

Topics in Gravitation - Numerical Simulations of Event Horizons and Parameter Estimation for LISA

Thesis by
Michael I. Cohen

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2011
(Defended 26 July 2010)

Acknowledgments

It is a pleasure to thank those who made this thesis possible, and I would like to start by thanking my formal professorial advisor, Yanbei Chen. Although not present from the start of my work, he contributed many useful discussions and suggestions on the general direction of my research. My work was primarily mentored by my co-advisors Mark Scheel and Curt Cutler, who have been a constant support, and wonderfully patient with me through every setback and success. Without them, all this would not have been possible.

For the first part of this thesis, I would like to thank Harald Pfeiffer and Mark Scheel for the invaluable help, advice and suggestions over the years. In the long research work that I have done with event horizons, Harald and Mark have been my mentors, and have helped guide and shape not only my scientific skills, but my programming abilities as well. Almost invariably, it has been to Mark or Harald whom I have turned time and time again when beset by difficulties, or just looking for a few useful ideas. Their patience, time, and unstinting good graces have been the solid bedrock of the first part of this thesis.

For their role in the second part of this thesis, I am likewise forever grateful to Curt Cutler and Michele Vallisneri. They generously took me under their wing and guided me through my work on parameter estimation for LISA, helping me make this work a reality. They have been a joy to work with, and I am grateful for their forbearance.

I would also like to thank my former formal advisor, Kip Thorne, now hopefully enjoying his retirement, who I have always looked up to with deep respect. I would like to thank Chris Mach, for all the help he has given me over the years. From all the weird software I have needed to install to the odd requests for this or that piece of hardware, Chris has been an often under-appreciated reason that my research has continued to tick along. I would like to thank Matt Robbins and Joerg Wachner, who were kind enough to design for me the beautiful figures in Section 6.3

I would like to thank my family. My mother and father, who have taught me all through my life how to think, how to act, and how to succeed; my siblings, Jonathan, Dafna, Sheira and Tali, who have never faltered in their task of keeping me honest; my darling girl Maia, who has taught me that all storms can be weathered, and that blue skies are just around every corner.

Finally, and most of all, I would like to thank my wife, Rebecca. She is my better half and the foundation on which my life rests, and I dedicate this thesis to her.

Abstract

In Part I, we consider numerical simulations of event horizons. Event horizons are the defining physical features of black hole spacetimes, and are of considerable interest in studying black hole dynamics. Here, we reconsider three techniques to find event horizons in numerical spacetimes, and find that straightforward integration of geodesics backward in time is most robust. We apply this method to various systems, from a highly spinning Kerr hole through to an asymmetric binary black hole inspiral. We find that the exponential rate at which outgoing null geodesics diverge from the event horizon of a Kerr black hole is the surface gravity of the hole. In head-on mergers we are able to track quasi-normal ringing of the merged black hole through seven oscillations, covering a dynamic range of about 10^5 . In the head-on “kick” merger, we find that computing the Landau-Lifshitz velocity of the event horizon is very useful for an improved understanding of the kick behaviour. Finally, in the inspiral simulations, we find that the topological structure of the black holes does not produce an intermediate toroidal phase, though the structure is consistent with a potential re-slicing of the spacetime in order to introduce such a phase. We further discuss the topological structure of non-axisymmetric collisions.

In Part II, we consider parameter estimation of cosmic string burst gravitational waves in Mock LISA data. A network of observable, macroscopic cosmic (super-)strings may well have formed in the early Universe. If so, the cusps that generically develop on cosmic-string loops emit bursts of gravitational radiation that could be detectable by gravitational-wave interferometers, such as the ground-based LIGO/Virgo detectors and the planned, space-based LISA detector. We develop two versions of a LISA-oriented string-burst search pipeline within the context of the Mock LISA Data Challenges, which rely on the publicly available MultiNest and PyMC software packages, respectively. We use the F -statistic to analytically maximize over the signal’s amplitude and polarization, \mathcal{A} and ψ , and use the FFT to search quickly over burst arrival times t_C . We also demonstrate an approximate, Bayesian version of the F -statistic that incorporates realistic priors on \mathcal{A} and ψ . We calculate how accurately LISA can expect to measure the physical parameters of string-burst sources, and compare to results based on the Fisher-matrix approximation. To understand LISA’s angular resolution for string-burst sources, we draw maps of the waveform fitting factor [maximized over (\mathcal{A}, ψ, t_C)] as a function of sky position; these maps dramatically illustrate why (for LISA) inferring the correct sky location of the emitting string loop will often be practically impossible. In addition, we identify and elucidate several symmetries that are embedded in this search problem, and we derive the distribution of cut-off frequencies f_{\max} for observable bursts.

Contents

Acknowledgments	iii
Abstract	iv
1 Introduction	1
1.1 History of gravitational theory	1
1.2 Gravitational Waves	2
Bibliography	6
I Simulating Event Horizons in Numerical Spacetimes	7
2 Introducing Event Horizons	8
Bibliography	12
3 Background and Basic Techniques	15
3.1 Notation	15
3.2 Basic Techniques for Finding Event Horizons	15
3.2.1 Geodesic Method	16
3.2.2 Surface Method	17
3.2.3 Level-set Method	20
3.3 Proof that points in the Surface Method follow geodesics	21
3.4 An Introduction to SpEC	22
3.5 Numerical Implementation	24
Bibliography	27
4 Application to Kerr Spacetime	30
4.1 Kerr spacetime using the geodesic method	32
4.2 Proof of Surface Gravity conjecture	33
4.3 Kerr spacetime using the surface method	34
Bibliography	37

5	Head-on Mergers	38
5.1	Head-on Nonspinning Binary Black Hole Merger	38
5.1.1	Details of BBH evolution	38
5.1.2	EH finder behaviour	39
5.1.3	Quasinormal Modes during Ringdown	41
5.1.4	Treatment of Merger	44
5.1.5	Analysis of Merger Phase	48
5.2	Head-on Spinning BBH Merger	51
5.2.1	Configuration	51
5.2.2	Overview of 4-Momentum Conservation in the Landau-Lifshitz Formalism . .	52
5.2.3	Landau-Lifshitz Velocities of Event and Apparent Horizons	55
	Bibliography	57
6	Binary Black Hole Inspirals and Mergers	61
6.1	Introduction to Binary Black Hole Inspirals	61
6.2	Collision Detection	61
6.3	Topological structure of the Event Horizon for an Inspiral-Merger	64
6.4	Topological Structure of Simulated Event Horizons	68
	Bibliography	73
7	Conclusion	74
	Bibliography	76
II Searches for Cosmic-String Gravitational-Wave Bursts in Mock LISA Data		78
8	Introduction	79
8.1	LISA and the Mock LISA Data Challenges	79
8.2	Cosmic Strings	80
	Bibliography	84
9	Theoretical Background	86
9.1	The gravitational waveform from cosmic-string bursts	86
9.2	Maximization over the extrinsic parameters	87
9.3	Bayesian version of the F -statistic	89

9.4	Distribution of f_{\max} for detected bursts	91
Bibliography		93
10	Near-symmetries and overlap maps	94
10.1	Sky-position reflection across the LISA plane	94
10.2	Broad F -statistic quasi-degeneracy across the sky	95
10.2.1	Proof of the fourth FF-map symmetry	99
10.3	Effects of degeneracies on searches	100
Bibliography		102
11	Search Methods	103
11.1	Markov Chain Monte Carlo	103
11.2	MultiNest	104
11.3	High-SNR limit and the Fisher-Matrix formalism	105
Bibliography		107
12	Results from the Mock LISA Data Challenges	110
Bibliography		114
13	Conclusions	115
A	Momentum flow in black-hole binaries: II. Numerical simulations of equal-mass, head-on mergers with antiparallel spins.	117
A.1	Abstract	117
A.2	Introduction	118
A.2.1	Motivation	118
A.2.2	Linear momentum flow in BBHs and gauge dependence	118
A.2.3	BBH mergers with recoil	120
A.2.4	Overview and summary	121
A.3	4-Momentum Conservation in the Landau-Lifshitz Formalism	124
A.4	Simulations of head-on BBH collisions with anti-aligned spins	127
A.4.1	Generalized harmonic	127
A.4.2	BSSN-moving-puncture	132
A.5	Momentum flow	135
A.5.1	Generalized harmonic results	135
A.5.2	BSSN-moving-puncture results and gauge	143

A.5.3	Comparison with post-Newtonian predictions	146
A.6	Conclusion	152
A.7	Acknowledgments	153
A.8	Appendix A: Excision initial data	153
A.8.1	Superposed-Kerr-Schild (SKS) initial data	153
A.8.2	Superposed-Harmonic-Kerr (SHK) initial data	155
A.9	Appendix B: Numerical methods for evolutions	157
A.9.1	Generalized harmonic evolutions	157
A.9.2	BSSN-moving-puncture evolutions	160

