer different perspectives on the spatial representation of the fields and the treatment of the particle source.

It is convenient to choose a particle size ℓ that is very small, i.e. $\ell \ll h$. This simplifies evaluating or approximating the expressions for the source terms, and we will do so below.

3.4.1 Finite Volume Method

We use the Finite Volume Method (FVM) [189, 138] for spatial discretization due to the nature of the source terms. The variational approach (weak solutions) of FVM is suitable for treating source terms that are integrable but not continuous, such as our case. (Although the source terms are continuous in our case, they have very large gradients, and thus "appear" discontinuous at the discrete level.) We follow the Method of Lines (MoL) approach [266], discretizing the spatial domain while keeping time continuous. This results in a system of ordinary differential equations (ODEs) which we evolve using a Strongly Stability Preserving Runge-Kutta (SSPRK) method [261].

Standard FVM methods have been designed for conservation laws where it is important to handle discontinuities well. Our case is different – we have source terms that have low regularity, i.e. which are "almost singular" – and we will therefore spend some time discussing the fundamentals of FVM, justifying our choice.

In this subsection we will describe the FVM method in 1D for brevity. The discussion would extend to higher dimensions in a straightforward manner.

First-order formulation

We rewrite the system (3.37) above in a fully first-order form by introducing a new variable $\Psi_i := \partial_i \Phi$:

$$\partial_t \Phi_{\rm r} = \Pi_{\rm r} - S_\Phi \tag{3.43}$$

$$\partial_t \Pi_{\rm r} = \delta^{ij} \partial_i \Psi_{{\rm r},j} - S_\Pi \tag{3.44}$$

$$\partial_t \Psi_{{\rm r},i} = \partial_i \Pi_{\rm r} - S_\Psi \tag{3.45}$$

where

$$S_\Psi := \partial_i S_\Phi \,. \tag{3.46}$$

Characteristics

We represent the field equations as a first-order hyperbolic system. We define the field state-vector U_f and the evolution system as:

$$\partial_t \bar{U}_f + A\partial_x \bar{U}_f = S \tag{3.47}$$

where

$$U_f = \begin{pmatrix} \Pi_{\mathrm{r}} \\ \Psi_{\mathrm{r}} \\ \Phi_{\mathrm{r}} \end{pmatrix}, \quad A = \begin{pmatrix} 0 & -1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad S = \begin{pmatrix} -a^i \partial_{v^i} \Pi_{\mathrm{lab}} \\ -a^i \partial_{v^i} \Psi_{\mathrm{lab}} \\ -a^i \partial_{v^i} \Phi_{\mathrm{lab}} \end{pmatrix}, \tag{3.48}$$

where U_f is the field part of the state-vector, A is the principal part and S is the source.

The principal part A admits an eigendecomposition

$$A = R\Lambda R^{-1} \tag{3.49}$$

where R contains the right eigenvectors r^m

$$R = [r^1|r^2|...|r^m] \tag{3.50}$$

and the matrix Λ contains its eigenvalues λ^m

$$\Lambda = \mathrm{diag}(\lambda^1, \lambda^2, ..., \lambda^m). \tag{3.51}$$

Computational grid

We discretize the spatial domain $[-L, L]$ into N cells of width h. We employ a staggered grid where the i-th cell is defined by the interval

$$\Omega_i = (x_{i-1/2}, x_{i+1/2}). \tag{3.52}$$

where $x_{i\pm1/2} = x_i \pm h/2$. The discrete variable $\bar{U}_i^n$ describes the average value

$$\bar{U}_i^n = \frac{1}{h} \int_{\Omega_i} U(t_n, x)\, \mathrm{d}x. \tag{3.53}$$

By working with cell averages we can ensure the numerical method is conservative even in the presence of discontinuities. Averaging is equivalent to computing the coefficients of an L^2 orthogonal projection onto the space of piecewise constant functions.

The spatial discretization of the field $U(\cdot, x)$ can be described as projection onto a subspace of piecewise constant functions. Let the domain be partitioned into disjoint cells Ω_i of width h. We define the computational basis as the set of characteristic functions $\{b_i(x)\}$ with

$$b_i(x) = \begin{cases} 1 & \text{if } x \in \Omega_i, \\ 0 & \text{otherwise.} \end{cases} \tag{3.54}$$

These functions form an orthogonal basis for the subspace $V_h \subset L^2(\mathbb{R})$. The approximation of the field, denoted $U_h \in V_h$, is constructed via the orthogonal projection operator $\mathcal{P}_h : L^2 \to V_h$. This is defined by the condition that the projection error is orthogonal to the subspace,

$$\langle U - \mathcal{P}_h U,\, b_i \rangle = 0 \quad \forall i, \tag{3.55}$$

where $\langle f, g \rangle = \int f(x) g(x)\, \mathrm{d}x$ denotes the standard inner product. Expanding the projection in terms of the basis, $\mathcal{P}_h U(\cdot, x) = \sum_j \bar{U}_j b_j(x)$, and substituting into the orthogonality condition yields

$$\sum_j \bar{U}_j \langle b_j, b_i \rangle = \langle U, b_i \rangle. \tag{3.56}$$

Due to the disjoint support of the basis functions, the expansion coefficients $\bar{U}_j$ are determined as

$$\bar{U}_j = \frac{1}{h} \int_{\Omega_j} U(x)\, \mathrm{d}x. \tag{3.57}$$

Thus, $\bar{U}_j$ denotes the coefficient of basis function b_j in the L^2 projection; for the piecewise-constant basis used here, it is the cell average over Ω_j. The standard cell-averaging operation is equivalent to the

coefficient extraction for the L^2 projection of the continuum field onto the piecewise constant basis functions.

The above is relevant because our source terms are not regular and need to be properly projected into our function space. For smooth source functions, it is often possible (and often done) to evaluate them at the centre of the cells as approximation of this integral. This approach does not work here because the respective approximation error would be unbounded. Instead, (3.57) needs to be used for the source term, as we describe in more detail below.

Weak formulation

Since (3.47) contains low regularity source terms, we interpret the evolution in a weak sense. Specifically, we require the residual of the evolution equation to be orthogonal to our basis functions $b_i(x)$:

$$\langle \partial_t U + \partial_x(AU) - S,\, b_i \rangle = 0. \tag{3.58}$$

Expanding the inner product using the definition of $b_i(x)$ yields:

$$\int_{\Omega_i} \partial_t U \,\mathrm{d}x + \int_{\Omega_i} \partial_x(AU)\,\mathrm{d}x = \int_{\Omega_i} S\,\mathrm{d}x. \tag{3.59}$$

Applying the fundamental theorem of calculus to the flux term and substituting the cell average $\bar{U}_i$, we obtain the semi-discrete scheme:

$$h\frac{d\bar{U}_i}{dt} + \big[AU(t, x_{i+1/2}) - AU(t, x_{i-1/2})\big] = \int_{\Omega_i} S\,\mathrm{d}x. \tag{3.60}$$

By defining the physical flux $F_{i\pm1/2} = AU(t, x_{i\pm1/2})$ and the cell-averaged source $S_i = \frac{1}{h}\int_{\Omega_i} S\,\mathrm{d}x$, we arrive at the standard FVM form:

$$\frac{d\bar{U}_i}{dt} = -\frac{1}{h}(F_{i+1/2} - F_{i-1/2}) + S_i. \tag{3.61}$$

This formulation demonstrates that the FVM is mathematically equivalent to a weighted residual method [127], where the test functions are

the same piecewise constant functions used for the L^2 projection. This choice is also known as a Bubnov-Galerkin method [81, 137]. For higher order schemes where the reconstruction is done with linear interpolation, instead of piecewise constant, the trial functions are now piecewise linear while the test functions remain piecewise constant. This is now called a Petrov-Galerkin method [222]. Note that the equation is in semi-discrete form, suitable for the MoL method.

Numerical fluxes

Since the flux calculations are being performed at the cell faces, we need to introduce a numerical flux to approximate the fluxes at the faces based on the values of U at the cell centers. We use the Godunov flux [138], which is the same as the Rusanov flux [260] for the linear wave equation, since the local Lax-Friedrichs speed [134] is given exactly by the eigenvalues of the system,

$$F^n_{i-1/2} = A^+ U^{\mathrm{R}}_{i-1} + A^- U^{\mathrm{L}}_i \tag{3.62}$$

where A^+ and A^- are defined as

$$A^+ = R\Lambda^+ R^{-1} \,\text{and}\, A^- = R\Lambda^- R^{-1} \tag{3.63}$$

where $\Lambda^{\pm}$ are the diagonal matrices containing only the positive or negative eigenvalues respectively. Note that the values U^{R}_{i-1} and U^{L}_i are evaluated on the common face at $x_{i-1/2}$ between the cells Ω_{i-1} and Ω_i. As the state vector quantities "live" in the cell centers, we have to reconstruct the values at the faces using conservative interpolation. First order schemes use piecewise constant interpolation, while for second order schemes, piecewise linear interpolation is needed.

Second-order approximation

To resolve discontinuous data at second order, a standard practice is to employ slope limiters and linear reconstruction in characteristic

space. These methods are also known as Monotonic Upstream-centered Scheme for Conservation Laws (MUSCL) [294, 295]. The homogeneous part is thus ensured to be Total Variation Diminishing (TVD) [153]. It is important to use a *monotone* flux [189], such as the Godunov or Rusanov, in order to be compatible with the MoL.

Source treatment

We employ an unsplit method for the time evolution of the field equations, wherein the source term is incorporated directly into the right-hand side of the homogeneous evolution equations. The source contribution is computed analytically by evaluating the cell-averaging integral over each computational cell. To accurately capture the behavior at the particle location, we split the integral of the cell that contains the particle at the particle location. This places the particle at an integration boundary, which simplifies evaluating the integral:

$$S_i = \frac{1}{h}\left(\int_{x_{i-1/2}}^{z_\mathrm{p}} S_\mathrm{left}\,\mathrm{d}x + \int_{z_\mathrm{p}}^{x_{i+1/2}} S_\mathrm{right}\,\mathrm{d}x\right). \tag{3.64}$$

Convergence analysis

While smooth solutions allow for pointwise convergence analysis in the L_∞ (max) norm, discontinuous data typically undergoes "smearing" over a few grid cells due to numerical diffusion. In such cases, the error at the jump does not decrease in the infinity norm regardless of the grid refinement. Therefore, we quantify convergence using the L_1 norm, which represents the integral of the absolute error (see 8.1.1 from [189]). To calculate the convergence order p without an analytical solution, we compare three successive resolutions $(h, h/2, h/4)$:

$$p = \log_2 \frac{\|\bar{U}_h - \mathcal{R}\bar{U}_{h/2}\|_1}{\|\mathcal{R}\bar{U}_{h/2} - \mathcal{R}\mathcal{R}\bar{U}_{h/4}\|_1}, \tag{3.65}$$

where $\mathcal{R}$ denotes a restriction operator that averages the higher-resolution cells onto the coarser grid. This averaging ensures the comparison is conservative and consistent with the finite volume projection.

The convergence behavior for non-smooth data is fundamentally limited. As shown by LeVeque [189] (pg 156), for a first-order upwind method applied to the scalar advection of a Heaviside function, the L_1 error behaves as $\mathcal{O}(h^{1/2})$. This is because the numerical width of the discontinuity—the region where the solution is smeared—typically grows as $\sqrt{th}$. Consequently, even if the scheme is formally first-order for smooth data, the presence of a jump discontinuity causes the L_1 error to decay only as $\sqrt{h}$.

By using the L_1 norm and a second-order MUSCL scheme [294, 295], we aim to recover a convergence rate closer to $\mathcal{O}(h^1)$ for discontinuous solutions, but we acknowledge that the $\mathcal{O}(h^{1/2})$ limit provides a theoretical baseline for the "worst-case" performance near the particle.

3.4.2 Spectral method

We also discretize the equations using a spectral method. We do this for two reasons: First, spectral methods are perceived as very "mathematical" and "difficult to get right". By showing that spectral methods are working we hope to increase acceptance of our approach in the community. And second, spectral methods are very simple to implement for a linear system, and it is straightforward to apply a filter to the source terms.

We choose a one-dimensional domain of size $[-L, +L]$ with $L = 1$. For simplicity we use Dirichlet boundary conditions where we set $\Phi_\mathrm{r} = 0$ and $\Pi_\mathrm{r} = 0$. We restrict our evolution to the time range $[0, 0.4L]$ and only analyze the results in the interval $[-0.5L, +0.5L]$ so that information from the boundary conditions does not influence our results.

We expand Φ_r and Π_r into Fourier modes. Given the boundary

conditions, all cos terms vanish, and our basis functions are

$$b^n(x) \quad := \quad \sin \pi n \frac{x+1}{2} \tag{3.66}$$

for $n \in 1 \dots N$. These basis function automatically impose the boundary conditions.

In this basis, the Laplace operator is diagonal and is given by

$$\Delta \quad := \quad \mathrm{diag}\left[-\left(n\frac{\pi}{2} \right)^2 \right] \tag{3.67}$$

When projecting the source terms onto our bases functions we use a Gauss-Legendre quadrature [2] with 3200 nodes. These are many more nodes than basis functions, and we assume that our numerical quadrature is essentially exact up to floating-point round-off error for our algorithm.

Our state vector consists of

$$\Phi_{\mathrm{r}}^n \quad : \quad \text{Coefficients for remainder field} \tag{3.68}$$

$$\Pi_{\mathrm{r}}^n \quad : \quad \text{Coefficients for time derivative of remainder field} \tag{3.69}$$

$$z_{\mathrm{p}} \quad : \quad \text{: Particle position} \tag{3.70}$$

$$v_{\mathrm{p}} \quad : \quad \text{Particle velocity} \tag{3.71}$$

To evolve the coupled particle/field system we use the method of lines (MoL) (see 3.4.1). We evaluate the right hand side (RHS) of the state vector in two steps:

1. We evaluate the fields Φ_{r} and Π_{r} and their gradients at the particle position z. We do this exactly, summing the spectral basis functions and their derivatives. We then evaluate the terms in the particle's equations of motion (3.18). We also add a forcing term for some test cases (see below).

2. We apply the Laplace operator to Φ_{r}^n. Since we already calculated the particle acceleration $a = d/dt\, v$ we can evaluate the source

terms at each Gauss-Legendre quadrature node. This allows us to explicitly project the source terms onto each basis function, filtering out all high-frequency noise.

It well known that specral methods converge "spectrally" (essentially exponentially) for smooth functions, i.e. that the error decays as $\exp(-\alpha N)$ for N spectral collocation points. The clause "for smooth functions" does the heavy lifting in this statement. In our case, the source terms are not smooth in this sense, and we thus only expect first order convergence, i.e. an error proportional to $1/N$. (A Fourier expansion converges for a large class of functions, including discontinuous functions, and the coefficients fall of as $1/N$.)

There exist well-known methods to improve the accuracy and convergence order in the presence of non-smooth term. In one dimension one can split the domain into two, with a moving domain boundary at the particle location; this would restore the spectral convergence. Similar approaches are possible in higher dimensions. We do not pursue these approaches here to keep our presentation simple.

3.5 Numerical results

For simplicity we consider a scalar field in $1+1$ dimensions. While this greatly simplifies the numerical implementation, it is also more difficult to understand intuitively because we are trained to think of a particle as having a $1/r$ potential. This is not the case in $1+1$ dimensions where the potential is proportional to $|x|$.

To allow for a nontrivial scenario we consider a single scalar particle which, in addition to the scalar field, also sees a static harmonic potential. This exerts a periodic external force onto the particle where the particle radiates and loses energy, leading to a shrinking oscillation amplitude.

3.5.1 Finite Volume Method results

For the numerical implementation, we employ a five-stage, fourth-order SSPRK [261] scheme for time integration. Spatial discretization is achieved via a second-order MUSCL reconstruction [294, 295] using the Monotonized Central (MC) slope limiter [294] in characteristic space.

We validate our framework by simulating a harmonically oscillating particle, evaluating the convergence of the numerical solution against both the analytically calculated retarded field and its own self-convergence. We forgo explicit boundary conditions for the scalar field; instead, we allow the field to evolve freely and truncate the computational domain to exclude any regions contaminated by numerical artifacts from the boundaries.

Forced motion

We simulate a particle with charge $q = 1$ that oscillates with maximum velocity $v_{max} = 0.1$ and amplitude $A = 0.01$. We evolve for a total time $t_f = 2.5$ during which the particle oscillates 4 times, thus we choose the domain $[-L, L]$ with $L = 6$. This domain size is chosen specifically so that, during the time frame of the four oscillations, any numerical artifacts or "junk" information from the boundaries cannot propagate into the region of interest. This allows us to safely truncate the domain as previously described, without loss of accuracy. Our baseline resolution is $dx_0 = 1$, which we refine while keeping a constant Courant-Friedrichs-Lewy factor of $\texttt{cfl} = 0.1$.

We find that our numerical fields show agreement with the analytic retarded solutions for both Π_r and Ψ_r, as demonstrated in Figure 3.1.

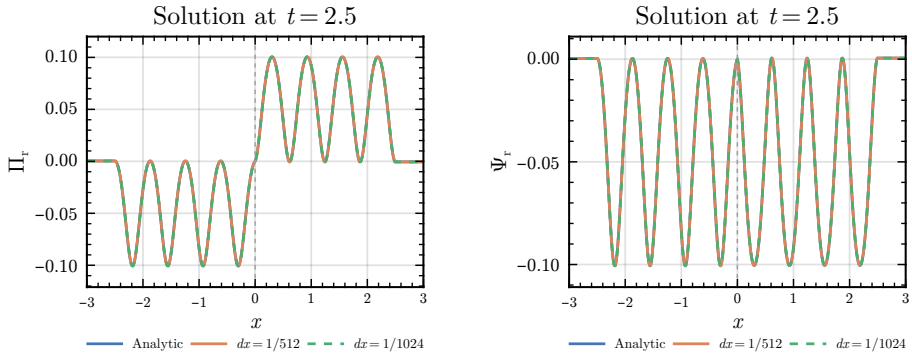


Figure 3.1: Field Π_r (left) and Ψ_r (right) at a late time $t = 2.5$ for two resolutions $dx = 1/512$ and $dx = 1/1024$, and comparing to the analytic retarded solution. The vertical dashed line indicates the position of the particle at that time.

We also show the evolution of the fields in spacetime Figure 3.2, and zoom onto the particle in Figure 3.3 to see the particle trajectory as well, for a resolution of $h = 1/512$.

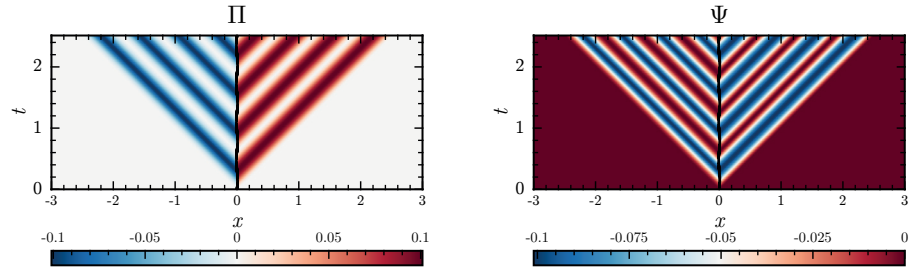


Figure 3.2: Space-time evolution of Π_r (left) and Ψ_r (right) for a particle in a harmonic potential with $q = 1$ and $v_{max} = 0.1$ for $N_{osc} = 4$ periods for $h = 1/512$. The field is initially zero everywhere, and the oscillating particle generates waves that propagate outwards. See Figure 3.3 for a zoom onto the particle's trajectory.

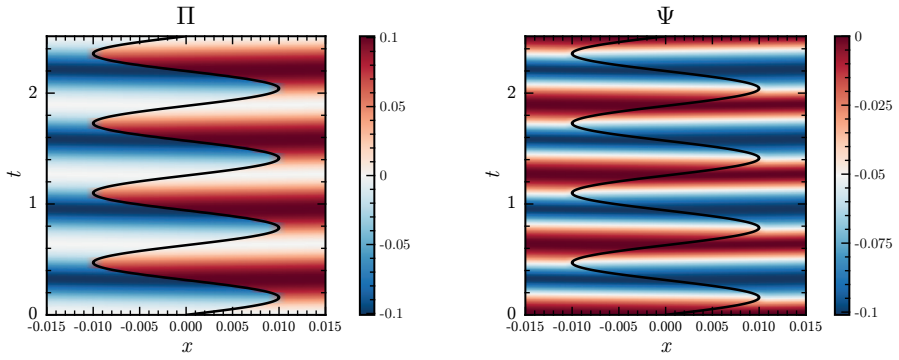


Figure 3.3: Zoom onto the particle: space-time evolution of Π_{r} (left) and Ψ_{r} (right) for a particle in a harmonic potential with $q = 1$ and $v_{max} = 0.1$ for $N_{osc} = 4$ periods for $h = 1/1024$. The waves emanating from the oscillating particle are clearly visible.

Since we can calculate the retarded field analytically, we can evaluate the convergence rate of our numerical scheme by comparing the computed fields against the analytic retarded fields throughout the entire evolution and investigate further at late times, such as $t = 2.5$.

In Figure 3.4 we show the convergence order for Π_{r} (left) and Ψ_{r} (right) fields against the analytic retarded fields. The convergence order approaches 2 for late times, while at earlier times the fields are only converging as 1st order. This is expected and will be explained in the end of this section.

We further investigate the convergence order of the numerical fields against the analytic retarded solution for late times as $t = 2.5$ in Figure 3.5, by plotting the scaled errors, scaled for 2nd order convergence. We observe the presence of spikes in the errors that do not align on top of each other after scaling for 2nd order. The location of these spikes are at the peaks of the waveforms (see Figure 3.1) and discontinuities. At a peak the slope changes sign, and as a result the slope limiter is choosing the slope calculation corresponding to an upwind calculation

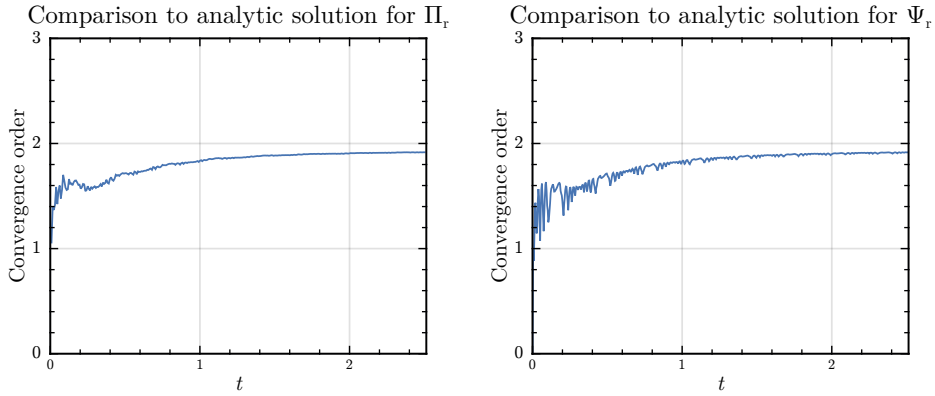


Figure 3.4: Convergence order for Π_r (left) and Ψ_r (right) fields against the analytic retarded fields, with resolutions $h = 1/512$ and $h = 1/1024$. As expected, our calculations are overall second order convergent. At early times the convergence order is reduced to about 1.5 for two reasons: (a) near the particle we only expect a convergence order of 1/2 (as explained by LeVeque, see main text), and (b) the error is zero almost everywhere else in the domain because the radiation from the particle has not propagated yet. At late times, the small region near the particle does not affect the overall convergence much, as can be seen in figure 3.5.

(1st order) instead of a 2nd order slope. This is a well-known artifact of slope-limited schemes when applied to non-smooth data. Away from these regions, the errors align on top of each other after scaling, showing second order convergence as expected for smooth data.

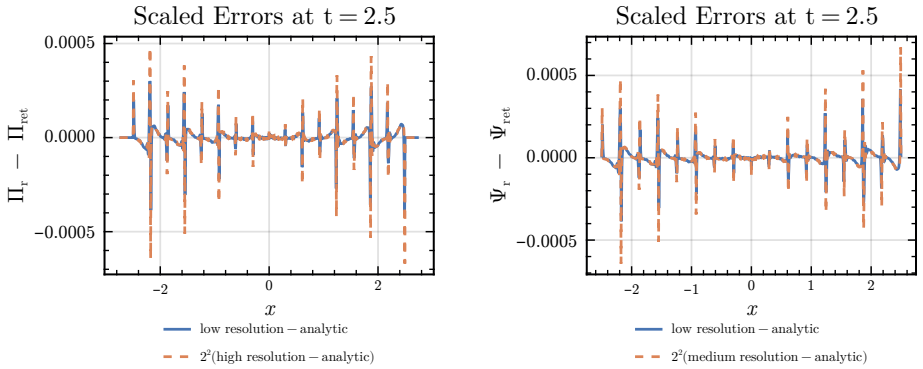
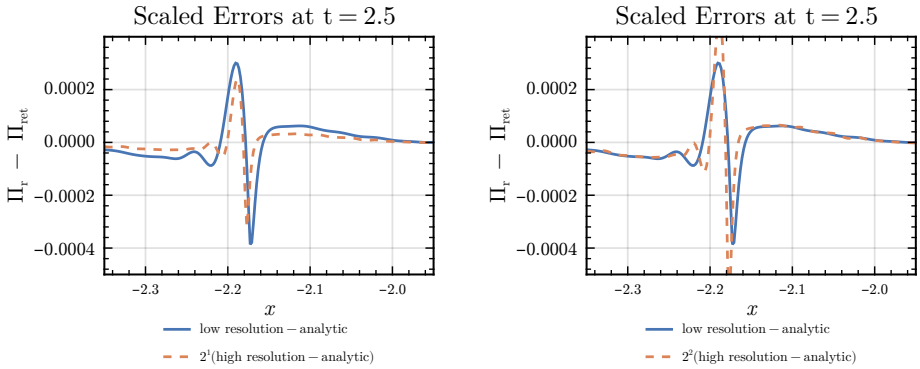


Figure 3.5: Scaled errors of the numerical fields Π_r (left) and Ψ_r (right) at $t = 2.5$ for resolutions $h_{low} = 1/512$ and $h_{high} = 1/1024$, also comparing to the analytic retarded solution. We see an overall second order convergence, overlaid by spikes for first order convergence (which do not affect the overall convergence order as shown in figure 3.4. The spikes are caused by the slope limiter near extrema of the solution, as explained in the main text, and are a general feature of slope limiters.

To illustrate this more explicitly we zoom in on one particular spike for Π_r at $r = 2.5$ and show the errors, scaled for both 1st and 2nd order convergence, in figure 3.6. We see the errors generally align on top of each other for the smooth region in the 2nd order (right) subfigure 3.6b, while the peak is not aligning well. The scaled high resolution error overshoots, indicating a lower convergence order there. The 1st order (left) subfigure 3.6a shows that the peak aligns reasonably well with a scaling for a 1st order error.



(a) Scaled errors for Π_r for 1st order convergence.

(b) Scaled errors for Π_r for 2nd order convergence.

Figure 3.6: Scaled errors for Π_r scaled for 1st (left) and 2nd (right) order near a spike location.

Having understood the effect of the slope limiters and how they affect the convergence order in non-smooth regions and near extrema, we now explain why the overall convergence order is reduced at early times, as promised when discussing figure 3.4. The solution starts with the initial conditions $\Pi_r = \Psi_r = 0$. At the first time step, the flux terms will be zero, and the source terms will be added to the RHS as is. These are discontinuous for a 1D particle field. In the next time step, this non-smoothness makes the slope limiters limit the slopes, reducing the convergence order. Figure 3.7 shows the solution for the remainder field Π_r at a very early time where the discontinuity at the particle location is clearly visible. At a later time $t = 0.3$, Π_r has a non-trivial shape for $|x| \leq t$, and the discontinuity at the particle does not dominate the numerical error any more. For more extensive discussion about high-resolution FVM with slope limiters, we refer the reader to [189].

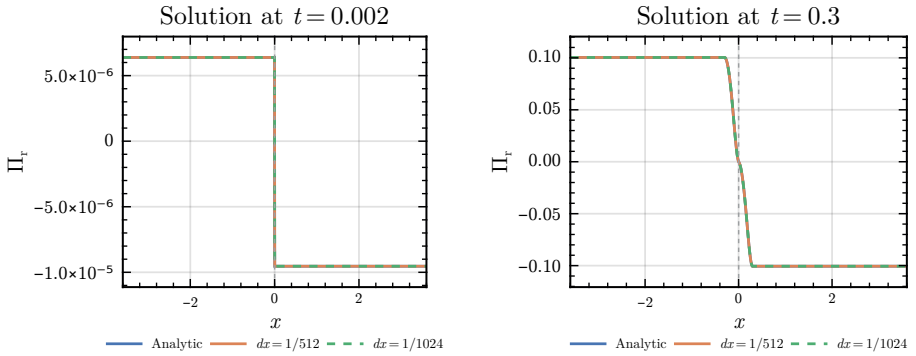


Figure 3.7: Resolutions for Π_r at $t = 0.002$ and a bit later at $t = 0.3$.

Overall, the finite volume discretization is working well, in line with the expected behaviour for discontinuous source terms and slope limiters. This implementation of our method converges overall to second-order as expected. The trajectory of the particle is given a priori, i.e. the particle does not lose energy or momentum while radiating. The focus of this subsection is to demonstrate that the non-regular source terms can be handled well by numerical methods.

Unfortunately, our attempts to let the particle be accelerated by the gradient of the remainder field failed, likely because our (ultimately piecewise constant) FVM discretization did not provide sufficiently accurate approximation of the field gradient near the particle. Further work is required to make self-consistent particle/field calculations work with finite volume discretizations.

We show a fully interacting particle/field system in the next subsection.

3.5.2 Spectral method results

We choose a domain $[-L, L]$ with $L = 1$. We expand the fields Φ_r and Π_r into Fourier modes and impose Dirichlet boundary conditions $\Phi_r = \Pi_r = 0$. We choose basis functions $b_n(x) = \sin(\pi n(x + 1)/2)$ which

strongly impose these boundary conditions. To avoid contamination of our results by boundary effects, we restrict our analysis to the time interval $[0, 0.4]$ and consider only the spatial region $[-0.5, 0.5]$ when evaluating the results. We choose a particle with $Q = 1$ and $M_0 = 1$.

We also use a non-relativistic approximation, i.e. we remove all relativistic terms from the equations of motion. We do so to be able to better compare to analytic results. The heart of our method is not affected by using Newtonian physics.

We choose an oscillation period P of $1/10$. We consider two different oscillation amplitudes A of $P/10$ and $P/100$. We use four different resolutions N with 200, 400, 800, and 1600 expansion coefficients.

We illustrate a particle's self-field and its effective source in figure 3.8. As mentioned earlier, the particle's field Φ grows linearly with distance. Φ corresponds to the potential U in electrodynamics. Π (not shown here), which corresponds to the electric field E in electrodynamics, is constant in space. The effective source changes over a length scale given by the particle size ℓ, i.e. it "looks" discontinuous.

These properties hold only in one spatial dimension. In three dimensions, the particle field falls off as $1/r$, and the effective source would fall off as $1/r^2$ near the particle. (Even in the limit $\ell \to 0$, the particle field and its effective source remain integrable in a neighbourhood of the particle.)

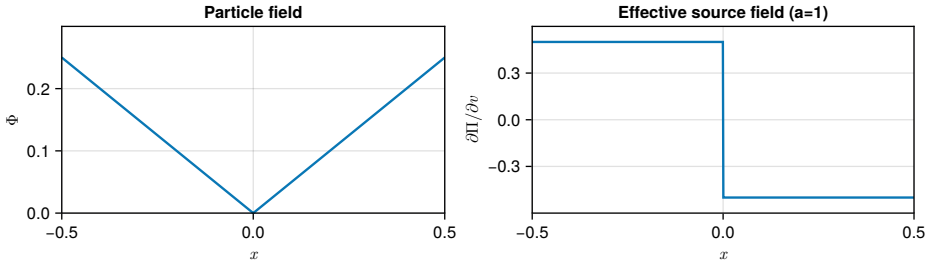


Figure 3.8: A particle's self-field and its effective source. In one space dimension, a particle's field is proportional to $|x|$ instead of $1/r$. The radiation field of a particle is given by its effective source $\partial \Pi_{\text{lab}}/\partial v$, multiplied by the particle's acceleration $\partial_t v_{\text{p}}$. (A non-accelerating particle does not radiate, and its effective source is zero.) (We show the effective source for $a = 1$.) This figure neglects relativistic effects such as the Lorentz contraction.

We show position and velocity of the oscillating particle in figure 3.9 and show a respective convergence test in figure 3.10. We compare two cases – a particle where the radiation reaction has been turned off (the particle sees only the harmonic potential, i.e. it moves on a pre-specified trajectory) and a particle that sees the radiation field it generates, and which thus loses energy. This energy loss is visible as a reduction in the oscillation amplitide and peak velocity. The convergence tests shows the rescaled error in the position of the radiating particle, demonstrating first order convergence (as expected).

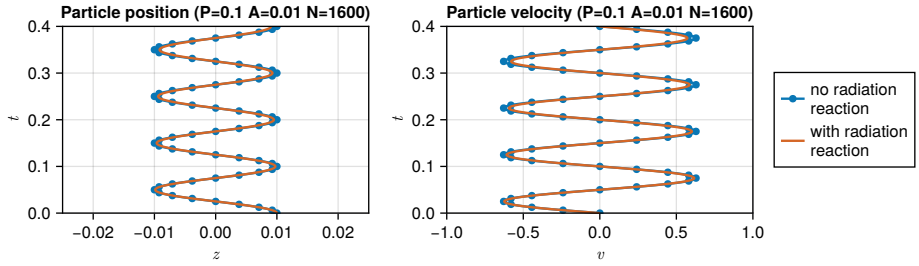


Figure 3.9: Particle position and velocity vs. time both without and with radiation reaction. Both graphs are very similar. The self-interacting particle's amplitude and maximum velocity decrease slightly over time because the particle loses energy as it radiates.

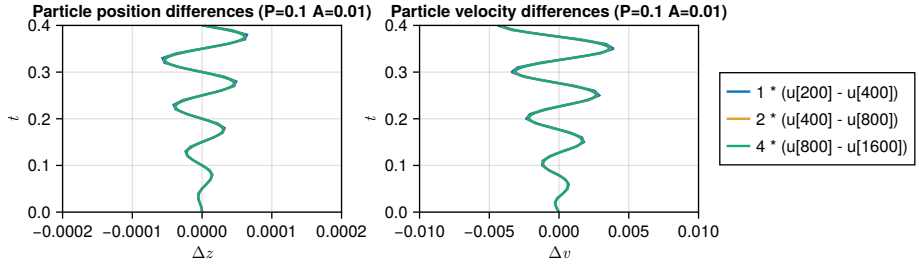


Figure 3.10: A convergence test showing the scaled particle position and velocity differences for four resolutions. As expected our discretization is first-order convergent. This low convergence order is expected for a spectral representation of the discontinuous effective source.

We show the radiation field Π in figure 3.11 and a respective convergence test in figure 3.12. We subtract the initial field, i.e. $\Pi(t = 0)$, to clarify the presentation. This shows the radiation field at a late time ($t = 0.4$); at this time, the particle is located at $z = 0.01$, where the radiation field crosses zero.

The waves to the right of the particle propagate to the right, the waves on the left propagate to the left. The radiation field has in general a discontinuity at the particle location (although there is no discontinuity at the show time $t = 0.4$).

The field has a clearly visible asymmetry – the maxima and minima of the field have a visibly different shape. This is caused by the Doppler effect; when the particle is moving to the right, the right-moving wave is squeezed, and when the particle is moving to the left, the right-moving wave is stretched. The maximum particle velocity is about 0.6, which is a significant fraction of the wave speed 1, although the simulation is non-relativistic.

The Liénard-Wiechert approximation used here evaluates the field for a prescribed particle trajectory and therefore neglects the self-force. It assumes that the particle is on a force trajectory, i.e. it ignores the self-force. At early times (i.e. near $|z| = 0.4$), it agrees well with our numerical calculation. At late times (i.e. closer to the origin), our numerical calculations produce less radiation because the particle has lost energy and its oscillation amplitude has shrunk. Outside of $|z| = 0.4$ the radiation field is zero, because the radiation has not yet reached this region.

We see first-order convergence of the radiation field as well.

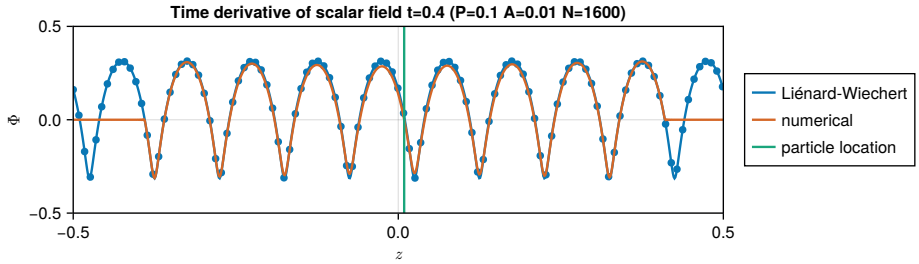


Figure 3.11: Generated radiation field at the final time, compared to the Liénard-Wiechert approximation which ignores the radiation reaction. The numerical calculation initially agrees well with the Liénard-Wiechert approximation (near $x = \pm 0.4$), but later produces less radiation because the oscillating particle has lost energy. The radiation field is lopsided because the particle oscillates fast ($v > 0.5$), leading to a significant Doppler effect.

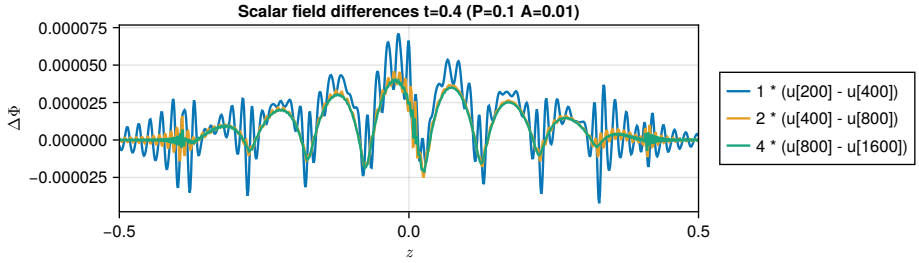


Figure 3.12: A convergence test showing the scaled radiation field differences at the final time. As expected, our discretization is first-order convergent. This low convergence order is expected for a spectral representation of the discontinuous effective source.

We examine the particle's radiative energy loss in figures 3.13 and 3.14 as function of time.

In figure 3.13, we estimate the particle's radiated power via a Liénard-Wiechert approximation. We show the left-propagating, the right-propagating, as well as the total emitted power. This shows that the particle radiates in bursts; it radiates most when it accelerates most (when it is near the extremum of its oscillation), and does not radiate when it is at the origin. The left/right asymmetry, caused by the Doppler effect, is also clearly visible.

In figure 3.14 we show the potential energy of the particle for two cases, a non-self-interacting and a self-interacting particle. The self-interacting particle loses energy to radiation; it has lost about 20% of its energy at the end of the simulation.

We also show the total energy of the particle for both cases. The total energy is the sum of the potential and kinetic energy. The non-self-interacting particle's energy is constant in time, the self-interacting particle's energy decreases monotonically over time. The energy loss is not linear because the particle only radiates when it accelerates in its oscillatory motion. We estimate the energy loss due to radiation via a Liénard-Wiechert approximation ("approximately corrected for radiation"). This approximation is not exact because it over-estimates the radiation at late times, it does not account for the reduction in oscillation amplitude over time.

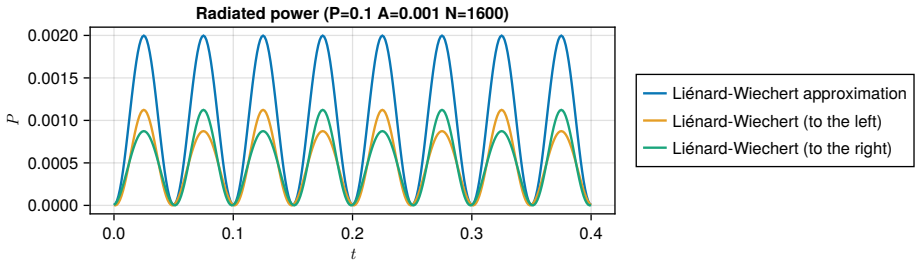


Figure 3.13: Particle estimated radiation power vs. time, calculated via a Liénard-Wiechert potential.

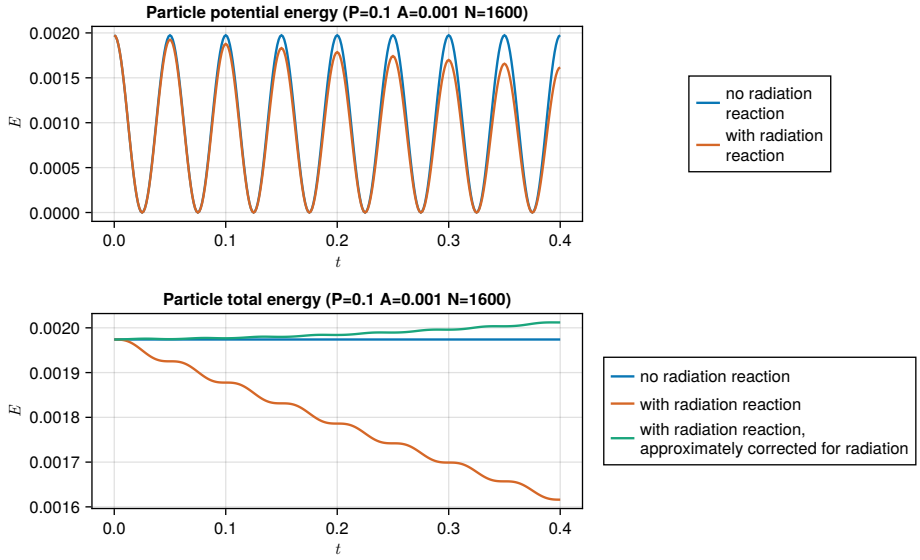


Figure 3.14: Particle energy vs. time. The particle energy decreases over time. It does so in bursts because the particle is oscillating forth and back and not in a circular motion. The particle's acceleration is largest when it reverses direction and is zero in between, and the radiated power is proportional to the particle's acceleration. This is well modeled by the (averaged) expected radiation power calculated via a Liénard-Wiechert potential. At late times the Liénard-Wiechert approximation overestimates the radiated power because the particle has lost energy.

3.6 Summary

We have presented a new approach for modeling small (unresolved) particles in numerical calculations. Our approach treats particles as small but finite-sized, and models the particle field effectively as an additional basis function on top of a regular remainder field. This

sidesteps any issues arising from handling infinities, and subsumes accuracy and convergence into well-known mechanisms for numerical methods.

We showed numerical calculations using two different discretization methods (finite volumes and spectral methods) with excellent agreement to analytic solutions. The fully self-interacting case agrees well with analytic approximations. We emphasize that it is necessary to ensure that the effective sources created by the particle acceleration are properly projected into the discrete function space. While this can be done analytically, we show that this can also be done numerically, which may be preferable because the source terms are complicated.

While our example calculations were performed in one dimension, our approach generalizes to any number of dimensions. In three dimensions, the effective source terms fall off as $1/r^2$ but still contain directional discontinuities. To improve convergence one could move to more advanced numerical methods that are better suited to handle such sources. One possible approach would be a finite element method where the element containing the particle is split into sub-elements so that the particle can be placed into element corners. Different from our one-dimensional calculations, where the solution of the Poisson equation grows as $|x|$, outer boundaries will not be an issue in three dimensions.

To limit the region affected by the effective sources, one can employ a window function in the definition of the particle field Φ_{p}. This leads to additional terms in the effective sources, but can reduce the overall computational cost. Indeed the sole purpose of such a window function is to reduce computational cost; a particle field Φ_{p} with infinite support is fully self-consistent.

Scalar particles and electrodynamic particles are described by their charge, mass, position, and velocity. Particles describing a compact object in relativity have an additional degree of freedom, namely their angular momentum. This could be implemented as a boosted Kerr black hole in some coordinates, or by an ansatz that only models

the leading order terms of the singularity. In our approach, only the combined particle and remainder fields have a physical meaning, and the particle field by itself does not need to satisfy any constraints or gauge conditions.

Our approach to limit the physical system to a finite resolution $h \gtrsim 0$ makes sense for the physical scenarios we described in the introduction, e.g. a small compact object with mass m orbiting a large black hole of mass M. The practically interesting range for the resolution h is $M \gg h \gg m$, with the numerical resolution h sandwiched between the sizes of the large and small objects. Our method is designed for this, and describes the small object as particle of size $\ell > 0$. In the limit $h \to 0$, the resolution is at some point finer than the small object, and we then have $M \gg m \gg h$. The small object can be fully resolved, and case we do not need to treat it as a particle any more. Thus we never need to take the limit $\ell \to 0$. This idea is at the core of our approach.


4

Chapter 4

Cusp Formation in Merging Black Hole Horizons

Publication status

This chapter is based on the following paper:

> S. Kastha, **S. Vretinaris**, D. Pook-Kolb, B. Krishnan, "Cusp Formation in Merging Black Hole Horizons", arXiv: 2605.10874, Submitted to Physical Review D.

Contributions

I contributed to the implementation, numerical modeling of horizons, production of figures, interpretation of results, and manuscript writing.

Abstract

An important question in binary black hole mergers is to connect properties of the remnant black hole to those of the two initial black holes. These properties include not only the final mass and spin of the remnant, but also higher multipoles and answers to other questions such as, for a given initial configuration, which quasi-normal modes of the final black hole are excited, and what are the amplitudes of these modes? Such questions have thus far been primarily addressed through a study of the emitted gravitational wave signal. In this chapter we consider a different alternative, namely using quasi-local black hole horizons themselves to establish the link between the initial and final states. Recent work has elucidated the behavior of black hole horizons in a merger. Cusps forming in such otherwise smoothly evolving horizons have been shown to play a central role in connecting the two initially separate black holes with the final remnant. In the present work, we will discuss from a numerical perspective how such cusps form in detail for the head-on collision of two non-spinning black holes. We show how the mass and higher mass multipole moments behave at the cusp and suggest a phenomenological model.

4.1 Introduction

For almost its entire lifetime, any black hole in our universe is described with great accuracy by a particular solution to the Einstein equations, namely the Kerr solution. Thus, given that the Kerr solutions are a 2-parameter family of solutions, black holes and the spacetime geometry in their vicinity are determined by just two parameters. These are, of course, the mass M and spin angular momentum $\mathbf{S}$ (often expressed in terms of the dimensionless spin parameter $\chi := |\mathbf{S}|/M^2$ which satisfies $\chi < 1$). All higher multipole moments are determined by these two parameters. Even for black holes with external perturbations such such as accretion disks or a distant binary companion, the black holes are still well approximated as Kerr black holes. If necessary, the higher order effects of such external influences can be accurately modeled by linear perturbations of the Kerr solution. All of these

features make black holes rather distinct from other stars where the higher multipole moments are not determined uniquely by the mass and angular momentum.

There is however one clear case where the Kerr solution does not suffice, namely when two black holes get sufficiently close to each other and coalesce to yield a single remnant black hole.[1] For the merger of two roughly comparable mass black holes, non-perturbative effects must necessarily be taken into account. Even though the merger lasts for a very short duration compared to the lifetime of a typical black hole, the black hole mass and spin generally change appreciably during the merger. One may therefore expect the distribution of black hole masses and spins in our universe to be significantly impacted by binary mergers. It is currently not possible to obtain a full analytic solution of the Einstein equations for a binary black hole merger, but the problem can be addressed numerically. Over the past two decades, numerical relativity has matured and it is now possible to obtain accurate numerical solutions to the binary black hole problem (see, e.g., [276, 85]).

One can view a binary black hole merger as a process whereby two Kerr black holes with masses (M_1, M_2), spin angular momenta $(\mathbf{S}_1, \mathbf{S}_2)$, and linear momenta $(\mathbf{P}_1, \mathbf{P}_2)$ merge to yield a remnant Kerr black hole with mass M_f, spin $\mathbf{S}_f$ and momentum $\mathbf{P}_f$. If the two initial black holes are gravitationally bound in a binary system, then $\mathbf{P}_1$ and $\mathbf{P}_2$ are not independent and would be determined by properties of the binary system such as the initial separation, eccentricity etc. We can encapsulate this process as a mapping from an initial to a final configuration:

$$\left\{\begin{matrix}(M_1, \mathbf{S}_1, \mathbf{P}_1)\\(M_2, \mathbf{S}_2, \mathbf{P}_2)\end{matrix}\right\} \to (M_f, \mathbf{S}_f, \mathbf{P}_f)\,. \tag{4.1}$$

Determining the remnant black hole parameters following such a mapping is useful for several reasons. In the case of supermassive black holes, the final parameters would be linked to several astrophysical observables such as the location of the central black hole within the host galaxy, the orientation of jets in the case of active galactic nuclei, the velocity dispersion of stars

[1] Even for black hole mergers, black hole perturbation theory retains a lot of its utility when the binary system is highly asymmetric [123].

within the central galactic bulge etc; see, e.g., [202]. For stellar mass black holes such as those observed by the LIGO, Virgo and KAGRA observatories, the distribution of black hole parameters and our inferences regarding the various formation channels of binary black hole systems would depend on the above mapping; see, e.g., [256]. Finally, the spectrum of quasi-normal modes emitted by the remnant black holes, and the associated tests of general relativity [115, 58] is also determined by the mass and spin of the remnant black hole. Moreover, a detailed understanding of the *approach* to the final state will enable us to determine precisely which quasi-normal modes will be excited for a given initial configuration.

Several approaches have been previously employed in the literature to determine the above map from the initial to final states, all relying on numerical relativity results to varying degrees. The first is a phenomenological approach based on a large number of numerical relativity simulations of binary black hole mergers with a wide variety of initial conditions [317, 156, 195, 159, 290, 42, 252, 253, 296]. A systematic approach which exploits the underlying symmetries in the problem and employs a power series expansion in the spins is given in [72, 73, 82]. A third approach relies on employing suitably accurate models for the emitted gravitational wave signal. Given a gravitational waveform, one can calculate the radiated fluxes of energy and angular momentum which, given the total energy and angular momentum at any given time, can be used to compute the remnant mass and spin. This approach has been used thus far within the Effective-One-Body formalism [94, 95].

In this chapter we initiate an entirely different approach to this problem, namely by tracking the black hole horizon dynamics and multipole moments all the way through the merger. As is conventional and practically useful in numerical relativity, we shall use the notion of quasi-local horizons (QLHs), which are 3-surfaces obtained by the time evolution of "apparent horizons" or more correctly marginally outer trapped surfaces (MOTSs) [24, 180, 23, 25, 67]. These notions will be defined more precisely below; for now, we note that QLHs differ significantly from event horizons only in the dynamical merger regime where they lie behind the event horizon. We apply this framework specifically to the head-on collision of two non-spinning black holes (though we expect the method to be generalizable for general configurations). This will

give us the masses, and higher multipole moments of the two individual black holes as functions of time up to the merger, followed by a discontinuous jump at the merger, and then finally the subsequent evolution of the parameters of the remnant black hole as it settles down to its final state. The dominant effect, as we shall see, is the discontinuous jump at the merger. It turns out (as verified in numerical simulations), that masses and spin *magnitudes* do not change appreciably during the inspiral regime. If the initial spins happen to be mis-aligned with the orbital angular momentum, then the system can exhibit precession which modifies the direction of the two spins, but the spin magnitudes themselves do not vary appreciably (see, e.g., [17]). Similarly, the late ringdown is well modeled as a perturbed Kerr black hole, whence the mass and spin of the remnant do not change appreciably in this regime as well. It is in fact in the comparatively short merger regime where the largest variations in mass and spin occur. Consistent with this observation, it is here that the black hole horizons absorb the most infalling radiation leading to significant area increase and the consequent variations in the physical parameters.

Horizon measures for black hole mass and angular momentum are used routinely in numerical simulations [116]. However, this has so far been typically used either for the two initial black holes (often in the initial data) or for the remnant black hole at late times.[2] [3] The challenge here is to do this at the merger.

As we shall describe in detail below, this involves highly distorted horizons and the formation of cusps which are challenging to locate numerically. Over the past decade, accurate numerical methods have been developed and applied in binary black hole simulations for addressing this challenge [228, 231, 230]. The case of head-on collisions has been studied in great detail and the process by which two horizons merge and eventually form a single horizon is now well understood [229, 226, 227, 69, 233, 234]. Though we shall

[2]These horizon measures are in fact used in most of the works referenced above for constructing the map of Eq. (4.1).

[3]While there are some subtleties in comparing the remnant parameters obtained from the horizon measures versus those obtained from the waveform and other asymptotic data [167], the mappings obtained from the two methods should, in principle, be equivalent.

not pursue this question in this chapter, among the physical applications of this work will be a quantitative understanding of how the remnant black hole approaches equilibrium and its relation to quasi-normal modes [149, 212, 176]. Can we predict, based on the initial configuration, specifically which quasi-normal modes are excited, and what are the amplitudes of these modes? As indicated earlier, such a prediction would be useful for the program of black hole spectroscopy and for other tests of general relativity [58]. There are several works which address this question through a detailed study of the emitted gravitational wave signal. The present work suggests an alternative based upon the horizons themselves. How is it possible to relate the gravitational wave signal to the horizons which are, after all, inside the event horizon and thus causally disconnected from the wave zone where gravitational wave observations are made? The reasoning rests on the following three considerations:

a) First, there is now growing evidence that the gravitational wave signal observed in the wave zone is closely connected to the *infalling* radiation at the horizon. This was proposed in [254, 170, 171]. The heuristic explanation is that both the infalling flux at the horizon and outgoing waves that we observe have the same source, namely by the complicated non-linear dynamical spacetime outside the horizons. Thus, while the horizon dynamics is itself not the source of observed signals, the two are strongly correlated. Since this initial suggestion, there are several numerical studies providing detailed evidence for the existence of such correlations [149, 176, 212, 176, 88].

b) Second, it is understood how a given quasi-local horizon (QLH), i.e. a 3-surface foliated by Marginally Outer Trapped Surfaces (MOTSs), evolves in time. This includes fluxes across the QLH which leads to changes in mass, angular momentum and higher multipoles. This includes "physical process" versions of the laws of black hole mechanics [23, 24, 140, 22], and the development of geometric notions of time evolution on a QLH [20, 88].[4]

[4]There remain still some gaps in this formulation. Examples are the technical details of how this extends when we transition between a spacelike QLH to one with a mixed signature.

c) Point (b) above only applies to the smooth part of the evolution; it does *not* apply to the singular behavior at the merger. There is so far no study, either analytical or numerical, in this direction. The goal of this chapter is to address this missing link in the case of a head-on merger of two non-spinning black holes.

It should now be clear that if all three points enumerated above are achieved, then we can establish a link between the progenitor and remnant black holes. We would thus have an alternative, and complementary formalism to the gravitational wave signal itself. In this chapter, we will consider the same simulation as in Refs. [229, 231], i.e. head-on collisions of two non-spinning black holes. We expect similar considerations should hold also for more generic initial data with say spinning black holes.

The rest of this chapter is structured as follows. Basic notions and definitions, and our numerical set-up are introduced in Sec. 4.2. Sec. 4.3 calculates geometrical quantities numerically on various horizons and elucidates its properties. Sec. 4.5 studies the jump in the mass and higher mass multipoles across the merger. Finally Sec. 4.6 presents a summary and concluding remarks.

4.2 Basic Notions: Quasi-local horizons, their mergers, and multipole moments

4.2.1 Quasi-local horizons

Let $(\mathcal{M}, g_{ab})$ be a 4-dimensional spacetime foliated by spacelike Cauchy surfaces (Σ, h_{ij}, K_{ij}) with Riemannian 3-metric h_{ij} and extrinsic curvature K_{ij}. A smooth closed spacelike 2-surface $\mathcal{S} \subset \Sigma$ with a Riemannian 2-metric q_{ab} is equipped with two future pointing null normal directions. We assume it is possible to assign an outward direction and let ℓ^a and n^a be future pointing outgoing and ingoing null normals, respectively. These can be rescaled with positive definite functions and we cross normalize these via $\ell \cdot n = -1$. Congruences of null geodesics starting in the ℓ^a or n^a directions

then have expansions

$$\Theta_{(\ell)} = q^{ab}\nabla_a \ell_b \,, \qquad \Theta_{(n)} = q^{ab}\nabla_a n_b \,, \tag{4.2}$$

respectively. $\mathcal{S}$ is called a *trapped surface* if both expansions are negative, a *marginally trapped surface* if $\Theta_{(\ell)} = 0$ and $\Theta_{(n)} < 0$, and a *marginally outer trapped surface* (MOTS) if $\Theta_{(\ell)} = 0$ with no condition on $\Theta_{(n)}$. Note that there is still a freedom to scale the null normals by some positive function $f > 0$, i.e. $\ell^a \to f\ell^a$, $n^a \to f^{-1}n^a$, without affecting the cross normalization. This, however, leaves the signs of the two expansions, and in particular the condition $\Theta_{(\ell)} = 0$, invariant. Thus the above definitions do not depend on the scaling of the null-normals.

A hypersurface $\mathcal{H}$ will here be called a *Quasi-Local Horizon* (QLH) if it admits a foliation of MOTSs $\mathcal{S}$. Starting from a MOTS on a Cauchy surface, QLHs arise by a time evolution. The question of whether the time evolution leads to a smooth QLH is studied in [13, 14, 68]. It is shown that the time evolution leads to a smooth $\mathcal{H}$ if the MOTS satisfies a stability condition which turns out to be equivalent to the spectral properties of an elliptic operator on $\mathcal{S}$. The properties of $\mathcal{H}$ have also been studied elsewhere; see [67, 24, 25] for reviews. The focus of this chapter will not be $\mathcal{H}$ but rather the MOTSs themselves.

The work presented here is made possible by the high accuracy MOTS finder which allows us to locate highly distorted MOTSs [230]. This is an adaptation of a previous method developed by Thornburg [288, 289] which is restricted to "star-shaped" surfaces. If $\mathcal{S}$ is a star shaped surface, then there exists a point p such that every ray from p intersects $\mathcal{S}$ exactly once. Thus, every point on $\mathcal{S}$ can be uniquely parametrized in terms of distance from p. The method of [230] removes this restriction and instead parametrizes $\mathcal{S}$ using distance from a *reference surface* which can itself be arbitrarily distorted.

With this MOTS finder at hand, it has been possible to understand how, beginning with two distinct and widely separated MOTSs, we end up eventually with a single MOTS corresponding to the remnant black hole. When the two MOTSs get sufficiently close to each other, a common MOTS appears which encloses the two progenitors BHs. This common MOTS, initially highly distorted, moves outwards and settles down to the final

remnant black hole, *absorbing* gravitational radiation in the process. The recent progress made in Refs. [229, 231] shows that a connected sequence of MOTSs does in fact exist which connects the two initial MOTSs to the final MOTS. Let us summarize the general behavior (see Fig. 4.1): In a head-on merger of two non-spinning black holes, the two individual horizons $\mathcal{S}_1$ and $\mathcal{S}_2$ were found to approach each other until they touch (at a time denoted t_{touch}) and subsequently start to intersect. A little before t_{touch}, a common MOTS surrounding $\mathcal{S}_1$ and $\mathcal{S}_2$ appears which immediately bifurcates into two, an inner and outer branch, denoted $\mathcal{S}_{\text{inner}}$ and $\mathcal{S}_{\text{outer}}$, respectively. We call the time of bifurcation $t_{\text{bifurcate}}$. The outer branch $\mathcal{S}_{\text{outer}}$ expands outwards and becomes more and more symmetric as it approaches the event horizon of the final Schwarzschild black hole. The inner branch $\mathcal{S}_{\text{inner}}$ on the other hand, becomes increasingly distorted and at t_{touch}, the union $\mathcal{S}_1 \cup \mathcal{S}_2$ coincides with $\mathcal{S}_{\text{inner}}$, with a cusp at the common point. It was found that $\mathcal{S}_{\text{inner}}$ continues to evolve from this point to the future, where it has a self-intersection. From t_{touch} to the past, $\mathcal{S}_{\text{inner}}$ evolves smoothly, first to the past and then turning to the future at $t_{\text{bifurcate}}$, to eventually asymptote to the final Schwarzschild horizon. As $\mathcal{S}_{\text{inner}}$ turns to the future, we call it $\mathcal{S}_{\text{outer}}$ and $(\mathcal{S}_{\text{outer}}, \mathcal{S}_{\text{inner}})$ can be seen as the two branches of the common horizon. With the exception of the MOTS at t_{touch}, all other MOTSs shown here are seen to be smooth.[5] Several mathematical results are known in the scenario described above; we mention one which will be useful for us later. At t_{touch}, it turns out that $\mathcal{S}_1$ and $\mathcal{S}_2$ have the same mean curvature at the point of contact [211]. Thus, at the point of contact, the ingoing expansions of $\mathcal{S}_1$ and $\mathcal{S}_2$ are identical. Equivalently, considering $\mathcal{S}_1$ and $\mathcal{S}_2$ to be embedded in a Cauchy surface Σ, the mean curvatures (i.e. the trace of the second fundamental forms) will coincide at the point of contact. It is also shown in [12] that when $\mathcal{S}_1$ and $\mathcal{S}_2$ get sufficiently close, then they must be enclosed by a common MOTS.

Our goal here is to elucidate details of the cusp formation on $\mathcal{S}_{\text{inner}}$, and to also understand the geometry of $\mathcal{S}_1$ and $\mathcal{S}_2$ in this process. This whole picture hinges on the fact that $\mathcal{S}_{\text{inner}}$ tends to $\mathcal{S}_1 \cup \mathcal{S}_2$ as $t \to t_{\text{touch}}$ from both sides, and hence necessarily develops a *cusp* during its evolution. However, this has not yet been fully understood. First, consider $\mathcal{S}_{\text{inner}}$ at t_{touch} where

[5] Self-intersecting MOTSs are smooth in the sense that they are smoothly immersed 2-spheres.

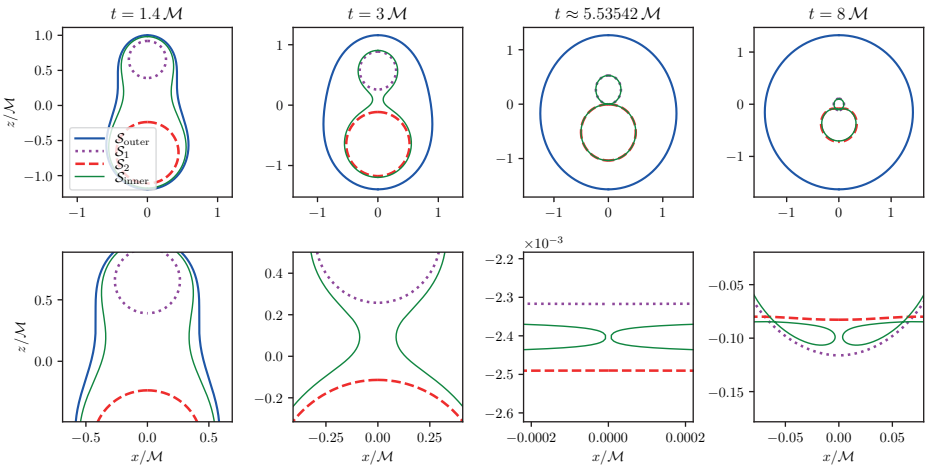


Figure 4.1: MOTS structure during the head-on collision. The second row shows close-ups of the configuration in the first row. The third column is very close before $t_{\text{touch}} \approx 5.5378\,\mathcal{M}$. The blue curve is $\mathcal{S}_{\text{outer}}$, $\mathcal{S}_{\text{inner}}$ is in green, $\mathcal{S}_1$ is the dashed magenta curve, and finally $\mathcal{S}_2$ is the dashed-red curve.

it consists of $\mathcal{S}_1 \cup \mathcal{S}_2$ and has a cusp. Given that the radius of curvature of $\mathcal{S}_{\text{inner}}$ is necessarily zero at the cusp, how does it *approach* zero and still be equal to the radius of curvature of $\mathcal{S}_1$ and $\mathcal{S}_2$, which clearly do *not* go to zero? In other words, will the radius of curvature jump discontinuously at the cusp? Furthermore, on either $\mathcal{S}_1$ or $\mathcal{S}_2$, how do their radii of curvature (and their angular derivatives) behave near the point of touching? How are these radii of curvature related to each other while still maintaining the Gauss-Bonnet theorem? In this chapter we will give a detailed answer to these questions.

4.2.2 Axisymmetric MOTSs and multipole moments

The main analytical tool we use here are suitable geometric multipole moments of a distorted horizon. Since the simulation considered in the present chapter is manifestly axisymmetric, we are able to simplify the representation of axisymmetric surfaces. Without loss of generality, let the symmetry axis be the z axis of our numerical coordinates. Any such surface will then be considered the surface of revolution around the z axis of a *curve* γ in the (x, z) half plane.

On each axisymmetric MOTS $\mathcal{S}$, we will construct uniquely defined coordinates (ζ, φ) using the procedure presented in Ref. [21] which is briefly summarized here; this construction applies in fact to any axisymmetric 2-surface. Let ϕ^a be the axial Killing vector on $\mathcal{S}$ and we assume that it vanishes at exactly 2 points on $\mathcal{S}$, these are the poles on $\mathcal{S}$, and also that it has closed integral curves. Let $A_{\mathcal{S}}$ be the area of $\mathcal{S}$ and $R_{\mathcal{S}} = \sqrt{A_{\mathcal{S}}/4\pi}$ the area radius. The coordinate φ is the affine parameter along the integral curves of ϕ^a, normalized so that $0 \leq \varphi < 2\pi$. The other coordinate ζ is defined via

$$\mathcal{D}_a \zeta = \frac{4\pi}{A_{\mathcal{S}}} \widetilde{\epsilon}_{ba} \varphi^b \,, \quad \oint_{\mathcal{S}} \zeta \, dA = 0 \,. \tag{4.3}$$

Here D_a is the covariant derivative on $\mathcal{S}$ compatible with q_{ab}, and ϵ_{ab} is the geometric volume 2-form on $\mathcal{S}$. The 2-metric on $\mathcal{S}$ has the simple form

$$ds_q^2 = R_{\mathcal{S}}^2 \left(\frac{d\zeta^2}{F(\zeta)} + F(\zeta) d\phi^2 \right) , \tag{4.4}$$

where

$$F(\zeta) = \frac{4\pi \varphi_a \varphi^a}{A_{\mathcal{S}}} \,. \tag{4.5}$$

The invariant coordinate $\zeta \in [-1, 1]$ thus becomes a parameter along γ. For the standard "round" 2-sphere in Euclidean space with spherical coordinates (θ, ϕ), it is straightforward to verify that $F(\zeta) = 1 - \zeta^2$ and $\zeta = \cos\theta$. We shall refer to F on the "round" 2-sphere as $F_o = 1 - \zeta^2$. We will often define $\theta = \cos^{-1}\zeta$ in this way even on distorted spheres and show results as functions of θ; we shall usually suppress the dependence on the orthogonal

angular coordinate φ. Note that we could in principle choose a different range of values for ζ such that $\zeta_{min} \leq \zeta \leq \zeta_{max}$. However, the total area will still be $2\pi R_{\mathcal{S}}^2(\zeta_{max} - \zeta_{min}) = 4\pi R_{\mathcal{S}}^2$. It follows that $\zeta_{max} - \zeta_{min} = 2$, and it is convenient to choose $\zeta_{max} = 1$ and $\zeta_{min} = -1$. The point $\zeta = +1$ is referred to as the north pole, and $\zeta = -1$ as the south pole.

In several places, it will be useful to parametrize the curve γ using its normalized proper length parameter, which is related to ζ via

$$\tilde{s}(\zeta) := \frac{\int_\zeta^1 F(\zeta')^{-1/2}\, d\zeta'}{\int_{-1}^1 F(\zeta')^{-1/2}\, d\zeta'}\,. \tag{4.6}$$

This is zero at the north pole and one at the south pole.

A few details about the canonical metric Eq. (4.4) will be useful for us later. First, since the axial Killing vector ϕ^a vanishes at the poles, we see from Eq. (4.5) that

$$\lim_{\zeta\to\pm1} F(\zeta) = 0\,. \tag{4.7}$$

Second, in order to avoid conical singularities at the pole, we must have

$$\lim_{\zeta\to\pm1} F'(\zeta) = \mp 2 \tag{4.8}$$

These conditions also hold if ζ_{max} and ζ_{min} differ from $+1$ and -1 respectively (as long as F vanishes at these points).

The scalar curvature $\mathcal{R}$ calculated from the metric Eq. (4.4) is

$$\mathcal{R}(\zeta) = -\frac{1}{R_{\mathcal{S}}^2} F''(\zeta)\,. \tag{4.9}$$

Note that the geometric volume element corresponding to the metric Eq. (4.4) is $d^2V = R_{\mathcal{S}}^2 d\zeta\, d\phi$ and is thus independent of F, and therefore the same as on a "round" 2-sphere. This fact is important when discussing the orthogonality of spherical harmonics. We can define the spherical harmonics $Y_{\ell m}(\theta, \phi)$ (with $\theta = \cos^{-1}\zeta$) on any distorted axisymmetric sphere. These would then satisfy the same orthogonality relationship with the geometric volume element on $\mathcal{S}$ as for the "round" 2-sphere. We will follow the normalization

$$\oint_{\mathcal{S}} Y_\ell^m Y_{\ell'}^{m'\star} d^2V = R_{\mathcal{S}}^2 \delta_{\ell\ell'}\delta_{mm'}\,. \tag{4.10}$$

In this chapter we will use the $m = 0$ harmonics $Y_{\ell 0}$. In terms of the Legendre polynomials $P_\ell(\zeta)$, these are

$$Y^0_\ell(\theta, \phi) = \sqrt{\frac{(2\ell + 1)}{4\pi}} P_\ell(\zeta)\,. \tag{4.11}$$

The above coordinate system can be employed to construct invariant geometric multipole moments on an axisymmetric MOTS. For a non-spinning axisymmetric QLH, the key geometric quantity of interest is the intrinsic scalar curvature $\mathcal{R}$ on a MOTS. First, we recall the well known Gauss-Bonnet theorem relating $\mathcal{R}$ on a compact 2-manifold $\mathcal{S}$ (with volume element d^2V) without boundary to its genus g:

$$\oint_{\mathcal{S}} \mathcal{R}\, d^2V = 8\pi(1 - g)\,. \tag{4.12}$$

For a sphere $g = 0$ so that

$$\oint_{\mathcal{S}} \mathcal{R}\, d^2V = 8\pi\,. \tag{4.13}$$

Here we shall deal only with topological spheres; even the MOTSs with self-intersections are, topologically speaking, spheres. In the invariant coordinate system, it is easy to verify the Gauss-Bonnet theorem by an explicit integration using Eqs. (4.9) and (4.8).

We define then geometric multipole moments on $\mathcal{S}$ by decomposing $\mathcal{R}$ using the spherical harmonics [21, 20]

$$\begin{aligned} I_\ell &= \frac{1}{4} \oint_{\mathcal{S}} \mathcal{R} Y_{\ell 0}(\zeta, \phi)\, d^2V && (4.14)\\ &= \pi R^2_{\mathcal{S}} \sqrt{\frac{(2\ell + 1)}{16\pi}} \int_{-1}^{1} \mathcal{R}(\zeta) P_\ell(\zeta)\, d\zeta && (4.15)\\ &= -\pi \sqrt{\frac{(2\ell + 1)}{16\pi}} \int_{-1}^{1} F''(\zeta) P_\ell(\zeta) d\zeta\,. && (4.16) \end{aligned}$$

In principle, if we were to use different angular coordinates on $\mathcal{S}$, the values of the multipole moments would differ. Using the geometrically defined (ζ, ϕ) removes this ambiguity. Also, since we shall only consider axisymmetric

surfaces, $\mathcal{R}$ will have no ϕ dependence and we do not need to consider $Y_{\ell m}(\zeta,\phi)$ (and thus $I_{\ell m}$) with non-vanishing m. Using again the Gauss-Bonnet theorem, we see that $I_0 = \sqrt{\pi}$. Finally, from Eq. (4.14) it can be shown that $I_1 = 0$ since

$$I_1 = -\sqrt{\frac{3\pi}{16}} \int_{-1}^{1} F''(\zeta)\zeta \, d\zeta \tag{4.17}$$

$$= -\sqrt{\frac{3\pi}{16}} \left\{ (F'\zeta - F)\big|_{-1}^{1} \right\} = 0 \,. \tag{4.18}$$

We have used Eqs. (4.7) and (4.8).

Since I_1 vanishes identically, we need only consider $\ell \geq 2$. The above *geometric* multipoles are dimensionless, and must be rescaled to get the physical horizon mass multipole moments:

$$\mathscr{M}_\ell = \sqrt{\frac{4\pi}{2\ell+1}} \frac{M R_{\mathcal{S}}^{\ell}}{2\pi} I_\ell \,. \tag{4.19}$$

Here $M = \frac{R_S}{2}$ is the horizon mass (for a non-spinning BH). In this case

$$\mathscr{M}_\ell = \sqrt{\frac{4\pi}{2\ell+1}} \frac{R_{\mathcal{S}}^{\ell+1}}{4\pi} I_\ell \,. \tag{4.20}$$

The vanishing of the mass dipole $\mathscr{M}_1$ can be taken to imply that we are in the center of mass frame. There are also appropriate definitions for angular momentum and spin multipole moments, but we shall not focus on them in this chapter. See [141] for a comparison of different notions of multipole moments for a Kerr black hole.

4.2.3 The numerical set-up

The initial data is time-symmetric Brill-Lindquist data [79] for two non-spinning and uncharged black holes. Thus, we consider Euclidean space $\mathbb{R}^3$ with two points (the "punctures") $\mathbf{x}_{1,2}$ excluded. Let δ_{ij} be the Euclidean metric on $\mathbb{R}^3$, and d the Euclidean distance between the punctures: $d = |\mathbf{x}_1 - \mathbf{x}_2|$. The extrinsic curvature is trivial: $K_{ij} = 0$, and the 3-metric

is conformally flat: $h_{ij} = \Phi^4 \delta_{ij}$. The conformal factor which satisfies the constraint equations of vacuum general relativity, and $\Phi \to 1$ at spatial infinity, is

$$\Phi(\mathbf{x}) = 1 + \frac{m_1}{2\,|\mathbf{x} - \mathbf{x}_1|} + \frac{m_2}{2\,|\mathbf{x} - \mathbf{x}_2|}\,. \tag{4.21}$$

Despite the simplicity, this data contains a rather complicated set of MOTSs. For a sequence of initial data with $m_{1,2}$ fixed and d decreasing, it is seen that for large d, the initial data contains two separate MOTSs corresponding to the two progenitor BHs. As d is decreased, a common horizon appears at a particular value of d which in turn bifurcates into an inner and outer branch, rather similar to what is seen during time evolution; see [228] for details.

As mentioned earlier, we use the same simulation as in Refs. [229, 231]. The bare masses are $m_1 = 0.5$ and $m_2 = 0.8$. The total mass of our simulation is hence $m_1 + m_2 = 1.3$ and our plots and numerical values use this mass, which we denote with $\mathcal{M}$. The punctures are located on the z axis, initially at $\pm 0.65\,\mathcal{M}$. The exact parameter files are available in the repository [232]. No mesh refinement or multiple grid resolutions are used. The simulations are performed with the *Einstein Toolkit* [194, 122] and *TwoPunctures* [15] for the initial conditions, and *McLachlan* and *Kranc* [80, 163, 177] to solve the Einstein equations.

4.3 The area and mass at the merger

With the necessary background and the tools outlined above, we are now ready to start studying the cusp at the merger. There are three MOTS sequences of interest to us: $\mathcal{S}_{\text{inner}}$, $\mathcal{S}_1$ and $\mathcal{S}_2$. The behavior of $\mathcal{R}$ on $\mathcal{S}_1$ and $\mathcal{S}_2$ is relatively straightforward: it remains smooth through the merger even when they touch and subsequently intersect each other. On the other hand, $\mathcal{S}_{\text{inner}}$ is more complicated since it develops a narrow "neck" leading to a cusp at the time t_{touch}. $\mathcal{R}_{\text{inner}}$ takes increasingly large negative values at the neck and then becomes singular at the cusp exactly at t_{touch}. Subsequently, when $\mathcal{S}_{\text{inner}}$ develops self-intersections, $\mathcal{R}$ is again smooth. Therefore, we will focus here on the behavior of $\mathcal{R}_{\text{inner}}$ near t_{touch}. Consistent with the

scenario sketched in Sec. 4.2.1, the mean curvatures of $\mathcal{S}_1$ and $\mathcal{S}_2$ coincide at the point of contact; see Fig. 4.2. The mean curvature of $\mathcal{S}_{\rm inner}$ is seen to have a spike around the neck and around $t_{\rm touch}$; the region with positive mean curvature shrinks near $t_{\rm touch}$.

At $t_{\rm touch}$ we would like to relate the masses of $\mathcal{S}_1$ and $\mathcal{S}_2$ with the mass of $\mathcal{S}_{\rm inner}$. This is straightforward because at $t_{\rm touch}$, we must have

$$A_{in} = A_1 + A_2\,, \tag{4.22}$$

which implies, $R_{in}^2 = R_1^2 + R_2^2$. Thus, since these are non-spinning black holes, at $t_{\rm touch}$, the mass of $\mathcal{S}_{\rm inner}$ is obtained from the masses of $\mathcal{S}_1$ and $\mathcal{S}_2$ as

$$M_{in} = \sqrt{M_1^2 + M_2^2}\,. \tag{4.23}$$

This value of M_{in} can then be used as an initial value as we follow $\mathcal{S}_{\rm inner}$, initially backwards in time, through $t_{\rm bifurcate}$ and then eventually towards the final black hole. It follows that $M_{in} < M_1 + M_2$ and moreover

$$\begin{aligned} M_{in} &= (M_1+M_2)\sqrt{1 - \frac{2M_1M_2}{(M_1+M_2)^2}} \\ &= (M_1+M_2)\sqrt{1-2\eta} \end{aligned} \tag{4.24}$$

with $\eta = M_1M_2/(M_1+M_2)^2$ being the symmetric mass ratio. Since $0 \leq \eta \leq 0.25$, we must have $M_{in} \leq M_1 + M_2$. For equal masses, $\eta = 0.25$ which leads to the largest value of $M_1 + M_2 - M_{in} = (M_1+M_2)(1-1/\sqrt{2})$ (this is the same limit as obtained in [155]). For highly asymmetric systems, i.e. small η, we will have $M_{in} \approx M_1 + M_2$. When we connect this to the final remnant mass (by tracing $\mathcal{S}_{\rm inner}$ *backwards* in time till we reach $t_{\rm bifurcate}$, and then forward in time till we reach the final state of the remnant) the final mass will be larger than the above value of M_{in}. In greater detail, if we trace the mass of the system using the horizon masses of the above sequence of MOTSs, we get three distinct effects:

1. The behavior of $\mathcal{S}_1$ and $\mathcal{S}_2$: the mass is $M_1 + M_2$ which increases monotonically for $-\infty < t < t_{\rm touch}$, as the two horizons absorb infalling radiation.