Bibliography		161
---------------------	--	------------

List of Figures

- 3.1 Cross-sections through event and apparent horizons during a BBH merger. Before the merger, $t < t_{\text{CEH}}$, the surface includes the set of generators that will merge onto the event horizons through the cusps in the individual event horizons (green dashed curves). The point C is the point of symmetry for the head-on merger, which will be used in Section 3.2.3. 16
- 3.2 A slice of the event horizon (\mathcal{S}_t), produced by the intersection of a spatial hypersurface Σ_t with the world tube of the event horizon \mathcal{N} , showing n^μ (the timelike normal to Σ_t), s^μ (the spatial normal of \mathcal{S}_t), and ℓ^μ (the null normal to the EH world-tube \mathcal{N}). 18
- 4.1 **Geodesic method** applied to a Kerr black hole. The **top panels** show the area difference between the computed and exact solution, normalized by the area of the exact solution. The **bottom panels** show the difference between the computed and exact location of the EH, as measured by (4.6). These data are shown for two series of runs: In the **left panels** we keep the dimensionless spin of the black hole fixed at $a/M = 0.6$ and vary the resolution L of the EH finder. In the **right panels** we vary the spin parameter a/M at fixed resolution. In all cases, the EH finder starts at $t = 0$ and the geodesics are evolved backward in time. 31
- 4.2 Approach of the tracked null surface onto the event horizon of Kerr black holes with various spins. The symbols show the numerical data (the same data as in the lower right panel of Figure 4.1), and the solid lines are representative least-squares fits. Table 4.1 compares the numerically computed e-folding time to the surface gravity of the black hole. 32
- 4.3 Effect of filtering using (4.19) and (4.20) for the **surface method**. Shown are evolutions with the same angular resolution $L = 18$, but for different numbers F of truncated modes in (4.19). **Left panel:** For a Schwarzschild black hole, the surface method is stable with or without this additional filtering. **Right panel:** For a Kerr black hole with $a/M = 0.6$, $F = 7$ performs best. The EH finder starts at $t = 0$ and the surface is evolved backward in time. 35

- 4.4 **Surface Method**, applied to a Kerr black hole. The **top panels** show the normalized area difference between the computed and exact solution. The **bottom panels** show the difference between the computed and exact location of the EH, as measured by (4.6). These data are shown for two series of runs: In the **left panels** we keep the dimensionless spin of the black hole fixed at $a/M = 0.6$ and vary the resolution L of the EH finder. In the **right panels** we vary the spin a at fixed resolution. The value F denotes the number of truncated modes during filtering according to (4.19). For each case, we show the value of F that provides the most accurate evolution. Also, in all cases, the EH finder starts at $t_{\text{end}} = 0$ and the surface is evolved backward in time. Compare to Figure 4.1. 36
- 5.1 Evolution of a head-on BBH merger: **normalized constraint violations**. The left panel shows the complete evolution. The right panel enlarges the time around merger, with formation of a common apparent horizon and time of regriding indicated by ‘CAH’ and ‘regrid,’ respectively. The discontinuity at t_{CAH} arises because the constraints are computed only outside the common AH for $t > t_{\text{CAH}}$. At t_{regrid} , the constraints jump because of the different numerical truncation error of the ringdown domain decomposition. 39
- 5.2 Effect of changing the resolution of the EH finder when applied to the BBH evolution at fixed high resolution. Shown are relative differences in the area $A(t)$ of the tracked surface. The label “7-15” denotes the difference between simulations with $L_1 = 7$ and $L_2 = 15$, normalized by $A(t)$ of L_2 . Vertical lines on the graph denote the formation of common event and apparent horizons. Note that the time scale of both plots change at $t/M = 25$ 40
- 5.3 Error estimates for the **surface** and **geodesic** methods, with surface resolution $L = 47$. The left panels show the root-mean-square pointwise deviation between the different runs, whereas the right panels show the differences in the surface area. The lines labelled “G:N#-N#” (“S:N#-N#”) in the upper panels show the difference between the geodesic method (surface method) when applied to merger simulations of different resolution N. The lines labelled “N#:S-G” in the lower panels show the differences between the surface and geodesic methods for a given N (where N0, N1, and N2 are resolutions of the merger simulation). Note that the time scale of all plots change at $t = 25M$ 42
- 5.4 Surface area differences between the EH, AH, and the final area, normalized by the final area. Also plotted are error estimates for $A_{\text{EH}}(t)$ and $A_{\text{AH}}(t)$ 43

5.5	Left panel: Area element \sqrt{h} along a few representative geodesics during the head-on merger. Each geodesic is labelled by the angle α between the initial location of the geodesic (at t_{end}) and the axis of symmetry. Three types of behaviour are apparent: Geodesics entering the horizon from \mathcal{I}^- ($\alpha = 85^\circ$); geodesics entering the horizon from an area in the vicinity of the individual event horizons before merger ($\alpha = 72^\circ$ or 76°), and geodesics remaining on the horizon throughout. The right panels show the time derivative of \sqrt{h} , highlighting the clear signature when a geodesic enters the horizon.	45
5.6	Convergence of the surface area of the event horizon during merger. The lower plot shows results for placement of the geodesic pole parallel to the axis of symmetry (i.e. consistent with axisymmetry), the upper plot has a geodesic axis perpendicular to the axis of symmetry of the merger. In both cases, geodesics are tracked using the geodesic method; derivatives for \sqrt{h} (cf. (3.28)) are computed with finite-differences; geodesics are removed from the event horizon based on (5.5). Lines are the difference between each resolution and the next highest.	47
5.7	Area of the event horizon and of the apparent horizons before merger and during merger. The vertical dotted lines indicate formation of a common event horizon and appearance of a common apparent horizon; the inset shows an enlargement for early time.	49
5.8	Difference between event-horizon area and the sum of the individual apparent horizon areas. The vertical dotted lines indicate formation of common event horizon and appearance of a common apparent horizon. This difference is computed using finite-difference derivatives on the event horizon, and we see a clear improvement in the exponential change in the area at higher order finite-differencing.	49
5.9	Spacetime diagram of the head-on merger. The pale lines denote geodesics that will join the event horizon. Some of these geodesics come from past null infinity, but others come from a region close to the individual event horizons (cf. the arrow and the circled geodesics on the far black hole).	50
5.10	Initial configuration of the head-on BBH spinning merger. The holes move primarily along the x axis, but they also accelerate in the $-y$ (downward) direction due to frame dragging. For the particular simulation on which the event horizon was evaluated, $d = 7.804/M_{\text{ADM}}$. This Figure appears in Appendix A as Figure A.1.	51
5.11	The regions of space around and inside a binary-black-hole system. This Figure appears in Appendix A as Figure A.4.	54

5.12 The velocity of the individual and merged black holes. The Landau-Lifshitz velocity $v_{LL}^y := p_{LL}^y/p_{LL}^t$, where p_{LL}^μ is the Landau-Lifshitz 4-momentum enclosed, is measured on the individual and common apparent horizons (labeled AH and AHC, respectively) and also on the event horizon (labeled EH). For comparison, the coordinate velocities v_{coord}^y of the apparent horizons are also shown. This Figure appears in Appendix A as Figure A.8. 56

5.13 The effective velocity v_{LL}^y calculated on the event horizon surface, with the specified snapshots in Fig. 5.14 of the event horizon surface marked: a,b, $t = 27.7M_{\text{ADM}}$; c, $t = 30.8M_{\text{ADM}}$; d, $t = 31.6M_{\text{ADM}}$; e, $t = 35.5M_{\text{ADM}}$; f, $t = 40.8M_{\text{ADM}}$. This Figure appears in Appendix A as Figure A.13. 56

5.14 Snapshots of the event horizons at the times indicated in Fig. 5.13: a,b, $t = 27.7M_{\text{ADM}}$; c, $t = 30.8M_{\text{ADM}}$; d, $t = 31.6M_{\text{ADM}}$; e, $t = 35.5M_{\text{ADM}}$; f, $t = 40.8M_{\text{ADM}}$. All snapshots are looking down the z-axis to the x-y plane, except for shot a, which is slightly skewed (slightly rotated about the y axis) to better see the geodesic structure. In shot a, the future generators of the horizon are visible as small blue dots. Note how the future generators map out a surface that meets the event horizon at the event horizon’s cusps; this is where the future generators join the horizon. The data shown are from the high-resolution evolution N2.C. This Figure appears in Appendix A as Figure A.14. 58

6.1 Diagram of collision detection between a moving triangle and a moving point. . . . 63

6.2 Slices through the event horizon hypersurface at the exact point of merger (to within numerical accuracy). **Upper panel:** Equal mass non-spinning 16-orbit inspiral at $t/M = 3932.779$, the precise point of merger to within $\pm 0.006t/M$. Here M is the sum of the ADM masses. **Lower panel:** Generic 2:1 mass ratio random spin-direction 1.5 orbit inspiral at $t/M = 117.147$, the precise point of merger to within to within $\pm 0.005t/M$. Here M is the sum of the Christodolou masses, since the black holes have spin. At prior times the two black hole horizons are disjoint. At later times, the two horizons have merged into a single S2 surface. No toroids are evident in the limit of our accuracy. 65

- 6.3 Diagrams of the event horizon null hypersurface in axisymmetric and non-axisymmetric mergers. The merger is along the z -axis. In both panels, the regions $\mathcal{C} \cup \mathcal{X}$ are spacelike. **Left panel:** In the axisymmetric case, the caustic/crossover set is reduced to a single line of caustic points, the “inseam” of the “pair of pants,” labeled \mathcal{C} . The x direction is suppressed but, since the x and y directions are identical for axisymmetry, the identical diagram would obtain if we were to suppress y in favour of x . **Right panel:** In the non-axisymmetric case, such as an inspiral (where we have “unwound” the legs of the “pair of pants”), the set of crossover points \mathcal{X} is evident, bounded on both sides by “inseams” \mathcal{C} . Unlike the axisymmetric case, here the x and y directions are not identical. Since the caustic/crossover set of points is a 2-surface, the diagram we would obtain by suppressing y in favour of x would look identical to the left panel case, except that the single caustic “inseam” would be composed of crossover points. 66
- 6.4 This figure shows a 2-dimensional slice through the event horizon null hypersurface in an axisymmetric merger. The horizontal direction in the right panel could be either x or y . We attempt to construct a slice \mathcal{S}_1 in x (or y) from point \mathcal{P} that intersects the black hole. This slice is clearly not spacelike. Since \mathcal{N} is spacelike only at \mathcal{C} , only a non-intersecting slice such as \mathcal{S}_0 can be spacelike. 67
- 6.5 This figure shows a 2-dimensional slice through the event horizon null hypersurface in a non-axisymmetric merger. Unlike the previous figure, the horizontal direction in the right panel is not interchangeable between x and y . We construct three slices $\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2$ from the starting point \mathcal{P} which all intersect the event horizon in different ways. Since $\mathcal{C} \cup \mathcal{X}$ is spacelike, all these slices are spacelike. 68
- 6.6 A 3-dimensional representation of slice \mathcal{S}_1 from Figure 6.5 Here we clearly see the continuation of the slice in the z direction, which creates a toroidal intermediate black hole. The toroidal region of the slice is that part of the slice that has dipped through the crossover region \mathcal{X} , with the center of the torus at \mathcal{P} 69
- 6.7 A snapshot of the geodesics being followed by the event horizon finder at time $t/M = t_{\text{merger}}/M - 0.067$, for the equal-mass inspiral. The small dots are geodesics currently on the event horizon. The larger points represent geodesics in the process of merging onto the event horizon. Crosses represent points merging through caustic points, while circles represent points merging through crossovers. In this slice, it is clear that the cusp on the black hole is linear, composed of crossover points with caustics at the end-points. 70

6.8	A snapshot of the geodesics being followed by the event horizon finder at time $t/M = t_{\text{merger}}/M$, the exact point of merger in the equal-mass inspiral simulation. The small dots are geodesics currently on the event horizon. The larger points represent geodesics in the process of merging onto the event horizon. Crosses represent points merging through caustics, while circles represent points merging through crossovers. Although finding the exact point of merger is difficult given limited numerical time-accuracy, we can extrapolate the “X”-shape of the cusps to see that the merger point is clearly a crossover point.	71
6.9	A snapshot of the geodesics being followed by the event horizon finder at time $t/M = t_{\text{merger}}/M + 0.039$, shortly after merger, for the equal-mass inspiral. The small dots are geodesics currently on the event horizon. The larger points represent geodesics in the process of merging onto the event horizon. Crosses represent points merging through caustics, while circles represent points merging through crossovers. The “bridge” between the two black holes has two lines of merger points running on either side of it, with the majority being crossover points anchored by caustics at either end. . . .	72
8.1	Expected LIGO and LISA sensitivity ranges, and the sources that may be found within those ranges. Sourced from the LISA Science Case (www.lisascience.org).	80
9.1	Expected distribution $dN/d(\log(f_{\text{max}}))$ of the maximum burst frequency f_{max} for the string bursts detectable by LISA.	92
10.1	FF maps for high- (top) and low-frequency (bottom) bursts: $-\log_{10}(1 - \text{FF})$ contours are computed between (β, λ, ψ) target sources (with $\beta = 0, \pi/3, \pi/2, \lambda \in [0, \pi/3], \psi \in [0, \pi/4]$) and (β', λ') templates across the sky ($\beta' \in [-\pi/2, \pi/2], \lambda' \in [-\pi, \pi]$, each small square). Because of the symmetries discussed in Section 10.2, these λ and ψ ranges exhaust the variety of maps seen across their entire ranges. The target-source latitude $\beta = \pi/3$ is also representative of latitudes intermediate between the equator $\beta = 0$ and the pole $\beta = \pi/2$. At the equator, ψ has no effect on the maps (except for $\psi = \pi/4$, where there is no LISA response); at the pole, λ is degenerate, and ψ is defined consistently with the $\lambda = 0$ meridian.	97
10.2	Symmetries between FF maps, as explained in the main text, exemplified for the case of $\beta = \pi/6, \lambda = \pi/9, \psi = \pi/6$	98

- 10.3 Fraction of the sky with $\text{FF}(A + E + T) > 1 - 10^{-x}$, for target-source $\beta = \pi/3$, and uniformly distributed (λ, ψ) , where each pair corresponds to one of the superimposed curves. The curves were obtained by generating 40×40 maps as for Fig. 10.1, assigning a weight to each pixel corresponding to its area in the sky, sorting the resulting sequence by increasing FF, and computing normalized cumulative weights. 99
- 10.4 Posterior-probability structure for the strongest source (#3) in the noiseless training data set from MLDC 3.4. **Left:** in this sky map, the density of dots (MultiNest equal-weight “resamples”) is proportional to the posterior probability, maximized over \mathcal{A} , ψ and t_C , and marginalized over f_{\max} . Crosses mark the true location of the source, and its LISA-plane-reflected counterpart. The map is plotted in the area-preserving Mollweide projection, which we adopt throughout the rest of this paper. **Right:** F -statistic as a function of ecliptic latitude and longitude, for the same sky locations as in the left panel. Here F is offset by a constant $\simeq 3,029$, and it is only slightly higher for the neighborhoods of the true and reflected sky locations than for the arcs connecting them. 100
- 10.5 Effect of different noise realizations on the posterior-probability structure for the strongest source (#3) in the noisy MLDC 3.4 training data set, and in five more data sets with the same source and different noise realizations. 101
- 10.6 **Left:** the best-fit value for the template amplitude, as computed by the F -statistic, increases dramatically for sky positions close to the instantaneous LISA plane, as shown here for source #3 in the noisy MLDC 3.4 training data set. **Right:** the Bayesian F_B -statistic shifts the best-fit sky locations away from the instantaneous LISA plane, as seen here for the six data sets of Fig. 10.5. In some cases, the best-fit location moves to the other side of the sky; this is not significant, given that reflected points have essentially the same posterior probability against the same source. . . . 102
- 11.1 Comparison of MultiNest posterior distributions with Fisher-matrix estimates, in the case of the strongest source (#3) of the MLDC 3.4 noiseless training data. The top plot shows that the posterior distribution on the sky is more “banana-shaped” than ellipsoidal. The next six plots compare the true posterior distribution (restricted to the neighborhood of the “true” mode) with Gaussian distributions of variance $\sigma_\mu^2 = (\Gamma^{-1})^{\mu\mu}$. The variance factor, defined as $\sigma_{\text{fit}}^2 / \sigma_{\text{Fisher}}^2$, ranges between 0.6 and 8.8. 108
- 11.2 Same as Fig. 11.1, except that the source’s SNR is now 1,000. In this case, the posterior is fit very well by the Fisher-matrix prediction. Even at this high SNR, a secondary maximum is present around the reflected location, but it is not shown in this plot. 109

12.1	MultiNest sky-location posteriors for sources 0–2 in the MLDC 3.4 challenge data set. The density of the dots is proportional to the posterior probability (including the F_B prior correction described in Section 9.3), maximized over t_C , and marginalized over \mathcal{A} , ψ , and f_{\max} . Crosses and circles indicate the true and best-fit locations, respectively. For source 1, the stars indicate the location of <i>mirage</i> best fits discarded by F_B	113
A.1	Initial configuration of the head-on BBH considered in this paper. The holes move primarily along the x axis, but they also accelerate in the $-y$ (downward) direction due to frame dragging. See Table A.1 for the value of $d = 2x_0$	121
A.2	The effective velocity v_{LL}^y for the individual (red dotted line) and common (green dashed line) apparent horizons and for the event horizon (black solid line). The inset shows the velocity of the common apparent horizon at late times.	122
A.3	A contour plot of the y component of the momentum density at the moment when the common apparent horizon forms. The common horizon encloses the momentum inside the individual horizons and also the momentum in the gravitational field. The grey-shaded region and solid, red contours indicate positive momentum density, while the white-shaded region and blue, dashed contours indicate negative momentum density. The individual apparent horizons are shaded black, and the common apparent horizon is shown as a thick black line.	123
A.4	The regions of space around and inside a binary-black-hole system.	126
A.5	Constraint violation at different numerical resolutions for the generalized harmonic evolutions S1. The common apparent horizon forms at time $t = 34.73M_{\text{ADM}}$. Labels of the form $Nx.y$ indicate the grid resolution, where the pre-merger resolution is labeled (from coarse to fine) by $x = 0, 1, 2$ and the post-merger resolution is labeled by $y = A, B, C$. The constraints decrease exponentially with higher resolution; the convergence rate is smaller near merger.	130
A.6	Recoil velocity for initial data set S1 inferred from the gravitational wave signal Ψ_4 extracted at $r_{\text{extr}} = 160M_{\text{ADM}}$ at the highest resolution (upper panel). Differences between several coarser resolutions and the highest resolution result are plotted in the lower panel.	131
A.7	Gravitational recoil for model P1 as estimated from the gravitational wave signal Ψ_4 extracted at $r_{\text{ex}} = 73.5 M_{\text{ADM}}$ using the highest resolution (upper panel). Differences in the recoil obtained at coarse, medium and fine resolution rescaled for second order convergence (lower panel).	134