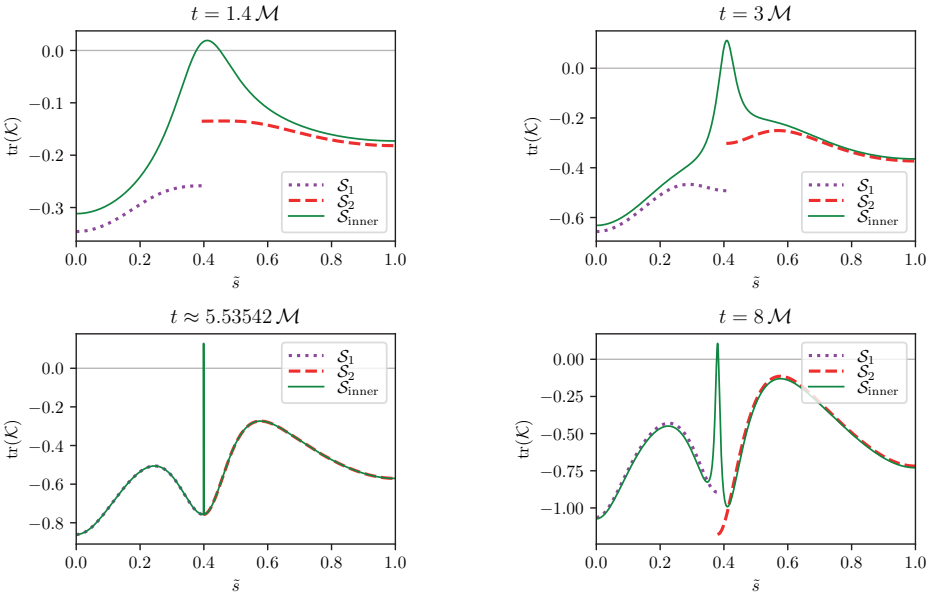


Figure 4.2: Mean curvature $\mathrm{tr}(\mathcal{K})$ of $\mathcal{S}_{\mathrm{inner}}$ (green) and the two individual MOTSs (dashed purple and red). The parameter $\tilde{s}$ on the x-axis here is the proper length of a constant-ϕ curve on $\mathcal{S}_1$, $\mathcal{S}_2$; these parameters on $\mathcal{S}_1$ and $\mathcal{S}_2$ have been transformed linearly so that they respectively end or start at the location $\tilde{s}_\star$ (where $\mathcal{S}_{\mathrm{inner}}$ has a locally minimal circumference away from its poles). This allows us to compare the various mean curvatures on the same plot even away from t_{touch} where $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_{\mathrm{inner}}$ are distinct surfaces.

2. At t_{touch}: A discontinuous decrease in the mass when we go from $M_1 + M_2 \to M_{in}$.

3. Approach to equilibrium: M_{in} generally increases till we reach the final equilibrium state (this increase is however not always monotonic [229]).

The discontinuous jump in (2) above is larger (in absolute value) than the mass increases in (1) and (3), so that the final remnant mass is less than the sum of the initial masses. This is consistent with the corresponding scenario at null infinity, where the flux of gravitational waves carries away energy from the system leading to a mass-loss.

4.4 The geometry of $\mathcal{S}_{\text{inner}}$ near t_{touch}

Going beyond the mass to the higher multipole moments requires a more detailed analysis of the geometry of $\mathcal{S}_{\text{inner}}$ (and its relation to $\mathcal{S}_1, \mathcal{S}_2$) near t_{touch}. Since $\mathcal{S}_{\text{inner}}$ (like the other MOTSs) is axisymmetric, we can use the invariant coordinates defined earlier. Let (ζ, ϕ) (with $-1 \leq \zeta < 1$ and $0 \leq \phi < 2\pi$) be the invariant coordinate system on $\mathcal{S}_{\text{inner}}$ and let $\mathcal{R}(\zeta)$ be the scalar curvature; it is independent of ϕ due to axisymmetry. Define the angle $\theta = \cos^{-1}\zeta$. Let $\zeta_\star$ be the location of the neck, i.e. the narrowest point of $\mathcal{S}_{\text{inner}}$. Let further $\theta_\star$ and $\tilde{s}_\star$ be the corresponding values of θ and $\tilde{s}$. The circumference of a constant-ζ curve is seen to be $2\pi R\sqrt{F(\zeta)}$, whence the narrowest point $\theta_\star$ corresponds to a minimum of $F(\zeta)$. When considering the behavior of the Ricci scalar $\mathcal{R}_{\text{inner}}$ (and F) of $\mathcal{S}_{\text{inner}}$, we need to analyze two limits, the one of $\theta \to \theta_\star$ and the limit $t \to t_{\text{touch}}$. Fixing a $\theta \neq \theta_\star$, the Ricci scalar converges to that of $\mathcal{S}_1$ or $\mathcal{S}_2$ (depending on whether $\theta < \theta_\star$ or $\theta > \theta_\star$).

Fig. 4.3 shows the Ricci scalars of $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_{\text{inner}}$ with the curve parameters suitably rescaled and shifted so that the values can be compared with $\mathcal{S}_{\text{inner}}$; for $\mathcal{S}_1$ we will have $0 < \theta < \theta_\star$, while for $\mathcal{S}_2$ we will have $\theta_\star < \theta < \pi$. We see that $\mathcal{R}_1$ and $\mathcal{R}_2$ are everywhere positive while $\mathcal{R}_{\text{inner}}$ is negative around the neck. Moreover, from the bottom left panel of Fig. 4.3 it is evident that the curvatures on $\mathcal{S}_1$ and $\mathcal{S}_2$ agree at the point of contact.

This issue, along with the question of the coordinate singularity at the poles, is discussed further in Appendix 4.7 and in Appendix 4.8.

Let $\mathcal{R}_{12} : [0, \pi] \setminus \{\theta_\star\} \to \mathbb{R}$ be the values of the Ricci scalar $\mathcal{R}$ along first $\mathcal{S}_1$ and then $\mathcal{S}_2$, with the transition happening at the neck of $\mathcal{S}_{\text{inner}}$, i.e.

$$\mathcal{R}_{12}(\theta) := \begin{cases} \mathcal{R}_1\left(\pi \, \dfrac{\theta}{\theta_\star}\right) , & 0 \leq \theta < \theta_\star \\ \mathcal{R}_2\left(\pi \, \dfrac{\theta - \theta_\star}{\pi - \theta_\star}\right) , & \theta_\star < \theta \leq \pi \end{cases} . \tag{4.25}$$

We see that at t_{touch}, if the negative spike were to be excluded, the resulting function for $\mathcal{R}_{\text{inner}}$ would be everywhere continuous and differentiable, and would in fact agree with $\mathcal{R}_{12}$. However, $\mathcal{R}_{12}$ cannot, by itself, be the correct model for $\mathcal{R}_{\text{inner}}$ since it would give the incorrect topology as far as the Gauss-Bonnet theorem is concerned. At the neck, the plot of the maximum of $|\mathcal{R}_{\text{inner}}|$ shown in Fig. 4.4 indicates that $\mathcal{R}$ diverges to $-\infty$ as $t \to t_{\text{touch}}$. Thus, at t_{touch}, the difference between $\mathcal{R}_{\text{inner}}$ and $\mathcal{R}_{12}$ is a δ-function normalized precisely so that the Gauss-Bonnet formula is satisfied

If we write $\mathcal{R}_{\text{inner}} = \mathcal{R}_{12} + \Delta\mathcal{R}$, then Eq. (4.13) implies

$$\lim_{t \to t_{\text{touch}}} \int_{\mathcal{S}_{\text{inner}}} \Delta\mathcal{R} \; dA = -8\pi \, , \tag{4.26}$$

provided $\mathcal{S}_{\text{inner}}$ merges with $\mathcal{S}_{1,2}$ at t_{touch}. The above integral is shown in Fig. 4.5 at times around t_{touch}. This shows that we have strong numerical support for the validity of Eq. (4.26). Additionally, we have mentioned earlier that the facing points on the poles of $\mathcal{S}_{1,2}$ take on the same value of the Ricci scalar at exactly t_{touch}, which is the time when these two points meet. Fig. 4.3 already suggests this possibility and Fig. 4.6 shows that we indeed find a strong numerical indication for this fact.

It is also instructive to see what the above implies for the behavior of F at t_{touch}. Since $\mathcal{R}_{\text{inner}}$ is the second derivative of F, the δ-function singularity of $\mathcal{R}_{\text{inner}}$ at $\zeta_\star$ implies a discontinuity of F' at $\zeta_\star$. At t_{touch}, we thus deduce the following facts about F near $\zeta_\star$. The neck has vanishing circumference:

$$\lim_{\zeta \to \zeta_\star} F(\zeta) = 0 \, . \tag{4.27}$$

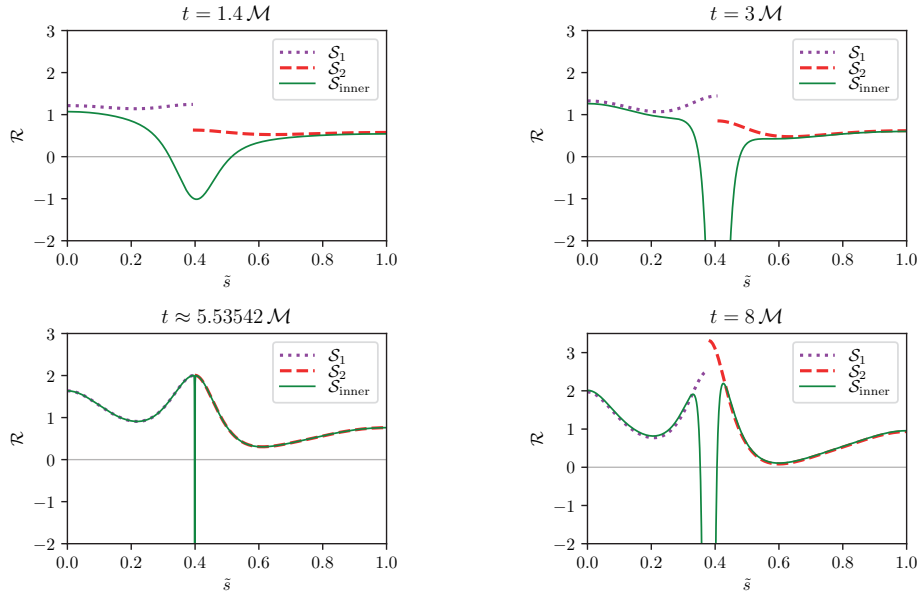


Figure 4.3: Ricci scalar $\mathcal{R}$ of $\mathcal{S}_{\text{inner}}$ and the two individual MOTSs. The curve parameter of $\mathcal{S}_1$, $\mathcal{S}_2$ has been transformed linearly to end or start at $\tilde{s}_\star$, respectively.

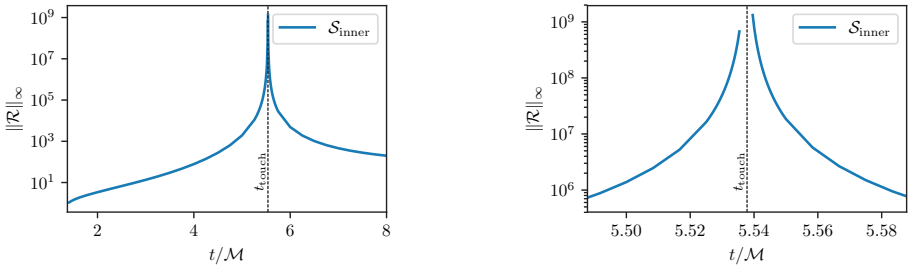


Figure 4.4: Maximum value of the Ricci scalar $|\mathcal{R}|$ along $\mathcal{S}_{\text{inner}}$ as function of time. The right panel shows a close-up near t_{touch}. We find numerical indication that the Ricci scalar diverges as $t \to t_{\text{touch}}$.

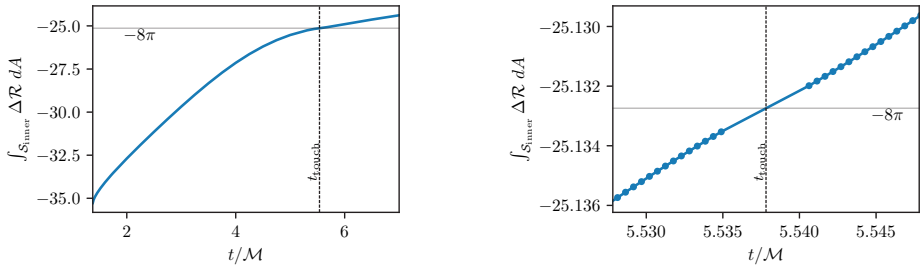


Figure 4.5: Integral of $\Delta\mathcal{R}$ as function of simulation time. This shows our numerical support for Eq. (4.26).

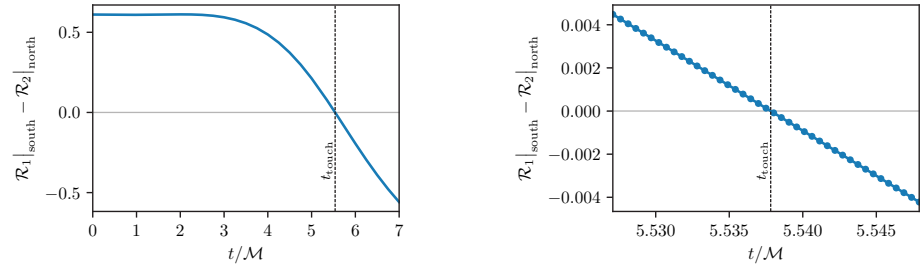


Figure 4.6: Difference of the Ricci scalar at the south pole of $\mathcal{S}_1$ and north pole of $\mathcal{S}_2$. This shows that the Ricci scalars coincide at the two points that touch at t_{touch}.

Note that F is positive-definite away from the poles and the neck. The absence of conical singularities implies (as in Eq. (4.8)) that

$$\lim_{\zeta\to\zeta_\star^-} F'(\zeta) = -2\,, \qquad \lim_{\zeta\to\zeta_\star^+} F'(\zeta) = 2\,. \tag{4.28}$$

Thus, near the neck, F'' has a δ-function singularity: $F'' \approx 4\delta(\zeta-\zeta_\star)$. Using the proper area element on $\mathcal{S}_{\rm inner}$, this is seen to be consistent with Eq. (4.26) as it should be. It is thus natural to say that the neck of $\mathcal{S}_{\rm inner}$ pinches off into the "north-pole" of $\mathcal{S}_2$ ($\zeta\to\zeta_\star^-$) and the "south-pole" of $\mathcal{S}_1$ ($\zeta\to\zeta_\star^+$).

Finally, the value of $\zeta_\star$ at $t_{\rm touch}$ can be calculated as follows. Since $\mathcal{S}_{\rm inner} = \mathcal{S}_1 \bigcup \mathcal{S}_2$, the total area of $\mathcal{S}_{\rm inner}$ must equal the sum of the areas of $\mathcal{S}_1$ and $\mathcal{S}_2$. The area of $\mathcal{S}_{\rm inner}$ from the south-pole $\zeta_{in}=-1$ up to $\zeta=\zeta_\star$ is $2\pi R_{in}^2(\zeta_\star+1)$. But this should also be equal to the area of $\mathcal{S}_1$, i.e. $4\pi R_1^2$. Thus:

$$\begin{aligned}\zeta_\star &= \frac{2R_1^2}{R_{in}^2} - 1 = 2\frac{M_1^2}{M_{in}^2} - 1 \\ &= \frac{M_1^2 - M_2^2}{M_1^2 + M_2^2}\,.\end{aligned} \tag{4.29}$$

4.5 The mass multipole moments at the merger

In this section we present analytical results regarding the mass multipole moments of $\mathcal{S}_{\rm inner}$ at the merger, i.e. at $t = t_{\rm touch}$.

4.5.1 The mass dipole moment at the merger

The mass dipole moment $\mathscr{M}_1$ (or alternatively the geometric moment I_1) of $\mathcal{S}_{\rm inner}$ has some features of interest as we shall now discuss. As mentioned earlier, $\mathscr{M}_1$ vanishes identically. Following [21], this is interpreted to say that the coordinates (ζ,ϕ) place us in the center of mass of the horizon. Previously, we proved $I_1 = 0$ in Eq. (4.18). However that derivation assumed

that F vanishes only at the poles, but now it vanishes additionally at $\zeta_\star$ and we must therefore revisit the vanishing of I_1 for $\mathcal{S}_{\rm inner}$ at $t_{\rm touch}$.

Let us break up the integral for I_1 at the neck:

$$\begin{aligned} I_1 &= -\sqrt{\frac{3\pi}{16}} \int_{-1}^{1} F''(\zeta)\zeta \, d\zeta \\ &= -\sqrt{\frac{3\pi}{16}} \left\{ \int_{-1}^{\zeta_\star} F''(\zeta)\zeta \, d\zeta + \int_{\zeta_\star}^{1} F''(\zeta)\zeta \, d\zeta \right\} \end{aligned} \tag{4.30}$$

At first glance, one might interpret the two individual integrals appearing here with the dipole moments of $\mathcal{S}_1$ and $\mathcal{S}_2$. This is however not correct and these two terms do not individually vanish[6]. A simple calculation yields (using Eq. 4.28):

$$\int_{-1}^{\zeta_\star} F''(\zeta)\zeta \, d\zeta = (F'\zeta - F)\big|_{-1}^{\zeta_\star} = 2 - 2\zeta_\star \,, \tag{4.31}$$

$$\int_{\zeta_\star}^{1} F''(\zeta)\zeta \, d\zeta = (F'\zeta - F)\big|_{-1}^{\zeta_\star} = -2 - 2\zeta_\star \,. \tag{4.32}$$

The sum therefore does not vanish:

$$I_1 = \sqrt{\frac{3\pi}{16}} \times 4\zeta_\star = \sqrt{3\pi}\frac{M_1^2 - M_2^2}{M_1^2 + M_2^2} \,. \tag{4.33}$$

We conclude that the mass-dipole moment of $\mathcal{S}_{\rm inner}$ does *not* vanish at $t_{\rm touch}$. Immediately away from $t_{\rm touch}$ however I_1 does vanish. We can interpret this behavior as a jump in the center of mass corresponding to the value

$$\Delta I_1 = \sqrt{3\pi}\frac{M_1^2 - M_2^2}{M_1^2 + M_2^2} \,. \tag{4.34}$$

This shift vanishes for an equal mass binary, and is otherwise in the direction of the heavier black hole; it approaches $\pm\sqrt{3\pi}$ in the limit of extreme

[6]In fact, we have here three different coordinate systems. On $\mathcal{S}_{\rm inner}$ we have the coordinates (ζ, ϕ). These are however not the same as the coordinates on $\mathcal{S}_1, \mathcal{S}_2$. If (ζ_1, ϕ) and (ζ_2, ϕ) are coordinates on $\mathcal{S}_1$ and $\mathcal{S}_2$ respectively, then the mass dipole moments vanish in these coordinates. However, this does not hold in the coordinates (ζ, ϕ).

mass ratios. The same results can be obtained also directly from the scalar curvature. Modeling $\mathcal{R}_{\rm inner}$ as the sum of a smooth part and a δ-function (as we shall discuss shortly), the δ-function is seen to be responsible for the above shift in the center of mass.

4.5.2 The mass quadrupole and higher moments

Let us now turn to the geometric multipole moments I_ℓ for $\ell \geq 2$. It will be easier to work directly with the scalar curvature on $\mathcal{S}_{\rm inner}$ with its δ-function singularity. Thus, we take at $t_{\rm touch}$

$$\mathcal{R}_{\rm inner} = \mathcal{R}_{12} - \frac{4}{R_{in}^2}\delta(\zeta - \zeta_\star)\,. \tag{4.35}$$

Here $\mathcal{R}_{12}$ is a smooth function whose integral over $\mathcal{S}_{\rm inner}$ is 16π; it coincides with $\mathcal{R}_1$ on the portion $\mathcal{S}_1$ and with $\mathcal{R}_2$ on $\mathcal{S}_2$. Applying Eq. (4.14) then leads to three contributions for $\mathcal{S}_{\rm inner}$:

$$I_\ell^{(in)} = I_\ell^{(in,1)} + I_\ell^{(in,2)} + I_\ell^{(cusp)} \tag{4.36}$$

where

$$I_\ell^{(in,1)} = \pi R_{in}^2 C_\ell \int_{-1}^{\zeta_\star} \mathcal{R}_1(\zeta) P_\ell(\zeta) d\zeta\,, \tag{4.37}$$

$$I_\ell^{(in,2)} = \pi R_{in}^2 C_\ell \int_{\zeta_\star}^{1} \mathcal{R}_2(\zeta) P_\ell(\zeta) d\zeta\,, \tag{4.38}$$

$$I_\ell^{(cusp)} = -8\pi^2 C_\ell \int_{-1}^{1} \delta(\zeta - \zeta_\star) P_\ell(\zeta) d\zeta\,. \tag{4.39}$$

We have defined

$$C_\ell = \sqrt{\frac{(2\ell+1)}{16\pi}}\,. \tag{4.40}$$

At $t_{\rm touch}$ we obtain

$$I_\ell^{(cusp)} = -8\pi^2 \sqrt{\frac{(2\ell+1)}{16\pi}} P_\ell(\zeta_\star) \tag{4.41}$$

We see that the contribution of the cusp will generally be non-vanishing unless $\zeta_\star$ happens to lie at one of the zeros of $P_\ell(\zeta)$. As we have seen in Eq. (4.29), the value of $\zeta_\star$ itself depends on the relative size of the two black holes, and thus eventually on the mass ratio $q = M_2/M_1$ of the binary system. Thus, there should be values of q, depending on ℓ, for which this contribution should vanish at the cusp. As an example, consider the quadrupole $\ell = 2$. The zeros of $P_2(\zeta)$ occur at $\zeta = \pm(1/\sqrt{3})$. Taking then $\zeta_\star = 1/\sqrt{3}$ leads to $q^2 = (\sqrt{3}-1)/(\sqrt{3}+1)$, i.e. $q \approx 0.268$. Thus, for configurations with mass ratios close to this value, we should find correspondingly smaller deviations in the quadrupole moment across the merger. Similar considerations apply to the higher multipoles.

Consider now the other contributions to the multipole moments due to $\mathcal{R}^{(1)}_{\rm inner}$ and $\mathcal{R}^{(2)}_{\rm inner}$ at $t_{\rm touch}$. As before, we define $\mathcal{R}^{(1)}_{\rm inner}$, $\mathcal{R}^{(2)}_{\rm inner}$ as having support on $\zeta \in [-1, \zeta_\star)$ and $(\zeta_\star, 1]$, respectively, such that away from $\zeta_\star$, $\mathcal{R}^{(1)}_{\rm inner} + \mathcal{R}^{(2)}_{\rm inner} = \mathcal{R}_{12}$. However, at $t_{\rm touch}$ itself, the two will differ by the δ-function at the cusp. Since $\mathcal{S}_1 \bigcup \mathcal{S}_2 = \mathcal{S}_{\rm inner}$ at $t_{\rm touch}$, we must be careful to distinguish between the different ζ coordinates on $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_{\rm inner}$. On these surfaces, we will have coordinates $\zeta_1, \zeta_2, \zeta_{in}$ each of which takes values within the range $[-1, 1]$. Thus, even though $\mathcal{R}^{(1)}_{\rm inner}$ agrees exactly with $\mathcal{R}_1$, we need to consider that $\mathcal{R}_1(\zeta_1)$ is defined for $-1 \leq \zeta_1 \leq 1$ while for $\mathcal{R}^{(1)}_{\rm inner}(\zeta_{in})$ we will have $-1 \leq \zeta_{in} < \zeta_\star$. Note that

$$\int_{-1}^{1} \mathcal{R}_{12} P_\ell d\zeta = \int_{-1}^{\zeta_\star} \mathcal{R}^{(1)}_{\rm inner} P_\ell d\zeta + \int_{\zeta_\star}^{1} \mathcal{R}^{(2)}_{\rm inner} P_\ell d\zeta \,. \tag{4.42}$$

At $t_{\rm touch}$, each of the terms on the right-hand-side should be closely related to the multipole moments on $\mathcal{S}_1$ and $\mathcal{S}_2$ respectively. For clarity, at $t_{\rm touch}$, we shall distinguish between ζ_{in}, ζ_1 and ζ_2. Thus, for example, by changing

variables from ζ_{in} to ζ_1, the first term becomes

$$\begin{aligned}
&\int_{-1}^{\zeta_\star} \mathcal{R}^{(1)}_{\rm inner}(\zeta_{in}) P_\ell(\zeta_{in}) d\zeta_{in} \\
&= \int_{-1}^{1} \mathcal{R}^{(1)}_{\rm inner}(\zeta_{in}(\zeta_1)) P_\ell(\zeta_{in}(\zeta_1)) \frac{d\zeta_{in}}{d\zeta_1} d\zeta_1 \\
&= \int_{-1}^{1} \mathcal{R}_1(\zeta_1) P_\ell(A\zeta_1 + B) \frac{d\zeta_{in}}{d\zeta_1} d\zeta_1 \\
&= A \int_{-1}^{1} \mathcal{R}_1(\zeta_1) P_\ell(A\zeta_1 + B) d\zeta_1 \,. \qquad (4.43)
\end{aligned}$$

In the third line we have assumed that (at $t_{\rm touch}$), at points of $\mathcal{S}_1$, $\mathcal{R}^{(1)}_{\rm inner}$ is identical to the scalar curvature $\mathcal{R}_1$. To obtain the last line, note that the transformation $\zeta_{in}(\zeta_1)$ is linear

$$\zeta_{in} = A\zeta_1 + B \qquad (4.44)$$

where

$$A = \frac{M_2^2}{M_1^2 + M_2^2} = \frac{1}{1+q^2} \,, \qquad (4.45)$$

$$B = -\frac{M_1^2}{M_1^2 + M_2^2} = -\frac{q^2}{1+q^2} \,. \qquad (4.46)$$

This transformation gives the expected results: $\zeta_{in}(-1) = -1$ and $\zeta_{in}(1) = \zeta_\star$ consistent with Eq. (4.29). Thus, we have the following result: The multipole moments of $\mathcal{R}^{(1)}_{\rm inner}(\zeta_{in})$ viewed as the curvature of $\mathcal{S}_{\rm inner}$, are *not* the same as the multipole moments on $\mathcal{S}_1$, even though the curvatures agree. Since the transformation $\zeta_{in} \to \zeta_1$ is linear, we will obtain the ℓ^{th} and also lower multipoles of $\mathcal{S}_1$. Similar considerations apply for the integral over $\mathcal{S}_2$ as well.

Illustrating the above for the quadrupole moment, we will clearly obtain a combination of the dipole terms (which vanish for the individual black holes) and the monopoles. Putting all of this together, a straightforward

calculation yields:

$$I_2^{(in)} = \frac{1}{(1+q^2)^2} I_2^{(1)} + \frac{1}{q^2(1+q^2)^2} I_2^{(2)} + \sqrt{\frac{5\pi}{4}} \frac{(1-q^2)}{(1+q^2)^2} \left(q^2 + \frac{2}{q^2} \right) - \sqrt{20\pi^3} P_2(\zeta_\star) \,. \tag{4.47}$$

The two terms in the first line are the contributions from the individual quadrupole moments. In the second line, the first term consists of the monopole contributions from the two black holes, while the second term is the contribution from the cusp. It is straightforward, if somewhat tedious, to obtain analogous equations for the higher mass multipole moments.

The considerations thus far are valid only at $t_{\rm touch}$ where the scalar curvature has a δ-function singularity. As the inner horizon $\mathcal{S}_{\rm inner}$ evolves away from the cusp, the δ-function is smeared out and eventually the contributions of the cusp decay. This however requires a numerical study which is presented in the next section.

4.5.3 Evolution of the mass multipole moments

The above considerations suggest that $\mathcal{R}_{\rm inner}$ can be modeled as the sum of a smooth function plus αG which limits to a δ-function at $t_{\rm touch}$; the parameter α is the scaling parameter while G is a function which limits to a δ-function with unit normalization. Our strategy will be to first estimate α and G near the cusp. Then, $\mathcal{R}_{\rm inner} - \alpha G$ will be a smooth function on $\mathcal{S}_{\rm inner}$ which is meant to consist of contributions from the two individial horizons. We now describe the details of this fitting procedure.

The neck $\zeta_\star$ is a minimum of the scalar curvature:

$$\left.\frac{d\mathcal{R}_{\rm inner}}{d\zeta}\right|_{\zeta=\zeta_\star} = 0 \qquad \left.\frac{d^2\mathcal{R}_{\rm inner}}{d\zeta^2}\right|_{\zeta=\zeta_\star} > 0 \,. \tag{4.48}$$

There could be multiple minima of $\mathcal{R}_{\rm inner}$ so this requires an inspection of $\mathcal{S}_{\rm inner}$ empirically. If we restrict our attention sufficiently close to $t_{\rm touch}$, this is unambiguously the global minimum of $\mathcal{R}_{\rm inner}$ over $\mathcal{S}_{\rm inner}$.

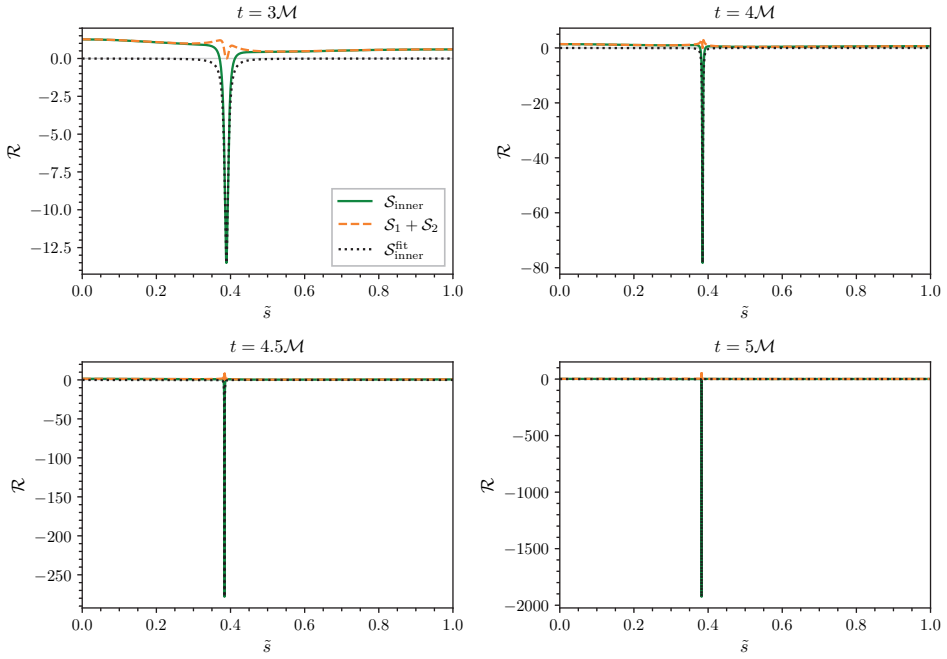


Figure 4.7: Ricci scalar $\mathcal{R}_{\text{inner}}$ of $\mathcal{S}_{\text{inner}}$ (green), the two individual MOTSs (orange dashed) and the numerically fitted function (including a wrapped Cauchy distribution) to $\mathcal{S}_{\text{inner}}$ (black dotted). As $t \to t_{\text{touch}}$ the width of the Cauchy distribution decreases and approaches a δ-function. The model is meant to capture the behavior of $\mathcal{R}_{\text{inner}}$ around the neck, and for times near t_{touch}. The behavior near $\zeta_\star$ is shown in Fig. 4.8, which resolves the spike more clearly.

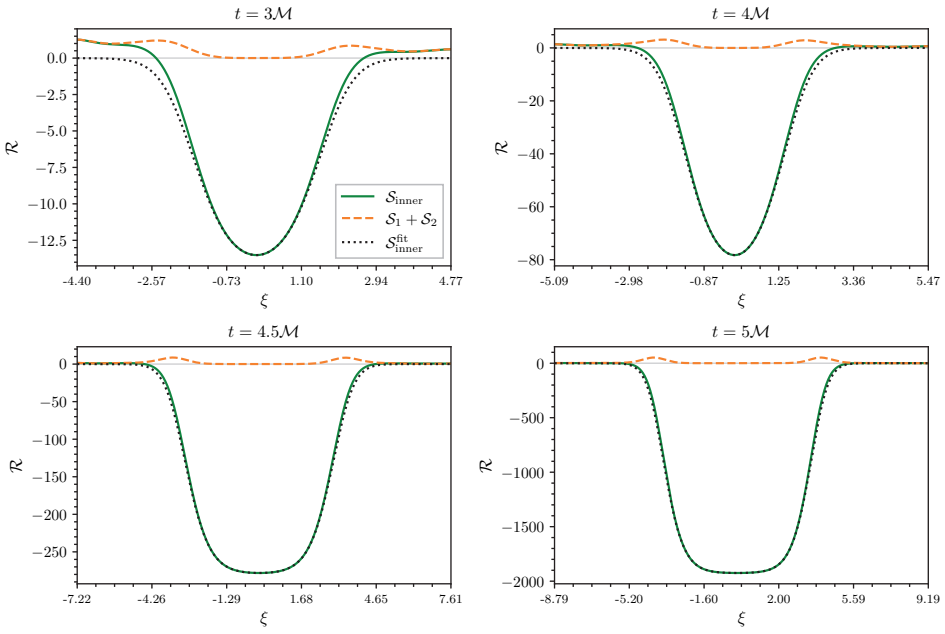


Figure 4.8: Ricci scalar $\mathcal{R}_{\text{inner}}$ of $\mathcal{S}_{\text{inner}}$ (green), the two individual MOTSs (orange dashed) and the numerically fitted function (with a wrapped Cauchy distribution) to $\mathcal{S}_{\text{inner}}$ (black dotted). As $t \to t_{\text{touch}}$ the width of the Cauchy distribution decreases and approaches a delta function. In this version of the figure the horizontal axis is the transformed coordinate $\xi = \text{asinh}\left(\frac{\tilde{s}-\tilde{s}_{\text{min}}(t)}{w(t)}\right)$, where $\tilde{s}_{\text{min}}(t)$ is the location of the minimum of $\mathcal{R}$ on $\mathcal{S}_{\text{inner}}$ for the given time slice and $w(t)$ is the width parameter controlling the amount of horizontal zoom. Thus, ξ is a recentered and stretched version of $\tilde{s}$, introduced to resolve the increasingly narrow region near the cusp as $t \to t_{\text{touch}}$.

We use the wrapped Cauchy distribution to model the singular part of $\mathcal{R}_{\text{inner}}$:

$$G(\theta;\theta_\star,\sigma) := \frac{1}{\pi}\frac{\sinh\sigma}{\cosh\sigma-\cos(2\theta-2\theta_\star)} \tag{4.49}$$

here scaled to the appropriate range of θ. The distribution is normalized to unity

$$\int_0^\pi G(\theta;\theta_\star,\sigma)d\theta = 1 \tag{4.50}$$

Furthermore, the peak of the distribution is located at $\theta_\star$ and σ is related to the width around $\theta_\star$. We obtain a δ-function when $\sigma \to 0$

$$\lim_{\sigma\to 0} G(\theta;\theta_\star,\sigma) = \delta(\theta-\theta_\star)\,. \tag{4.51}$$

We fit a wrapped Cauchy profile, multiplied by a scaling factor, to the neighborhood of the peak of the data, i.e. to the region near the neck. Since the wrapped Cauchy model is intended only as a local description, we do not fit it over the full angular domain. For each time slice, we first locate the peak and construct a small symmetric fitting window around it. We then enlarge this window one point at a time on each side and solve the resulting nonlinear least-squares problem with a Levenberg–Marquardt-type procedure. To improve conditioning in the sharpening regime, we fit in the variables (q,α), where $q=\log\sigma$, rather than in (σ,α) directly, thereby enforcing $\sigma>0$ automatically. We also use an algebraically equivalent but numerically more stable form of the wrapped Cauchy kernel, which avoids loss of significance when σ becomes very small.

The accepted fitting window is chosen adaptively. We accept the first converged local fit as a baseline and then compare each larger candidate window to the previously accepted one. A larger window is retained only if the fit remains sufficiently stable according to three diagnostics: the change in the fitted width parameter, the relative mean-square residual, and the predictive error at the newly added edge points. In this way, the fit is allowed to grow while it continues to represent the local peak accurately, but it is prevented from absorbing regions where the data are no longer well described by a wrapped Cauchy profile. Near t_{touch}, where the problem becomes very stiff, we allow a small number of rejected enlargements before terminating

the growth of the window, so that a single marginal step does not end the procedure prematurely.

Because of the increasingly singular behavior as $t \to t_{\text{touch}}$, arbitrary-precision arithmetic is used throughout the fitting stage. The fitted values of σ decrease rapidly and are consistent with $\sigma \to 0$ as $t \to t_{\text{touch}}$, as can be seen in Fig. 4.9. The associated scaling parameter is tracked separately as a function of time, as shown in Fig. 4.10.

Once the local wrapped Cauchy model has been fitted, we subtract the fitted inner profile from the full data to define the remainder, which we interpret as $\mathcal{R}_1 + \mathcal{R}_2$. We then compute the multipole moments I_ℓ of this residual field. For the required angular integrals, we interpolate the sampled data using a Piecewise Cubic Hermite Interpolating Polynomial (PCHIP), as implemented in `DataInterpolations.jl` [62], and integrate the interpolant over the sampled angular interval. This choice preserves monotonicity near sharp features and avoids the overshoot often produced by higher-order spline interpolants when the data are not fully smooth. The resulting fits for different times are shown in Fig. 4.7, and in greater detail around the neck in Fig. 4.8. These figures show both the singular and smooth parts of the fit. Fig. 4.9 shows the width of the Cauchy distribution, which vanishes in the limit $t \to t_{\text{touch}}$ as expected, and Fig. 4.10 shows the scaling parameter α. Finally, with the singular part αG determined by the fit, we can subtract it and obtain the smooth part corresponding to $\mathcal{R}_{12}$ and thus it's multipole moments. This is shown in Fig. 4.11. These moments would eventually asymptote to the multipole moments of the remnant black hole, and in this case since the final remnant is a Schwarzschild black hole, it means these will eventually decay away. This is studied in greater detail in [149, 212]. The moments we have calculated on $\mathcal{S}_{\text{inner}}$ need to be matched with those on the apparent horizon to obtain a complete description. This will take us away from the main topic of this work, which is the merger itself, and will be presented elsewhere.

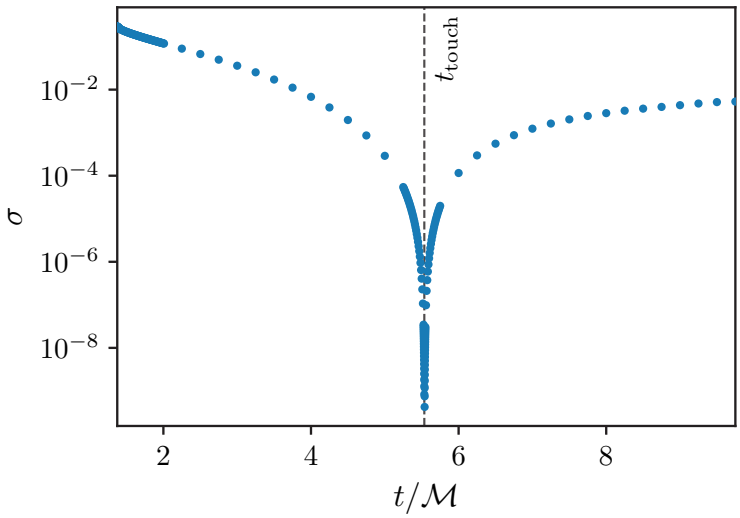


Figure 4.9: Fitted values for σ parameter of the wrapped Cauchy distribution. The dashed line denotes the time instance where $t = t_{\text{touch}}$.

4.6 Conclusions

A binary black hole merger can be understood as a merger of MOTSs. During the merger, the two MOTSs $\mathcal{S}_1$ and $\mathcal{S}_2$ associated with the two progenitor black holes approach each other and eventually merge with the inner MOTS $\mathcal{S}_{\text{inner}}$ associated with the remnant black hole. At the merger time, $\mathcal{S}_{\text{inner}}$ is seen to have a cusp. This can be used not only to connect properties such as mass and multipole moments of the progenitors with those of the remnant black hole, but to obtain a detailed record of the time evolution of these parameters.