A.8 The velocity of the individual and merged black holes. The Landau-Lifshitz velocity $v_{LL}^y := p_{LL}^y/p_{LL}^t$, where p_{LL}^μ is the Landau-Lifshitz 4-momentum enclosed, is measured on the individual and common apparent horizons (labeled AH and AHC, respectively) and also on the event horizon (labeled EH). For comparison, the coordinate velocities v_{coord}^y of the apparent horizons are also shown. The data shown are from the high-resolution evolution N2.C. 136

A.9 Convergence of v_{LL}^y with resolution. Specifically, differences between v_{LL}^y at the highest resolution N2.C and at various lower resolutions are shown. Labels of the form Nx.y indicate the grid resolution, where the pre-merger resolution is labeled (from coarse to fine) by $x = 0, 1, 2$ and the post-merger resolution is labeled by $y = A, B, C$. The difference between the second-highest and highest resolution is below 0.1km/s except near merger, when it grows as large as 1 km/s. 136

A.10 Contour plots of the y (up-down) component of the momentum density, which points along or opposite of the holes' motion due to frame dragging. Adjacent contours correspond to a factor of 10 difference in the magnitude of the momentum density. Contours of positive y momentum density are shown as solid red lines, while contours of negative y momentum density are shown as dashed blue lines. The region containing positive y momentum density is shaded grey. The regions inside the apparent horizons are shaded black, except for the upper right panel, where the region inside the *individual* horizons is shaded black, while the common apparent horizon is indicated by a thick black line. The data shown are from the high-resolution evolution N2.C. 138

A.11 *Left:* A comparison of the common apparent horizon's effective velocity and the horizon's shape and area. The top panel shows the horizon's effective velocity v_{LL}^y . The middle panel shows the pointwise minimum and maximum of the horizon's dimensionless intrinsic scalar curvature; both $M_{\text{Chr}}^2 \min(R)$ and $M_{\text{Chr}}^2 \max(R)$ relax to the Schwarzschild value of 1/2 as the horizon rings down. (The first four local minima of $M_{\text{Chr}}^2 \min(R)$ are indicated by vertical dashed lines.) The bottom panel shows the area A of the common apparent horizon normalized by the total area of the individual horizons at $t = 0$. The data shown are from the high-resolution evolution N2.C. *Right:* The dimensionless intrinsic scalar curvature $M_{\text{Chr}}^2 R$ of the common apparent horizon at the times labeled A–F in the left panel. The horizon begins peanut-shaped, then rings down, eventually settling down to a sphere with a constant curvature $M_{\text{chr}}^2 R = 0.5$ 139

- A.12 A comparison of various measures of the final velocity of the merged hole in simulation S1. The kick velocity inferred from the gravitational-wave flux (thin dashed lines) and the Landau-Lifshitz effective velocities v_{LL}^y (thin solid lines) are measured on spheres of radius $100M_{\text{ADM}}$, $120M_{\text{ADM}}$, $140M_{\text{ADM}}$, and $160M_{\text{ADM}}$. The value of the kick velocity at the final time is extrapolated to $r = \infty$ (black cross). The effective velocity measured on the common apparent horizon (thick solid line) and the coordinate velocity (thick dashed line) are also shown. The data shown are from the high-resolution evolution N2.C. 140
- A.13 The effective velocity v_{LL}^y calculated on the event horizon surface, with the specified snapshots in Fig. A.14 of the event horizon surface marked: a,b, $t = 27.7M_{\text{ADM}}$; c, $t = 30.8M_{\text{ADM}}$; d, $t = 31.6M_{\text{ADM}}$; e, $t = 35.5M_{\text{ADM}}$; f, $t = 40.8M_{\text{ADM}}$ 141
- A.14 Snapshots of the event horizons at the times indicated in Fig. A.13: a,b, $t = 27.7M_{\text{ADM}}$; c, $t = 30.8M_{\text{ADM}}$; d, $t = 31.6M_{\text{ADM}}$; e, $t = 35.5M_{\text{ADM}}$; f, $t = 40.8M_{\text{ADM}}$. All snapshots are looking down the z-axis to the x-y plane, except for shot a, which is slightly skewed (slightly rotated about the y axis) to better see the geodesic structure. In shot a, the future generators of the horizon are visible as small blue dots. Note how the future generators map out a surface that meets the event horizon at the event horizon's cusps; this is where the future generators join the horizon. The data shown are from the high-resolution evolution N2.C. 142
- A.15 Velocity perpendicular to the line of sight associated with the horizon integrals of the the Landau-Lifshitz tensor obtained for models P1 and S1. The shaded area represents the numerical uncertainty for P1. During the pre-merger phase, in each simulation the velocities of the individual horizons are identical. 144
- A.16 Same as Fig. A.15 for model P2 of Table A.1. 145
- A.17 A comparison of numerical and post-Newtonian longitudinal velocities (i.e., v^x/c) versus time. The predicted coordinate velocities at several post-Newtonian orders are shown as broken curves. *Left:* A comparison of S1 numerical data and post-Newtonian predictions. The numerical and post-Newtonian curves agree qualitatively. When the 2.5 PN curve is offset by a certain amount, it agrees quantitatively with the coordinate velocity v_{coord}^x and the effective velocity v_{LL}^x . *Right:* A comparison of H1 numerical data and PN predictions. The effective velocity v_{LL}^x (thick black line) closely tracks the coordinate velocity v_{coord}^x ; both numerical curves also agree well with the 2.0 PN and 2.5 PN curves. 147

- A.18 A comparison of numerical and post-Newtonian transverse velocities (i.e., v^y in km/s) versus time. The left panel shows numerical results from simulation S1, while the right panel shows numerical results from simulation H1. The predicted coordinate velocity at several post-Newtonian orders are shown as broken curves. The effective velocity is shown in black; it has been rescaled by a factor of $3/2$ in order to aid comparison with the post-Newtonian point-particle velocities, as discussed in Sec. A.5.3. The turn around in the 2.5 PN curves is due to a 2.5 order spin-orbit term becoming quite large at a separation of roughly $2M_{\text{ADM}}$. One can argue this is due to the post-Newtonian approximation breaking down at this small separation. 148
- A.19 A comparison of numerical and post-Newtonian velocities. In the figure, v^y in km/s is plotted against v^x/c . The effective velocity v_{LL}^y of the highest-resolution (N2.C) evolution of initial data S1 (Table A.1) on the left and of the evolution of initial data H1 (Table A.1) on the right are shown as a thick black line. The predicted coordinate velocity at several post-Newtonian orders are shown as broken curves. The transverse effective velocities only agree qualitatively with post-Newtonian predictions; however, the coordinate velocity agrees very well with post-Newtonian predictions. In the left panel, the coordinate velocity has been artificially truncated shortly before merger because at that point we do not have a good measure of the coordinate velocity. The effective velocity has been rescaled by a factor of $3/2$ to aid comparison with the post-Newtonian point-particle velocities, as discussed in Sec. A.5.3. 149
- A.20 Comparison of numerical (top row) and post-Newtonian (bottom row) y momentum density. The numerical data comes from the harmonic evolution H1 described in Appendix A.8.2. The 1.5 PN momentum density is computed from Eqs. (A2a)–(A2c) in paper I using the numerical hole trajectories. As in Fig. A.10, contours represent powers of 10 in y momentum density. The positive y momentum density contours are shown in red, negative in blue. The region of positive y momentum density is shaded grey. In the numerical plots the apparent horizons are shown in black. 151

List of Tables

4.1	Exponential approach of the null surface to the correct event horizon location. Mg_H represents the (dimensionless) surface-gravity of a Kerr black hole with spin a/M . M/τ is the numerical rate of approach as determined by fits to the data shown in Figure 4.2.	33
12.1	True source parameter values and MCMC and MultiNest best fits for the MLDC 3.4 challenge data set. When the estimated f_{\max} is larger than the 0.5 Hz Nyquist frequency.	112
12.2	Differences between true source parameter values and MCMC and MultiNest best fits, for the MLDC 3.4 challenge data set. The Δ_{sky} error is measured in radians along the geodesic arc between the true and best-fit sky positions.	112
12.3	Parameter accuracy achieved by MultiNest for source #3 in the MLDC 3.4 training data set, with the original and boosted SNR.	113
A.1	Parameters of the initial data configurations studied in this work. Model S1 (see Sec. A.4.1) gives the parameters used to construct a set of Superposed-Kerr-Schild quasiequilibrium excision initial data. Model H1 (see Appendix A.8.2) gives the parameters for the larger separation Superposed-Harmonic-Kerr initial data set. Both S1 and H1 were used in generalized harmonic, pseudospectral evolutions. P1 and P2 provide the Bowen-York parameters for the two systems evolved with the BSSN-moving-puncture method. The holes are initially separated by a coordinate distance $d = 2x_0$ and are located at coordinates $(x, y, z) = (\pm x_0, 0, 0)$. For clarity, only 4 significant figures are shown.	128