We have investigated this for head-on collisions of two non-spinning black holes. Generally, the time evolution of the mass and mass multipole moments is smooth except at the time t_{touch} where the two individual MOTSs merge to yield the remnant. There is a discontinuous jump at t_{touch} corresponding to a δ-function singularity of the curvature at the cusp. The magnitude of these discontinuities is closely related to basic topology of surfaces. The

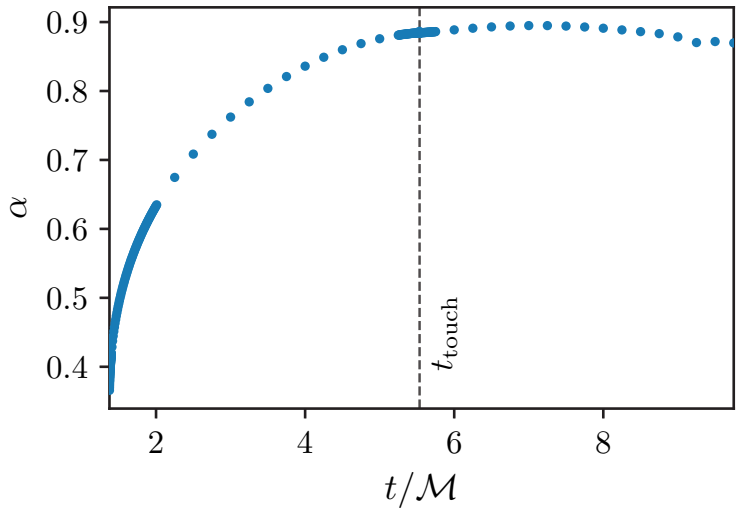


Figure 4.10: Fitted values for the scaling parameter α of the wrapped Cauchy distribution. The dashed line denotes the time instance where $t = t_{\text{touch}}$.

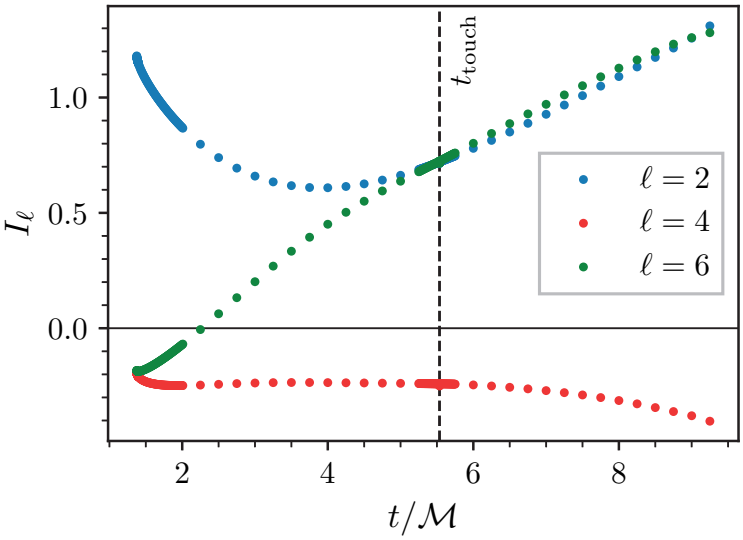


Figure 4.11: Multipole moments of $\mathcal{R}_{12}$ after subtraction of our model. The dashed line denotes the time instance where $t = t_{\rm touch}$.

jump in the mass results just from the additivity of the area, while the jump in the higher multipole moments is determined by the Gauss invariant (which controls the strength of the δ-function singularity in the curvature). Away from t_{touch}, this δ-function is rapidly "smeared" over $\mathcal{S}_{\text{inner}}$, and is not detectable at late times. In this sense, accurate numerical calculations of the MOTSs at t_{touch} are essential for understanding this phenomenon. While these jumps are only part of the total change in these quantities starting from the progenitors at early times to the final remnant black hole, these are often the bigger effects. Consider for example the area of the black holes. The area, and therefore the irreducible mass generally increases gradually during the smooth evolution of the MOTSs. On the other hand the final mass is less than the sum of the initial masses. This discrepancy is explained by the larger discontinuous decrease at t_{touch}.

Looking ahead, the most interesting aspect is to investigate the effects of angular momentum (and its higher multipoles), and to go away from axisymmetric configurations. To illustrate some of these issues, let us briefly consider angular momentum. For spinning black holes, the angular momentum is determined by a 1-form $\widetilde{\omega}_a$ on the MOTS. Integrals of $\widetilde{\omega}_a$ contracted with suitable axial symmetry vector fields ϕ^a yields the components of the angular momentum vector $\mathbf{J}$. At t_{touch}, we would therefore have the vectors $\mathbf{J}_1, \mathbf{J}_2$ for the individual horizons $\mathcal{S}_1, \mathcal{S}_2$ respectively, and $\mathbf{J}$ for $\mathcal{S}_{\text{inner}}$. How would $\mathbf{J}_{1,2}$ be related with $\mathbf{J}$ at t_{touch}? Are there any topological restrictions on $\widetilde{\omega}_a$ such as the Gauss invariant? An important topological property of closed surfaces is that the space of harmonic 1-forms is connected with the genus. Thus on spheres, there exist no harmonic 1-forms. Does this fact lead to restrictions on $\widetilde{\omega}$? Moreover, the notion of angular momentum in classical mechanics depends on the center of mass; there are additional terms in the angular momentum when the reference point differs from the center of mass. We have seen that the mass-dipole moment at t_{touch} is non-vanishing, and this might be interpreted as a jump in the center of mass. Does this have any implications for angular momentum? How would the direction of $\mathbf{J}$ be related with the directions of $\mathbf{J}_{1,2}$ when they are not aligned?

Next, in the case we have studied in this chapter, the presence of symmetries leads to several simplifications. For example, $\mathcal{S}_1$ and $\mathcal{S}_2$ remain exactly axisymmetric and they touch at the poles. This will not hold for

more general configurations. Finally, an important physical effect which has been extensively studied is the recoil velocity produced by the merger. This has important astrophysical implications, and very large recoil velocities have been found for certain configurations. Would it be possible to obtain a recoil velocity from these considerations?

Finally we mention possible connections with aspects of binary black hole waveforms. Analyses of numerical relativity waveforms have shown intriguing properties at the merger. Of relevance to us are the results in [70] which analyze different modes of the gravitational wave signal. In the inspiral regime, post-Newtonian theory predicts amplitudes of the various modes in terms of the binary parameters. Interestingly, it is shown in [70] that these amplitudes are often preserved across the merger. This fact then can be used to predict the amplitude of ringdown modes in terms of the progenitor binary system parameters. It is plausible that there is a link between this behavior of the waveforms with those of the source multipole moments that we have studied here. The horizon dynamics of $\mathcal{S}_{\text{inner}}$ and especially those of the outermost horizon also contain imprints of the ringdown modes [128, 212, 176]. The connection between these various complementary aspects of the merger regime could be useful in gravitational wave astronomy and tests of general relativity, especially regarding black hole spectroscopy [57, 58, 115].

4.7 Appendix: Ricci scalar at the poles

We compute here the Ricci scalar of an axisymmetric MOTS $\mathcal{S}$ at its two poles. Dealing with the coordinate singularity requires some work, which we describe here. Note that we assume $\mathcal{S}$ to be an immersed sphere such that it necessarily possesses two poles on the z axis (as opposed to MOTSs of toroidal topology which we do not consider here). The steps to arrive at the final expression are very similar to those taken in Ref. [69] and we shall repeat just the cornerstones.

We work in coordinates (ρ, φ, z) related to our Cartesian coordinates via $\rho^2 = x^2 + y^2$ and $\tan\varphi = y/x$ such that $\partial_\varphi = -y\partial_x + x\partial_y$ is the rotational Killing field. The (ρ, z) half-plane for constant φ can then simply be taken

as the $(x > 0, z)$ half-plane. Let $\gamma(s)$ be a curve in this (ρ, z) half-plane such that $\mathcal{S}$ is the surface of revolution of γ around the z axis. We take the parameter s to be the proper length of γ measured from one of the two poles. The tangent $T := \dot{\gamma} := \frac{\partial \gamma}{\partial s}$ then has unit length, $T^a T_a = 1$. Let further R be the circumferential radius defined by

$$2\pi R = \int_0^{2\pi} \sqrt{h(\partial_\varphi, \partial_\varphi)}\, d\varphi\,,$$

i.e. $R^2 = h_{\varphi\varphi}$. Taking now (s, φ) as coordinates on $\mathcal{S}$, the induced metric on $\mathcal{S}$ becomes

$$q_{AB} = \begin{pmatrix} 1 & 0 \\ 0 & R^2 \end{pmatrix}. \tag{4.52}$$

This form can be used to find a simple expression for the intrinsic scalar curvature $\mathcal{R}$ of $\mathcal{S}$, i.e. with the Ricci tensor $R_{AB} := R^C{}_{ACB}$ we find

$$\begin{aligned} \mathcal{R} &= q^{ss} R_{ss} + q^{\varphi\varphi} R_{\varphi\varphi} \\ &= 2 q^{ss} R^{\varphi}{}_{s\,\varphi s} \\ &= 2\left(\partial_s^2 \ln R + (\partial_s \ln R)^2\right) \\ &= -2\frac{\ddot{R}}{R}\,, \end{aligned} \tag{4.53}$$

where as before the dot represents differentiation with respect to s.

The (normalized) left-hand normal to γ can be written as (see again [69])

$$\begin{aligned} N^\flat &= \sqrt{\bar{h}}(-T^z\, d\rho + T^\rho\, dz) \\ N &= \frac{1}{\sqrt{\bar{h}}}\left(-(h_{z\rho}T^\rho + h_{zz}T^z)\partial_\rho + (h_{\rho\rho}T^\rho + h_{\rho z}T^z)\partial_z\right), \end{aligned}$$

where $(N^\flat)_a := N_a = \bar{h}_{ab} N^b$, $\bar{h}_{ab}$ is the 2-metric on the half-plane and $\bar{h}$ its determinant. We assume γ to be parametrized such that N is pointing in the outside direction and we will call $\gamma(s = 0)$ the *north pole* of $\mathcal{S}$. With $\bar{D}_a$ denoting the covariant derivative compatible with $\bar{h}_{ab}$, we write the acceleration

$$\bar{D}_T T = \kappa N\,.$$

The expansion of the outgoing and ingoing null normals can then be expressed as, respectively,

$$\Theta_{(\ell)} = k_u + (-\kappa + \bar{D}_N \ln R)$$
$$\Theta_{(n)} = k_u - (-\kappa + \bar{D}_N \ln R)\,.$$

Here, $k_u = q^{ij}K_{ij}$ is the trace of the extrinsic curvature of $\mathcal{S}$ with respect to its timelike normal u. For $\mathcal{S}$ to have vanishing expansion, we must hence have

$$\kappa = \kappa^{\pm} = k_u \pm \bar{D}_N \ln R\,. \tag{4.54}$$

If $\kappa = \kappa^+$, then $\mathcal{S}$ is a MOTS. Note that we can reverse the direction of γ, i.e. $\gamma(s=0)$ is the *south pole* of $\mathcal{S}$ and the outward normal is $-N$. Then, $\mathcal{S}$ is a MOTS if $\kappa = \kappa^-$.

To compute the value of $\mathcal{R}$ at the poles, we let $s \to 0$ and use $\kappa = \kappa^+$ or $\kappa = \kappa^-$ for the north and south pole, respectively. To this end, first let $R_T := T^a \bar{D}_a R$, $R_{TT} := T^a T^b \bar{D}_a \bar{D}_b R$, and so on. Then

$$\begin{aligned}\ddot{R} = \dot{R}_T &= T^a \partial_a (T^b \bar{D}_b R) \\ &= (T^a \bar{D}_a T^b)\bar{D}_b R + T^a T^b \bar{D}_a \bar{D}_b R \\ &= \kappa R_N + R_{TT}\end{aligned}$$

As shown in Ref. [69], κ remains finite and nonzero in the limit $s \to 0$. However, R_N, R_{TT} and R vanish. We hence evaluate this limit with L'Hospital's rule using

$$\dot{R}_N = -\kappa R_T + R_{TN}$$
$$\dot{R}_{TT} = 2\kappa R_{TN} + R_{TTT}$$

and noting that $R_{TN} = 0$ for $s = 0$. With these identities and using [69]

$$\kappa_0^{\pm} = \lim_{s\to 0} \kappa^{\pm} = \frac{1}{2}\left(\frac{R_{TN}}{R_T} \pm k_u\right)\,, \tag{4.55}$$

we finally have

$$\begin{aligned}\mathcal{R}_0 := \lim_{s\to 0}\mathcal{R} &= \lim_{s\to 0}\left(-2\frac{\ddot{R}}{R}\right)\\ &= -\kappa_0^{\pm}\left(\frac{R_{TN}}{R_T}\mp k_u\right) - 2\frac{2\kappa^{\pm}R_{TN}+R_{TTT}}{R_T}\\ &= \frac{1}{2}k_u^2 - 2\frac{R_{TTT}}{R_T}\,.\end{aligned} \tag{4.56}$$

Note that $\mathcal{R}_0$ does not depend on the direction of the normal, i.e. it takes on the same value whether $\Theta_{(\ell)} = 0$ or $\Theta_{(n)} = 0$. In other words, if two arbitrary axisymmetric MOTSs $\mathcal{S}_1$ and $\mathcal{S}_2$ share a common point $p \in \mathcal{S}_1 \cap \mathcal{S}_2$ and p lies on the symmetry axis, then the scalar curvature $\mathcal{R}$ of $\mathcal{S}_1$ and $\mathcal{S}_2$ is equal at p.

4.8 Appendix: Derivatives of the Ricci scalar near the poles

In the main text, we have seen that the mean curvature and the Ricci scalar of $\mathcal{S}_1$ and $\mathcal{S}_2$ agree at the point of contact, i.e. at the north pole of $\mathcal{S}_2$ and the south pole of $\mathcal{S}_1$. Furthermore, the curvatures of $\mathcal{S}_1$ and $\mathcal{S}_2$ remain smooth at all times, even at t_{touch}. Evidence for this is presented in Fig. 4.12. In these plots, the Ricci scalars of $\mathcal{S}_1$ and $\mathcal{S}_2$ and their first three non-vanishing even derivatives with respect to proper length s are shown as function of time. They are evaluated at the facing poles, which coincide exactly at t_{touch}. The top-left plot shows again the agreement of the Ricci scalars at the point of contact. The remaining three plots are respectively the second, fourth and sixth derivatives, showing the smoothness of the Ricci scalars. Note that the odd derivatives of these functions must vanish identically due to symmetry. Finally, Fig. 4.13 shows $\mathcal{R}_{\text{inner}}$ as function of normalized proper length $\tilde{s}$ and its first two derivatives.

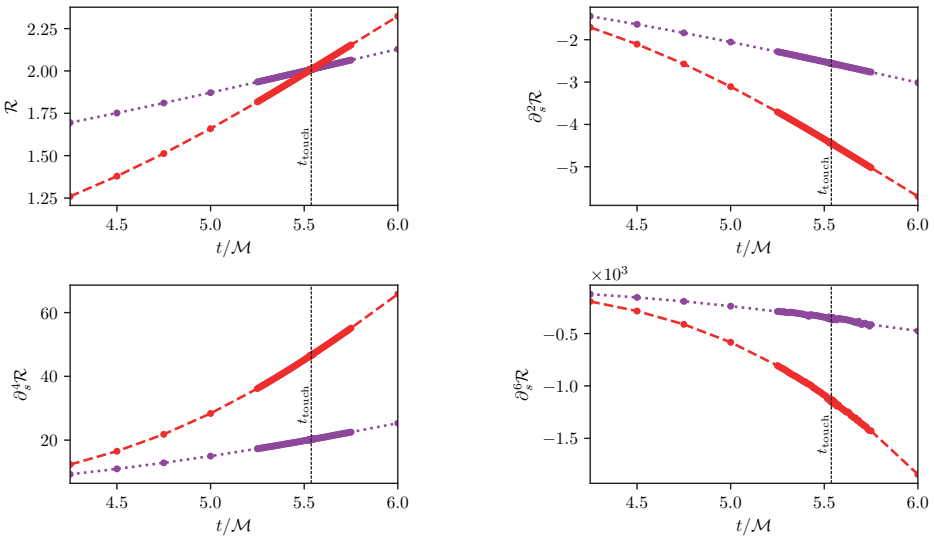


Figure 4.12: Even derivatives of the Ricci scalar $\mathcal{R}$ of $\mathcal{S}_{1,2}$ evaluated at the facing poles and plotted as a function of time. The derivative is taken with respect to the proper length parameter s. The top-left panel shows that the values of $\mathcal{R}$ coincide at the time when $\mathcal{S}_{1,2}$ touch. Compare also with Fig. 4.6.

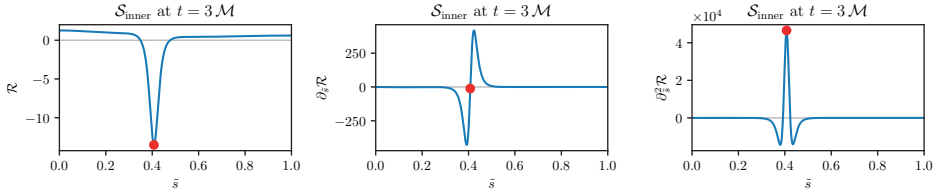


Figure 4.13: Ricci scalar $\mathcal{R}$ of $\mathcal{S}_{\text{inner}}$ as function of the proper length parameter $\tilde{s}$ and its first two derivatives. The red dot marks the location of the neck.


5

Chapter 5

Robust and Fast Parameter Estimation for Gravitational Waves from Binary Neutron Star Merger Remnants

Publication status

This chapter is based on the following paper:

> **S. Vretinaris**, G. Vretinaris, C. Mermigkas, M. Karamanis, and N. Stergioulas, "Robust and fast parameter estimation for gravitational waves from binary neutron star merger remnants," arXiv: 2501.11518, *Physical Review D* **113**, 024012, 2026. DOI: 10.1103/g1qs-j74x.

Contributions

I contributed to the formulation of the problem, development of the method, implementation of the code, simulations, data analysis, production of figures, interpretation of results, and manuscript writing.

Abstract

We present a robust and efficient methodology for parameter estimation of gravitational waves generated during the post-merger phase of binary neutron star mergers. Our approach leverages an analytic waveform model combined with empirical relations to predict prior ranges for the post-merger frequencies based on measurements of the chirp mass and effective tidal deformability in the inspiral phase. This enables robust inference of the main features of the post-merger spectrum, avoiding possible multi-modality induced by wide priors. Using waveforms derived from numerical relativity, we systematically validate our model across a broad spectrum of neutron star equations of state and mass configurations, demonstrating high fitting factors. Our method can be applied in future detections of gravitational waves from the post-merger phase with third-generation gravitational wave observatories. Furthermore, by integrating the Preconditioned Monte Carlo sampling method within the pocoMC framework, we achieve substantial computational acceleration compared to conventional Bayesian techniques.

5.1 Introduction

Gravitational Waves (GWs) have revolutionized our exploration of physical phenomena in the universe since their first direct observation in 2015 [4]. Propagating as ripples in spacetime, these waves carry invaluable information about their astrophysical sources, providing a new way to study cosmic processes. Among the various sources of GWs, Binary Neutron Star (BNS) mergers hold a special place due to their connection to multi-messenger astronomy [264, 257, 164, 19, 91] and the rich physics they reveal about the properties of neutron stars and their remnants [52, 31, 130, 264, 32, 18]. Throughout the initial three observation runs, the LIGO-Virgo-KAGRA (LVK) collaboration [3, 8, 27, 28] has documented 90 detections, predominantly involving binary black hole systems, with only two being BNS systems [7]. The ongoing O4 phase is seeing an increase in the number of GW candidate events. As upgrades are made to existing detectors and third-generation detectors, such as Cosmic Explorer [278, 125] and Einstein

Telescope [158, 74], are being developed, it is anticipated that a larger volume of signals will be detected, including more BNS systems [6, 29, 18]. Here, we focus on parameter estimation of post-merger signals of BNS systems, which could provide new insights into the Equation of State (EOS) of neutron star matter. New designs for dedicated observatories like NEMO [9] and the High Frequency design proposed in [315] are motivated by the significance of post-merger waveform signals.

The study of BNS mergers is a rapidly evolving field, with recent advancements in numerical simulations and astrophysical modeling shedding light on their various properties. The inspiral part of the GW signal provides important constraints about the EOS up to the central densities of the individual neutron stars involved in the merger (e.g. [83, 166, 87, 165]), while the post-merger part will allow us to set EOS constraints closer to the maximum allowed densities, which cannot be probed by the inspiral phase, see [33, 46, 30, 54, 154, 110, 76, 175, 50, 75, 93, 75, 121] and references therein.

Gravitational-wave spectra obtained from numerical-relativity simulations consistently exhibit distinct, characteristic peaks, which originate from oscillatory or transient dynamics of the surviving remnant and correlate with the bulk properties of the BNS merger remnant [281, 160, 45, 287, 52, 46, 220, 90, 244, 223, 126, 174]. The most important of these peaks, can be labeled as f_{2-0}, $f_{\rm spiral}$, $f_{\rm peak}$, f_{2+0}, in increasing order of frequency [281, 45]. The dominant peak, denoted as $f_{\rm peak}$ (or sometimes f_2), arises from the fundamental quadrupole oscillation excited in the remnant at the moment of merger. The peaks at frequencies f_{2-0} and f_{2+0} are combination tones, that arise due to the quasi-linear coupling between quadrupole and the quasi-radial oscillation modes of the remnant. The peak at the frequency $f_{\rm spiral}$ is due to transient spiral arms formed immediately after merger, that have a slower rotation rate than the core of the remnant. For further details on the various characteristics of the post-merger spectrum, see the reviews [53, 49, 47, 264, 131, 245].

The structure of the remnant produced in a binary neutron star merger is tightly related to the bulk properties of the individual stars before merger. Thus, the post-merger frequency spectrum does not only carry the imprint of the properties of the remnant, but also of the properties of the individual

stars before merger and a number of empirical relations for the dominant post-merger frequencies have been constructed [268, 271, 215, 43, 44, 160, 286, 45, 287, 46, 130, 126, 56, 71, 46, 77, 302, 75], connecting them to NS properties in the inspiral phase.

Analytical time-domain models of post-merger GW emission were presented, e.g., in [52, 71, 77, 119, 273]. In [120] a hierarchical model was used to predict postmerger spectra in the frequency domain. The training set was extended in [221], where a model for the amplitude of the post-merger spectrum was constructed, using artificial neural networks. Frequency domain models have been developed to encompass both the inspiral and post-merger parts of the signal [309, 78, 241]. Faithful models of the post-merger GW spectrum could enable more accurate parameter estimation for the post-merger part of the GW signal, revealing important information about the microphysics involved in these events [168, 299, 217, 255], including thermal effects [126, 249, 126, 65, 206, 250], and the nature of phase transitions that may occur at extreme densities [136, 124, 237, 64, 210, 48, 209, 238].

The detectability of post-merger signals relies on a signal-to-noise ratio (SNR) greater than about 8, considering high-frequency noise artefacts and false alarm criteria [219]. Current detectors are not sufficiently sensitive at high frequency to allow for a direct detection of the post-merger phase [5, 181, 142]. However, future detectors are expected to allow for the first observations of post-merger signals [316], enabling a deeper understanding of the physics behind these events. The classification of post-merger signals is also an area of interest, as it allows for the differentiation between systems that experience prompt collapse or the formation of a neutron star remnant [10, 291, 240, 169]. New numerical codes using e.g. a moving-mesh approach [193] are expected to allow for more accurate determination of the post-merger dynamics.

An important aspect of extracting information from the post-merger phase is the parameter estimation of the source properties, see e.g. [119]. The conventional approach to Bayesian inference typically adopts an impartial stance toward the subject being studied, relying on statistical methods like Markov Chain Monte Carlo (MCMC) algorithms to highlight prominent features [119]. However, this approach encounters a limitation when there is a lack of substantial prior knowledge.

Here, we aim to address this knowledge gap, by taking advantage of observational properties that can be extracted from the inspiral phase and using empirical relations (connecting pre- and post-merger phases) to constrain the priors in the post-merger phase, leading to a more robust parameter estimation. In this way, we have successfully achieved accurate and robust parameter estimation, even when resampling is performed. This accomplishment has led to a significant reduction in computational time for the sampling procedure, by utilizing the Preconditioned Monte Carlo method.

This chapter is organised as follows. In Section 5.2 we review the numerical relativity data used in this work, and we present the analytic post-merger waveform model used in the parameter estimation. Furthermore we provide new empirical relations, and present our prior distributions and the preconditioned Monte Carlo sampling technique. In Section 5.3 we present the results of our parameter estimations for various cases of signal-to-noise ratio (SNR) and for post-merger spectral types. The chapter concludes in Section 5.5, which summarizes our findings.

Conventions

To refer to a particular post-merger waveform we will use the format "EOS-Mx", where "EOS" is the equation-of-state name and "x" stands for the average mass of the binary (the mass of either component when they are isolated), measured in solar masses.

5.2 Methods

5.2.1 Numerical Waveforms and Injections

In this study we focus mainly on *equal mass* BNS mergers and utilize a total of 15 waveforms. Of these, 13 are waveforms from the publicly available numerical relativity catalog `CoRe` [139, 1] and the remaining two waveforms are from [273]. The waveforms are provided as $h_+(t)$ and $h_\times(t)$, the plus

Table 5.1: Different post-merger waveforms included in this study. The first column lists their label in the CoRe database [139, 1] (except for the last two entries). The remaining columns list the EOS name, the mass ratio and either the mass of individual components (when $q = 1$) or their average mass (when $q \neq 1$).

Label	**EOS**	**q**	**(Average) Mass [$M_\odot$]**	**References**
THC:0036:R03	SLy	1.0	1.350	[244]
THC:0019:R05	LS220	1.0	1.350	[246, 248]
BAM:0088:R01	MS1b	1.0	1.500	[139, 1]
THC:0002:R01	BHBlp	1.0	1.300	[246, 248]
THC:0011:R01	DD2	1.0	1.250	[246, 248]
BAM:0070:R01	MS1b	1.0	1.375	[109]
BAM:0065:R03	MS1b	1.0	1.350	[55]
THC:0010:R01	DD2	1.0	1.200	[246, 248]
BAM:0002:R02	2H	1.0	1.350	[55]
BAM:0053:R01	H4	1.5	1.375	[111]
BAM:0124:R01	SLy	1.5	1.250	[109]
BAM:0090:R02	MS1b	1.0	1.600	[139, 1]
BAM:0092:R02	MS1b	1.0	1.700	[139, 1]
Soultanis et al.	MPA1	1.0	1.200	[273]
Soultanis et al.	MPA1	1.0	1.550	[273]

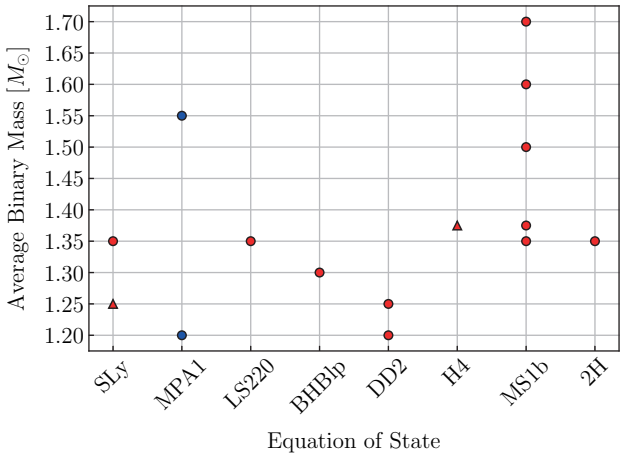


Figure 5.1: EOS and average binary mass configurations used in this work. Red markers correspond to waveforms from the CoRe database [1, 139] and blue markers correspond to waveforms from [273]. Circles indicate equal-mass ($q = 1$) mergers, whereas triangles indicate the two cases with $q = 1.5$.

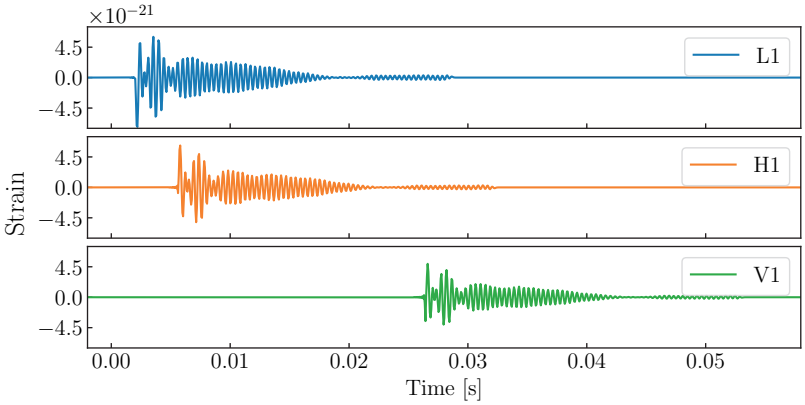


Figure 5.2: An example of a projected waveform onto the L1, H1 and V1 detector network. The EOS is BHBlp and the average binary mass is $1.3M_{\odot}$. The $t = 0s$ marks the time of first detection. In this case it happens for the Livingston detector.

(+) and cross (×) polarizations, respectively, as extracted in the numerical-relativity simulations.

Notice that our set of models is an expanded set of the 9 waveforms that were used in [119] and includes those. Table 5.1 lists the detailed properties of each numerical waveform, which are their label in the CoRe database (otherwise the label is Soultanis et al.), the EOS name, the mass ratio $q \geq 1$, the average mass of the two binary components and a reference to a publication where more details on the particular simulation were presented[1] Two of our chosen waveforms do not have equal mass, but have a relatively large mass ratio of $q = 1.5$. These were included to test the robustness of our results to the mass ratio. Fig. 5.1 shows the set of our selected numerical waves in the configuration space of average binary mass vs. EOS. Equal-mass cases are shown as circles, while the two $q = 1.5$ cases are shown as triangles.

We interpolated the numerical waveforms to a uniform time step, which corresponds to a sampling frequency of 16384 Hz. To isolate the post-merger

[1]For two of the CoRe waveforms, no particular publication exists and we cite the CoRe database.

part of the signal, we truncate the waveforms at the *merger time*, defined as

$$t_c = \max\{h_+^2(t) + h_\times^2(t)\}. \tag{5.1}$$

To facilitate a direct comparison with the previous results of [119], we zero-pad both sides of all signals to a total duration of 0.125s, with the merger time roughly centered in the middle[2]

Next, we project the post-merger parts of the numerical waveforms onto a *reference* three-detector network, consisting of the advanced Livingston (L1), the advanced LIGO Hanford (H1) and the advanced Virgo (V1) detectors. Fig. 5.2 displays, as an example, the projected waveforms onto the L1, H1 and V1 detectors for the BHBlp EOS with $M = 1.3M_\odot$. The resulting strain $h(t)$ for each detector is then *injected* into Gaussian noise, assuming that the detectors are operating at their design sensitivities. For the injections, we are using the `BILBY` [26] software and assume a fixed sky position[3].

We stress that we do not expect realistic detection rates for the BNS post-merger phase with such a detector network, comprising only second-generation (2G) detectors at their design sensitivity. Instead, we use it only as a well-defined reference noise level. Then, we vary the distance of the source to achieve SNR values that are in the range of 8-50, which are expected to be achieved with different designs of third-generation (3G) detectors at distances of $\sim$ 200Mpc, depending on the assumed mass and EOS, see [278] for a detailed study. *Therefore, our results are relevant for realistic searches with 3G detector networks*[4].

5.2.2 Analytic Post-merger Waveform Model

To perform parameter estimation for post-merger injections, we will use an analytic waveform model, which is an extension of the model presented

[2]P. Easter, personal communication.

[3]We chose the same fixed sky position and other extrinsic characteristics as in [119] (P. Easter, personal communication.), namely $\psi = 0°$, right ascension $= 0°$ and declination $= 0°$.

[4]Except for the somewhat different frequency dependence of the noise curves of 3G detectors, compared to 2G detectors, in the frequency range between 1.5 and 3.5 kHz.

in [119] and includes four damped oscillators and a linear drift for each frequency. Our analytic post-merger waveform model reads

$$\begin{aligned} h(\boldsymbol{\theta}, t) &= h_{+}(\boldsymbol{\theta}, t) - i h_{\times}(\boldsymbol{\theta}, t) \\ &= \sum_{j=1}^{4} \left[h_{j,+}(\boldsymbol{\theta}, t) - i h_{j,\times}(\boldsymbol{\theta}, t) \right], \end{aligned} \tag{5.2}$$

with the individual (+) polarization components written as

$$h_{j,+}(\boldsymbol{\theta}, t) = A_j \exp\left[-\frac{t}{T_j} \right] \cos\left[2\pi f_j t (1 + a_j t) + \psi_j \right]. \tag{5.3}$$

In Eq. 5.3, $\boldsymbol{\theta}$ is a vector containing the intrinsic parameters of the model

$$\boldsymbol{\theta} = \{A_j, T_j, f_j, a_j, \psi_j\} \quad \text{for} \quad j \in [1, 4],$$

where A_j are the individual amplitudes, f_j the frequencies, T_j the damping times, ψ_j the phases, and a_j the frequency drift coefficients. The (×) polarization components are then produced by applying a $\pi/2$ phase shift on the respective (+) polarization components.

Compared to the analytic model used in [119], our model has two main differences:

- we have added a fourth oscillator, f_4 at frequencies $> f_{\text{peak}}$, which could correspond to e.g. f_{2+0} [281] when this combination tone is present with sufficient amplitude, or to some other high-frequency peak in the post-merger spectrum, and
- we are using the amplitudes A_j as free parameters, whereas in [119] an additional constraint of $A_1 + A_2 + A_3 = H$ was imposed, where H was some global amplitude.

In our investigation, we found that adopting the $A_1 + A_2 + A_3 = H$ constraint imposes a *simplex* geometry that could potentially introduce *degeneracies* [59] in the posterior distribution of parameter estimation.

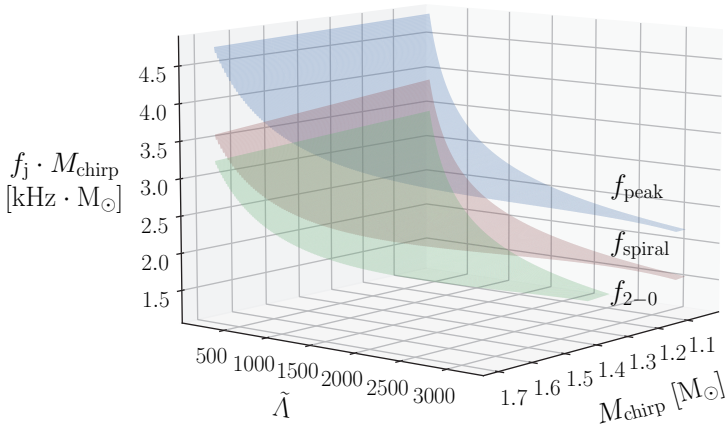


Figure 5.3: Empirical surfaces for the three most important post-merger frequencies vs. $M_{\rm chirp}$ and $\tilde{\Lambda}$. The blue surface corresponds to $f_{\rm peak}$, the red to $f_{\rm spiral}$ and the green to f_{2-0}.

5.2.3 Empirical Relations and Classification

Here, we use empirical relations for some of the post-merger frequencies to tighten the prior distribution in the Bayesian parameter estimation. As we show in Sec. 5.3, this leads to a robust parameter estimation, overcoming the problem of multi-peak distributions obtained when a flat prior distribution over a wide frequency range is assumed for the post-merger peaks (see, e.g., Figs. C1, C5 in [119]).

For the primary post-merger frequency, a useful empirical relation for $f_{\rm peak}$ was derived in [302]:

$$f_{\rm peak} M_{\rm chirp} = 1.392 - 0.108 M_{\rm chirp} + 51.70 \tilde{\Lambda}^{-1/2}, \tag{5.4}$$

where $M_{\rm chirp} = (M_A M_B)^{3/5}/(M_A + M_B)^{1/5}$, is the chirp mass of the binary,

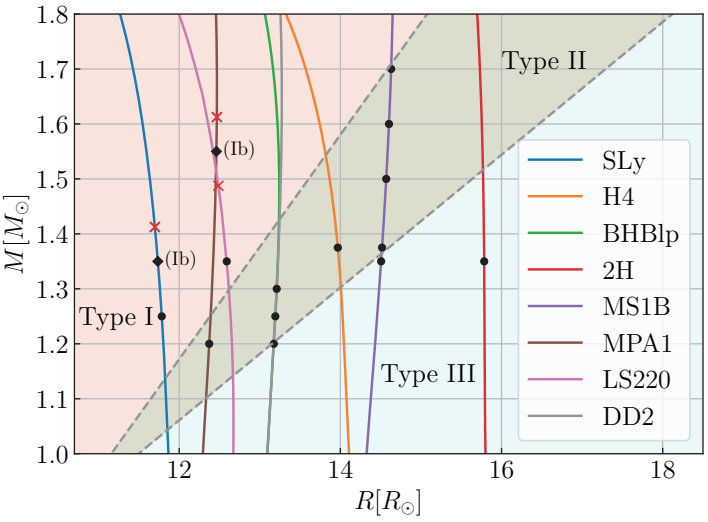


Figure 5.4: Mass-Radius diagram of equilibrium sequences constructed with the different EOS in Table 5.1. For each model BNS system used in this study, the average mass of its two components is shown with black markers. We use the classification scheme of [45, 302] to indicate models that give rise to post-merger spectra of Type I (light red region), Type II (light green region) and Type III (light blue region), see text for details. Red crosses denote 1/2 of the threshold mass to prompt collapse. Models close to the threshold mass (shown here as Type Ib) exhibit a different behavior, see [273].

with M_A, M_B being the individual masses and

$$\tilde{\Lambda} = \frac{16}{13} \frac{(M_A + 12M_B)\, M_A^4 \Lambda_A + (M_B + 12M_A)\, M_B^4 \Lambda_B}{(M_A + M_B)^5} \tag{5.5}$$

is the mass-weighted tidal deformability, with Λ_A, Λ_B being the individual dimensionless tidal deformabilities, respectively. This empirical relation has a maximum residual of 0.302kHz and $R^2 = 0.985$. It was constructed using a large set of both equal-mass and un-equal-mass models and is thus relatively insensitive to the mass ratio q.

We use the particular empirical relation in Eq. (5.4), because both $M_{\rm chirp}$ and $\tilde{\Lambda}$ can be constrained by GW observations in the inspiral phase of a BNS merger. Using the data in Table II of [302] and a standard Levenberg-Marquardt nonlinear least-squares minimization algorithm, we construct similar empirical relations for two additional post-merger peaks:

$$f_{2-0} M_{\rm chirp} = 0.558 - 0.406 M_{\rm chirp} + 48.696 \tilde{\Lambda}^{-1/2}, \tag{5.6}$$

with a maximum residual of 0.362kHz and $R^2 = 0.941$, and

$$f_{\rm spiral} M_{\rm chirp} = 1.2 - 0.442 M_{\rm chirp} + 45.819 \tilde{\Lambda}^{-1/2}, \tag{5.7}$$

with a maximum residual of 0.461kHz and $R^2 = 0.944$. Fig. 5.3 displays all three empirical relations in Eqs. (5.4), (5.7) and (5.6) as two-dimensional, non-intersecting surfaces.

Post-merger spectra can be classified into three main types (I, II, and III), based on the relative strength of the secondary frequencies f_{2-0} and $f_{\rm spiral}$ [45, 302]. Fig. 5.4 shows the mass-radius diagram of equilibrium sequences constructed with the different EOS in Table 5.1. For each model BNS system used in this study, the average mass of its two components is shown with black markers. The region of models with Type I post-merger spectra is shown in light red, with Type II spectra in light green, and with Type III spectra in light blue color.

Notice that the boundaries between the regions of different post-merger types in Fig. 5.4 are not sharp. Rather, there is a continuous transition between the different types. In addition, in [273] a different behavior was noticed for merger remnants that come close to the threshold mass for

prompt collapse[5] (see the red crosses in Fig. 5.4). For these models, both f_{2-0} and $f_{\rm spiral}$ are important, and there is partial overlap of the peaks in the frequency spectrum. This is because the quasi-radial frequency decreases rapidly, as the quasi-radial instability is approached; see [129]. We indicate such models as Type Ib in Fig. 5.4.

Depending on the particular BNS merger, fewer than four damped oscillators may be meaningful for inclusion in the analytic post-merger model of Eq. (5.2). In our Bayesian parameter estimation, we use the classification scheme of [45, 302], as applied in Fig. 5.4, to decide which secondary frequencies will be constrained through their respective empirical relations and which will be assumed to have a uniform prior in a wide frequency range (see Sec. 5.2.5).