Chapter 1 Introduction

1.1 History of gravitational theory

In the study of the macroscopic universe, the most important force to reckon with is the force of gravity. The universe is strongly shaped by gravity which, although the weakest of the forces, exerts a universal attraction on all matter. The force of gravity governs the majority of the structure of our universe, from the orbits of the planets in our solar system, through to the dynamics of stars and nebulae, black holes and galaxies. The earliest known theory of gravity was formulated by the English scientist, Sir Isaac Newton in his *Philosophiæ Naturalis Principia Mathematica*. In this work Newton defines gravity as a force emanating radially from the centers of mass of objects, with an inverse square nature (i.e. the strength of the gravitational attraction between two objects decreases as the square of the distance between their respective centers of mass). This simple theory of gravity adequately describes weak gravitational phenomena, and is enough for many useful applications of gravitational analysis, such as the motion of the solar system, or the ballistic behavior of objects in the presence of a massive stellar body. Along with his theory of gravity, Newton devised his famous laws of motion

When Scottish physicist James Clerk Maxwell formulated the laws of electromagnetism in the 1860's, he showed that light was merely an oscillation in the electromagnetic field, and that these oscillations propagated at a set speed of approximately 300,000 km/s, a number which depended on the known physical constants ϵ_0 , the “permittivity of free space,” and μ_0 , the “permeability of free space.” Given that Newton's laws and the associated so-called “Galilean” transform were used to compute motion, the obvious question was “300,000 km/s relative to what frame of reference?” Initially, this led to theories of the “ether,” a field that provided a medium for light to travel. In the old Galilean/Newtonian point of view, this ether provided a preferred reference frame for the universe. A famous experiment performed by American physicists Albert Michelson and Edward Morley in 1887 disproved the existence of the ether, showing that the speed of light was 300,000 km/s in any reference frame. As a result, it was clear that Maxwell's equations were not compatible with Newton's laws of motion, and that it was Newton's laws that needed changing, not Maxwell's.

It was in order to resolve this fundamental difficulty that Albert Einstein in 1905 proposed his Special Theory of Relativity. In this theory, the “Galilean” transformation (which codified simple assumptions about the addition of velocities and the universal nature of time) was discarded in favor of a transformation between reference frames devised by Dutch physicist Hendrik Lorentz, who was attempting to explain the Michelson-Morley experiment. By transforming physical laws between

non-accelerating reference frames using the Lorentz transformation, Einstein was able to present a theory of physical motion consistent with Michelson and Morley's discovery that light propagated at the same velocity in all reference frames. Lorentz's transformation also paved the way for the unification of space and time into spacetime. No longer were space and time distinct quantities, for in the Lorentz frame, observers traveling at different speeds measure time at different rates.

Unfortunately, one of the main casualties of special relativity was Newton's theory of gravity. The two were simply incompatible, and much subsequent effort over the next decade was devoted to reconciling special relativity and Newton's gravity. This effort culminated in 1915 with the publication of Einstein's General Theory of Relativity. General relativity postulated that the effect of gravity was indistinguishable from the effect of being stationary in an accelerating reference frame. Known as the *Equivalence Principle*, this idea lies at the heart of general relativity. In general relativity, the presence of mass warps space and time, much as a bowling ball warps a rubber sheet it is dropped onto. In order to describe exactly how spacetime is warped by the presence of matter, general relativity applies the tools of differential geometry, a branch of mathematics that describes the properties of curved and distorted n-dimensional surfaces.

1.2 Gravitational Waves

General relativity gives rise to a number of interesting phenomena. Several of these provided early evidence for relativity, such as the precession of the perihelion of the orbit of Mercury, and the gravitational bending of starlight around the sun. One major phenomenon predicted by general relativity is the existence of gravitational waves (GWs). Much as light is a propagating wave disturbance in the electric and magnetic fields, so the gravitational field admits wavelike disturbances. The existence of these waves has not until now been directly observed.

Indirect evidence of gravitational waves has been found in observations of a binary pulsar system known as the Hulse-Taylor binary, which was discovered in 1974. Pulsars are neutron stars with very strong magnetic fields, spinning at high speeds. The magnetic field of the star causes beam-like emissions to occur from the magnetic poles. These beams of electromagnetic radiation are easily observed by astronomers, and pulsars spin at frequencies that are extremely consistent. Functioning as a kind of cosmic lighthouse, pulsar beams allow astronomers to, among other things, measure rotational frequency to extremely high accuracy. The Hulse-Taylor binary was the earliest binary pulsar system to be discovered, and its rotational energy has been carefully measured for over 35 years. This careful measurement has led to the realization that the Hulse-Taylor binary is losing rotational energy at precisely the rate one would expect, if that energy were being carried away by gravitational waves as predicted by general relativity.

Direct evidence of gravitational waves is expected to be discovered within the next decade through

the use of extremely precise laser interferometry. Several major experiments around the globe, notably the LIGO experiment¹ in the United States, Virgo in Italy, and GEO 600 in Germany, are refining their sensitivity in the hope of observing GWs. LIGO is composed of a series of Michelson interferometers, situated at Hanford, Washington, and Livingston, Louisiana. Each site has an L-shaped vacuum system, 4 kilometers on a side, within which operate one or more interferometers, the longest being the full length of the vacuum system. The arms of the interferometer are Fabry-Perot cavities that recycle the laser light, effectively extending the arm length by a factor of ~ 100 . This arm length results in LIGO's sensitivity being highest in the frequency range of 10 to 1000 Hz. LIGO's sensitivity has progressively increased over the years, with improvements such as Enhanced LIGO and Advanced LIGO modifying the configuration of the system to greatly increase LIGO's range. It is expected that Advanced LIGO, currently being installed and expected to be operational in 2014, will improve LIGO's sensitivity to the point where detection of a gravitational wave will become almost a daily occurrence.

Direct experimental evidence of GWs, it is hoped, will further confirm the validity of Einstein's general relativity, and provide a stepping-stone towards observing the universe in the gravitational spectrum, just as we currently do in the electromagnetic spectrum. However, it is also possible that the lack of detection, or detection of waves not in line with the predictions of general relativity, will result in direct experimental evidence which contradicts relativity. Either way, evidence, or lack thereof, of gravitational waves is a key part of confirming or disproving relativity.

The challenges for GW detection are even today very daunting. In current or foreseen experiments, any expected GW signal is buried beneath significant noise. The LIGO mirrors, which function as the end-points of the interferometer, and motion of these mirrors must be detected at the level of 10^{-19} m. Such small distances are extremely hard to measure, and matched filtering techniques are typically used to amplify the signal above the noise. As such, a good idea of what these signals can be expected to look like is needed.

For any given type of event, such as a binary black hole merger, there is a large space of possible waveforms, arising from differences in the physical parameters of the emitting system. For example, in the case of binary black hole mergers, differences between the masses of the two black holes can result in significantly different mergers, with different signals. Since the two-body problem does not have a known analytic solution, numerical simulations and analytic approximations of potential sources can be used to lay down a series of "templates" - sample simulations or approximate analytic formulae that cover a region of the parameter space of all possible events of the given type. These have in the past been primarily performed using the post-Newtonian expansion of general relativity - a simplified theory of gravity which applies only in situations of weaker gravitational forces. Increasingly, templates are being constructed using more accurate numerical simulations of

¹www.ligo.caltech.edu

the full theory of relativity, with the increased accuracy this brings.

In order to observe direct evidence using laser interferometric experiments such as LIGO, three main areas of expertise must be developed. Firstly, the equipment itself must be designed, understood, characterized and improved. These experiments require highly developed apparatus in order to isolate the systems as much as possible from sources of noise. Major developments include improved seismic isolation and suspension systems for the mirrors that form the ends of the interferometer cavities, improvements in thermal coating of mirrors to reduce noise from internal thermal fluctuations, and advances in interferometer design, such as increased laser power, to improve sensitivity.

Secondly, the scientific community must determine which astrophysical events present the highest likelihood for detection, and what waveforms those events would produce. Because of the conservation of momentum, in a multipolar analysis the lowest-order (and strongest) GW emission comes from a varying mass quadrupole. Linear motion is not sufficient, but any two bodies orbiting each other are primarily engaged in quadrupole motion. However, the gravitational waves emitted by typical quadrupolar motions visible to astronomers (such as the orbit of a planet around a star) are not nearly powerful enough to be detected. Only the most violent events in the universe produce gravitational waves sufficiently powerful to stand any chance of detection. Currently, LIGO is expected to be able to detect a pair of solar mass neutron stars orbiting each other, if they are within 60 million light years. When emitting detectable waves, these neutron stars would orbit each other approximately 100 times per second at a distance of about 100km between them. Key aspects of detecting such violent events include numerical simulations of events in the strong gravity regime, where full general relativity must be applied, without simplifying approximations that would be valid in weaker gravity. These simulations can be used to extract gravitational waveforms, the study of which can be used to develop generalized waveforms which are functions of the astrophysical parameters of the event.

Thirdly, techniques must be developed to search for these signals in the noise. Since even comparatively simple waveforms may have upwards of 5 parameters (such as the relative masses, sky position, polarization and time of arrival), finding a waveform in experimental data requires searching over large parameter spaces. Techniques of statistical analysis must be developed and tested, so that we can search for real events in detector output data.

This thesis addresses issues in the second and third parts of the quest to detect gravitational waves. In Part I, I discuss my work as part of the “sXs” collaboration, a group of researchers at Caltech, Cornell, and CITA who develop and use “SpEC” software. “SpEC” is a code built to solve Einstein’s equations, specializing in the simulation of binary black hole mergers. My work has centered on investigating the numerical simulation of event horizons of black holes. I have implemented simulations to track the motion and merger of the event horizons of binary black holes

in head-on and inspiral mergers, building on existing simulations of the bulk gravitational fields. I report on results concerning the topology of event horizon mergers, among other results. In order to perform this work, I have designed and implemented code modules within the framework of SpEC.

In Part II I discuss my work on parameter estimation of gravitational wave templates. These templates are models of the expected detector response to a gravitational wave emitted from a given source, depending on parameters such as the position of the source on the sky, its polarization, etc. This work was done in the context of the planned LISA mission, a collaboration between NASA and the European Space Agency. LISA will be a configuration of three spacecraft in an equilateral triangle, which will orbit the sun at 1 AU distance. LISA is intended to complement the sensitivity of earthbound observatories like LIGO by attempting to detect gravitational waves in the frequency range of 10^{-4} to 1 Hz. It is impractical for earthbound detectors to be sensitive at these low frequencies due to seismic activity. LISA therefore, is planned to operate in space, and will have arm lengths of approximately five million kilometers. LISA has not yet been fully approved, so a launch date is at least a decade away. In order to improve the science case for LISA, and to advance the science of parameter estimation in GW models, the LISA International Science Team has launched a series of “Mock LISA Data Challenges” (MLDCs). The MLDCs are competitions where simulated LISA data sets are produced, based on theoretical models of the behavior of the LISA constellation. Signals from potential sources are embedded into these data sets, and scientists around the world are challenged to locate these signals within the noise. This section of the thesis focuses on my work for MLDC challenge 3.4, which for the first time includes signals from cosmic strings as potential sources of GWs for LISA.

Not all work in this thesis is my own. In Part I, the underlying numerical simulations of Einstein’s equations are performed by various members of the SpEC group. Significant portions (primarily the latter half) of Chapter 2 are quoted from Cohen, Pfeiffer & Scheel [1]. All of Chapters 3 and 4 are quoted from [1] except Section 3.4, which is newly written for this thesis. In Chapter 5, Section 5.1 is quoted from [1] (except for Figure 5.8, which has been updated), while Section 5.2 is quoted from Lovelace et. al. [2], which is reproduced in its entirety as Appendix A. Cohen, Pfeiffer & Scheel [1] was written collaboratively by the authors. Lovelace et al. [2] was written primarily by Geoffrey Lovelace, not myself.

All of Part II except Section 8.1 is quoted from Cohen, Cutler & Vallisneri [3]. In Part II, the waveform generator for cosmic-string burst gravitational waveforms was originally written by Michele Vallisneri. The F-statistic and Fourier transform approach was originally devised by Curt Cutler and Michele Vallisneri. All sections of the paper were written collaboratively, with Curt taking the lead on those sections reproduced in Chapters 8 and 9, Michele on Chapter 10, and myself on Chapters 11, 12 and 13.

Bibliography

- [1] Michael Cohen, Harald P. Pfeiffer, and Mark A. Scheel. Revisiting event horizon finders. *Class. Quantum Grav.*, 26:035005, 2009.
<http://iopscience.iop.org/0264-9381/26/3/035005/>, ©2009 IOP Publishing Ltd.
- [2] Geoffrey Lovelace, Yanbei Chen, Michael Cohen, Jeffrey D. Kaplan, Drew Keppel, Keith D. Matthews, David A. Nichols, Mark A. Scheel, and Ulrich Sperhake. Momentum flow in black-hole binaries: II. Numerical simulations of equal-mass, head-on mergers with antiparallel spins. arXiv:0907.0869 (gr-qc), 2009.
- [3] M. Cohen, C. Cutler, and M. Vallisneri. Searches for cosmic-string gravitational-wave bursts in mock lisa data. *Class. Quantum Grav.*, 27:185012, 2010.
<http://iopscience.iop.org/0264-9381/27/18/185012/>, ©2010 IOP Publishing Ltd.

Part I

Simulating Event Horizons in Numerical Spacetimes

Chapter 2 Introducing Event Horizons

The idea that a body could be sufficiently massive that light could not escape its pull dates back to an 18th century geologist named John Michell, who postulated the idea based on Newtonian principles. In the modern sense, however, black holes were conceived in the few years following the publication of Einstein's Theory of General Relativity in 1915. Several months following publication, Karl Schwarzschild published a solution to Einstein's equations for a spherically symmetric point mass. This solution showed that if the mass at the center of the field was dense enough, there would be a region outside the mass which was causally disconnected from the rest of the universe. No information or energy from inside this region could escape to the outside. This is because the gravitational field inside this region is so strong, not even light can escape its pull. The boundary of this region, known as the Schwarzschild radius, is linearly dependent on the mass of the object. Only an object whose entire mass lies inside its Schwarzschild radius creates this causally disconnected region. By way of example, you would need to force the entire mass of the earth into a sphere roughly 9 millimeters in size to create a black hole. Most black holes, however, are not spherically symmetric. In 1963, Roy Kerr discovered the general solution for a rotating black hole, known as a Kerr black hole. Additional modifications need to be made to this metric to include the effects of charge, but astrophysically, charge is not typically a dominant feature of a black hole. Rotation, however, is a dominant feature, and most results that deal with a single black hole tend to spring from the general Kerr metric. In general, the Kerr and Schwarzschild metrics establish a surface which, if all matter of the object is contained within it, creates a causally disconnected region of spacetime. The boundary between this disconnected region and the rest of the universe is known as the *event horizon* of the black hole. One important feature of black holes, known as the Second Law of Black Hole Thermodynamics, is that as long as the weak energy condition holds (i.e. no negative energy), the surface area of the event horizon (and correspondingly the mass of the black hole) can only increase in time.

A precise definition of an event horizon is given in [1], and can be paraphrased thusly:

An asymptotically flat (and strongly asymptotically predictable) spacetime M is said to contain a black hole if not every point of M is contained in the causal past of future null infinity. The black hole region, B , of such a spacetime is defined to be the points of M not contained in the causal past of future null infinity. The boundary of B in M is called the event horizon.

The event horizon is a null hypersurface, defined as the sum of the paths of the null generators of the horizon. These generators are a specific set of null geodesics that have no future end-point, and, once they are on the event horizon, do not intersect any other generator. In non-stationary

spacetimes, not all generators of the black hole are on the event horizon at all times. Some generators of the event horizon have merged onto it at some specific point in time. At this point, they intersected other generators of the event horizon. Prior to this point, these generators were not on the event horizon at all. Excepting these merger points, one and only one generator passes through each point on the event horizon.

Inside the event horizon, at the center of the black hole, a point of infinite curvature is predicted by general relativity. This point, known as a singularity, is a point where the theory of relativity breaks down. It is postulated that all such singularities must be surrounded by an event horizon, and be thus causally disconnected from the rest of the universe. This idea, known as cosmic censorship, is of profound importance to our understanding of the limits of general relativity.

Given that theoretically black holes can exist, one must still find a process by which they may be reliably created. In the 1930's it was established that collapse of a large amount of matter could form a black hole. This collapse would occur within a star which, having reached the end of its lifespan, has run out of materials to sustain its fusion reaction and is no longer capable of sustaining itself against its own self-gravity. Small stars, when they reach this point, expand to become red giants and throw off large quantities of stellar matter, after which they will collapse to form a white dwarf if the remaining mass is $< 1.44M_{\odot}$, where M_{\odot} is the mass of the sun (this limit is known as the Chandrasekhar limit). A white dwarf is a stellar remnant, composed of matter which is held up by the pressure of electron degeneracy (the tendency of electrons to not wish to be in the same state and location as another electron), with a mass approximately like the Sun and a volume approximately like the Earth. Larger stars will typically undergo a much more energetic explosion known as a supernova. For these stars, if a larger mass remains (up to approximately $3M_{\odot}$), the collapse continues through the white dwarf stage until the star stabilizes as a neutron star. A neutron star has collapsed further than a white dwarf, and is held up by the force of neutron degeneracy (neutrons, just like electrons, do not want to be in the same state and location as another neutron). However for stars with more than about $3M_{\odot}$ remaining after their supernova, the forces of neutron degeneracy which hold up a neutron star are not sufficient to halt the collapse. No known force can hold up the mass, and the star collapses into a black hole.

Given that the process of stellar collapse can form a black hole if the star is sufficiently massive, the population of large stars in the universe suggests that black holes are quite common-place. Extremely massive compact objects, potentially supermassive black holes, are postulated to reside at the centers of galaxies like the Milky Way, based on measurements of the motion of stars and other visible masses around the galactic center. Globular clusters may also harbor large black holes at their centers. There is as yet no widely agreed mechanism by which these supermassive black holes are formed. Current theories include formation by slow accretion of nearby matter, the collapse of extremely large gas clouds, or multiple mergers of smaller black holes. Additionally, stellar-mass

black holes (or objects very much like them) have been observed in binary systems, often through x-ray emissions thought to originate from the accretion disk around the black hole, which is fed by matter pulled from the black hole’s more visible companion.

Black holes are relatively simple astrophysical objects. A number of “no hair” theorems (see e.g. p. 876 of [2]) show that black holes at rest can be uniquely characterized by only three parameters, their mass, spin and charge. Since significant charge is unlikely to build up around astrophysical objects (due to the presence of both positively and negatively charged particles in a roughly 1:1 ratio), most discussions of black holes omit charge in favor of concentrating on mass and spin alone. Since black holes are such simple objects, binary black hole systems are the configuration of choice for studying the two-body problem in general relativity, for which no analytic solution exists.