5.2.4 Bayesian inference

Our Bayesian parameter estimation is based on the likelihood function

$$\mathcal{L}(d \mid \boldsymbol{\theta}) \propto \exp\big[-\big\langle d(t) - h(\boldsymbol{\theta}, t), d(t) - h(\boldsymbol{\theta}, t)\big\rangle\big], \tag{5.8}$$

where $d(t) = h(t) + n(t)$ is the *injected waveform* in each detector, that is, the sum of the projected numerical waveform $h(t)$ and the noise $n(t)$ into which it is injected. In Eq. (5.8), the *noise-weighted inner product* is

$$\big\langle h, g\big\rangle = 4\mathfrak{Re}\int_0^\infty \frac{\tilde{h}(f)\tilde{g}^*(f)}{S(f)} df, \tag{5.9}$$

where a tilde denotes the Fourier transform, an asterisk denotes complex conjugations and $S(f)$ is the detector's noise power spectral density.

To evaluate the effectiveness of our parameter estimation, we compute the *noise-weighted fitting factor* $\mathcal{F}$ [16] between an injected waveform $d(t)$ and a *posterior* waveform $h(\boldsymbol{\theta}, t)$

$$\mathcal{F}(d(t), h(\boldsymbol{\theta}, t)) \equiv \frac{\big\langle d(t)|h(\boldsymbol{\theta}, t)\big\rangle}{\sqrt{\big\langle d(t)|d(t)\big\rangle\big\langle h(\boldsymbol{\theta}, t)|h(\boldsymbol{\theta}, t)\big\rangle}}. \tag{5.10}$$

[5]The values of the threshold mass for prompt collapse can be found in Table 1 of [51].

For the projected numerical waveforms $h(t)$, we use `pyCBC` [214] to calculate their optimal SNR for a single detector as

$$\rho_{\mathrm{opt},i} = \left\langle h, h \right\rangle^{1/2}, \tag{5.11}$$

and the network SNR as

$$\rho_{\mathrm{opt}} = \left(\sum_{i \in \mathrm{HLV}} \rho^2_{\mathrm{opt},i} \right)^{1/2}. \tag{5.12}$$

As already mentioned in Sect. 5.2, we scale the source distance to obtain a desired network SNR in the range of 8 to 50, corresponding to the expected event rates for a network of 3G detectors.

5.2.5 Informed Gaussian Priors

Gravitational waves from the post-merger phase of individual BNS mergers are expected to be detected for systems that are loud enough to have a post-merger network SNR larger than about 8. This means that the inspiral part of the network will have a significantly larger SNR, allowing for a very accurate determination of the chirp mass M_{chirp}, as well as fairly good constraints on the mass-weighted tidal deformability $\tilde{\Lambda}$ (see [166, 165] for recent estimates).

Since all of the empirical relations we use in this work have the same form

$$f_j M_{\mathrm{chirp}} = c_1 + c_2 M_{\mathrm{chirp}} + c_3 \tilde{\Lambda}^{-1/2}, \tag{5.13}$$

one can derive the general form of the error propagation, i.e. how observational uncertainties in the determination of M_{chirp} and $\tilde{\Lambda}$ result in uncertainties δf_j in the frequencies. Because M_{chirp} will be measured to extremely high precision with 3G detectors, the error in the frequency will be dominated by the error in the determination of $\tilde{\Lambda}$ and thus, to good approximation

$$\delta f_j \simeq \frac{c_3}{2 M_{\mathrm{chirp}} \tilde{\Lambda}^{3/2}} \delta\tilde{\Lambda}. \tag{5.14}$$

For a GW170817-like event, assuming nominal values of $M_{\mathrm{chirp}} = 1.2 M_\odot$, $\tilde{\Lambda} = 300$ and an error of $\delta\tilde{\Lambda} \simeq 20$ for a 3G detector [86], the uncertainties

in determining the three frequencies f_j through their empirical relations are $\simeq 80$ Hz, which is several times smaller than the maximum error in the empirical relations.

Hence, as a proof of principle, we will assume that the binary system's $M_{\rm chirp}$ and $\tilde{\Lambda}$ have been obtained from the inspiral part of the waveform with accuracy better than the residuals of the empirical relations in Eqs. (5.4), (5.6) and (5.7). In future work, realistic distributions for the measured $M_{\rm chirp}$ and $\tilde{\Lambda}$, depending on the expected SNR for next-generation detectors, should be incorporated in this analysis.

Given $M_{\rm chirp}$ and $\tilde{\Lambda}$ from the inspiral phase, we can calculate the expected frequencies for $f_{\rm peak}$, f_{2-0} and $f_{\rm spiral}$ from the empirical relations in Eqs. (5.4), (5.6) and (5.7), respectively. These can then be used to set informed priors for some of the frequencies of our analytic model in the Bayesian parameter estimation. Specifically, we assume a Gaussian normal distribution with mean value equal to the prediction of the respective empirical relation and a standard deviation (σ), chosen such that the maximum residual in the empirical relations in Eqs. (5.4), (5.6) and (5.7) is equal to 3σ.

We always assume a Gaussian normal prior distribution for the $f_{\rm peak}$ frequency and a uniform prior in the range $(f_{\rm min}, 5{\rm kHz})$, where $f_{\rm min} = f_{\rm peak} + 0.3{\rm kHz}$ for the frequency f_4 of the fourth damped oscillator (with frequency $> f_{\rm peak}$) in our analytic model (5.2). For the other two frequencies, we use the classification scheme discussed in Sec. 5.2.3 to set priors, according to the type of post-merger spectrum, as follows:

- Type I: Gaussian priors, $\mathcal{N}(f_{2-0}, \sigma^2)$ for f_{2-0} and uniform priors $\mathcal{U}(1,5)$[kHz] for $f_{\rm spiral}$.
- Type II: Gaussian priors, $\mathcal{N}(f_{2-0}, \sigma^2)$ for f_{2-0} and $\mathcal{N}(f_{\rm spiral}, \sigma^2)$ for $f_{\rm spiral}$.
- Type III: Gaussian priors, $\mathcal{N}(f_{\rm spiral}, \sigma^2)$ for $f_{\rm spiral}$. For f_{2-0}, we use uniform priors $\mathcal{U}(1,5)$[kHz].

Notice that when we use uniform priors for f_{2-0} or $f_{\rm spiral}$, we adopt the same frequency range as in [119]. The reason we adopt uniform priors in the range 1-5kHz for f_{2-0} or $f_{\rm spiral}$, when these peaks are expected to be

very weak (based on the post-merge type), is not that we expect them to have frequencies outside the range predicted by the corresponding empirical relation, but to give the sampler the opportunity to explore other features of the spectrum, where it may be more meaningful to fit a damped sinusoid than to use it to fit a very weak f_{2-0} or $f_{\rm spiral}$.

The priors for the amplitudes and frequency drift terms in the analytic model are set to $A_j \sim \mathcal{U}(-24, -19)$ and to $\alpha_j \sim \mathcal{U}(-6.4, 6.4)$, respectively. The priors for the remaining intrinsic parameters in $\boldsymbol{\theta}$ are set as in [119].

In Appendix 5.6, we demonstrate, in a specific example, the improvements one can obtain in the parameter estimation when using informed priors, as opposed to the uninformative flat priors used in the methodology of [119].

5.2.6 Preconditioned Monte Carlo Sampling

Because the waveform depends nonlinearly on its frequency-drift and damping parameters, weak post-merger signals can produce non-Gaussian and, in some cases, multimodal posterior distributions. These deviations from Gaussianity can impede the effectiveness of most sampling techniques, including many MCMC and nested sampling variants, to produce independent samples from the posterior distribution. To overcome this challenge, we used the recently developed method of Preconditioned Monte Carlo (PMC) as implemented in the open source `pocoMC` Python package [172, 173]. *PMC utilizes a normalizing flow to iteratively decorrelate the parameters of the posterior distribution*, effectively Gaussianizing it.

PMC targets an annealed version of the posterior, with density given by

$$p_t(\boldsymbol{\theta}|d) \propto \mathcal{L}^{\beta_t}(\boldsymbol{\theta}|d)\pi(\boldsymbol{\theta}), \tag{5.15}$$

where β_t is the parameter controlling the annealing process. A collection of particles gradually transitions from the prior (i.e. $\beta_0 = 0$) to the posterior distribution (i.e. $\beta_T = 1$) through a sequence of reweighting, resampling, and mutation steps. After each iteration, the normalizing flow is fitted to the distribution of the particles and the sampling proceeds in the uncorrelated latent space of the normalizing flow. Since the transformed annealed posterior distribution exhibits almost no correlations in the latent space, the sampling

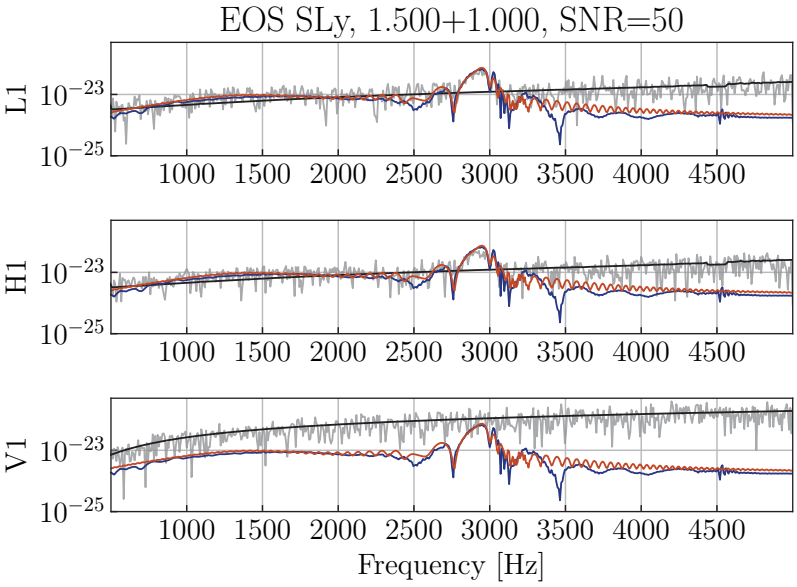


Figure 5.5: Amplitude spectral density (ASD) in $1/\sqrt{\text{Hz}}$ of the injected waveforms (blue lines) in Gaussian noise around the ASD of each of the three detectors comprising the network. Grey curves represent the sum of the noise plus the injected waveforms. Red curves are the maximum likelihood reconstructed waveforms.

performance is substantially increased. The sampling procedure ends when the annealing parameter β reaches the value of unity, which corresponds to the actual posterior distribution.

In our application, `pocoMC` generates the required number of samples almost *an order of magnitude faster* in terms of the total number of likelihood evaluations and the wall clock time, compared to nested sampling using `dynesty` [275].

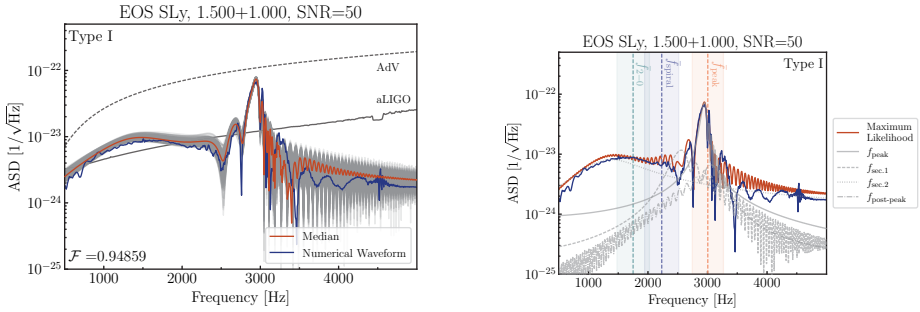


Figure 5.6: Amplitude spectral density (ASD) of gravitational wave signals from a binary neutron star merger with EOS SLy and equal component masses of 1.5 $M_\odot$. The left panel compares the ASD to the sensitivity of aLIGO and AdV detectors, while the right panel analyzes the spectral features of the signal with respect to various characteristic frequencies, at an SNR of 50. The fitting factor $\mathcal{F}$ is noted, suggesting the closeness of the numerical waveform to the median signal.

5.3 Parameter Estimation

The cases studied in this work have a network SNR of 50, 16 and 8. This is motivated by 3G detector designs that can reach post-merger network SNR's up to O(50) at distances of 200Mpc [278]. We start with the case of SNR = 50 and discuss the parameter estimation obtained for four representative cases, one for each of the four types I, Ib, II and III, according to the classification discussed in Sec. 5.2.3.

5.3.1 SNR = 50

Type I (EOS SLy, 1.5+1.0 $M_\odot$)

As a first example, we discuss the case of a 1.5+1.0 $M_\odot$ merger with EOS SLy. This EOS is very soft and the model lies well in the Type I region of Fig. 5.4, i.e. f_{2-0} is the dominant secondary peak in the post-merger

spectrum[6]. Consequently, our scheme leads to Gaussian priors $\mathcal{N}(f_{\rm peak}, \sigma^2)$, $\mathcal{N}(f_{2-0}, \sigma^2)$ centered on the predictions of the empirical relations (5.4) and (5.6) for $f_{\rm peak}$ and f_{2-0}, correspondingly, and uniform priors $\mathcal{U}(1, 5)$[kHz] for $f_{\rm spiral}$. The fourth oscillator, f_4, in our analytic model (5.2) is initialized with uniform priors in the range $(f_{\rm min}, 5{\rm kHz})$, where $f_{\rm min} = f_{\rm peak} + 0.3{\rm kHz}$.

In Fig. 5.5, we show the amplitude spectral density (ASD) of the injected waveforms (blue lines) in Gaussian noise around the ASD of each of the three detectors comprising the network [191, 192]. The gray curves in this plot represent the sum of the noise plus the injected waveforms.

For each run, we identify the posterior sample with the highest value of the log-likelihood returned by `pocoMC` (we the call the corresponding waveform "maximum likelihood reconstructed waveform" in the remainder of the text)[7].

We identify the set of parameters of the model that maximize the likelihood, Eq. (5.8), and display the maximum likelihood reconstructed waveform as a red curve. These agree to a high degree with the injected waveforms.

In the left panel of Fig. 5.6, we display the ASD of 1000 randomly sampled waveforms (gray lines), constructed using the posterior distribution of the inferred parameters of our analytic model of Eq. (5.2). The ASDs of the random reconstructed waveforms cluster around their median (red line), in good agreement with the ASD of the numerical waveform (blue line). Some oscillations can be seen at high frequencies, which have only a small impact on the overall agreement between the reconstruction and the numerical waveform. The fitting factor between the median reconstructed waveform and the numerical waveform is $\mathcal{F} = 0.94859$.

The right panel of Fig. 5.6 shows, in more detail, the individual contributions of the different damped oscillators included in the analytic model of Eq. (5.2). Here, the red line represents the maximum likelihood reconstructed

[6]Because the mass ratio is $q = 1.5$, the amplitude of the $f_{\rm spiral}$ contribution is even weaker than it would have been for an equal mass case, since the $f_{\rm spiral}$ contribution is created by antipodal tidal bulges that are strong emitters when they are symmetric [45].

[7]For a more accurate determination, we would start from this point, and perform a local optimization of the likelihood surface.

waveform, whereas the different gray lines correspond to the individual damped oscillators. Vertical dashed lines represent the predicted frequencies using the empirical relations (5.4), (5.6) and (5.7), for $f_{\rm peak}$, f_{2-0} and $f_{\rm spiral}$, respectively. The colored bands indicate the maximum error for each of the three empirical relations.

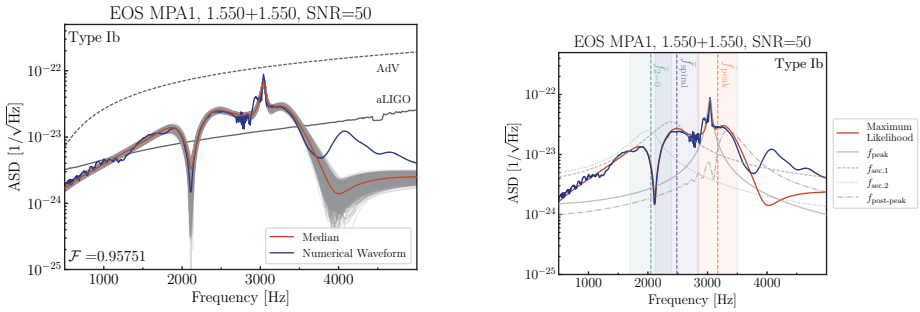


Figure 5.7: Same as Fig. 5.6, but for a 1.55+1.55 $M_\odot$ model with EOS MPA1 (Type Ib post-merger spectrum).

We observe in the right panel of Fig. 5.6 that the damped oscillator corresponding to $f_{\rm peak}$ describes the shape of $f_{\rm peak}$ with high accuracy. On the other hand, $f_{\rm spiral}$ can barely be seen in this spectrum and the damped oscillator corresponding to $f_{\rm spiral}$ (shown as $f_{\rm sec.1}$ in the figure), has converged to a frequency at the upper edge of the frequency band of the $f_{\rm spiral}$ empirical relation, to accommodate another peak in between[8] $f_{\rm spiral}$ and $f_{\rm peak}$. The f_{2-0} appears prominently in this Type I spectrum, and the corresponding damped oscillator (shown as $f_{\rm sec.2}$ in the figure) converges at the left edge of the corresponding empirical frequency band.

At frequencies higher than $f_{\rm peak}$, the strongest peak is very close to $f_{\rm peak}$. In fact, the two features at frequencies somewhat smaller and somewhat higher than $f_{\rm peak}$ seem to form a triplet with the latter. This could be due to modulations induced in $f_{\rm peak}$ by the evolution of the remnant created by

[8]Notice that the classification scheme of [45] has not yet been fully extended to the case of large mass ratios and such additional features still need to be fully explained and taken into account in analytic models.

the unequal mass merger. The fourth oscillator (shown as $f_{\rm post-peak}$ in the figure) converges to a frequency in between the main $f_{\rm peak}$ and the feature at somewhat higher frequency.

Overall, for this Type I model, the $f_{\rm peak}$ and f_{2-0} frequency peaks are well described by the corresponding damped oscillators, whereas the other two oscillators in the analytic model converge close to other prominent features, in a way that the sum of all four (shown in red as the maximum likelihood waveform) achieves a high fitting factor.

Type Ib (EOS MPA1, 1.55+1.55 $M_\odot$)

Next, we discuss a representative case of a Type Ib spectrum. This is a $1.55+1.55 M_\odot$ merger using the MPA1 EOS. Its average mass is close to the threshold mass for prompt collapse to black holes (see Fig. 5.8). It was already discussed in [273] that in such cases, the f_{2-0} peak appears closer to $f_{\rm peak}$ and practically merges with the $f_{\rm spiral}$ frequency, because the quasi-radial frequency f_0 becomes smaller, as the threshold to prompt black hole formation is approached. These are outlier cases, which are at the edge of the maximum error of the empirical relation for f_{2-0}. A second prominent feature of this particular model is the appearance of a late-time ($t-t_{\rm merger} > 10$ms) rotational instability; see Fig. 27 in [273] and the detailed analysis in [265]. In the post-merger GW spectrum, this adds a characteristic narrow peak to the broader $f_{\rm peak}$ of the first 10-15ms after merger.

In the left panel of Fig. 5.7, the 1000 randomly sampled waveforms (gray lines), constructed using the posterior distribution of the inferred parameters of our analytic model of Eq. (5.2) cluster tightly around their median (red line), in good agreement with the ASD of the numerical waveform (blue line), except for very high frequencies, where the detectors are not sensitive. The fitting factor between the median reconstructed waveform and the numerical waveform is $\mathcal{F} = 0.95751$.

In the right panel of Fig. 5.7, we observe that the damped oscillator corresponding to $f_{\rm peak}$ has converged specifically to the narrow $t - t_{\rm merger} > 10$ms contribution of the rotational instability. The broader $f_{\rm peak}$ contribution of the first 10ms after the merger is described mainly by the fourth oscillator, shown as $f_{\rm post-peak}$ in this figure. On the other hand, the first secondary

oscillator ($f_{\rm sec.1}$) describes relatively well the merged f_{2-0} and $f_{\rm spiral}$ frequencies[9]. Finally, the other secondary oscillator ($f_{\rm sec.2}$) has converged to a low-frequency peak that is not part of the classification of [45] (for soft EOSs, $f_{\rm peak}$ can be so high, that additional features of the post-merger GW spectrum can appear at frequencies higher than the highest frequency of the inspiral phase).

Overall, we see that the analytic model of Eq. (5.2), combined with our scheme for choosing the Gaussian priors based on the empirical relations in Eqs. (5.4), (5.6) and (5.7), and the classification scheme of [45] has sufficient flexibility to capture all essential features of the post-merger GW spectrum of Type Ib models, achieving a high fitting factor.

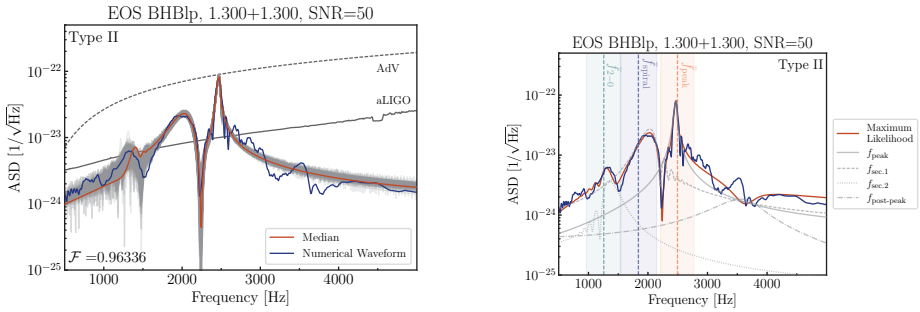


Figure 5.8: Same as Fig. 5.6, but for a 1.3+1.3 $M_\odot$ model with EOS BHBlp (Type II post-merger spectrum).

Type II (EOS BHBlp, 1.3+1.3 $M_\odot$)

Our third example is a 1.3+1.3 $M_\odot$ merger using the BHBlp EOS, which is well inside the Type II region in Fig. 5.4 and, therefore, both secondary peaks $f_{\rm spiral}$ and f_{2-0} are important. In the left panel of Fig. 5.8, the 1000 randomly sampled waveforms (gray lines), constructed using the posterior distribution of the inferred parameters of our analytic model of Eq. (5.2)

[9]We stress that the f_{2-0} peak is not at the center of the empirical band for f_{2-0} in Fig. 5.7, but at the right edge of this band and has merged with $f_{\rm spiral}$.

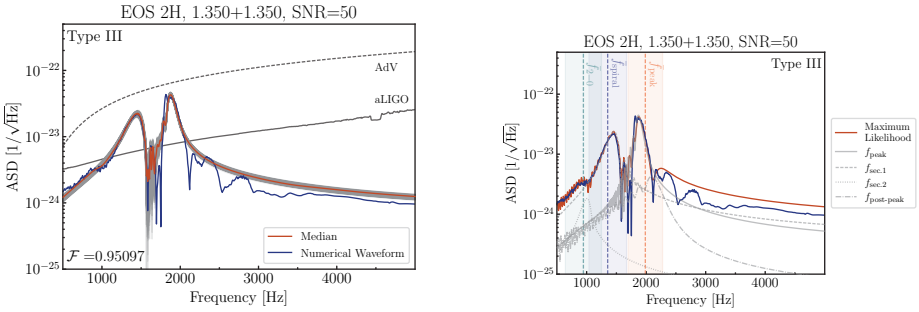


Figure 5.9: Same as Fig. 5.6, but for a 1.35+1.35 $M_\odot$ model with EOS 2H (Type III post-merger spectrum).

cluster tightly around their median (red line), in good agreement with the ASD of the numerical waveform (blue line), with larger variations only at low frequencies. The fitting factor between the median reconstructed waveform and the numerical waveform is $\mathcal{F} = 0.96336$.

In the right panel of Fig. 5.8, all three main post-merger peaks, $f_{\rm peak}$, $f_{\rm spiral}$ and f_{2-0} have frequencies close to the central frequency predicted by each corresponding empirical relation. The damped oscillator corresponding to $f_{\rm peak}$ has converged very well to the main $f_{\rm peak}$ contribution. The first secondary oscillator ($f_{\rm sec.1}$) describes well the $f_{\rm spiral}$ contribution and the other secondary oscillator ($f_{\rm sec.2}$) has converged to a frequency corresponding to the f_{2-0} contribution. On the other hand, there is no strong feature at high frequencies and $f_{\rm post-peak}$ simply contributes a very broad peak (in practice the fourth oscillator could be left out of the analytic model in this case).

Overall, the analytic model of Eq. (5.2), when combined with the empirical relations and classification scheme, describes very well all the main features of the post-merger GW spectrum of type II models, and high fitting factors can be obtained.

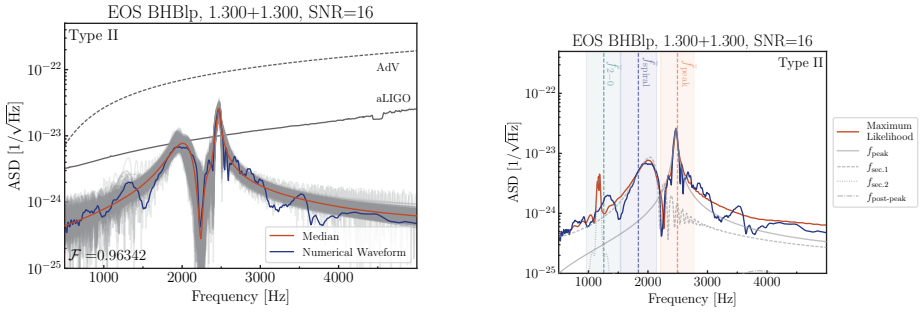


Figure 5.10: Same as Fig. 5.8, but for signal-to-noise ratio of 16.

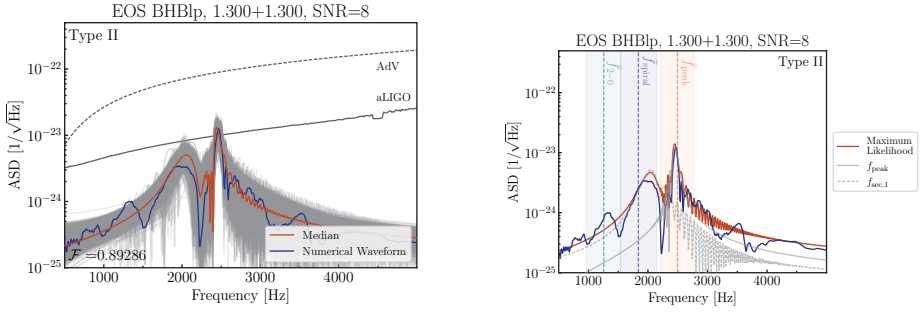


Figure 5.11: Same as Fig. 5.8, but for signal-to-noise ratio of 8.

Type III (EOS 2H, 1.35+1.35 $M_\odot$)

Our last example is a 1.35+1.35 $M_\odot$ merger with EOS 2H. This EOS is very stiff and the model lies deep inside the Type III region of Fig. 5.4, i.e. $f_{\rm spiral}$ is the dominant secondary peak in the post-merger spectrum. Consequently, our scheme leads to Gaussian priors $\mathcal{N}(f_{\rm peak}, \sigma^2)$, $\mathcal{N}(f_{\rm spiral}, \sigma^2)$ centered on the predictions of the empirical relations (5.4) and (5.7) for $f_{\rm peak}$ and $f_{\rm spiral}$, correspondingly, and uniform priors $\mathcal{U}(1, 5)$[kHz] for f_{2-0}. The fourth oscillator, $f_{\rm post-peak}$, in our analytic model (5.2) is initialized with uniform priors in the range $(f_{\rm min}, 5{\rm kHz})$, where $f_{\rm min} = f_{\rm peak} + 0.3{\rm kHz}$.

In the left panel of Fig. 5.9, we display the ASD of 1000 randomly sampled posterior waveforms (gray lines), which cluster very tightly around their median (red line) at all frequencies, in excellent agreement with the ASD of the numerical waveform (blue line). The fitting factor between the median reconstructed waveform and the numerical waveform is $\mathcal{F} = 0.95097$.

The right panel of Fig. 5.9 shows, in more detail, the individual contributions of the different damped oscillators included in the analytic model (5.2). We observe that the damped oscillator corresponding to $f_{\rm peak}$ describes the shape of $f_{\rm peak}$, which lies well inside the band predicted by the empirical relation (5.4) with high accuracy. The same can be said for the damped oscillator corresponding to $f_{\rm spiral}$. Although the f_{2-0} peak is barely distinguishable in this Type III spectrum, the corresponding damped oscillator still converges to the correct frequency and amplitude (notice that the maximum likelihood waveform is the *sum* of the individual contributions). At high frequencies, a strong f_{2+0} does not exist in the Type III spectrum, but several smaller peaks can be observed. Because we gave a wide frequency range as prior to the fourth oscillator, $f_{\rm post-peak}$, it does not converge to a particular secondary peak at high frequencies, but it contributes to the reconstructed waveform over a wide frequency range.

Overall, the analytic model of Eq. (5.2), when combined with the empirical relations and classification scheme, describes very well the main features of the post-merger GW spectrum of type III models, and high fitting factors are obtained.

5.3.2 SNR = 16

When the source is at a larger distance than considered in Sec. 5.3.1, the signal to noise ratio drops. Here, we examine the impact of a lower SNR = 16 on the reconstruction of the post-merger spectrum for the 1.3+1.3 $M_\odot$ type II merger using the BHBlp EOS. We chose this model because both f_{2-0} and $f_{\rm spiral}$ are important and could be reconstructed at SNR = 50 (although f_{2-0} had greater uncertainty).

In the left panel of Fig. 5.10, the 1000 randomly sampled waveforms (gray lines) have a significantly wider distribution, compared to the corresponding SNR = 50 case in Sec. 5.3.1. The median reconstructed waveform only informs about $f_{\rm peak}$ and $f_{\rm spiral}$, while the f_{2-0} is no longer visible in the reconstruction. The fitting factor is still high, $\mathcal{F} = 0.96342$, due to the good reconstruction of the main peak.

The right panel of Fig. 5.10, confirms that only two post-merger peaks, $f_{\rm peak}$ and $f_{\rm spiral}$ are reconstructed, with frequencies close to the central frequency predicted by each corresponding empirical relation. The damped oscillator corresponding to $f_{\rm peak}$ has converged very well to the main $f_{\rm peak}$ contribution. The first secondary oscillator ($f_{\rm sec.1}$) describes well the $f_{\rm spiral}$ contribution, but the other secondary oscillator ($f_{\rm sec.2}$) has converged to a frequency peak close to f_{2-0}, but with an incorrect amplitude and damping timescale. The fourth oscillator in the model $f_{\rm post-peak}$ does not converge. In practice, we find that for SNR = 16, a simpler analytic model employing only two oscillators, corresponding to $f_{\rm peak}$ and $f_{\rm spiral}$ would suffice to describe the post-merger spectrum for this merger.

5.3.3 SNR = 8

For even larger source distances, when SNR = 8, the uncertainties in the reconstruction of the post-merger spectrum for the 1.3+1.3 $M_\odot$ type II merger using the BHBlp EOS increase significantly. Because of this, there is no point in using an analytic model with four damped oscillators. Instead, we attempt to reconstruct the post-merger waveform using only two damped oscillators, corresponding to $f_{\rm peak}$ and $f_{\rm spiral}$.

In the left panel of Fig. 5.11, the 1000 randomly sampled waveforms

(gray lines) have a much wider distribution than for the SNR = 16 case and the fitting factor has dropped to $\mathcal{F} = 0.89286$. However, the median reconstructed waveform still informs about both the $f_{\rm peak}$ and $f_{\rm spiral}$ peaks.

The right panel of Fig. 5.11, confirms that both post-merger peaks, $f_{\rm peak}$ and $f_{\rm spiral}$ are reconstructed fairly accurately, with frequencies close to the central frequency predicted by each corresponding empirical relation. The damped oscillator corresponding to $f_{\rm peak}$ has converged very well to the main $f_{\rm peak}$ contribution and the secondary oscillator ($f_{\rm sec.1}$) describes well the $f_{\rm spiral}$ contribution.

5.4 Fitting Factors and Reconstruction

We can evaluate the effectiveness of our methodology by computing the noise-weighted fitting factor of Eq. (5.10) for each of the 15 cases we consider. Fig. 5.12 shows the distribution of the fitting factors in our sample of models. In each case, the upper and lower horizontal lines show the 99.7% confidence intervals, based on 10 different noise realizations, with an intermediate line denoting the median value. In addition, a vertical line indicates the 95% confidence interval. The median fitting factors are in the range of 0.935 to 0.980, similar to the results in [119].

A representative case of the reconstruction of the post-merger waveform using the maximum likelihood waveform is shown in Fig. 5.13, which is employed for the 1.3+1.3 BHBlp data. The reconstruction agrees well with the numerical waveform up to about 20ms after merger. The differences after that can be explained as follows: The reconstructed waveform is consistent with the assumption of the analytic model that the signal comprised several damped oscillators. On the other hand, the numerical waveform shows signs of the revival of the quadrupole frequency after about 25ms from merger. This behaviour is not built into the analytic model. The revival could be due to e.g. rotational instabilities developing in the remnant (see [265] and references therein).

At a post-merger network SNR of 50, the posterior distributions show no multi-modal behaviour, and the typical 1σ error for extracting the $f_{\rm peak}$ frequency is $\sim 0.22\%$. For the dominant secondary peak $f_{\rm sec.1}$, the 1σ error is

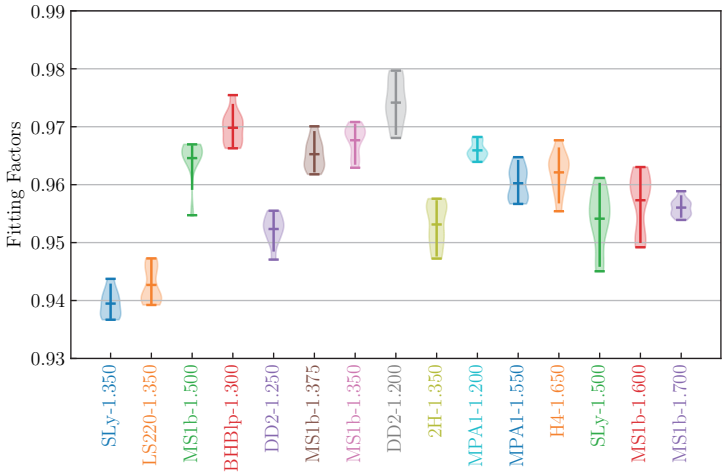


Figure 5.12: Fitting-factor distributions for each post-merger numerical-relativity waveform. 99.7% confidence intervals are represented by the upper and lower horizontal bars. The median value appears as the horizontal bar in the centre. The 95% confidence intervals are indicated by the thick vertical line. The range of the median fitting factors is 0.935 to 0.980.

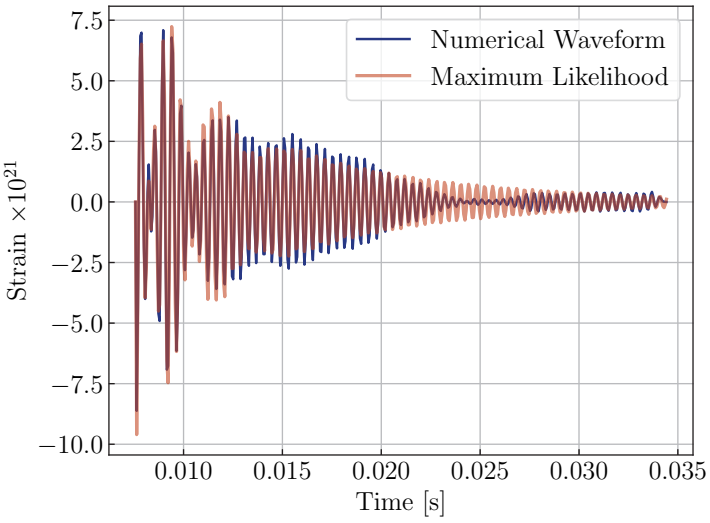


Figure 5.13: Scaled data alongside the maximum likelihood reconstruction time domain plot for the 1.3+1.3 BHBlp EOS.

$\sim 1.4\%$ and for the weaker $f_{\rm sec.2}$ and $f_{\rm post-peak}$ frequencies it is $\sim 6.2\%$. At lower post-merger network SNR values, we still typically find no multi-modal behavior in the posterior distributions. The 1σ errors are 0.72%, 2.6%, 7.3% and 18% for $f_{\rm peak}$, $f_{\rm sec.1}$, $f_{\rm sec.2}$ and $f_{\rm post-peak}$ respectively, when SNR=16. For the weakest case of SNR=8, we find that $f_{\rm peak}$ can still be extracted with 1σ of 2.4%, which increases to 8.5% for the dominant secondary frequency.

5.5 Summary and Discussion

We propose an improved method for parameter estimation in the post-merger phase of binary neutron star mergers. The improvements come from the application of empirical relations for the three main post-merger frequencies to the choice of priors, as well as from the application of a classification of the different post-merger spectra. Our dataset consists of 15 waveforms of

binary neutron star mergers that form differentially rotating remnants, 13 of which are from the CoRe catalog and two from [273]. Most of our waveforms correspond to equal-mass cases, with the exception of two cases that have a mass ratio of 1.5. The individual neutron star radii in our dataset cover a range of $\sim$ 11.5km to 16km and masses from $\sim 1.2M_\odot$ to $1.7M_\odot$. This range accommodates all three waveform types of the classification scheme introduced in [45, 302], as can be seen in Fig. 5.4.

The analytic model we use to describe gravitational waves in the post-merger phase is an extension of the analytical model in [119]. Apart from the dominant $f_{\rm peak}$ frequency and the two subdominant frequencies $f_{\rm spiral}$, f_{2-0} modes we also allow for the existence of a higher subdominant frequency.

The main objective of our study is to develop a robust framework that systematically leads to converged posterior distributions of the main parameters involved in the analytic model, even when choosing different noise realizations. In addition, our objective was to resolve the appearance of multimodal posterior distributions, seen previously in some cases studied in [119].

Following [302], we derived additional empirical relations for post-merger frequencies, in terms of the tidal deformability ($\tilde{\Lambda}$) and the chirp mass ($\mathcal{M}_{\rm chirp}$), which are observables that can be determined from the inspiral phase. This allows us to set significantly smaller prior ranges for several parameters, which avoids convergence to wrong or multiple solutions. The classification scheme of [45, 302] is used to first determine the type of post-merger spectrum and then choose an appropriate strategy for choosing the priors or leaving out certain components of the analytic model. Depending on the type of the post-merger spectrum, we use Gaussian priors for the dominant and main subdominant frequency peaks and uniform priors only for the weakest oscillators in the analytic model. In this way, the parameter estimation is optimized, taking into account the available knowledge from the inspiral phase and our current understanding of the post-merger GW spectrum. We make detailed corner plots of the posterior distributions for all model parameters for the cases discussed in this work available in the accompanying repository [303].

Using different components in the analytic model, depending on type of the post-merger waveform, allowed us to identify the resolved peaks with

particular physical processes in the merger remnant. This is important for testing the detailed physical assumptions that are made when simulating binary neutron star mergers.

Overall, we demonstrate that the analytic model of Eq. (5.2), when combined with the empirical relations, the choice of priors and the classification scheme, describes well the main features of the post-merger GW spectra of the three main types I, II and III and high fitting factors are obtained. In addition, we show that our scheme has sufficient flexibility to capture all essential features of the post-merger GW spectrum of Type Ib models, also reaching high fitting factors in this case.

Transitioning from `dynesty` to `pocoMC` as our primary sampler resulted in an acceleration of nearly tenfold in terms of walltime. The enhanced sampling performance offered by `pocoMC` arises from several contributing factors. Firstly, the efficiency of sampling is increased in the latent space of the normalizing flow due to the lack of strong correlations between the parameters of the posterior distribution. This eliminates the necessity of segregating the posterior into distinct sections - often accomplished via an ellipsoid mixture within the nested sampling context - and the subsequent need to sample each region individually. Additionally, the application of broad priors on certain parameters means that the prior volume substantially exceeds the posterior volume. This disparity significantly impairs the sampling performance of nested sampling methods, though the impact on PMC implementations such as `pocoMC` is less pronounced. Lastly, the inherent parallelizability of `pocoMC` enabled us to harness the power of multiple CPUs to further boost the sampling speed. This increased pace allowed for the exploration of a wider range of cases and a large number of noise realizations. We have thus demonstrated that `pocoMC` is an outstanding tool for parameter estimation of gravitational wave signals.