The two body problem in general relativity has been the focus of extensive work for many years and, because there is no analytic solution, it must be solved numerically. Binary black hole mergers are expected to be one of the most astrophysically common sources of gravitational radiation for detectors such as LIGO [3, 4]. Recent advances in simulating binary black hole mergers include the development of the generalized harmonic evolution system [5] and the moving punctures technique [6, 7]. In the last several years the field has reached a stage where binary black hole simulations are becoming routine. Numerical simulations have been remarkably successful in expanding our understanding of binary black holes, but challenges remain.

One particular challenge is to be able to more accurately locate the holes during the merger. There are two useful concepts to describe the location of black holes in a spacetime, *apparent horizons* (AH) and *event horizons* (EH). An EH is the true surface of a black hole: it is defined as the boundary of the region of the spacetime that is causally connected to future null infinity. Because the definition of the EH involves global properties of the spacetime, one must know the *full* future evolution of the spacetime before the EH can be determined exactly. This difficulty has led researchers to instead identify black holes with apparent horizons, which are defined in terms of the expansion of null congruences¹. Indeed, AH finders are highly developed and have been the subject of extensive work (see, e.g. the review [8]) Unlike an EH, an AH can be located from data on a single spacelike hypersurface, i.e. on each timestep of a numerical evolution, without knowing the future evolution of the spacetime. The AH is often an effective substitute for the EH for several reasons. First, according to the cosmic censorship conjecture, if an AH is present, it must be surrounded by an EH. Second, if an AH is present on a spacelike hypersurface through a stationary spacetime, it coincides with the EH. Finally, in numerical simulations, apparent horizons generally show behaviour attributed to event horizons: For instance, the area of the AH typically does not decrease and it is usually almost constant whenever the spacetime is only mildly dynamic.

¹More precisely, we define AH as the outermost marginally outer-trapped surface, where an outer-trapped surface is a topological 2-sphere with zero expansion along outgoing null normals.

In fact, apparent horizons have motivated the development of “isolated” and “dynamical” horizons (see [9] for a review). These surfaces satisfy analogues of the laws of black hole thermodynamics, although they are defined quasi-locally, rather than globally.

However, using the AH to locate the holes is not always appropriate. For instance, the AH is slicing dependent, while the EH is not. Indeed, the Schwarzschild spacetime can be sliced in such a way that no AH exists [10]. Furthermore, even on slicings on which an AH is present, there are few precise mathematical statements about how “close” AH and EH are. Finally, AH and EH behave qualitatively differently during a black hole merger: The EH² around each black hole expands continuously until the two components of the EH join into one, whereas a common apparent horizon appears discontinuously quite some time after the EHs have merged. The common AH encompasses the two individual AHs, which continue to exist as surfaces of zero outgoing null expansion for some time after the merger.

Early EH finders [11, 12] followed null geodesics forward in time and determined whether or not each geodesic eventually escapes to infinity. Following geodesics forward in time is unstable in that slightly perturbed geodesics will diverge from the EH and either escape to infinity or fall into the singularity. Furthermore, a large number of geodesics with different directions must be sampled at each point and at each time step to determine if one of these succeeds in escaping to infinity [12]. To reduce the number of sampling points, the EH search in [12] was performed on a series of time slices proceeding backward from late to early times; to find the EH on each time slice, they integrated geodesics forward in time, using the already-located EH at the later time as an initial guess.

Since outgoing null geodesics diverge from the event horizon when going forward in time, when going *backward* in time they will converge onto the event horizon [13, 14]. All recent EH finders use this observation, and follow null geodesics or null surfaces backward in time [8, 13, 14, 15, 16, 17, 18, 19, 20, 21].

Several algorithms have been developed to follow null geodesics backward in time. These can be divided into three types, which we shall refer to as the “geodesic method,” the “surface method” and the “level-set method.” The geodesic method works by simply integrating the geodesic equation, as done by Libson et. al. [14]. Libson et al. express concerns that the geodesic method may be susceptible to tangential “drifting” of the geodesics. However, this is not evident when the method is applied to the science applications in that paper, nor do we find tangential drifting in our simulations. To avoid any issues with drifting, Libson et al. introduced the surface method: a complete null surface (rather than individual geodesics) is evolved backward in time. In [13, 14] this surface was parameterized based on axisymmetry (although the parameterization of [13, 14] cannot handle *generic* axisymmetric situation, cf. Section 3.2.2 below), and many interesting results on the structure of caustics and the geometry of the horizon for axisymmetric spacetimes were obtained in

²More precisely, the 2-surface formed by the intersection of the spatial slice and the EH.

[15, 17]. Diener [21] and Caveny et al. [18, 19, 20] independently introduced the level-set method by recasting the surface method in a way that does not assume symmetry: rather than evolving a single 2-D surface, they evolve a volume-filling series of surfaces given as the level-sets of a space-time function $f(t, x^i)$. To avoid exponentially steepening gradients of f , Caveny et al. introduce an artificial diffusive term, whereas Diener reinitializes f whenever necessary.

This thesis re-examines these techniques for event horizon finding in the context of the Caltech/Cornell Spectral Einstein Code (SpEC), which provides an infrastructure for highly accurate simulations of Einstein’s equations for single and binary black holes. Recent work includes highly accurate computations of gravitational waveforms from inspiraling binaries [22, 23, 24]. The availability of high accuracy binary evolutions motivates the development of very precise event horizon finding techniques in order to extract all possible physics from these simulations. Therefore, this paper reconsiders the three techniques mentioned above in the context of general binary black hole mergers without any symmetries.

We implement the geodesic method, and generalize the surface method to arbitrary situations without symmetries. Both methods are then applied to single Kerr black holes, and a head-on binary black hole merger. In both cases, the geodesic method is found to be more robust. We encounter two fundamental problems with the level-set method, and therefore halted our efforts to implement it in SpEC.

This part of the thesis is organized as follows: In Chapter 3, we explain the three methods in more detail and give details of our numerical implementation. Chapter 4 presents results for a single Kerr black hole, and in Chapter 5, we apply the techniques to a head-on BBH merger, where we extract ringdown behaviour and the behaviour of the individual event horizons before merger. We also consider an application of event horizon finding to gaining physical insight into the “maximum kick” head-on merger. In Chapter 6 we compute event horizons of spinning and non-spinning binary black hole inspirals, and examine in detail the topology of these mergers. We close with a conclusion in Chapter 7. Note that in subsequent chapters I assume knowledge of general relativity at the senior undergraduate or junior graduate level at the least. For a good introductory text, see “An Introduction to General Relativity, Spacetime and geometry” by Sean M. Carroll.

Bibliography

- [1] Robert M. Wald. *General Relativity*. University of Chicago Press, Chicago and London, 1984.
- [2] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. Freeman, New York, New York, 1973.

- [3] Barry C. Barish and Rainer Weiss. LIGO and the detection of gravitational waves. *Phys. Today*, 52(10):44–50, Oct 1999.
- [4] Samuel J. Waldman. Status of LIGO at the start of the fifth science run. *Class. Quantum Grav.*, 23(19):S653–S660, 2006.
- [5] Frans Pretorius. Simulation of binary black hole spacetimes with a harmonic evolution scheme. *Class. Quantum Grav.*, 23(16):S529–S552, 2006.
- [6] Manuela Campanelli, Carlos O. Lousto, Pedro Marronetti, and Yosef Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Phys. Rev. Lett.*, 96(11):111101, 2006.
- [7] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter. Gravitational-wave extraction from an inspiraling configuration of merging black holes. *Phys. Rev. Lett.*, 96(11):111102, 2006.
- [8] Jonathan Thornburg. Event and apparent horizon finders for 3+1 numerical relativity. *Living Rev. Rel.*, 10, June 2007. 3.
- [9] Abhay Ashtekar and Badri Krishnan. Isolated and dynamical horizons and their applications. *Living Rev. Rel.*, 7, December 2004. 10.
- [10] R. M. Wald and V. Iyer. Trapped surfaces in the Schwarzschild geometry and cosmic censorship. *Phys. Rev. D*, 44:R3719, 1991.
- [11] Stuart L. Shapiro and Saul A. Teukolsky. Gravitational collapse to neutron stars and black holes: Computer generation of spherical spacetimes. *Astrophys. J.*, 235:199–215, 1980.
- [12] Scott A. Hughes, Charles R. Keeton, II, Paul Walker, Kevin T. Walsh, Stuart L. Shapiro, and Saul A. Teukolsky. Finding black holes in numerical spacetimes. *Phys. Rev. D*, 49(8):4004–4015, 1994.
- [13] Peter Anninos, David Bernstein, Steven Brandt, Joseph Libson, Joan Massó, Edward Seidel, Larry Smarr, Wai-Mo Suen, and Paul Walker. Dynamics of apparent and event horizons. *Phys. Rev. Lett.*, 74(5):630, 1995.
- [14] J. Libson, J. Massó, , E. Seidel, W.-M. Suen, and P. Walker. Event horizons in numerical relativity: Methods and tests. *Phys. Rev.*, D53:4335–4350, 1996.
- [15] R. A. Matzner, E. Seidel, S. L. Shapiro, L. L. Smarr, W.-M. Suen, S. A. Teukolsky, and J. Winicour. Geometry of a black hole collision. *Science*, 270(5238):941–947, 1995.