Here, we presented a robust and fast parameter estimation method of post-merger GW signals, based on a specific analytic model for the post-merger phase and on numerical waveforms that were obtained through numerical-relativity simulations where only the general-relativistic hydrodynamics was evolved. An alternative approach to analytic models is the use of machine-learning representations of the post-merger GW spectrum [308, 221, 274].

A number of additional physical effects are expected to impact the waveform produced in the post-merger phase, e.g. turbulence modeling [247], magnetic fields [284, 34, 292, 35], neutrino transport [312, 284, 217], bulk viscosity [208], dark matter imprints [283], etc. These effects are expected to have an impact on the onset (or suppression) of late-time convective or rotational instabilities in the remnant [223, 97, 265, 205]. Furthermore, deviations from general relativity are also expected to influence the structure of the post-merger remnant and affect the detected waveforms, see e.g. [40, 269, 262, 200, 117, 60, 279, 182, 41, 218, 270, 263, 201, 118, 184, 61, 280, 183, 186, 293]. In future studies, our current method could be tested for robustness against the impact of such effects.

5.6 Appendix: Robustness of parameter estimation

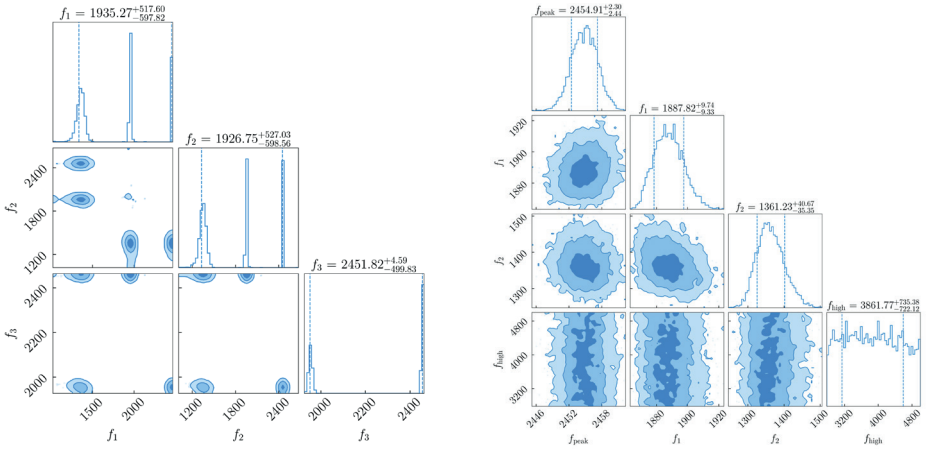
The baseline for the work presented here was the publication by Easter et al. [119], and it is instructive to examine how the changes and augmentations we implemented result in a better quantitative outcome. While in some cases, the two methods can yield comparable results, in several cases there is a clear improvement in the robustness of the parameter estimation. One such case, which we chose to illustrate this difference, is the BHBlp-1.3 Type II model, discussed in Sec. 5.3.1.

The left panel of Fig. 5.14 displays the posterior distributions for the three frequencies f_1, f_2, f_3 participating in the analytic model used in [119], when we performed parameter estimation with the same uninformative flat prior choices as in [119]. It is evident, that the parameter estimation has not converged to a unique solution, since the posteriors are a mixture of difference frequencies.

On the other hand, the right panel of Fig. 5.14 displays the corresponding posterior distributions of the four frequencies $f_{\rm peak}, f_1, f_2, f_{\rm high}$ participating in our extended analytic model, where the informed priors were used, as discussed in Sec. 5.2.5. Clearly, the frequencies have converged to a unique solution, with nearly Gaussian posterior distributions, except for the 4th

oscillator $f_{\rm high}$, which remains uninformative (see Appendix 5.7 for a case, where the 4th oscillator does make a difference in the outcome).

Notice that in the above comparison, the same seed was used for the generation of the noise and the sampling.



(a) Using the methodology in Easter et al. [119].

(b) Using the methodology presented in this work.

Figure 5.14: Truncated corner plots for the BHBlp-1.3 case, showing only the frequencies involved. The informed priors used in this work (b) result in robust parameter estimation.

5.7 Appendix: Contribution of the 4th oscillator

As discussed in Sec. 5.2.2, there are four main frequencies that we consider in our analytic model of the post-merger phase. Of these, the fourth oscillator with frequency f_4, could fit the dominant feature at frequencies higher than $f_{\rm peak}$, hence it is labeled $f_{\rm post-peak}$ in Figs. 5.6 - 5.10. Examining these figures closely, we see that in most cases the features at frequencies higher than $f_{\rm peak}$

have amplitudes well below the detector noise level. Hence, the inclusion of f_4 ($f_{\rm post-peak}$) has an almost negligible impact on the reconstruction of the post-merger waveform.

However, in the case of the Type Ib 1.55+1.55 $M_\odot$ model with EOS MPA1, displayed in Fig. 5.7, the main peak comprises a narrow, late-time component and a second, broader, early-time component at somewhat higher frequency, see the discussion in Sec. 5.3.1. In this case, the inclusion of f_4 in our analytic model, allows us to capture both of these components in Fig. 5.7, resulting in an accurate reconstruction of the post-merger waveform. On the other hand, we demonstrate in Fig. 5.15 that omitting f_4 for this model, results in a significantly worse reconstruction, with the fitting factor dropping from 0.95751 in Fig. 5.7 to 0.9105 in Fig. 5.15. Thus, at least for Type Ib cases, the inclusion of f_4 makes a real difference in the accuracy of the reconstruction.

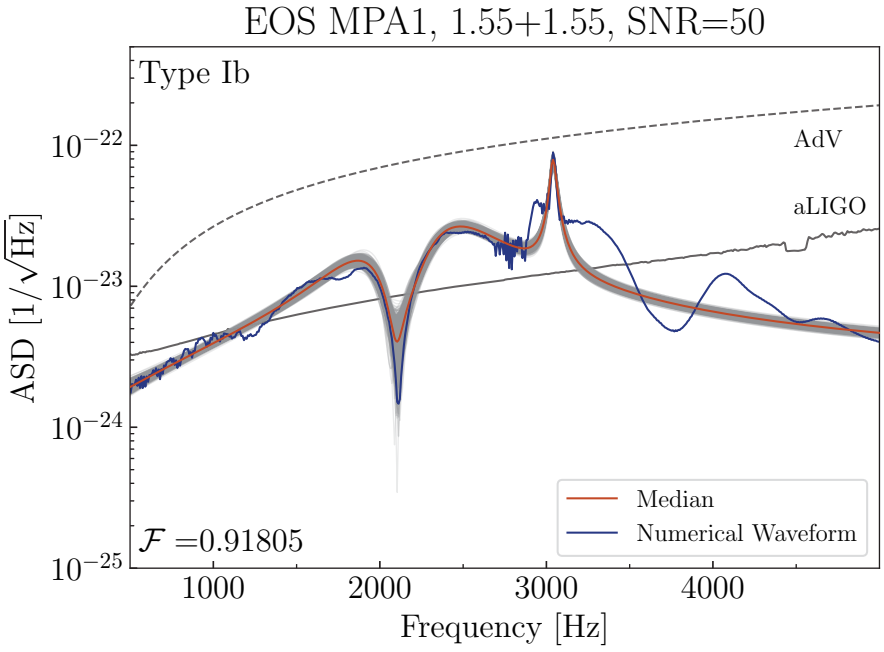


Figure 5.15: Reconstruction of the waveform using the medians of all parameters from the posterior distribution. In this case, the fourth oscillator, f_4, was omitted, to demonstrate that its absence for Type Ib models can result in a significantly smaller fitting factor than when it is included, as was the case in Fig. 5.7.


7

Chapter 6

Conclusions

The four articles collected in this thesis address compact-binary modeling at different points along the chain from numerical formulation to physical interpretation. Taken individually, each chapter solves a specific problem. Taken together, they show that progress in compact-binary science depends on more than improved astrophysical modeling alone. It also requires stable discretizations in geometries adapted to the problem, reduced descriptions that remain faithful to the structure of the underlying equations, geometric diagnostics that resolve the strongly nonlinear merger regime, and inference methods that convert modeled signals into usable physical constraints [148, 224, 229, 304]. The principal lesson of the thesis is therefore methodological as much as phenomenological: robust scientific conclusions in gravitational physics depend on controlling how the problem is posed numerically, how singular or multiscale structure is represented, and how simulated signals are interpreted.

6.1 Stable Discretization as Infrastructure

The first chapter contributes at the level of numerical infrastructure. Its main result is not the treatment of a particular astrophysical source, but the construction of summation-by-parts operators that remain accurate and stable in curvilinear coordinate systems with singularities. This is signifi-

cant because many problems in relativistic astrophysics are most naturally expressed in symmetry-adapted coordinates, yet coordinate singularities at the origin or axis often force compromises in accuracy or in the clarity of the discrete stability argument [148, 239, 216]. The chapter shows that this compromise is not inevitable.

The contribution of this work lies in providing a systematic construction where ad hoc regularization strategies are often used. We derived for the first time higher order operators for collocated grids. We also showed that previously used methods for staggering are inheritantly unstable at the origin, and must be replaced with a stable staggered scheme, which we provide. While the chapter does not provide an 8th order operator, nor does it show evolutions of vector and tensor wave equations, it does provide a principled route to high-order stable finite differencing in a class of coordinates that is pervasive in compact-object physics.

6.2 Self-Force Formulated in the Lab Frame

The second chapter turns from discretization infrastructure to modeling in the extreme-mass-ratio regime. Its central insight is that formal regularization alone is not enough to make the self-force problem numerically useful. The field equations and equations of motion must also retain the structure of a viable initial-value problem. This is the point at which the chapter departs from a purely formal treatment of singular fields and addresses the practical demands of self-consistent evolution [224, 101, 36].

The chapter is most convincing in its identification of a specific obstruction: acceleration-dependent singular-field constructions can introduce self-referential higher-derivative dependence into the dynamics [157, 105, 311, 310]. That observation matters because it pinpoints a failure mode that is easy to hide inside an otherwise correct perturbative framework. The proposed remedy, namely a formulation that incorporates acceleration effects without requiring higher derivatives of the acceleration as additional effective data, is therefore conceptually important. At the same time, the chapter is deliberately modest in scope. It treats a scalar model in a simplified setting rather than the full gravitational problem. We do not address curved

spacetime and we resort to 1D evolutions instead of 3D evolutions. This limits our capability to compare to other works that use different methods. Although we showed we recover the analytic solution, it is essential to go to 3D and have a comparison with previous results with the Effective Source Method [298, 107].

6.3 Horizon Geometry as a Merger Diagnostic

The third chapter shifts from methodological development to the nonlinear phenomenology of black-hole coalescence. Its contribution is to treat the merger itself through the behavior of marginally trapped surfaces and their multipolar structure, rather than only through waveforms or remnant fitting formulae. The resulting picture is geometrically direct: the merger appears as a singular transition in the horizon structure, with discontinuous jumps in mass-related quantities tied to the cusp formation on the inner remnant surface [229, 231, 69, 233, 234].

What emerges from this analysis is a more precise view of where the major changes in black-hole parameters occur. The work shows that the merger is not merely a brief interpolation between inspiral and ringdown, but the stage at which the strongest geometric reorganization takes place. This is an important corrective to approaches that focus primarily on early inspiral information or late-time remnant properties. The chapter is also clear about its limitations. The analysis is restricted to head-on collisions of non-spinning black holes in a highly symmetric setting. As a consequence, many questions of greatest astrophysical interest, especially those involving orbital angular momentum, precession, recoil, and generic asymmetry, remain open. Even so, the chapter establishes that horizon dynamics can carry physically interpretable information across the merger itself. That is a nontrivial result, and it suggests that quasi-local geometry can serve as a complementary diagnostic to waveform-based descriptions of strong-field dynamics [254, 170, 212, 176, 70]. In more general mergers, horizon asymmetries may also help connect strong-field geometry to the recoil generated by anisotropic

gravitational-wave emission.

6.4 From Post-Merger Signals to Inference

The fourth chapter addresses a different part of the compact-binary problem: not how to evolve the system, but how to extract physical information from a modeled signal once the merger has occurred. This regime is rich in information about dense matter, yet it is observationally difficult because the signal is short, high-frequency, and structurally more complicated than the inspiral waveform [31, 264, 46]. The chapter's main contribution is to reduce the inference problem by importing physically motivated prior information from the inspiral and from empirical relations between pre-merger and post-merger properties [119, 302, 304].

The value of this work lies in its pragmatism. It does not attempt to solve the dense-matter problem at the level of first principles, nor does it claim that a compact analytic model captures all possible post-merger physics. Instead, it shows that carefully constrained models can transform a difficult and sometimes multimodal inference problem into one that is substantially more robust and computationally tractable. The acceleration gained through the sampling strategy is important, but the more consequential point is that the chapter treats inference itself as part of the modeling pipeline. A waveform model that cannot support reliable parameter estimation has limited scientific value, no matter how accurate its underlying simulations may be. The limitations are also evident: the study relies on a relatively small waveform set, on empirical relations calibrated to current numerical-relativity results, and on simulations that do not yet include the full range of microphysical effects that may shape real post-merger signals [46, 205, 247, 217]. These caveats do not undermine the chapter's contribution; rather, they define the conditions under which the method should be interpreted and extended.

6.5 Taken Together

Read together, the articles do not offer a single unified computational framework for compact binaries across all mass ratios. They identify several points at which compact-binary modeling can fail: at the level of discrete stability, at the level of singular source representation, at the level of interpreting the merger geometry, and at the level of extracting parameters from short and structured signals. The thesis argues, implicitly through these case studies, that no single approximation strategy is sufficient across the full problem. Different mass-ratio regimes and different scientific objectives demand different numerical languages.

The common thread is therefore not a single method but a standard of construction. In each chapter, the decisive issue is whether the formulation isolates the physically relevant structure of the problem without introducing avoidable numerical or interpretive ambiguity. The SBP chapter does this for stability in adapted coordinates. The self-force chapter does it for self-consistent reduced evolution. The horizon chapter does it for the nonlinear merger geometry. The BNS chapter does it for observational inference in the post-merger regime. This thesis should therefore be read as a contribution to the numerical modeling of gravitational physics: accurate science follows when the mathematical structure of the problem, the numerical representation of that structure, and the intended physical interpretation are kept in close alignment.



Summary

This thesis studied numerical modeling of compact binary systems across regimes in which the relevant physical and numerical difficulties are qualitatively different. Comparable-mass binaries require the nonlinear dynamics of spacetime and, for neutron stars, matter. Extreme-mass-ratio systems require a controlled treatment of an unresolved small body and its self-field. Across both regimes, the reliability of the calculation depends on preserving continuum structure at the discrete or statistical level: geometric quantities must retain their meaning, evolution systems must remain stable, singular or unresolved sources must be represented consistently, and inference problems must be constrained by the available physics.

The work presented here consists of four connected contributions. The first constructs high-order summation-by-parts discretizations for curvilinear coordinates with coordinate singularities. The second proposes a finite-resolution formulation for scalar self-force calculations. The third concerns the quasi-local geometry of binary black-hole mergers. The fourth develops a robust inference strategy for binary-neutron-star post-merger gravitational waves. Together, these results form a thesis on compact-binary modeling rather than on a single computational technique.

High-order SBP operators for curvilinear singularities

This thesis begins with the fundamental requirement of numerical relativity: that the discretized evolution equations should preserve the mathematical

structure of the underlying continuum physics. Coordinates adapted to symmetry, such as spherical coordinates, are natural in many relativistic calculations. They also contain coordinate singularities. The factors $1/r^p$ that appear at the origin are not physical singularities, but they can destabilize a finite-difference method if the discrete operators do not reproduce the correct integration-by-parts structure.

The work on summation-by-parts operators constructed high-order finite-difference schemes for first-order hyperbolic systems in such coordinates. The construction starts from Cartesian SBP operators but modifies the norm and derivative structure to account for the curvilinear volume measure and the distinction between scalar and vector-valued quantities. Non-diagonal vector-valued norm operators are allowed. Accuracy conditions near the origin determine the closure coefficients, and the covariant divergence operator is derived from the SBP identity itself.

Explicit operators were constructed up to sixth order. In spherical scalar-wave evolutions, the schemes conserve the discrete energy to floating-point accuracy under reflecting boundary conditions. A Gaussian wave packet can pass through the origin and re-emerge cleanly, with fourth- and sixth-order convergence observed for the corresponding operators. These tests demonstrate that the origin can be included on a collocated radial grid without sacrificing the energy stability expected from an SBP discretization.

This result supplies numerical infrastructure for relativistic evolution problems. It provides a systematic alternative to excising the origin or using staggered grids solely to avoid the coordinate singularity. Because the construction is based on weighted norms, compatible gradient and divergence operators, and the SBP identity, it is not restricted in principle to scalar waves. The method can be extended to arbitrary tensor rank and spacetime dimension, although its integration into full numerical-relativity systems remains future work.

Finite-resolution scalar self-force modeling

The scalar self-force work addressed the extreme-mass-ratio side of the thesis. In such systems, a small compact object moves in the field of a much larger

body and is influenced by its own field. A literal point-particle description introduces singular fields, which must then be regularized. In self-consistent evolutions, acceleration-dependent singular-field terms can lead to implicit or higher-derivative structures in the equations of motion. This is undesirable for an initial-value evolution scheme.

The approach developed here starts from a different premise: a numerical calculation has finite resolution. The relevant hierarchy is one in which the large scale is resolved, the small object is not, and the grid scale lies between them. In this regime the small object can be modeled as finite-sized below the resolution scale. The particle field is treated as an additional basis function, and the evolved field is decomposed into this particle field plus a regular remainder field. Only the sum has physical meaning.

The formulation was developed for a scalar particle coupled to a scalar field in flat spacetime. The equations were written explicitly in the lab frame, and the state vector was chosen before deriving the equations of motion. This leads to a self-consistent initial-value system that includes the radiation produced by particle acceleration without introducing infinities or requiring higher time derivatives of the self-acceleration as independent data.

Two discretizations were implemented: a finite-volume method and a spectral method. In one-dimensional tests, both showed excellent agreement with analytic solutions. The fully self-interacting case also agreed well with analytic approximations. The calculations emphasized that effective sources generated by particle acceleration must be projected properly into the discrete function space; this projection can be performed analytically or numerically, with the latter being attractive when the source expressions become complicated.

The present work is a model problem, but its conceptual implication is broader. It suggests that effective-source modeling can be formulated directly at finite resolution. In higher dimensions the effective sources will have stronger local structure, including directional discontinuities. This motivates discretizations adapted to the particle neighborhood, such as finite-element methods with local sub-element structure. Window functions can be used to localize the particle field and reduce computational cost. Extending the construction to electrodynamics and ultimately to general relativity would require additional degrees of freedom, including spin for

compact objects, but the central idea remains the same: the unresolved object should be represented in a way that is consistent with the numerical scale at which the field is evolved.

Black-hole horizon mergers and cusp formation

The study of cusp formation in merging black-hole horizons approached the binary black-hole merger through marginally outer trapped surfaces. In this quasi-local picture, the merger can be understood as the merger of the two individual MOTSs associated with the progenitor black holes with the inner MOTS associated with the remnant. For the head-on collision of two nonspinning black holes, the inner MOTS develops a cusp at the touching time.

The main result is that the time evolution of the horizon mass and mass multipole moments is smooth except at this touching time. At the cusp, the curvature contains a distributional contribution, and the quasi-local quantities exhibit discontinuous jumps. The jump in the mass follows from the additivity of the area at the merger. The jumps in the higher mass multipole moments are controlled by the Gauss invariant, which determines the strength of the distributional curvature at the cusp. Although these jumps are not the entire change between the initial progenitor horizons and the final remnant, they can be among the dominant contributions.

This result gives a source-side description of part of the merger dynamics. It shows that horizon multipoles contain information that is localized at the merger event and that can be lost rapidly in the subsequent smooth evolution of the common horizon. Accurate numerical resolution of the MOTSs near the touching time is therefore essential. The analysis also suggests a broader program: to connect horizon multipoles, angular momentum, recoil, and gravitational-wave mode amplitudes in generic black-hole mergers. The present work treated a symmetric head-on configuration. Extending the analysis to spinning, precessing, and non-axisymmetric binaries is the natural next step.

Post-merger inference for binary neutron stars

The binary-neutron-star post-merger work addressed a different aspect of comparable-mass dynamics: the extraction of physical parameters from the gravitational-wave signal of the remnant. The post-merger spectrum contains a dominant frequency $f_{\rm peak}$ and subdominant structures associated with the dynamics of the differentially rotating remnant. These features are potentially valuable probes of dense nuclear matter, but robustly inferring them is difficult because the signal is short, high-frequency, and multi-component.

The method developed in this thesis combines an analytic damped-oscillator model with empirical relations and a classification of post-merger spectra. Inspiral-measurable quantities, in particular the chirp mass $\mathcal{M}_{\rm chirp}$ and the effective tidal deformability $\tilde{\Lambda}$, are used to construct informed priors for the main post-merger frequencies. The spectral classification determines which components should be emphasized and which priors should be broad. This structure reduces the multimodality that can arise when broad uninformative priors allow different oscillators to exchange roles.

The method was validated on fifteen numerical-relativity waveforms of binary neutron-star mergers forming differentially rotating remnants. The set includes thirteen waveforms from the `CoRe` catalog and two additional waveforms, mostly equal-mass systems and two cases with mass ratio $q = 1.5$. The models span neutron-star radii of approximately 11.5 km to 16 km and masses from approximately $1.2M_\odot$ to $1.7M_\odot$, covering the relevant post-merger spectral types considered in the analysis.

The resulting fitting factors have median values in the range 0.935–0.980. At post-merger network signal-to-noise ratio 50, the typical one-sigma uncertainty in the dominant frequency $f_{\rm peak}$ is about 0.22%. For the dominant secondary peak it is about 1.4%, while for weaker secondary or post-peak frequencies it is about 6.2%. The informed-prior strategy gives converged posterior distributions in cases where less constrained analyses can produce ambiguous or multimodal results. Replacing `dynesty` with `pocoMC` yields nearly an order-of-magnitude speedup in wall time, making it possible to

explore more cases and noise realizations.

The broader significance of this result is that post-merger inference should not be treated as a blind fit of a flexible model. The inspiral and post-merger phases are physically connected, and the inference method should use that connection. Future work should test the robustness of the approach against additional physical effects that are absent or simplified in the waveform set considered here, including magnetic fields, neutrino transport, turbulence modeling, bulk viscosity, possible dark-matter imprints, and deviations from general relativity. It will also be important to compare analytic oscillator models with reduced-order or machine-learning representations of the post-merger spectrum.

Synthesis and outlook

The four parts of the thesis are connected by a common principle: numerical models should preserve the structures that make the continuum problem well defined. In the SBP study, the relevant structure is the energy estimate; preserving it discretely yields stable high-order evolution through a coordinate singularity. In the self-force study, the relevant structure is the scale hierarchy; the small object is unresolved and should be modeled at finite resolution rather than as a singularity that the discretization cannot represent. In the horizon study, the relevant structure is geometric and topological; the cusp carries a distributional curvature contribution that controls jumps in quasi-local quantities. In the post-merger study, the relevant structure is physical information shared between inspiral and post-merger observables; using that information in the priors turns an ambiguous inference problem into a robust one.

The comparable-mass direction opened by this work is to connect source-side diagnostics and detector-side observables more directly. For black-hole mergers, this means extending the cusp and multipole analysis to generic spinning binaries and investigating whether horizon multipoles can be quantitatively related to waveform modes and recoil. For binary neutron stars, it means applying robust post-merger inference to waveforms with increasingly realistic microphysics and to detector networks expected in the

third-generation era.

The extreme-mass-ratio direction is to move from scalar model problems to gravitational self-force calculations in curved spacetime. The finite-resolution viewpoint developed here should be tested in settings where long-time phase accuracy can be measured. It must also be combined with numerical methods capable of handling localized effective sources in multiple dimensions.

The numerical-methods direction is to incorporate the high-order curvilinear SBP operators into more complete evolution systems. This includes coupling to nonlinear equations, matter sources, gauge dynamics, and realistic boundary conditions. If successful, such operators would allow simulations to use symmetry-adapted coordinates while retaining the stability guarantees of an energy method.

Taken together, the results of this thesis advance compact-binary modeling across the mass-ratio spectrum. They do so by combining quasi-local geometry, gravitational-wave inference, finite-resolution source modeling, and stable high-order discretization. The unifying lesson is that robust numerical relativity is not only a matter of increasing resolution or computational cost. It is a matter of formulating the discrete problem so that the essential physics and mathematics of the continuum problem remain visible in the computation.

Samenvatting

Dit proefschrift bestudeerde numerieke modellering van compacte binaire systemen in regimes waarin de relevante fysische en numerieke moeilijkheden kwalitatief verschillend zijn. Binaire systemen met vergelijkbare massa vereisen de niet-lineaire dynamica van de ruimtetijd en, voor neutronensterren, van materie. Systemen met extreme massaverhouding vereisen een gecontroleerde behandeling van een onopgelost klein object en zijn zelfveld. In beide regimes hangt de betrouwbaarheid van de berekening af van het behouden van de continuümstructuur op discreet of statistisch niveau: geometrische grootheden moeten hun betekenis behouden, evolutiesystemen moeten stabiel blijven, singuliere of onopgeloste bronnen moeten consistent worden weergegeven, en inferentieproblemen moeten worden begrensd door de beschikbare fysica.

Het hier gepresenteerde werk bestaat uit vier samenhangende bijdragen. De eerste construeert hoog-orde summation-by-parts-discretisaties voor kromlijnige coördinaten met coördinaatsingulariteiten. De tweede stelt een eindige-resolutieformulering voor scalaire self-force-berekeningen voor. De derde betreft de quasi-lokale geometrie van fusies van binaire zwarte gaten. De vierde ontwikkelt een robuuste inferentiestrategie voor post-mergerzwaartekrachtsgolven van binaire neutronensterren. Samen vormen deze resultaten een proefschrift over het modelleren van compacte binaries en niet over slechts één computationele techniek.

Hoog-orde SBP-operatoren voor kromlijnige singulariteiten

Dit proefschrift begint met de fundamentele eis van numerieke relativiteit: dat de gediscretiseerde evolutievergelijkingen de wiskundige structuur van de onderliggende continuümfysica behouden. Coördinaten aangepast aan symmetrie, zoals bolcoördinaten, zijn natuurlijk in veel relativistische berekeningen. Zij bevatten ook coördinaatsingulariteiten. De factoren $1/r^p$ die in de oorsprong verschijnen zijn geen fysische singulariteiten, maar zij kunnen een eindige-differentiemethode destabiliseren als de discrete operatoren niet de juiste integratie-door-delen-structuur reproduceren.

Het werk aan summation-by-parts-operatoren construeerde hoog-orde eindige-differenceschema's voor eerste-orde hyperbolische systemen in zulke coördinaten. De constructie begint met Cartesische SBP-operatoren maar wijzigt de norm- en afgeleidestructuur om rekening te houden met de kromlijnige volumemaat en het onderscheid tussen scalaire en vectorwaardige grootheden. Niet-diagonale vectorwaardige normoperatoren zijn toegestaan. Nauwkeurigheidsvoorwaarden nabij de oorsprong bepalen de closure-coëfficiënten, en de covariante divergentie-operator wordt afgeleid uit de SBP-identiteit zelf.

Expliciete operatoren werden geconstrueerd tot en met zesde orde. In sferische scalaire-golfevoluties behouden de schema's de discrete energie tot op floating-point-nauwkeurigheid onder reflecterende randvoorwaarden. Een Gaussiaans golfpakket kan door de oorsprong gaan en weer schoon tevoorschijn komen, waarbij vierde- en zesde-orde convergentie wordt waargenomen voor de overeenkomstige operatoren. Deze testen tonen aan dat de oorsprong kan worden opgenomen op een collocated radiaal rooster zonder de energiestabiliteit op te offeren die van een SBP-discretisatie wordt verwacht.

Dit resultaat levert numerieke infrastructuur voor relativistische evolutieproblemen. Het biedt een systematisch alternatief voor het uitsnijden van de oorsprong of het gebruiken van staggered grids uitsluitend om de coördinaatsingulariteit te vermijden. Omdat de constructie is gebaseerd op gewogen normen, compatibele gradiënt- en divergentie-operatoren en de

SBP-identiteit, is zij in principe niet beperkt tot scalaire golven. De methode kan worden uitgebreid naar willekeurige tensorrang en ruimtetijddimensie, hoewel de integratie ervan in volledige numerieke-relativiteitssystemen toekomstig werk blijft.

Eindige-resolutie-modellering van scalaire self-force

Het werk over scalaire self-force behandelde de kant van extreme massaverhouding in dit proefschrift. In zulke systemen beweegt een klein compact object in het veld van een veel groter object en wordt het beïnvloed door zijn eigen veld. Een letterlijke puntdeeltjesbeschrijving introduceert singuliere velden, die vervolgens geregulariseerd moeten worden. In zelfconsistente evoluties kunnen versnellingafhankelijke termen in het singuliere veld leiden tot impliciete of hoger-afgeleide structuren in de bewegingsvergelijkingen. Dit is ongewenst voor een beginwaarde-evolutieschema.

De hier ontwikkelde benadering vertrekt van een andere premisse: een numerieke berekening heeft eindige resolutie. De relevante hiërarchie is er een waarin de grote schaal wordt opgelost, het kleine object niet, en de roosterschaal ertussenin ligt. In dit regime kan het kleine object worden gemodelleerd als eindig groot onder de resolutieschaal. Het deeltjesveld wordt behandeld als een extra basisfunctie, en het geëvolueerde veld wordt ontbonden in dit deeltjesveld plus een regulier restveld. Alleen de som heeft fysische betekenis.

De formulering werd ontwikkeld voor een scalair deeltje gekoppeld aan een scalair veld in vlakke ruimtetijd. De vergelijkingen werden expliciet opgeschreven in het lab frame, en de toestandsvector werd gekozen voordat de bewegingsvergelijkingen werden afgeleid. Dit leidt tot een zelfconsistent beginwaardesysteem dat de straling omvat die door deeltjesversnelling wordt geproduceerd zonder oneindigheden te introduceren of hogere tijdsafgeleiden van de zelfversnelling als onafhankelijke data te vereisen.

Twee discretisaties werden geïmplementeerd: een eindige-volumemethode en een spectrale methode. In eendimensionale testen vertoonden beide uit-

stekende overeenstemming met analytische oplossingen. Het volledig zelfinteragerende geval stemde eveneens goed overeen met analytische benaderingen. De berekeningen benadrukten dat effectieve bronnen die door deeltjesversnelling worden gegenereerd correct in de discrete functieruimte moeten worden geprojecteerd; deze projectie kan analytisch of numeriek worden uitgevoerd, waarbij het laatste aantrekkelijk is wanneer de bronuitdrukkingen gecompliceerd worden.

Het huidige werk is een modelprobleem, maar de conceptuele implicatie ervan is breder. Het suggereert dat effective-source-modellering direct op eindige resolutie kan worden geformuleerd. In hogere dimensies zullen de effectieve bronnen een sterkere lokale structuur hebben, inclusief directionele discontinuïteiten. Dit motiveert discretisaties die zijn aangepast aan de omgeving van het deeltje, zoals eindige-elementmethoden met lokale sub-elementstructuur. Windowfuncties kunnen worden gebruikt om het deeltjesveld te lokaliseren en de rekenkosten te verlagen. Het uitbreiden van de constructie naar elektrodynamica en uiteindelijk naar de algemene relativiteit zou extra vrijheidsgraden vereisen, waaronder spin voor compacte objecten, maar het centrale idee blijft hetzelfde: het onopgeloste object moet worden weergegeven op een manier die consistent is met de numerieke schaal waarop het veld wordt geëvolueerd.

Fusies van zwarte-gathorizonten en cuspvorming

De studie van cuspvorming in samensmeltende zwarte-gathorizonten benaderde de fusie van binaire zwarte gaten via marginaal buitenste gevangen oppervlakken. In dit quasi-lokale beeld kan de fusie worden begrepen als de fusie van de twee individuele MOTSs die geassocieerd zijn met de voorlopers van de zwarte gaten met de binnenste MOTS die geassocieerd is met het restant. Voor de frontale botsing van twee niet-spinnende zwarte gaten ontwikkelt de binnenste MOTS een cusp op het raaktijdstip.

Het belangrijkste resultaat is dat de tijdsevolutie van de horizonmassa en de massamultipoolmomenten glad is behalve op dit raaktijdstip. Bij de

cusp bevat de kromming een distributionele bijdrage, en de quasi-lokale grootheden vertonen discontinue sprongen. De sprong in de massa volgt uit de additiviteit van de oppervlakte bij de fusie. De sprongen in de hogere massamultipoolmomenten worden bepaald door de Gauss-invariant, die de sterkte van de distributionele kromming bij de cusp vastlegt. Hoewel deze sprongen niet de volledige verandering vormen tussen de initiële voorloper-horizonten en het uiteindelijke restant, kunnen zij tot de dominante bijdragen behoren.

Dit resultaat geeft een bronzijdige beschrijving van een deel van de fusiedynamica. Het laat zien dat horizonmultipolen informatie bevatten die gelokaliseerd is bij het fusiegebeurtenis en die snel verloren kan gaan in de daaropvolgende gladde evolutie van de gemeenschappelijke horizon. Nauwkeurige numerieke resolutie van de MOTSs nabij het raaktijdstip is daarom essentieel. De analyse suggereert ook een breder programma: horizon-multipolen, impulsmoment, recoil en amplituden van zwaartekrachtsgolfmodi met elkaar verbinden in generieke fusies van zwarte gaten. Het huidige werk behandelde een symmetrische frontale configuratie. Het uitbreiden van de analyse naar spinnende, precesserende en niet-axisymmetrische binaries is de natuurlijke volgende stap.

Post-mergerinferentie voor binaire neutronensterren

Het werk over de post-mergerfase van binaire neutronensterren behandelde een ander aspect van de dynamica van vergelijkbare massa's: het extraheren van fysische parameters uit het zwaartekrachtsgolfsignaal van het restant. Het post-mergerspectrum bevat een dominante frequentie f_{peak} en subdominante structuren die samenhangen met de dynamica van het differentieel roterende restant. Deze kenmerken zijn potentieel waardevolle sondes van dichte nucleaire materie, maar het robuust afleiden ervan is moeilijk omdat het signaal kort, hoogfrequent en uit meerdere componenten opgebouwd is.

De in dit proefschrift ontwikkelde methode combineert een analytisch model van gedempte oscillatoren met empirische relaties en een classificatie

van post-mergerspectra. Tijdens de inspiral meetbare grootheden, in het bijzonder de chirpmassa $\mathcal{M}_{\rm chirp}$ en de effectieve tidale vervormbaarheid $\tilde{\Lambda}$, worden gebruikt om geïnformeerde priors voor de belangrijkste post-mergerfrequenties te construeren. De spectrale classificatie bepaalt welke componenten benadrukt moeten worden en voor welke componenten de priors breed moeten zijn. Deze structuur vermindert de multimodaliteit die kan ontstaan wanneer brede, niet-informatieve priors verschillende oscillatoren van rol laten wisselen.

De methode werd gevalideerd op vijftien numeriekerelativiteitsgolfvormen van fusies van binaire neutronensterren die differentieel roterende restanten vormen. De verzameling omvat dertien golfvormen uit de `CoRe`-catalogus en twee aanvullende golfvormen, meestal systemen met gelijke massa en twee gevallen met massaverhouding $q = 1.5$. De modellen bestrijken neutronensterstralen van ongeveer 11.5 km tot 16 km en massa's van ongeveer $1.2M_\odot$ tot $1.7M_\odot$, en dekken de relevante post-mergerspectrale typen die in de analyse zijn beschouwd.

De resulterende fitting factors hebben mediaanwaarden in het bereik 0.935–0.980. Bij een netwerk-signaal-ruisverhouding van 50 in de post-mergerfase is de typische één-sigma-onzekerheid in de dominante frequentie $f_{\rm peak}$ ongeveer 0.22%. Voor de dominante secundaire piek is die ongeveer 1.4%, terwijl die voor zwakkere secundaire of post-piekfrequenties ongeveer 6.2% is. De strategie met geïnformeerde priors levert geconvergeerde posterioorverdelingen op in gevallen waarin minder beperkte analyses ambigue of multimodale resultaten kunnen opleveren. Het vervangen van `dynesty` door `pocoMC` levert bijna een orde van grootte snelheidswinst in wandkloktijd op, waardoor het mogelijk wordt meer gevallen en ruisrealisaties te onderzoeken.

De bredere betekenis van dit resultaat is dat post-mergerinferentie niet moet worden behandeld als een blinde fit van een flexibel model. De inspiral- en post-mergerfasen zijn fysisch met elkaar verbonden, en de inferentiemethode moet die verbinding gebruiken. Toekomstig werk moet de robuustheid van de aanpak testen tegen bijkomende fysische effecten die afwezig zijn of vereenvoudigd zijn in de hier beschouwde verzameling golfvormen, waaronder magnetische velden, neutrinotransport, turbulentie-modellering, bulkviscositeit, mogelijke afdrukken van donkere materie en afwijkingen van de algemene relativiteit. Het zal ook belangrijk zijn analytische oscillatormod-

ellen te vergelijken met reduced-order- of machine-learningrepresentaties van het post-mergerspectrum.

Synthese en vooruitblik

De vier delen van het proefschrift zijn verbonden door een gemeenschappelijk principe: numerieke modellen moeten de structuren behouden die het continuümprobleem goed gedefinieerd maken. In de SBP-studie is de relevante structuur de energieschatting; die discreet behouden levert stabiele hoog-orde evolutie door een coördinaatsingulariteit op. In de self-force-studie is de relevante structuur de schaalhiërarchie; het kleine object is onopgelost en moet op eindige resolutie worden gemodelleerd in plaats van als een singulariteit die de discretisatie niet kan weergeven. In de horizonstudie is de relevante structuur geometrisch en topologisch; de cusp draagt een distributionele krommingsbijdrage die sprongen in quasi-lokale grootheden bepaalt. In de post-mergerstudie is de relevante structuur fysische informatie die gedeeld wordt tussen inspiral- en post-merger-observabelen; het gebruik van die informatie in de priors verandert een ambigu inferentieprobleem in een robuust probleem.

De richting voor vergelijkbare massa's die door dit werk wordt geopend, is om diagnostiek aan de bronzijde en observabelen aan de detectorzijde directer met elkaar te verbinden. Voor fusies van zwarte gaten betekent dit het uitbreiden van de analyse van cusp en multipolen naar generieke spinnende binaries en het onderzoeken of horizonmultipolen kwantitatief gerelateerd kunnen worden aan golfvormmodi en recoil. Voor binaire neutronensterren betekent dit het toepassen van robuuste post-mergerinferentie op golfvormen met steeds realistischer microfysica en op detectornetwerken die in het tijdperk van de derde generatie worden verwacht.

De richting voor extreme massaverhouding is om van scalaire modelproblemen over te gaan naar gravitatie-self-force-berekeningen in gekromde ruimtetijd. Het hier ontwikkelde eindige-resolutieperspectief moet worden getest in situaties waarin langetermijnfase-nauwkeurigheid kan worden gemeten. Het moet ook worden gecombineerd met numerieke methoden die gelokaliseerde effectieve bronnen in meerdere dimensies aankunnen.

De richting voor numerieke methoden is om de hoog-orde kromlijnige SBP-operatoren in completere evolutiesystemen op te nemen. Dit omvat koppeling aan niet-lineaire vergelijkingen, materiebronnen, ijkdynamica en realistische randvoorwaarden. Indien succesvol zouden zulke operatoren simulaties in staat stellen symmetrie-aangepaste coördinaten te gebruiken met behoud van de stabiliteitsgaranties van een energiemethode.

Gezamenlijk genomen bevorderen de resultaten van dit proefschrift de modellering van compacte binaries over het volledige spectrum van massaverhoudingen. Zij doen dit door quasi-lokale geometrie, zwaartekrachtsgolfinferentie, bronmodellering op eindige resolutie en stabiele hoog-orde discretisatie te combineren. De verenigende les is dat robuuste numerieke relativiteit niet alleen een kwestie is van het verhogen van de resolutie of de rekenkosten. Het is een kwestie van het discrete probleem zo te formuleren dat de essentiële fysica en wiskunde van het continuümprobleem zichtbaar blijven in de berekening.

Curriculum Vitae

Stamatis Vretinaris was born on 19 January 1994 in Düsseldorf, Germany. Although it may sound like a cliche, my interest in the mysteries of the universe truly began in childhood. Looking up at the stars, I kept wondering what they were made of, how far away they were, and whether there might be life on other planets. That curiosity eventually led me to pursue a career in astrophysics.

In 2012 I started my BSc in Physics, followed in 2017 by an MSc in Computational Physics. I began working on gravitational waves in my BSc thesis in 2016, shortly before the official announcement of the first detection of gravitational waves by the LIGO collaboration [4].

For my MSc thesis, I worked on machine-learning methods for time series data of gravitational waves from binary neutron star mergers, with a focus on the post-merger regime.

During my master's studies I was fortunate to travel extensively. I was invited to the Kavli Institute for Astronomy and Astrophysics at Peking University in Beijing to present my work, and I visited the institute three times. I was also invited to CSIC in Barcelona for collaboration. In addition, I took part in many schools and conferences, where I learned a great deal about the LISA mission and the self-force program.

I had initially accepted a PhD position in Jena, Germany, to work on AI in gravitational wave data analysis. However, after attending an HPC school in Kobe, Japan, and speaking with professionals there, I realized that this was not the topic I wanted to pursue after all. After rejecting a second offer to work on mixed binaries, I eventually interviewed for a position at the Max Planck Institute for Gravitational Physics to work on the self-force

program. Fortunately, I received an offer, which I accepted, and I began working on the self-force program in February 2020.

I have always found it appealing to work on different topics, and this is reflected in the present thesis. Its four chapters are only loosely connected, but that also reflects my broader view of science: it should be multidisciplinary and not overly specialized. The universe is a complex system, and we should not restrict ourselves to studying only a tiny part of it with a single method.

During my PhD, I was fortunate to visit Perimeter Institute in Waterloo, Canada, every year.

Levensloop

Stamatis Vretinaris werd geboren op 19 januari 1994 in Düsseldorf, Duitsland. Hoewel het misschien als een cliche klinkt, begon mijn interesse in de mysteries van het universum al echt in mijn jeugd. Als ik naar de sterren keek, vroeg ik mij steeds af waaruit ze bestonden, hoe ver weg ze waren en of er misschien leven op andere planeten bestond. Die nieuwsgierigheid heeft mij er uiteindelijk toe gebracht een loopbaan in de astrofysica na te streven.

In 2012 begon ik aan mijn bacheloropleiding Natuurkunde, gevolgd in 2017 door een masteropleiding Computationele Fysica. In 2016 begon ik in het kader van mijn bachelorscriptie aan onderzoek naar zwaartekrachtsgolven, kort voor de officiële aankondiging van de eerste detectie van zwaartekrachtsgolven door de LIGO-samenwerking [4].

Voor mijn masterscriptie werkte ik aan machine-learningtoepassingen en de data-analyse van zwaartekrachtsgolven van binaire neutronensterfusies, met een focus op het post-mergerregime.

Tijdens mijn masterstudie had ik het geluk veel te kunnen reizen. Ik werd uitgenodigd door het Kavli Institute for Astronomy and Astrophysics aan Peking University in Beijing om mijn werk te presenteren, en ik heb het instituut drie keer bezocht. Ook werd ik uitgenodigd door CSIC in Barcelona voor samenwerking. Daarnaast nam ik deel aan veel scholen en conferenties, waar ik veel heb geleerd over de LISA-missie en het self-force-programma.

Aanvankelijk had ik een promotieplaats in Jena, Duitsland, geaccepteerd om te werken aan AI in zwaartekrachtsgolf-data-analyse. Na deelname aan een HPC-school in Kobe, Japan, en gesprekken met professionals daar, besefte ik echter dat dit toch niet het onderwerp was waarin ik verder wilde gaan. Nadat ik een tweede aanbod om aan mixed binaries te werken had

afgewezen, solliciteerde ik uiteindelijk op een positie bij het Max Planck Institute for Gravitational Physics om aan het self-force-programma te werken. Gelukkig kreeg ik een aanbod, dat ik accepteerde, en in februari 2020 begon ik aan mijn werk binnen het self-force-programma.

Ik heb het altijd aantrekkelijk gevonden om aan verschillende onderwerpen te werken, en dat komt ook in dit proefschrift tot uiting. De vier hoofdstukken hangen slechts losjes met elkaar samen, maar dat weerspiegelt ook mijn bredere visie op wetenschap: zij moet multidisciplinair zijn en niet te sterk gespecialiseerd. Het universum is een complex systeem, en we zouden ons niet moeten beperken tot het bestuderen van slechts een klein deel ervan met slechts één methode.

Tijdens mijn promotietraject had ik het geluk elk jaar het Perimeter Institute in Waterloo, Canada, te kunnen bezoeken.

Description of Research Data Management

This chapter describes the research data management relevant to this thesis in accordance with Article 72. It addresses the research data used in the thesis, the way these data were processed, how they were stored during the research, and how data and code are made available.

Scope of the research data

The thesis consists of four scientific chapters with different relationships to research data.

The chapter on marginally outer trapped surfaces (MOTSs) uses simulation data produced within the collaboration. The underlying numerical-relativity simulations were carried out by a collaborator, while the analysis in this thesis uses the resulting horizon data and derived quantities needed for the cusp-fitting and multipole analysis.

The chapter on binary-neutron-star (BNS) post-merger inference uses waveform data from the `CoRe` data set together with additional waveforms referred to in the chapter as the Soultanis waveforms. These waveforms form the research data on which the post-merger analysis and parameter-estimation study are based.

The scalar self-force chapter does not rely on an external research data set. The numerical results in that chapter were generated by simulations performed by the author. In this case, the research outputs consist primarily of

simulation results that are reproducible from the code and input parameters used in the study.

The chapter on summation-by-parts (SBP) operators likewise does not use an external research data set. The figures and numerical results shown there were generated from code written by the author, together with cited software infrastructure where relevant. The main research outputs are therefore the source code and the reproducible numerical results generated from it.

Processing of the research data

MOTS chapter

For the MOTS chapter, the starting point is the simulation output generated within the collaboration. From these outputs, the geometrical quantities relevant to the cusp-formation problem were extracted and post-processed. This included the computation and fitting of quantities such as horizon shape diagnostics, curvature quantities, and multipole-related observables near the touching time. The data used in the thesis are therefore derived analysis products constructed from the original simulation output rather than newly measured experimental data.

BNS chapter

For the BNS chapter, the waveform data were processed to produce frequency-domain and time-domain representations, posterior distributions, and derived parameter estimates. The analysis combines the waveform data with the damped-oscillator modelling and inference pipeline described in the chapter. Derived products include fitted post-merger frequencies, posterior samples, corner plots, violin plots, and the figures shown in the thesis. The scientific interpretation is based on these processed and derived products rather than on unmodified raw files alone.

Self-force and SBP chapters

For the self-force and SBP chapters, the numerical results were generated directly from source code developed by the author. In these cases, the relevant workflow is simulation followed by analysis and plot generation. The research outputs consist of source code, input parameters, numerical solutions, and derived figures. Since these chapters do not depend on external data sets, reproducibility is primarily ensured through the public availability of the analysis and simulation code.

During the course of the research, the data, simulation outputs, and derived analysis products were stored in digital form on research computing and local working systems used by the author and, where applicable, by collaborators. Working copies of manuscripts, source code, derived tables, and figures were maintained in structured project directories. Version control was used for the source material that underlies the thesis text and code repositories, and derived figures included in the thesis were kept together with the associated chapter material.

For chapters based on externally produced or collaboratively produced simulation outputs, only the data needed for the analysis presented in the thesis were retained locally in working form. For the chapters based on the author's own simulations, the numerical outputs can be recreated from the publicly available code and the documented analysis workflow.

Availability of data and code

Data availability

The BNS chapter uses waveform data from the publicly available `CoRe` data set as well as additional Soultanis waveforms used in the analysis. The MOTS chapter uses simulation data produced within the collaboration and derived analysis products based on those simulations. The self-force and SBP chapters do not rely on external data sets; their numerical results are reproducible from the code bases listed below.

Code availability

The software used to produce the analyses in this thesis is made available through the following public repositories:

- BNS chapter analysis code:
  `https://github.com/svretina/PythiaBNS`
- SBP chapter code:
  `https://github.com/svretina/SphericalSBPOperators.jl`
- MOTS chapter analysis code:
  `https://github.com/svretina/MOTSCuspFitting.jl`
- Self-force chapter code:
  `https://github.com/svretina/ScalarWaveFVM.jl`

These repositories provide the main analysis and simulation codes needed to reproduce the results presented in the thesis. For the SBP chapter and the self-force chapter, the code constitutes the primary reproducibility record. For the BNS and MOTS chapters, the code documents the analysis applied to the waveform or simulation data and thereby makes the processing pipeline transparent.

Field-specific compliance

The data-management approach in this thesis follows common standards in computational gravitational physics and numerical relativity. In particular, it distinguishes between externally available data, collaboratively produced simulation data, and numerical results that are reproducible from publicly released code. The key principle is that the scientific conclusions of the thesis are either directly traceable to publicly available data products or reproducible from publicly available analysis and simulation software.

Acknowledgements

Financing

The PhD research was financially supported through PhD positions at Radboud University Nijmegen and the Max Planck Institute for Gravitational Physics (Albert Einstein Institute).

I gratefully acknowledge the hospitality of the Perimeter Institute during repeated visits over the course of my PhD. The research environment there, the many stimulating discussions, and the people I had the opportunity to meet contributed significantly to the completion of this thesis and strengthened my motivation to continue research in this field.

I also acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), and the research support provided by the Aristotle University of Thessaloniki High Performance Computing Infrastructure and Resources.

Acknowledgements

I would like to thank my supervisors, Badri Krishnan and Erik Schnetter, for supporting me and guiding me through my PhD, especially during the hard times.

I thank Luis Lehner, Lars Andersson, Abhay Ashtekar, Ivan Booth, José Luis Jaramillo, Ricardo Uribe-Vargas, Andreas Bauswein, Nikolaos Karnesis, Theodoros Soultanis, Nikolaos Chatzarakis, and Sebastiano Bernuzzi for

helpful discussions, suggestions, and comments on the work.

I would also like to thank my family for their continuous support. Without their sacrifices I would not be here today. Thank you Mom and Dad. I would like to thank my brother Giorgos, who always believed in me and supported me emotionally but also scientifically.

Further I want to thank my partner Iris Papakonstantinou-Presvelou and all my friends who are scattered around the world. Just to name a few: Nikos Chatzarakis, Evgenia Vogiatzi, Stylianos Tzanetis, Dimitris Sioumpouras, Fani Alexandri, Athanasios Kogios, Aristomenis Yfantis, Armin Pornaki, Renan Assimos Martins, Thomas Koutalidis, Kostas Tzanavaris, Pau Amaro Seoane, Kostas Papanikolaou, Christos Mermigkas, Max, Adam, Gabriele Bozzola, Alvaro, Christina Kapsomani, Phillip Willians Leite, Vasilis Tzallas, Alexandros Emmanouilidis. I apologize if I forgot to mention you, but know that I love you all. Thank you for making the last years more enjoyable and interesting.

Bibliography

[1] http://www.computational-relativity.org.

[2] FastGaussQuadrature.jl Julia package, 2025. `https://github.com/JuliaApproximation/FastGaussQuadrature.jl`.

[3] Aasi J., et al. Advanced ligo. *Classical and Quantum Gravity*, 32(7):074001, mar 2015.

[4] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016.

[5] B. P. Abbott et al. Search for Post-merger Gravitational Waves from the Remnant of the Binary Neutron Star Merger GW170817. *Astrophys. J. Lett.*, 851(1):L16, 2017.

[6] B. P. Abbott et al. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Rev. Rel.*, 23(1):3, 2020.

[7] B.P. Abbott et al. GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run. *Phys. Rev. X*, 13:041039, Dec 2023.

[8] Acernese F., et al. Advanced virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, dec 2014.

[9] Ackley K., et al. Neutron star extreme matter observatory: A kilohertz-band gravitational-wave detector in the global network. *Publications of the Astronomical Society of Australia*, 37, 2020.

[10] Michalis Agathos, Francesco Zappa, Sebastiano Bernuzzi, Albino Perego, Matteo Breschi, and David Radice. Inferring prompt black-hole formation in neutron star mergers from gravitational-wave data. *Physical Review D*, 101(4), feb 2020.

[11] Warren G. Anderson and Alan G. Wiseman. A matched expansion approach to practical self-force calculations. *Class. Quant. Grav.*, 22:S783–S800, 2005.

[12] Lars Andersson, Marc Mars, Jan Metzger, and Walter Simon. The Time evolution of marginally trapped surfaces. *Class. Quant. Grav.*, 26:085018, 2009.

[13] Lars Andersson, Marc Mars, and Walter Simon. Local existence of dynamical and trapping horizons. *Phys.Rev.Lett.*, 95:111102, 2005.

[14] Lars Andersson, Marc Mars, and Walter Simon. Stability of marginally outer trapped surfaces and existence of marginally outer trapped tubes. *Adv.Theor.Math.Phys.*, 12, 2008.

[15] Marcus Ansorg, Bernd Brügmann, and Wolfgang Tichy. A single-domain spectral method for black hole puncture data. *Phys. Rev. D*, 70:064011, 2004.

[16] Theocharis A. Apostolatos. Search templates for gravitational waves from precessing, inspiraling binaries. *Phys. Rev. D*, 52(2):605–620, July 1995.

[17] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries. *Phys. Rev. D*, 49:6274–6297, 1994.

[18] M. Arimoto et al. Gravitational wave physics and astronomy in the nascent era. *Progress of Theoretical and Experimental Physics*, 2023(10):10A103, 04 2021.

[19] Stefano Ascenzi, Vanessa Graber, and Nanda Rea. Neutron-star measurements in the multi-messenger Era. *Astroparticle Physics*, 158:102935, June 2024.

[20] Abhay Ashtekar, Miguel Campiglia, and Samir Shah. Dynamical Black Holes: Approach to the Final State. *Phys. Rev.*, D88(6):064045, 2013.

[21] Abhay Ashtekar, Jonathan Engle, Tomasz Pawłowski, and Chris Van Den Broeck. Multipole moments of isolated horizons. *Class. Quant. Grav.*, 21:2549–2570, 2004.

[22] Abhay Ashtekar, Neev Khera, Maciej Kolanowski, and Jerzy Lewandowski. Charges and fluxes on (perturbed) non-expanding horizons. *JHEP*, 02:066, 2022.

[23] Abhay Ashtekar and Badri Krishnan. Dynamical horizons and their properties. *Phys. Rev.*, D68:104030, 2003.

[24] Abhay Ashtekar and Badri Krishnan. Isolated and dynamical horizons and their applications. *Living Rev. Rel.*, 7:10, 2004.

[25] Abhay Ashtekar and Badri Krishnan. Quasi-Local Black Hole Horizons: Recent Advances, 2 2025.

[26] Gregory Ashton, Moritz Hübner, Paul D. Lasky, Colm Talbot, Kendall Ackley, Sylvia Biscoveanu, Qi Chu, Atul Divakarla, Paul J. Easter, Boris Goncharov, Francisco Hernandez Vivanco, Jan Harms, Marcus E. Lower, Grant D. Meadors, Denyz Melchor, Ethan Payne, Matthew D. Pitkin, Jade Powell, Nikhil Sarin, Rory J. E. Smith, and Eric Thrane. Bilby: A user-friendly bayesian inference library for gravitational-wave astronomy. *The Astrophysical Journal Supplement Series*, 241(2):27, April 2019.

[27] Yoichi Aso, Yuta Michimura, Kentaro Somiya, Masaki Ando, Osamu Miyakawa, Takanori Sekiguchi, Daisuke Tatsumi, and Hiroaki Yamamoto. Interferometer design of the KAGRA gravitational wave detector. *Phys. Rev. D*, 88(4):043007, August 2013.

[28] Aso Y., et al. Interferometer design of the kagra gravitational wave detector. *Phys. Rev. D*, 88:043007, Aug 2013.

[29] M. Bailes et al. Gravitational-wave physics and astronomy in the 2020s and 2030s. *Nature Reviews Physics*, 3(5):344–366, April 2021.

[30] Luca Baiotti. Gravitational waves from neutron star mergers and their relation to the nuclear equation of state. *Progress in Particle and Nuclear Physics*, 109:103714, November 2019.

[31] Luca Baiotti. Gravitational waves from binary neutron stars. *Arabian Journal of Mathematics*, 11, 01 2022.

[32] Luca Baiotti. Gravitational waves from binary neutron stars. *Arabian Journal of Mathematics*, 11(1):105–118, 2022.

[33] Luca Baiotti and Luciano Rezzolla. Binary neutron star mergers: a review of einstein's richest laboratory. *Reports on Progress in Physics*, 80(9):096901, jul 2017.

[34] Jamie Bamber, Antonios Tsokaros, Milton Ruiz, and Stuart L. Shapiro. Jetlike structures in low-mass binary neutron star merger remnants. *Phys. Rev. D*, 110(2):024046, July 2024.

[35] Jamie Bamber, Antonios Tsokaros, Milton Ruiz, and Stuart L. Shapiro. Post-merger gravitational wave signals from binary neutron stars: Effect of the magnetic field. *arXiv e-prints*, page arXiv:2411.00943, November 2024.

[36] Leor Barack. Gravitational self-force in extreme mass-ratio inspirals. *Classical and Quantum Gravity*, 26(21):213001, oct 2009.

[37] Leor Barack and Darren A. Golbourn. Scalar-field perturbations from a particle orbiting a black hole using numerical evolution in 2+1 dimensions. *Phys. Rev. D*, 76:044020, 2007.

[38] Leor Barack, Yasushi Mino, Hiroyuki Nakano, Amos Ori, and Misao Sasaki. Calculating the gravitational self-force in schwarzschild spacetime. *Phys. Rev. Lett.*, 88:091101, 2002.

[39] Leor Barack and Amos Ori. Mode sum regularization approach for the self-force in black hole space-time. *Phys. Rev. D*, 61:061502, 2000.

[40] Enrico Barausse, Carlos Palenzuela, Marcelo Ponce, and Luis Lehner. Neutron-star mergers in scalar-tensor theories of gravity. *Phys. Rev. D*, 87(8):081506, April 2013.

[41] Enrico Barausse, Carlos Palenzuela, Marcelo Ponce, and Luis Lehner. Neutron-star mergers in scalar-tensor theories of gravity. *Phys. Rev. D*, 87:081506, Apr 2013.

[42] Enrico Barausse and Luciano Rezzolla. Predicting the direction of the final spin from the coalescence of two black holes. *Astrophys. J. Lett.*, 704:L40–L44, 2009.

[43] A. Bauswein and H.-T. Janka. Measuring neutron-star properties via gravitational waves from neutron-star mergers. *Phys. Rev. Lett.*, 108:011101, Jan 2012.

[44] A. Bauswein, H.-T. Janka, K. Hebeler, and A. Schwenk. Equation-of-state dependence of the gravitational-wave signal from the ring-down phase of neutron-star mergers. *Phys. Rev. D*, 86:063001, Sep 2012.

[45] A. Bauswein and N. Stergioulas. Unified picture of the post-merger dynamics and gravitational wave emission in neutron star mergers. *Phys. Rev. D*, 91:124056, Jun 2015.

[46] A Bauswein and N Stergioulas. Spectral classification of gravitational-wave emission and equation of state constraints in binary neutron star mergers. *Journal of Physics G: Nuclear and Particle Physics*, 46(11):113002, sep 2019.

[47] A. Bauswein and N. Stergioulas. Spectral classification of gravitational-wave emission and equation of state constraints in binary neutron star mergers. *Journal of Physics G Nuclear Physics*, 46(11):113002, November 2019.

[48] Andreas Bauswein, Niels-Uwe F. Bastian, David B. Blaschke, Katerina Chatziioannou, James A. Clark, Tobias Fischer, and Micaela Oertel. Identifying a first-order phase transition in neutron-star mergers through gravitational waves. *Phys. Rev. Lett.*, 122:061102, Feb 2019.

[49] Andreas Bauswein, Niels-Uwe Friedrich Bastian, David Blaschke, Katerina Chatziioannou, James Alexander Clark, Tobias Fischer, Hans-Thomas Janka, Oliver Just, Micaela Oertel, and Nikolaos Stergioulas. Equation-of-state constraints and the QCD phase transition in the era of gravitational-wave astronomy. In *Xiamen-CUSTIPEN Workshop on the Equation of State of Dense Neutron-Rich Matter in the Era of Gravitational Wave Astronomy*, volume 2127 of *American Institute of Physics Conference Series*, page 020013, July 2019.

[50] Andreas Bauswein, Sebastian Blacker, Georgios Lioutas, Theodoros Soultanis, Vimal Vijayan, and Nikolaos Stergioulas. Systematics of prompt black-hole formation in neutron star mergers. *Phys. Rev. D*, 103(12):123004, 2021.

[51] Andreas Bauswein, Sebastian Blacker, Vimal Vijayan, Nikolaos Stergioulas, Katerina Chatziioannou, James A. Clark, Niels-Uwe F. Bastian, David B. Blaschke, Mateusz Cierniak, and Tobias Fischer. Equation of state constraints from the threshold binary mass for prompt collapse of neutron star mergers. *Physical Review Letters*, 125(14), sep 2020.

[52] Andreas Bauswein, Nikolaos Stergioulas, and Hans-Thomas Janka. Exploring properties of high-density matter through remnants of neutron-star mergers. *The European Physical Journal A*, 52(3), March 2016.

[53] Andreas Bauswein, Nikolaos Stergioulas, and Hans-Thomas Janka. Exploring properties of high-density matter through remnants of neutron-star mergers. *European Physical Journal A*, 52:56, March 2016.

[54] Sebastiano Bernuzzi. Neutron star merger remnants. *General Relativity and Gravitation*, 52(11):108, November 2020.

[55] Sebastiano Bernuzzi and Tim Dietrich. Gravitational waveforms from binary neutron star mergers with high-order weighted-essentially-nonoscillatory schemes in numerical relativity. *Physical Review D*, 94(6), September 2016.

[56] Sebastiano Bernuzzi, Tim Dietrich, and Alessandro Nagar. Modeling the complete gravitational wave spectrum of neutron star mergers. *Physical Review Letters*, 115(9), aug 2015.

[57] Emanuele Berti et al. Black hole spectroscopy: from theory to experiment, 5 2025.

[58] Emanuele Berti, Alberto Sesana, Enrico Barausse, Vitor Cardoso, and Krzysztof Belczynski. Spectroscopy of Kerr black holes with Earth- and space-based interferometers. *Phys. Rev. Lett.*, 117(10):101102, 2016.

[59] Michael Betancourt. Cruising the simplex: Hamiltonian monte carlo and the dirichlet distribution. In *AIP Conference Proceedings*. AIP, 2012.

[60] Miguel Bezares, Ricard Aguilera-Miret, Lotte ter Haar, Marco Crisostomi, Carlos Palenzuela, and Enrico Barausse. No Evidence of Kinetic Screening in Simulations of Merging Binary Neutron Stars beyond General Relativity. *Phys. Rev. Lett.*, 128(9):091103, March 2022.

[61] Miguel Bezares, Ricard Aguilera-Miret, Lotte ter Haar, Marco Crisostomi, Carlos Palenzuela, and Enrico Barausse. No evidence of kinetic screening in simulations of merging binary neutron stars beyond general relativity. *Phys. Rev. Lett.*, 128:091103, Mar 2022.

[62] Sathvik Bhagavan, Bart de Koning, Shubham Maddhashiya, and Christopher Rackauckas. Datainterpolations.jl: Fast interpolations of 1d data. *Journal of Open Source Software*, 9(101):6917, 2024.

[63] Ashish Bhole, Boniface Nkonga, Stanislas Pamela, Guido Huijsmans, and Matthias Hoelzl. Treatment of polar grid singularities in the bi-cubic hermite-bézier approximations: Isoparametric finite element framework. *Journal of Computational Physics*, 471:111611, 2022.

[64] Sebastian Blacker, Niels-Uwe F. Bastian, Andreas Bauswein, David B. Blaschke, Tobias Fischer, Micaela Oertel, Theodoros Soultanis, and Stefan Typel. Constraining the onset density of the hadron-quark phase transition with gravitational-wave observations. *Phys. Rev. D*, 102(12):123023, December 2020.

[65] Sebastian Blacker, Hristijan Kochankovski, Andreas Bauswein, Angels Ramos, and Laura Tolos. Thermal behavior as indicator for hyperons in binary neutron star merger remnants. *Phys. Rev. D*, 109(4):043015, February 2024.

[66] Luc Blanchet, Alessandro Spallicci, and Bernard Whiting, editors. *Mass and motion in general relativity. Proceedings, School on Mass, Orleans, France, June 23-25, 2008*, volume 162, Dordrecht, 2011. Springer.

[67] Ivan Booth. Black hole boundaries. *Can. J. Phys.*, 83:1073–1099, 2005.

[68] Ivan Booth and Stephen Fairhurst. Isolated, slowly evolving, and dynamical trapping horizons: geometry and mechanics from surface deformations. *Phys. Rev.*, D75:084019, 2007.

[69] Ivan Booth, Robie A. Hennigar, and Daniel Pook-Kolb. Ultimate fate of apparent horizons during a binary black hole merger. I. Locating and understanding axisymmetric marginally outer trapped surfaces. *Phys. Rev. D*, 104(8):084083, 2021.

[70] Ssohrab Borhanian, K. G. Arun, Harald P. Pfeiffer, and B. S. Sathyaprakash. Comparison of post-Newtonian mode amplitudes with numerical relativity simulations of binary black holes. *Class. Quant. Grav.*, 37(6):065006, 2020.

[71] Sukanta Bose, Kabir Chakravarti, Luciano Rezzolla, B. S. Sathyaprakash, and Kentaro Takami. Neutron-star radius from a population of binary neutron star mergers. *Phys. Rev. Lett.*, 120:031102, Jan 2018.

[72] Latham Boyle and Michael Kesden. The spin expansion for binary black hole merger: new predictions and future directions. *Phys. Rev. D*, 78:024017, 2008.

[73] Latham Boyle, Michael Kesden, and Samaya Nissanke. Binary black hole merger: Symmetry and the spin expansion. *Phys. Rev. Lett.*, 100:151101, 2008.

[74] Branchesi M., et al. Science with the einstein telescope: a comparison of different designs. *Journal of Cosmology and Astroparticle Physics*, 2023(07):068, July 2023.

[75] Matteo Breschi, Sebastiano Bernuzzi, Kabir Chakravarti, Alessandro Camilletti, Aviral Prakash, and Albino Perego. Kilohertz gravitational waves from binary neutron star mergers: Numerical-relativity informed postmerger model. *Physical Review D*, 109(6), March 2024.

[76] Matteo Breschi, Sebastiano Bernuzzi, Daniel Godzieba, Albino Perego, and David Radice. Constraints on the maximum densities of neutron stars from postmerger gravitational waves with third-generation observations. *Physical Review Letters*, 128(16), apr 2022.

[77] Matteo Breschi, Sebastiano Bernuzzi, Francesco Zappa, Michalis Agathos, Albino Perego, David Radice, and Alessandro Nagar. Kilohertz gravitational waves from binary neutron star remnants: Time-domain model and constraints on extreme matter. *Phys. Rev. D*, 100(10):104029, November 2019.

[78] Matteo Breschi, Gregorio Carullo, and Sebastiano Bernuzzi. Pre/post-merger consistency test for gravitational signals from binary neutron star mergers. *Particles*, 6(3):731–738, July 2023.

[79] Dieter R. Brill and Richard W. Lindquist. Interaction energy in geometrostatics. *Phys. Rev.*, 131:471–476, 1963.

[80] J. David Brown, Peter Diener, Olivier Sarbach, Erik Schnetter, and Manuel Tiglio. Turduckening black holes: An Analytical and computational study. *Phys. Rev. D*, 79:044023, 2009.

[81] Ivan G Bubnov. Report on the works of professor timoshenko which were awarded the zhuranskyi prize. In *Symposium of the institute of communication engineers*, volume 81, 1913.

[82] Zhou-Jian Cao and Chen-Zhou Liu. Spin expansion for head-on collision of two black holes. *Int. J. Mod. Phys. D*, 20:43–57, 2011.

[83] Zack Carson, Andrew W. Steiner, and Kent Yagi. Future prospects for constraining nuclear matter parameters with gravitational waves. *Phys. Rev. D*, 100(2):023012, July 2019.

[84] Marc Casals, Sam Dolan, Adrian C. Ottewill, and Barry Wardell. Self-force and green function in schwarzschild spacetime via quasinormal modes and branch cut. *Phys. Rev. D*, 88:044022, 2013.

[85] Joan Centrella, John G. Baker, Bernard J. Kelly, and James R. van Meter. Black-hole binaries, gravitational waves, and numerical relativity. *Rev. Mod. Phys.*, 82:3069, 2010.

[86] Katerina Chatziioannou. Uncertainty limits on neutron star radius measurements with gravitational waves. *Physical Review D*, 105(8), April 2022.

[87] Katerina Chatziioannou, H. Thankful Cromartie, Stefano Gandolfi, Ingo Tews, David Radice, Andrew W. Steiner, and Anna L. Watts. Neutron stars and the dense matter equation of state: from microscopic theory to macroscopic observations. *arXiv e-prints*, page arXiv:2407.11153, July 2024.

[88] Yitian Chen, Prayush Kumar, Neev Khera, Nils Deppe, Arnab Dhani, Michael Boyle, Matthew Giesler, Lawrence E. Kidder, Harald P. Pfeiffer, Mark A. Scheel, and Saul A. Teukolsky. Multipole moments on the common horizon in a binary-black-hole simulation. *Physical Review D*, 106(12), December 2022.

[89] Matthew W. Choptuik. Universality and scaling in gravitational collapse of a massless scalar field. *Phys. Rev. Lett.*, 70:9–12, Jan 1993.

[90] J A Clark, A Bauswein, N Stergioulas, and D Shoemaker. Observing gravitational waves from the post-merger phase of binary neutron star coalescence. *Classical and Quantum Gravity*, 33(8):085003, mar 2016.

[91] Colombo A., et al. Multi-messenger observations of binary neutron star mergers in the o4 run. *The Astrophysical Journal*, 937(2):79, sep 2022.

[92] G.S. Constantinescu and S.K. Lele. A highly accurate technique for the treatment of flow equations at the polar axis in cylindrical coordinates using series expansions. *Journal of Computational Physics*, 183(1):165–186, 2002.

[93] Alexander W. Criswell, Jesse Miller, Noah Woldemariam, Theodoros Soultanis, Andreas Bauswein, Katerina Chatziioannou, Michael W. Coughlin, Galin Jones, and Vuk Mandic. Hierarchical bayesian method for constraining the neutron star equation of state with an ensemble of binary neutron star postmerger remnants. *Physical Review D*, 107(4), feb 2023.

[94] Thibault Damour and Alessandro Nagar. Final spin of a coalescing black-hole binary: An Effective-one-body approach. *Phys. Rev. D*, 76:044003, 2007.

[95] Thibault Damour, Alessandro Nagar, and Loïc Villain. Merger states and final states of black hole coalescences: a numerical-relativity-assisted effective-one-body approach. *Phys. Rev. D*, 89(2):024031, 2014.

[96] Simon Danisch and Julius Krumbiegel. Makie.jl: Flexible high-performance data visualization for Julia. *Journal of Open Source Software*, 6(65):3349, 2021.

[97] Roberto De Pietri, Alessandra Feo, José A. Font, Frank Löffler, Michele Pasquali, and Nikolaos Stergioulas. Numerical-relativity simulations of long-lived remnants of binary neutron star mergers. *Phys. Rev. D*, 101(6):064052, March 2020.

[98] David C. Del Rey Fernández, Pieter D. Boom, and David W. Zingg. A generalized framework for nodal first derivative summation-by-parts operators. *Journal of Computational Physics*, 266:214–239, 2014.

[99] David C. Del Rey Fernández, Jason E. Hicken, and David W. Zingg. Review of summation-by-parts operators with simultaneous approximation terms for the numerical solution of partial differential equations. *Computers & Fluids*, 95:171–196, 2014.

[100] David C. Del Rey Fernández and David W. Zingg. Generalized summation-by-parts operators for the second derivative. *SIAM Journal on Scientific Computing*, 37(6):A2840–A2864, 2015.

[101] Steven Detweiler. Perspective on gravitational self-force analyses. *Classical and Quantum Gravity*, 22(15):S681, jul 2005.

[102] Steven Detweiler. Gravitational radiation reaction and second order perturbation theory. *Phys. Rev. D*, 85:044048, 2012.

[103] Steven L. Detweiler and Bernard F. Whiting. Self-force via a green's function decomposition. *Phys. Rev. D*, 67:024025, 2003.

[104] Bryce S. DeWitt and Robert W. Brehme. Radiation damping in a gravitational field. *Annals Phys.*, 9:220–259, 1960.

[105] Peter Diener. Challenges in the self-consistent evolution of extreme mass ratio inspirals. Talk presented at the ICERM Workshop "Advances and Challenges in Computational Relativity", September 2020. September 16, 2020; in collaboration with Barry Wardell and Niels Warburton.

[106] Peter Diener, Ernst Nils Dorband, Erik Schnetter, and Manuel Tiglio. Optimized high-order derivative and dissipation operators satisfying summation by parts, and applications in three-dimensional multi-block evolutions. *J. Sci. Comput.*, 32:109–145, Jul 2007.

[107] Peter Diener, Ian Vega, Barry Wardell, and Steven Detweiler. Self-consistent orbital evolution of a particle around a schwarzschild black hole. *Phys. Rev. Lett.*, 108:191102, May 2012.

[108] Peter Diener, Ian Vega, Barry Wardell, and Steven Detweiler. Self-consistent orbital evolution of a particle around a schwarzschild black hole. *Phys. Rev. Lett.*, 108:191102, 2012.

[109] Tim Dietrich and Tanja Hinderer. Comprehensive comparison of numerical relativity and effective-one-body results to inform improvements in waveform models for binary neutron star systems. *Physical Review D*, 95(12), June 2017.

[110] Tim Dietrich, Tanja Hinderer, and Anuradha Samajdar. Interpreting binary neutron star mergers: describing the binary neutron star dynamics, modelling gravitational waveforms, and analyzing detections. *General Relativity and Gravitation*, 53(3):27, March 2021.

[111] Tim Dietrich, Maximiliano Ujevic, Wolfgang Tichy, Sebastiano Bernuzzi, and Bernd Bruegmann. *Phys. Rev. D*, 95(2):024029, 2017.

[112] Paul A. M. Dirac. Classical theory of radiating electrons. *Proc. Roy. Soc. Lond. A*, 167:148–169, 1938.

[113] Sam R. Dolan and Leor Barack. Self force via m-mode regularization and 2+1d evolution: Foundations and a scalar-field implementation on schwarzschild. *Phys. Rev. D*, 83:024019, 2011.

[114] Steve Drasco and Scott A. Hughes. Gravitational wave snapshots of generic extreme mass ratio inspirals. *Phys. Rev. D*, 73:024027, 2006.

[115] Olaf Dreyer, Bernard J. Kelly, Badri Krishnan, Lee Samuel Finn, David Garrison, and Ramon Lopez-Aleman. Black hole spectroscopy: Testing general relativity through gravitational wave observations. *Class. Quant. Grav.*, 21:787–804, 2004.

[116] Olaf Dreyer, Badri Krishnan, Deirdre Shoemaker, and Erik Schnetter. Introduction to isolated horizons in numerical relativity. *Phys. Rev. D*, 67:024018, 2003.

[117] William E. East and Frans Pretorius. Binary neutron star mergers in Einstein-scalar-Gauss-Bonnet gravity. *Phys. Rev. D*, 106(10):104055, November 2022.

[118] William E. East and Frans Pretorius. Binary neutron star mergers in einstein-scalar-gauss-bonnet gravity. *Phys. Rev. D*, 106:104055, Nov 2022.

[119] Paul J. Easter, Sudarshan Ghonge, Paul D. Lasky, Andrew R. Casey, James A. Clark, Francisco Hernandez Vivanco, and Katerina Chatziioannou. Detection and parameter estimation of binary neutron star merger remnants. *Phys. Rev. D*, 102:043011, Aug 2020.

[120] Paul J. Easter, Paul D. Lasky, Andrew R. Casey, Luciano Rezzolla, and Kentaro Takami. Computing fast and reliable gravitational waveforms of binary neutron star merger remnants. *Phys. Rev. D*, 100:043005, Aug 2019.

[121] Christian Ecker, Tyler Gorda, Aleksi Kurkela, and Luciano Rezzolla. Listening to the long ringdown: a novel way to pinpoint the equation of state in neutron-star cores. *arXiv e-prints*, page arXiv:2403.03246, March 2024.

[122] Einstein Toolkit: Open software for relativistic astrophysics. `http://einsteintoolkit.org/`.

[123] Roberto Emparan, Marina Martinez, and Miguel Zilhao. Black hole fusion in the extreme mass ratio limit. *Phys. Rev. D*, 97(4):044004, 2018.

[124] Pedro L. Espino, Aviral Prakash, David Radice, and Domenico Logoteta. Revealing phase transition in dense matter with gravitational wave spectroscopy of binary neutron star mergers. *Physical Review D*, 109(12), June 2024.

[125] Evans M., et al. A horizon study for cosmic explorer: Science, observatories, and community, 2021.

[126] Jacob Fields, Aviral Prakash, Matteo Breschi, David Radice, Sebastiano Bernuzzi, and André da Silva Schneider. Thermal Effects in Binary Neutron Star Mergers. *Astrophys. J. Lett.*, 952(2):L36, August 2023.

[127] Bruce A. Finlayson. *The Method of Weighted Residuals and Variational Principles*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2013.

[128] Xisco Jiménez Forteza and Pierre Mourier. High-overtone fits to numerical relativity ringdowns: Beyond the dismissed n=8 special tone. *Phys. Rev. D*, 104(12):124072, 2021.

[129] John L. Friedman and Nikolaos Stergioulas. *Rotating Relativistic Stars*. 2013.

[130] John L. Friedman and Nikolaos Stergioulas. Astrophysical implications of neutron star inspiral and coalescence. *International Journal of Modern Physics D*, 29(11):2041015, 2020.

[131] John L. Friedman and Nikolaos Stergioulas. Astrophysical implications of neutron star inspiral and coalescence. *International Journal of Modern Physics D*, 29(11):2041015, jul 2020.

[132] K. O. Friedrichs. Symmetric hyperbolic linear differential equations. *Communications on Pure and Applied Mathematics*, 7(2):345–392, 1954.

[133] K. O. Friedrichs. Symmetric positive linear differential equations. *Communications on Pure and Applied Mathematics*, 11(3):333–418, 1958.

[134] K. O. Friedrichs and P. D. Lax. Systems of conservation equations with a convex extension. *Proceedings of the National Academy of Sciences*, 68(8):1686–1688, 1971.

[135] Kurt Friedrichs and Hans Lewy. Über die eindeutigkeit und das abhängigkeitsgebiet der lösungen beim anfangswertproblem linearer hyperbolischer differentialgleichungen. *Mathematische Annalen*, 98(1):192–204, Mar 1928.

[136] Yuki Fujimoto, Kenji Fukushima, Kenta Hotokezaka, and Koutarou Kyutoku. Gravitational Wave Signal for Quark Matter with Realistic Phase Transition. *Phys. Rev. Lett.*, 130(9):091404, 2023.

[137] Boris Grigoryevich Galerkin. Series solution of some problems of elastic equilibrium of rods and plates. *Vestnik inzhenerov i tekhnikov*, 19(7):897–908, 1915.

[138] Sergei K. Godunov. A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Matematicheskii Sbornik*, 47(3):271–306, 1959. Original Russian title: *Raznostnyi metod chislennogo rascheta razryvnykh reshenii uravnenii gidrodinamiki*.

[139] Alejandra Gonzalez, Francesco Zappa, Matteo Breschi, Sebastiano Bernuzzi, David Radice, Ananya Adhikari, Alessandro Camilletti, Swami Vivekanandji Chaurasia, Georgios Doulis, Surendra Padamata, Alireza Rashti, Maximiliano Ujevic, Bernd Brügmann, William Cook, Tim Dietrich, Albino Perego, Amit Poudel, and Wolfgang Tichy. Second release of the core database of binary neutron star merger waveforms, 2022.

[140] Eric Gourgoulhon. A Generalized Damour-Navier-Stokes equation applied to trapping horizons. *Phys. Rev. D*, 72:104007, 2005.

[141] Eric Gourgoulhon, Alexandre Le Tiec, and Marc Casals. Horizon multipole moments of a kerr black hole. *Phys. Rev. D*, Apr 2026.

[142] Benjamin Grace, Karl Wette, and Susan M. Scott. Gravitational wave searches for postmerger remnants of GW170817 and GW190425. *Phys. Rev. D*, 110(8):083016, October 2024.

[143] Samuel E. Gralla. Second order gravitational self force. *Phys. Rev. D*, 85:124011, 2012.

[144] Samuel E. Gralla and Robert M. Wald. A rigorous derivation of gravitational self-force. *Class. Quant. Grav.*, 25:205009, 2008.

[145] Michael D Griffin, Everett Jones, and John D Anderson. A computational fluid dynamic technique valid at the centerline for non-axisymmetric problems in cylindrical coordinates. *Journal of Computational Physics*, 30(3):352–360, 1979.

[146] Carsten Gundlach, David Hilditch, and José M. Martín-García. Critical phenomena in gravitational collapse. *arXiv preprint arXiv:2507.07636*, 2025.

[147] Carsten Gundlach and José M. Martín-García. Critical phenomena in gravitational collapse. *Living Reviews in Relativity*, 10(1):5, Dec 2007.

[148] Carsten Gundlach, José M Martín-García, and David Garfinkle. Summation by parts methods for spherical harmonic decompositions of the wave equation in any dimensions. *Classical and Quantum Gravity*, 30(14):145003, jun 2013.

[149] Anshu Gupta, Badri Krishnan, Alex Nielsen, and Erik Schnetter. Dynamics of marginally trapped surfaces in a binary black hole merger: Growth and approach to equilibrium. *Phys. Rev.*, D97(8):084028, 2018.

[150] Bertil Gustafsson, Heinz-Otto Kreiss, and Joseph Oliger. *Time-Dependent Problems and Difference Methods*. Wiley, September 2013.

[151] Ernst Hairer, Syvert P. Nørsett, and Gerhard Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*, volume 8 of *Springer Series in Computational Mathematics*. Springer Berlin, Heidelberg, 2 edition, 1993.

[152] Ernst Hairer and Gerhard Wanner. *Solving Ordinary Differential Equations II*. Springer Berlin Heidelberg, 1996.

[153] Ami Harten. High resolution schemes for hyperbolic conservation laws. *Journal of Computational Physics*, 49(3):357–393, 1983.

[154] Carl-Johan Haster, Katerina Chatziioannou, Andreas Bauswein, and James Alexander Clark. Inference of the neutron star equation of state from cosmological distances. *Physical Review Letters*, 125(26), dec 2020.

[155] S. W. Hawking. Gravitational radiation from colliding black holes. *Phys. Rev. Lett.*, 26:1344–1346, May 1971.

[156] James Healy, Carlos O. Lousto, and Yosef Zlochower. Remnant mass, spin, and recoil from spin aligned black-hole binaries. *Phys. Rev. D*, 90(10):104004, 2014.

[157] Anna Heffernan, Adrian C Ottewill, Niels Warburton, Barry Wardell, and Peter Diener. Accelerated motion and the self-force in schwarzschild spacetime. *Classical and Quantum Gravity*, 35(19):194001, sep 2018.

[158] Hild S., et al. Sensitivity studies for third-generation gravitational wave observatories. *Classical and Quantum Gravity*, 28(9):094013, apr 2011.

[159] Fabian Hofmann, Enrico Barausse, and Luciano Rezzolla. The final spin from binary black holes in quasi-circular orbits. *Astrophys. J. Lett.*, 825(2):L19, 2016.

[160] Kenta Hotokezaka, Kenta Kiuchi, Koutarou Kyutoku, Takayuki Muranushi, Yu-ichiro Sekiguchi, Masaru Shibata, and Keisuke Taniguchi. Remnant massive neutron stars of binary neutron star mergers: Evolution process and gravitational waveform. *Phys. Rev. D*, 88:044026, Aug 2013.

[161] Nick Huggett. Zeno's Paradoxes. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2025 edition, 2025.

[162] Scott A. Hughes, Steve Drasco, Eanna E. Flanagan, and Joel Franklin. Gravitational radiation reaction and inspiral waveforms in the adiabatic limit. *Phys. Rev. Lett.*, 94:221101, 2005.

[163] Sascha Husa, Ian Hinder, and Christiane Lechner. Kranc: a Mathematica application to generate numerical codes for tensorial evolution equations. *Comput. Phys. Commun.*, 174:983–1004, 2006.

[164] Sabrina Huth, Peter T. H. Pang, Ingo Tews, Tim Dietrich, Arnaud Le Fèvre, Achim Schwenk, Wolfgang Trautmann, Kshitij Agarwal, Mattia

Bulla, Michael W. Coughlin, and Chris Van Den Broeck. Constraining neutron-star matter with microscopic and macroscopic collisions. *Nature*, 606(7913):276–280, June 2022.

[165] Rachael Huxford, Rahul Kashyap, Ssohrab Borhanian, Arnab Dhani, Ish Gupta, and B. S. Sathyaprakash. Accuracy of neutron star radius measurement with the next generation of terrestrial gravitational-wave observatories. *Phys. Rev. D*, 109(10):103035, May 2024.

[166] Francesco Iacovelli, Michele Mancarella, Chiranjib Mondal, Anna Puecher, Tim Dietrich, Francesca Gulminelli, Michele Maggiore, and Micaela Oertel. Nuclear physics constraints from binary neutron star mergers in the einstein telescope era. *Physical Review D*, 108(12), December 2023.

[167] Dante A. B. Iozzo et al. Comparing Remnant Properties from Horizon Data and Asymptotic Data in Numerical Relativity. *Phys. Rev. D*, 103(12):124029, 2021.

[168] M Jacobi, F M Guercilena, S Huth, G Ricigliano, A Arcones, and A Schwenk. Effects of nuclear matter properties in neutron star mergers. *Monthly Notices of the Royal Astronomical Society*, 527(3):8812–8828, November 2023.

[169] Tamanna Jain and Michalis Agathos. Improving inference on neutron star properties using information from binary merger remnants. *arXiv e-prints*, page arXiv:2404.12126, April 2024.

[170] J. L. Jaramillo, R. P. Macedo, P. Mösta, and L. Rezzolla. Towards a cross-correlation approach to strong-field dynamics in Black Hole spacetimes. *AIP Conf. Proc.*, 1458:158–173, 2011.

[171] Jose Luis Jaramillo, Rodrigo P. Macedo, Philipp Mösta, and Luciano Rezzolla. Black-hole horizons as probes of black-hole dynamics II: geometrical insights. *Phys. Rev.*, D85:084031, 2012.

[172] Minas Karamanis, Florian Beutler, John A Peacock, David Nabergoj, and Uroš Seljak. Accelerating astronomical and cosmological infer-

ence with preconditioned monte carlo. *Monthly Notices of the Royal Astronomical Society*, 516(2):1644–1653, August 2022.

[173] Minas Karamanis, David Nabergoj, Florian Beutler, John A. Peacock, and Uroš Seljak. pocomc: A python package for accelerated bayesian inference in astronomy and cosmology. *Journal of Open Source Software*, 7(79):4634, November 2022.

[174] Wolfgang Kastaun, Beatrix Willburger, and Kostas D. Kokkotas. On the saturation amplitude of the f-mode instability. *Phys. Rev. D*, 82:104036, 2010.

[175] Atul Kedia, Hee Il Kim, In-Saeng Suh, and Grant J. Mathews. Binary neutron star mergers as a probe of quark-hadron crossover equations of state. *Phys. Rev. D*, 106(10):103027, 2022.

[176] Neev Khera, Ariadna Ribes Metidieri, Béatrice Bonga, Xisco Jiménez Forteza, Badri Krishnan, Eric Poisson, Daniel Pook-Kolb, Erik Schnetter, and Huan Yang. Nonlinear Ringdown at the Black Hole Horizon. *Phys. Rev. Lett.*, 131(23):231401, 2023.

[177] Kranc: Kranc assembles numerical code.

[178] H. O. Kreiss and G. Scherer. *Mathematical Aspects of Finite Elements in Partial Differential Equations, edited by C. D. Boor.* Academica Press, 1974.

[179] H. O. Kreiss and G. Scherer. On the existence of energy estimates for difference approximations for hyperbolic systems. Technical report, Department of Scientific Computing, Uppsala University, 1977.

[180] Badri Krishnan. Fundamental properties and applications of quasi-local black hole horizons. *Class. Quant. Grav.*, 25:114005, 2008.

[181] Andrzej Królak, Piotr Jaranowski, Michał Bejger, Paweł Ciecieląg, Orest Dorosh, and Andrzej Pisarski. Search for postmerger gravitational waves from binary neutron star mergers using a matched-filtering statistic. *Classical and Quantum Gravity*, 40(21):215008, November 2023.

[182] Hao-Jui Kuan, Alan Tsz-Lok Lam, Daniela D. Doneva, Stoytcho S. Yazadjiev, Masaru Shibata, and Kenta Kiuchi. Dynamical scalarization during neutron star mergers in scalar-Gauss-Bonnet theory. *Phys. Rev. D*, 108(6):063033, September 2023.

[183] Hao-Jui Kuan, Alan Tsz-Lok Lam, Daniela D. Doneva, Stoytcho S. Yazadjiev, Masaru Shibata, and Kenta Kiuchi. Dynamical scalarization during neutron star mergers in scalar-gauss-bonnet theory. *Phys. Rev. D*, 108:063033, Sep 2023.

[184] Hao-Jui Kuan, Karim Van Aelst, Alan Tsz-Lok Lam, and Masaru Shibata. Binary neutron star mergers in massive scalar-tensor theory: Quasiequilibrium states and dynamical enhancement of the scalarization. *Phys. Rev. D*, 108:064057, Sep 2023.

[185] Diogenes Laertius. *Lives of Eminent Philosophers, Volume I: Books 1-5*, volume 184 of *Loeb Classical Library*. Harvard University Press, Cambridge, MA, 1925.

[186] Alan Tsz-Lok Lam, Hao-Jui Kuan, Masaru Shibata, Karim Van Aelst, and Kenta Kiuchi. Binary neutron star mergers in massive scalar-tensor theory: Properties of postmerger remnants. *Phys. Rev. D*, 110(10):104018, November 2024.

[187] Henry Desmond Pritchard Lee. *Zeno of Elea*. Cambridge University Press, 2015.

[188] Luis Lehner and Frans Pretorius. Black strings, low viscosity fluids, and violation of cosmic censorship. *Phys. Rev. Lett.*, 105:101102, Sep 2010.

[189] Randall J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2002.

[190] Steven L. Liebling and Carlos Palenzuela. Dynamical boson stars. *Living Reviews in Relativity*, 20(1):5, Nov 2017.

[191] LIGO Scientific Collaboration. Advanced LIGO anticipated sensitivity curves. `https://dcc.ligo.org/LIGO-T0900288/public`, 2010. ZERO_DET_high_P.txt.

[192] LIGO Scientific Collaboration. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. `https://dcc.ligo.org/LIGO-P1200087-v42/public`, 2016. fig1_adv_sensitivity.txt.

[193] Georgios Lioutas, Andreas Bauswein, Theodoros Soultanis, Rüdiger Pakmor, Volker Springel, and Friedrich K. Röpke. General relativistic moving-mesh hydrodynamic simulations with AREPO and applications to neutron star mergers. *Mon. Not. R. Astron. Soc.*, 528(2):1906–1929, February 2024.

[194] Frank Löffler, Joshua Faber, Eloisa Bentivegna, Tanja Bode, Peter Diener, Roland Haas, Ian Hinder, Bruno C. Mundim, Christian D. Ott, Erik Schnetter, Gabrielle Allen, Manuela Campanelli, and Pablo Laguna. The Einstein Toolkit: A Community Computational Infrastructure for Relativistic Astrophysics. *Class. Quantum Grav.*, 29(11):115001, 2012.

[195] Carlos O. Lousto and Yosef Zlochower. Black hole binary remnant mass and spin: A new phenomenological formula. *Phys. Rev. D*, 89(10):104052, 2014.

[196] K. Mattsson, M. Svärd, and M. Shoeybi. Stable and accurate schemes for the compressible navier–stokes equations. *Journal of Computational Physics*, 227(4):2293–2316, 2008.

[197] Ken Mattsson. Summation by parts operators for finite difference approximations of second-derivatives with variable coefficients. *J. Sci. Comput.*, 51:650–682, Jun 2012.

[198] Ken Mattsson and Jan Nordström. Summation by parts operators for finite difference approximations of second derivatives. *Journal of Computational Physics*, 199(2):503–540, 2004.

[199] Ken Mattsson, Magnus Svärd, and Jan Nordström. Stable and accurate artificial dissipation. *Journal of Scientific Computing*, 21(1):57–79, Aug 2004.

[200] Raissa F. P. Mendes, Néstor Ortiz, and Nikolaos Stergioulas. Nonlinear dynamics of oscillating neutron stars in scalar-tensor gravity. *Phys. Rev. D*, 104(10):104036, November 2021.

[201] Raissa F. P. Mendes, Néstor Ortiz, and Nikolaos Stergioulas. Nonlinear dynamics of oscillating neutron stars in scalar-tensor gravity. *Phys. Rev. D*, 104:104036, Nov 2021.

[202] David Merritt and Milos Milosavljević. Massive Black Hole Binary Evolution. *Living Reviews in Relativity*, 8:8, November 2005.

[203] Yasushi Mino. Perturbative approach to an orbital evolution around a supermassive black hole. *Phys. Rev. D*, 67:084027, 2003.

[204] Yasushi Mino, Misao Sasaki, and Takahiro Tanaka. Gravitational radiation reaction to a particle motion. *Phys. Rev. D*, 55:3457–3476, 1997.

[205] Miquel Miravet-Tenés, Florencia L. Castillo, Roberto De Pietri, Pablo Cerdá-Durán, and José A. Font. Prospects for the inference of inertial modes from hypermassive neutron stars with future gravitational-wave detectors. *Phys. Rev. D*, 107(10):103053, May 2023.

[206] Miquel Miravet-Tenés, Davide Guerra, Milton Ruiz, Pablo Cerdá-Durán, and José A. Font. Identifying thermal effects in neutron star merger remnants with model-agnostic waveform reconstructions and third-generation detectors. *arXiv e-prints*, page arXiv:2401.02493, January 2024.

[207] Kamran Mohseni and Tim Colonius. Numerical treatment of polar coordinate singularities. *Journal of Computational Physics*, 157(2):787–795, 2000.

[208] Elias R. Most, Alexander Haber, Steven P. Harris, Ziyuan Zhang, Mark G. Alford, and Jorge Noronha. Emergence of Microphysical Bulk

Viscosity in Binary Neutron Star Postmerger Dynamics. *Astrophys. J. Lett.*, 967(1):L14, May 2024.

[209] Elias R. Most, L. Jens Papenfort, Veronica Dexheimer, Matthias Hanauske, Horst Stoecker, and Luciano Rezzolla. On the deconfinement phase transition in neutron-star mergers. *The European Physical Journal A*, 56(2):59, 2020.

[210] Elias R. Most, L. Jens Papenfort, Veronica Dexheimer, Matthias Hanauske, Stefan Schramm, Horst Stöcker, and Luciano Rezzolla. Signatures of quark-hadron phase transitions in general-relativistic neutron-star mergers. *Phys. Rev. Lett.*, 122:061101, Feb 2019.

[211] P. Mösta, L. Andersson, J. Metzger, B. Szilágyi, and J. Winicour. The Merger of Small and Large Black Holes. *Class. Quant. Grav.*, 32(23):235003, 2015.

[212] Pierre Mourier, Xisco Jiménez Forteza, Daniel Pook-Kolb, Badri Krishnan, and Erik Schnetter. Quasinormal modes and their overtones at the common horizon in a binary black hole merger. *Phys. Rev. D*, 103(4):044054, 2021.

[213] Bernhard Müller. Hydrodynamics of core-collapse supernovae and their progenitors. *Living Reviews in Computational Astrophysics*, 6(1):3, Jun 2020.

[214] Alex Nitz, Ian Harry, Duncan Brown, Christopher M. Biwer, Josh Willis, Tito Dal Canton, Collin Capano, Thomas Dent, Larne Pekowsky, Soumi De, Miriam Cabero, Gareth S Cabourn Davies, Andrew R. Williamson, Duncan Macleod, Bernd Machenschalk, Francesco Pannarale, Prayush Kumar, Steven Reyes, dfinstad, Sumit Kumar, Shichao Wu, Márton Tápai, Leo Singer, veronica villa, Sebastian Khan, Stephen Fairhurst, Koustav Chandra, Alex Nielsen, Shashwat Singh, and Thomas Massinger. gwastro/pycbc: v2.1.2 release of pycbc, May 2023.

[215] R. Oechslin and H.-T. Janka. Gravitational waves from relativistic neutron-star mergers with microphysical equations of state. *Phys. Rev. Lett.*, 99:121102, Sep 2007.

[216] Ossian O'Reilly and N. Anders Petersson. Energy conservative SBP discretizations of the acoustic wave equation in covariant form on staggered curvilinear grids. *Journal of Computational Physics*, 411:109386, 2020.

[217] Michael A. Pajkos and Elias R. Most. The Influence of Muons, Pions, and Trapped Neutrinos on Neutron Star Mergers. *arXiv e-prints*, page arXiv:2409.09147, September 2024.

[218] Carlos Palenzuela, Enrico Barausse, Marcelo Ponce, and Luis Lehner. Dynamical scalarization of neutron stars in scalar-tensor gravity theories. *Phys. Rev. D*, 89:044024, Feb 2014.

[219] F H Panther and P D Lasky. The effect of noise artefacts on gravitational-wave searches for neutron star post-merger remnants. *Monthly Notices of the Royal Astronomical Society*, 523(2):2928–2933, may 2023.

[220] Vasileios Paschalidis, William E. East, Frans Pretorius, and Stuart L. Shapiro. One-arm spiral instability in hypermassive neutron stars formed by dynamical-capture binary neutron star mergers. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 92(12), December 2015. Publisher Copyright: © 2015 American Physical Society.

[221] Dimitrios Pesios, Ioannis Koutalios, Dimitris Kugiumtzis, and Nikolaos Stergioulas. Predicting binary neutron star postmerger spectra using artificial neural networks. *Phys. Rev. D*, 110(6):063008, September 2024.

[222] GI Petrov. Application of galerkin's method to the problem of stability of flow of a viscous fluid. *J. Appl. Math. Mech*, 4(3):3–12, 1940.

[223] Roberto De Pietri, Alessandra Feo, José A. Font, Frank Löffler, Francesco Maione, Michele Pasquali, and Nikolaos Stergioulas. Convective excitation of inertial modes in binary neutron star mergers. *Physical Review Letters*, 120(22), may 2018.

[224] Eric Poisson, Adam Pound, and Ian Vega. The motion of point particles in curved spacetime. *Living Reviews in Relativity*, 14(1):7, Sep 2011.

[225] Eric Poisson and Alan G. Wiseman. Suggestion at the 1st capra ranch meeting on radiation reaction, 1998.

[226] Daniel Pook-Kolb, Ofek Birnholtz, José Luis Jaramillo, Badri Krishnan, and Erik Schnetter. Horizons in a binary black hole merger I: Geometry and area increase, 6 2020.

[227] Daniel Pook-Kolb, Ofek Birnholtz, José Luis Jaramillo, Badri Krishnan, and Erik Schnetter. Horizons in a binary black hole merger II: Fluxes, multipole moments and stability, 6 2020.

[228] Daniel Pook-Kolb, Ofek Birnholtz, Badri Krishnan, and Erik Schnetter. Existence and stability of marginally trapped surfaces in black-hole spacetimes. *Phys. Rev. D*, 99(6):064005, 2019.

[229] Daniel Pook-Kolb, Ofek Birnholtz, Badri Krishnan, and Erik Schnetter. Interior of a binary black hole merger. *Phys. Rev. Lett.*, 123:171102, Oct 2019.

[230] Daniel Pook-Kolb, Ofek Birnholtz, Badri Krishnan, and Erik Schnetter. MOTS Finder version 1.1, March 2019. 10.5281/zenodo.2591105.

[231] Daniel Pook-Kolb, Ofek Birnholtz, Badri Krishnan, and Erik Schnetter. Self-intersecting marginally outer trapped surfaces. *Phys. Rev. D*, 100:084044, Oct 2019.

[232] Daniel Pook-Kolb, Ofek Birnholtz, Badri Krishnan, Erik Schnetter, and Victor Zhang. MOTS Finder version 1.3, July 2019.

[233] Daniel Pook-Kolb, Ivan Booth, and Robie A. Hennigar. Ultimate fate of apparent horizons during a binary black hole merger. II. The vanishing of apparent horizons. *Phys. Rev. D*, 104(8):084084, 2021.

[234] Daniel Pook-Kolb, Robie A. Hennigar, and Ivan Booth. What Happens to Apparent Horizons in a Binary Black Hole Merger? *Phys. Rev. Lett.*, 127(18):181101, 2021.

[235] Adam Pound. Self-consistent gravitational self-force. *Phys. Rev. D*, 81:024023, 2010.

[236] Adam Pound. Second-order gravitational self-force. *Phys. Rev. Lett.*, 109:051101, 2012.

[237] Aviral Prakash, Ish Gupta, Matteo Breschi, Rahul Kashyap, David Radice, Sebastiano Bernuzzi, Domenico Logoteta, and B. S. Sathyaprakash. Detectability of QCD phase transitions in binary neutron star mergers: Bayesian inference with the next generation gravitational wave detectors. *Phys. Rev. D*, 109(10):103008, May 2024.

[238] Aviral Prakash, David Radice, Domenico Logoteta, Albino Perego, Vsevolod Nedora, Ignazio Bombaci, Rahul Kashyap, Sebastiano Bernuzzi, and Andrea Endrizzi. Signatures of deconfined quark phases in binary neutron star mergers. *Phys. Rev. D*, 104:083029, Oct 2021.

[239] Bo Prochnow, Ossian O'Reilly, Eric M. Dunham, and N. Anders Petersson. Treatment of the polar coordinate singularity in axisymmetric wave propagation using high-order summation-by-parts operators on a staggered grid. *Computers & Fluids*, 149:138–149, 2017.

[240] Anna Puecher and Tim Dietrich. Machine-learning classifier for the postmerger remnant of binary neutron stars. *Physical Review D*, 110(12), December 2024.

[241] Anna Puecher, Tim Dietrich, Ka Wa Tsang, Chinmay Kalaghatgi, Soumen Roy, Yoshinta Setyawati, and Chris Van Den Broeck. Unraveling information about supranuclear-dense matter from the complete binary neutron star coalescence process using future gravitational-wave detector networks. *Phys. Rev. D*, 107(12):124009, June 2023.

[242] Theodore C. Quinn and Robert M. Wald. An axiomatic approach to electromagnetic and gravitational radiation reaction of particles in curved space-time. *Phys. Rev. D*, 56:3381–3394, 1997.

[243] Christopher Rackauckas and Qing Nie. DifferentialEquations.jl–a performant and feature-rich ecosystem for solving differential equations in Julia. *Journal of Open Research Software*, 5(1), 2017.

[244] David Radice, Sebastiano Bernuzzi, and Christian D. Ott. One-armed spiral instability in neutron star mergers and its detectability in gravitational waves. *Physical Review D*, 94(6), September 2016.

[245] David Radice, Sebastiano Bernuzzi, and Albino Perego. The dynamics of binary neutron star mergers and gw170817. *Annual Review of Nuclear and Particle Science*, 70(1):95–119, 2020.

[246] David Radice, Sebastiano Bernuzzi, Walter Del Pozzo, Luke F. Roberts, and Christian D. Ott. Probing extreme-density matter with gravitational-wave observations of binary neutron star merger remnants. *The Astrophysical Journal*, 842(2):L10, June 2017.

[247] David Radice and Ian Hawke. Turbulence modelling in neutron star merger simulations. *Living Reviews in Computational Astrophysics*, 10(1):1, December 2024.

[248] David Radice, Albino Perego, Francesco Zappa, and Sebastiano Bernuzzi. GW170817: Joint constraint on the neutron star equation of state from multimessenger observations. *The Astrophysical Journal*, 852(2):L29, January 2018.

[249] Carolyn A. Raithel and Vasileios Paschalidis. Influence of stellar compactness on finite-temperature effects in neutron star merger simulations. *Physical Review D*, 108(8), October 2023.

[250] Carolyn A. Raithel and Vasileios Paschalidis. Detectability of finite-temperature effects from neutron star mergers with next-generation gravitational wave detectors. *Phys. Rev. D*, 110(4):043002, August 2024.

[251] Hendrik Ranocha. SummationByPartsOperators.jl: A Julia library of provably stable semidiscretization techniques with mimetic properties. *Journal of Open Source Software*, 6(64):3454, 08 2021.

[252] Luciano Rezzolla. Modelling the final state from binary black-hole coalescences. *Class. Quant. Grav.*, 26:094023, 2009.

[253] Luciano Rezzolla, Peter Diener, Ernst Nils Dorband, Denis Pollney, Christian Reisswig, Erik Schnetter, and Jennifer Seiler. The Final spin from the coalescence of aligned-spin black-hole binaries. *Astrophys. J. Lett.*, 674:L29–L32, 2008.

[254] Luciano Rezzolla, Rodrigo P. Macedo, and Jose Luis Jaramillo. Understanding the 'anti-kick' in the merger of binary black holes. *Phys. Rev. Lett.*, 104:221101, 2010.

[255] Giuseppe Rivieccio, Davide Guerra, Milton Ruiz, and José A. Font. Gravitational-wave imprints of nonconvex dynamics in binary neutron star mergers. *Phys. Rev. D*, 109(6):064032, March 2024.

[256] Carl L. Rodriguez, Michael Zevin, Pau Amaro-Seoane, Sourav Chatterjee, Kyle Kremer, Frederic A. Rasio, and Claire S. Ye. Black holes: The next generation—repeated mergers in dense star clusters and their gravitational-wave properties. *Phys. Rev. D*, 100:043027, Aug 2019.

[257] S. Ronchini, M. Branchesi, G. Oganesyan, B. Banerjee, U. Dupletsa, G. Ghirlanda, J. Harms, M. Mapelli, and F. Santoliquido. Perspectives for multimessenger astronomy with the next generation of gravitational-wave detectors and high-energy satellites. *Astron. Astrophys.*, 665:A97, September 2022.

[258] Eran Rosenthal. Regularization of the second-order gravitational perturbations produced by a compact object. *Phys. Rev. D*, 72:121503, 2005.

[259] Eran Rosenthal. Second-order gravitational self-force. *Phys. Rev. D*, 74:084018, 2006.

[260] V.V Rusanov. The calculation of the interaction of non-stationary shock waves and obstacles. *USSR Computational Mathematics and Mathematical Physics*, 1(2):304–320, 1962.

[261] Steven J. Ruuth. Global optimization of explicit strong-stability-preserving runge-kutta methods. *Mathematics of Computation*, 75(253):183–207, 2006.

[262] Laura Sagunski, Jun Zhang, Matthew C. Johnson, Luis Lehner, Mairi Sakellariadou, Steven L. Liebling, Carlos Palenzuela, and David Neilsen. Neutron star mergers as a probe of modifications of general relativity with finite-range scalar forces. *Phys. Rev. D*, 97(6):064016, March 2018.

[263] Laura Sagunski, Jun Zhang, Matthew C. Johnson, Luis Lehner, Mairi Sakellariadou, Steven L. Liebling, Carlos Palenzuela, and David Neilsen. Neutron star mergers as a probe of modifications of general relativity with finite-range scalar forces. *Phys. Rev. D*, 97:064016, Mar 2018.

[264] Nikhil Sarin and Paul D. Lasky. The evolution of binary neutron star post-merger remnants: a review. *General Relativity and Gravitation*, 53(6), jun 2021.

[265] Argyro Sasli, Nikolaos Karnesis, and Nikolaos Stergioulas. Exploring the potential for detecting rotational instabilities in binary neutron star merger remnants with gravitational wave detectors. *Phys. Rev. D*, 109:043045, Feb 2024.

[266] William E. Schiesser. *The Numerical Method of Lines: Integration of Partial Differential Equations*. Elsevier Science, 2012.

[267] L.F. Shampine. Solving ODEs and DDEs with residual control. *Applied Numerical Mathematics*, 52(1):113–127, 2005.

[268] Masaru Shibata. Constraining nuclear equations of state using gravitational waves from hypermassive neutron stars. *Physical Review Letters*, 94(20), May 2005.

[269] Masaru Shibata, Keisuke Taniguchi, Hirotada Okawa, and Alessandra Buonanno. Coalescence of binary neutron stars in a scalar-tensor theory of gravity. *Phys. Rev. D*, 89(8):084005, April 2014.

[270] Masaru Shibata, Keisuke Taniguchi, Hirotada Okawa, and Alessandra Buonanno. Coalescence of binary neutron stars in a scalar-tensor theory of gravity. *Phys. Rev. D*, 89:084005, Apr 2014.

[271] Masaru Shibata, Keisuke Taniguchi, and K ōji Uryū. Merger of binary neutron stars with realistic equations of state in full general relativity. *Phys. Rev. D*, 71:084021, Apr 2005.

[272] Ralph A. Smith. GenericSchur.jl: Julia package for Schur decomposition of matrices with generic element types, 2025.

[273] Theodoros Soultanis, Andreas Bauswein, and Nikolaos Stergioulas. Analytic models of the spectral properties of gravitational waves from neutron star merger remnants. *Phys. Rev. D*, 105:043020, Feb 2022.

[274] Theodoros Soultanis, Kiril Maltsev, Andreas Bauswein, Katerina Chatziioannou, Friedrich K. Röpke, and Nikolaos Stergioulas. Gravitational-wave model for neutron star merger remnants with supervised learning. *Physical Review D*, 111(2), January 2025.

[275] Joshua S Speagle. dynesty: a dynamic nested sampling package for estimating bayesian posteriors and evidences. *Monthly Notices of the Royal Astronomical Society*, 493(3):3132–3158, feb 2020.

[276] Ulrich Sperhake. The numerical relativity breakthrough for binary black holes. *Class. Quant. Grav.*, 32(12):124011, 2015.

[277] Anand Srinivasan, Miguel Dumett, Christopher Paolini, Guillermo F. Miranda, and José E. Castillo. Mimetic finite difference operators and higher order quadratures. *GEM - International Journal on Geomathematics*, 14(1):19, Jul 2023.

[278] Varun Srivastava, Derek Davis, Kevin Kuns, Philippe Landry, Stefan Ballmer, Matthew Evans, Evan D. Hall, Jocelyn Read, and B. S. Sathyaprakash. Science-driven tunable design of cosmic explorer detectors. *The Astrophysical Journal*, 931(1):22, may 2022.

[279] Kalin V. Staykov, Daniela D. Doneva, Lavinia Heisenberg, Nikolaos Stergioulas, and Stoytcho S. Yazadjiev. Differentially rotating scalarized neutron stars with realistic postmerger profiles. *Phys. Rev. D*, 108(2):024058, July 2023.

[280] Kalin V. Staykov, Daniela D. Doneva, Lavinia Heisenberg, Nikolaos Stergioulas, and Stoytcho S. Yazadjiev. Differentially rotating scalarized neutron stars with realistic postmerger profiles. *Phys. Rev. D*, 108:024058, Jul 2023.

[281] Nikolaos Stergioulas, Andreas Bauswein, Kimon Zagkouris, and Hans-Thomas Janka. Gravitational waves and nonaxisymmetric oscillation modes in mergers of compact object binaries. *Mon. Not. Roy. Astron. Soc.*, 418:427, 2011.

[282] Bo Strand. Summation by parts for finite difference approximations for d/dx. *Journal of Computational Physics*, 110(1):47–67, 1994.

[283] D. Suárez-Fontanella, D. Barba-González, C. Albertus, and M. Ángeles Pérez-García. Gravitational Wave emission in Binary Neutron Star early post-merger within a dark environment. *arXiv e-prints*, page arXiv:2408.05226, July 2024.

[284] Lunan Sun, Milton Ruiz, Stuart L. Shapiro, and Antonios Tsokaros. Jet launching from binary neutron star mergers: Incorporating neutrino transport and magnetic fields. *Phys. Rev. D*, 105(10):104028, May 2022.

[285] Pranesh A. Sundararajan, Gaurav Khanna, and Scott A. Hughes. Towards adiabatic waveforms for inspiral into kerr black holes. i. a new model of the source for the time domain perturbation equation. *Phys. Rev. D*, 76:104005, 2007.

[286] Kentaro Takami, Luciano Rezzolla, and Luca Baiotti. Constraining the equation of state of neutron stars from binary mergers. *Phys. Rev. Lett.*, 113:091104, Aug 2014.

[287] Kentaro Takami, Luciano Rezzolla, and Luca Baiotti. Spectral properties of the post-merger gravitational-wave signal from binary neutron stars. *Phys. Rev. D*, 91:064001, Mar 2015.

[288] Jonathan Thornburg. A Fast apparent horizon finder for three-dimensional Cartesian grids in numerical relativity. *Class. Quant. Grav.*, 21:743–766, 2004.

[289] Jonathan Thornburg. Event and Apparent Horizon Finders for $3 + 1$ Numerical Relativity. *Living Rev. Rel.*, 10:3, 2007.

[290] Wolfgang Tichy and Pedro Marronetti. The Final mass and spin of black hole mergers. *Phys. Rev. D*, 78:081501, 2008.

[291] M. C. Tringali, A. Puecher, C. Lazzaro, R. Ciolfi, M. Drago, B. Giacomazzo, G. Vedovato, and G. A. Prodi. Morphology-independent characterization method of postmerger gravitational wave emission from binary neutron star coalescences. *Classical and Quantum Gravity*, 40(22):225008, November 2023.

[292] Antonios Tsokaros, Jamie Bamber, Milton Ruiz, and Stuart L. Shapiro. Masking equation of state effects in binary neutron star mergers. *arXiv e-prints*, page arXiv:2411.00939, November 2024.

[293] Alan Tsz-Lok Lam, Yong Gao, Hao-Jui Kuan, Masaru Shibata, Karim Van Aelst, and Kenta Kiuchi. Accessing universal relations of binary neutron star waveforms in massive scalar-tensor theory. *arXiv e-prints*, page arXiv:2410.00137, September 2024.

[294] Bram Van Leer. Towards the ultimate conservative difference scheme. IV. A new approach to numerical convection. *Journal of Computational Physics*, 23(3):276–299, 1977.

[295] Bram van Leer. Towards the ultimate conservative difference scheme. v. a second-order sequel to godunov's method. *Journal of Computational Physics*, 32(1):101–136, 1979.

[296] Vijay Varma, Davide Gerosa, Leo C. Stein, François Hébert, and Hao Zhang. High-accuracy mass, spin, and recoil predictions of generic black-hole merger remnants. *Phys. Rev. Lett.*, 122(1):011101, 2019.

[297] Ian Vega and Steven L. Detweiler. Regularization of fields for self-force problems in curved spacetime: Foundations and a time-domain application. *Phys. Rev. D*, 77:084008, 2008.

[298] Ian Vega, Peter Diener, Wolfgang Tichy, and Steven Detweiler. Self-force with $(3 + 1)$ codes: A primer for numerical relativists. *Phys. Rev. D*, 80:084021, Oct 2009.

[299] Vimal Vijayan, Ninoy Rahman, Andreas Bauswein, Gabriel Martínez-Pinedo, and Ignacio L. Arbina. Impact of pions on binary neutron star mergers. *Phys. Rev. D*, 108(2):023020, July 2023.

[300] Gregory Vlastos. A note on zeno's arrow. *Phronesis*, 11(1):3–18, 1966.

[301] Stamatis Vretinaris. SphericalSBPOperators.jl, May 2026.

[302] Stamatis Vretinaris, Nikolaos Stergioulas, and Andreas Bauswein. Empirical relations for gravitational-wave asteroseismology of binary neutron star mergers. *Phys. Rev. D*, 101:084039, Apr 2020.

[303] Stamatis Vretinaris, Georgios Vretinaris, Christos Mermigkas, Minas Karamanis, and Nikolaos Stergioulas. Extended figures for publication "Robust and fast parameter estimation for gravitational waves from binary neutron star merger remnants". `https://doi.org/10.5281/zenodo.1234567`, 2025. Zenodo repository.

[304] Stamatis Vretinaris, Georgios Vretinaris, Christos Mermigkas, Minas Karamanis, and Nikolaos Stergioulas. Robust and fast parameter estimation for gravitational waves from binary neutron star merger remnants. *Phys. Rev. D*, 113:024012, Jan 2026.

[305] Niels Warburton and Barry Wardell. Applying the effective-source approach to frequency-domain self-force calculations. *Phys. Rev. D*, 89:044046, 2014.

[306] Barry Wardell. *Self-force: Computational Strategies*, pages 487–522. Springer International Publishing, Cham, 2015.

[307] Barry Wardell, Chad R. Galley, Anil Zenginoglu, Marc Casals, Sam R. Dolan, and Adrian C. Ottewill. Self-force via green functions and worldline integration. *Phys. Rev. D*, 89:084021, 2014.

[308] Tim Whittaker, William E. East, Stephen R. Green, Luis Lehner, and Huan Yang. Using machine learning to parametrize postmerger signals from binary neutron stars. *Phys. Rev. D*, 105(12):124021, June 2022.

[309] Marcella Wijngaarden, Katerina Chatziioannou, Andreas Bauswein, James A. Clark, and Neil J. Cornish. Probing neutron stars with the full premerger and postmerger gravitational wave signal from binary coalescences. *Phys. Rev. D*, 105(10):104019, May 2022.

[310] Nikolas A. Wittek, Leor Barack, Harald P. Pfeiffer, Adam Pound, Nils Deppe, Lawrence E. Kidder, Alexandra Macedo, Kyle C. Nelli, William Throwe, and Nils L. Vu. Relieving scale disparity in binary black hole simulations. *Phys. Rev. Lett.*, 134:251402, Jun 2025.

[311] Nikolas A. Wittek, Adam Pound, Harald P. Pfeiffer, and Leor Barack. Worldtube excision method for intermediate-mass-ratio inspirals: Self-consistent evolution in a scalar-charge model. *Phys. Rev. D*, 110:084023, Oct 2024.

[312] Francesco Zappa, Sebastiano Bernuzzi, David Radice, and Albino Perego. Binary neutron star merger simulations with neutrino transport and turbulent viscosity: impact of different schemes and grid resolution. *Monthly Notices of the Royal Astronomical Society*, 520(1):1481–1503, jan 2023.

[313] Anil Zenginoglu and Chad R. Galley. Caustic echoes from a schwarzschild black hole. *Phys. Rev. D*, 86:064030, 2012.

[314] David Kai Zhang and Alex Aiken. High-performance branch-free algorithms for extended-precision floating-point arithmetic. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '25, page 695–710, New York, NY, USA, 2025. Association for Computing Machinery.

[315] Teng Zhang, Huan Yang, Denis Martynov, Patricia Schmidt, and Haixing Miao. A gravitational wave detector for post merger neutron stars: Beyond the quantum loss limit of michelson fabry perot interferometer, 2022.

[316] Teng Zhang, Huan Yang, Denis Martynov, Patricia Schmidt, and Haixing Miao. Gravitational-wave detector for postmerger neutron stars: Beyond the quantum loss limit of the fabry-perot-michelson interferometer. *Physical Review X*, 13(2), may 2023.

[317] Yosef Zlochower and Carlos O. Lousto. Modeling the remnant mass, spin, and recoil from unequal-mass, precessing black-hole binaries: The Intermediate Mass Ratio Regime. *Phys. Rev. D*, 92(2):024022, 2015. [Erratum: Phys.Rev.D 94, 029901 (2016)].

Radboud Universiteit