- [16] Paul Walker. *Horizons, Hyperbolic Systems, and Inner Boundary Conditions in Numerical Relativity*. PhD thesis, University of Illinois at Urbana-Champaign, 1998.
- [17] J Massó, E Seidel, W-M Suen, and P Walker. Event horizons in numerical relativity. II. analyzing the horizon. *Phys. Rev. D*, 59:064015, 1999.
- [18] Scott A. Caveny. *Tracking Black Holes in Numerical Relativity: Foundations and Applications*. PhD thesis, University of Texas at Austin, 2002.
- [19] Scott A. Caveny, Matthew Anderson, and Richard A. Matzner. Tracking black holes in numerical relativity. *Phys. Rev. D*, D68:104009, 2003.
- [20] Scott A. Caveny and Richard A. Matzner. Adaptive event horizon tracking and critical phenomena in binary black hole coalescence. *Phys. Rev. D*, D68:104003, 2003.
- [21] Peter Diener. A new general purpose event horizon finder for 3D numerical spacetimes. *Class. Quantum Grav.*, 20:4901–4917, 2003.
- [22] Harald P. Pfeiffer, Duncan A. Brown, Lawrence E. Kidder, Lee Lindblom, Geoffrey Lovelace, and Mark A. Scheel. Reducing orbital eccentricity in binary black hole simulations. *Class. Quantum Grav.*, 24(12):S59–S81, 2007.
- [23] Michael Boyle, Duncan A. Brown, Lawrence E. Kidder, Abdul H. Mroué, Harald P. Pfeiffer, Mark A. Scheel, Gregory B. Cook, and Saul A. Teukolsky. High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions. *Phys. Rev. D*, 76:124038, 2007.
- [24] Michael Boyle, Alessandra Buonanno, Lawrence E. Kidder, Abdul H. Mroué, Yi Pan, Harald P. Pfeiffer, and Mark A. Scheel. High-accuracy numerical simulation of black-hole binaries: Computation of the gravitational-wave energy flux and comparisons with post-Newtonian approximants. *Phys. Rev. D*, 78(12):104020, 2008.

Chapter 3 Background and Basic Techniques

3.1 Notation

Throughout this thesis, Greek indices ($\alpha, \beta, \gamma \dots$) will be used to denote sums over $\alpha = 0, 1, 2, 3$, while lower-case Latin indices ($a, b, c \dots$) denote sums over $a = 1, 2, 3$. Upper-case Latin indices ($A, B, C \dots$) will be occasionally used to denote sums over $A = 1, 2$. We use the Einstein summation convention throughout.

3.2 Basic Techniques for Finding Event Horizons

All EH-finding techniques considered here proceed backward in time and must therefore be performed after the numerical evolution of the spacetime has been completed. We assume that we have access to the spacetime metric in a 3+1 decomposition

$$ds^2 = -N^2 dt^2 + \gamma_{ij}(dx^i + \beta^i dt)(dx^j + \beta^j dt), \quad (3.1)$$

where N is the lapse, β^i is the shift, and γ_{ij} is the 3-metric on the slice. Latin indices $i, j, \dots = 1, 2, 3$ denote spatial dimensions; below we will use Greek indices to denote spacetime dimensions, $\alpha, \beta \dots = 0, 1, 2, 3$. The time t in (3.1) represents the coordinate time of the numerical evolution. Typically, the metric data γ_{ij} , β^i , and N are available at discrete times and at discrete spatial grid points. Evaluating the values of the metric components elsewhere requires interpolation.

A black-hole merger exhibits several characteristic features of relevance to EH finders, as illustrated in Figure 3.1¹ (cf. [1, 2, 3]). At times sufficiently far prior to merger, the EH and AH are expected to coincide closely (and indeed, we confirm this below for our simulation). The green dashed curves in Figure 3.1 represent future generators of the event horizon, i.e. null geodesics that will merge onto the event horizon through cusps in the individual event horizons. These cusps are clearly visible at time $t = 13.5M$ where the individual EHs have diverged significantly from their respective AHs. At $t_{\text{CEH}} = 14.6M$ the two previously disjoint components of the event horizon join. We shall refer to this time t_{CEH} as *the merger* of the black hole binary. After the merger, the event horizon of the merged black hole can be seen relaxing towards its final time-independent shape. The common apparent horizon appears at $t_{\text{CAH}} = 17.8M$, and approaches the event horizon as the evolution proceeds; at $t = 80M$, the AH coincides almost exactly with the event horizon.

¹While Figure 3.1 is meant as an illustration, it presents actual data from the head-on binary black hole merger discussed in Section 5.1. The time given at each frame of Figure 3.1 will aid in the discussion in Section 5.1.




Figure 3.1: Cross-sections through event and apparent horizons during a BBH merger. Before the merger, $t < t_{\text{CEH}}$, the surface includes the set of generators that will merge onto the event horizons through the cusps in the individual event horizons (green dashed curves). The point C is the point of symmetry for the head-on merger, which will be used in Section 3.2.3.

3.2.1 Geodesic Method

The most straightforward way to follow light rays is to simply integrate the geodesic equation [4, 5, 6, 1],

$$\frac{d^2 q^\mu}{d\lambda^2} + \Gamma_{\alpha\beta}^\mu \frac{dq^\alpha}{d\lambda} \frac{dq^\beta}{d\lambda} = 0, \quad (3.2)$$

where $q^\mu = q^\mu(\lambda)$ is the position of the photon on the geodesic, parameterized by an affine parameter λ , and $\Gamma_{\alpha\beta}^\mu$ are the spacetime Christoffel symbols.

Since we have access to our spacetime as a function of the evolution time coordinate t , it is convenient to rewrite (3.2), replacing λ by t along the geodesic. Writing $\dot{q}^\mu = dq^\mu/dt$, and $a = d\lambda/dt$, we find:

$$\frac{dq^\mu}{d\lambda} = \frac{1}{a} \dot{q}^\mu, \quad (3.3)$$

$$\frac{d^2 q^\mu}{d\lambda^2} = \frac{1}{a^2} \ddot{q}^\mu - \frac{\dot{a}}{a^3} \dot{q}^\mu. \quad (3.4)$$

Substituting into the geodesic equation we get

$$\ddot{q}^\mu = \frac{\dot{a}}{a} \dot{q}^\mu - \Gamma_{\alpha\beta}^\mu \dot{q}^\alpha \dot{q}^\beta. \quad (3.5)$$

The quantity a is determined by the requirement that $q^0 = t$, i.e. that at parameter value t along the geodesic, the geodesic is on the corresponding $t = \text{const}$ hypersurface of the evolution. This implies $\dot{q}^\mu = [1, \dot{q}^i]$ and $\ddot{q}^\mu = [0, \ddot{q}^i]$. Setting $\mu = 0$ in (3.5) gives $\frac{\dot{a}}{a} = \Gamma_{\alpha\beta}^0 \dot{q}^\alpha \dot{q}^\beta$. The spatial components of (3.5) are the desired evolution equation for the spatial coordinates as a function of coordinate time,

$$\ddot{q}^i = \Gamma_{\alpha\beta}^0 \dot{q}^\alpha \dot{q}^\beta \dot{q}^i - \Gamma_{\alpha\beta}^i \dot{q}^\alpha \dot{q}^\beta. \quad (3.6)$$

We convert this set of ordinary differential equations to first order form by defining $p^i \equiv \dot{q}^i$, which gives

$$\dot{q}^i = p^i, \quad (3.7a)$$

$$\dot{p}^i = \Gamma_{\alpha\beta}^0 p^\alpha p^\beta p^i - \Gamma_{\alpha\beta}^i p^\alpha p^\beta. \quad (3.7b)$$

This facilitates the use of standard ODE integrators like Runge-Kutta methods [7, 8].

While integrating geodesics is not new [5, 1], re-expressing the geodesic equation in terms of coordinate time in numerical integration seems to be new. It appears that the primary reason this technique has been phased out in favour of the two techniques described below is the concern that, in a full 3D implementation, slight tangential velocities may be imparted to the outgoing null geodesics through numerical inaccuracies, and that this tangential drift of geodesics could result in unphysical caustics. These concerns are discussed in detail in [1], where the idea of representing the whole surface, rather than individual geodesics, was introduced. This was justified on the basis that for a surface, tangential drift is irrelevant. However, while it is possible that tangential drift can be significant for very coarse, low-resolution simulations, we see no evidence that tangential drift affects our numerical tests of the geodesic method.

We finally like to point out that if one evolves $p_i = g_{i\mu} p^\mu$ instead of p^i (cf. [5]), then the evolution equations depend only on spatial derivatives of the spacetime metric. Evolving p_i therefore results in computational savings, because the time derivatives of the metric need not be stored or interpolated. This will be investigated in a future work.

3.2.2 Surface Method

The idea of the surface method dates back to Libson et al. [1], who used it in axisymmetry. The goal is to evolve a 2-dimensional surface \mathcal{S}_t backward in time such that it traces out a null hypersurface \mathcal{N} . The time coordinate t is inherited from the black hole simulation for which event horizons are to be determined, i.e. \mathcal{S}_t is the intersection of \mathcal{N} with the spatial hypersurfaces Σ_t of the evolution, as indicated in Figure 3.2. Before the black hole merger, $t < t_{\text{CEH}}$, the surface \mathcal{S}_t consists not only of the two disjoint parts of the event horizon, but also includes the future generators, which




Figure 3.2: A slice of the event horizon (\mathcal{S}_t), produced by the intersection of a spatial hypersurface Σ_t with the world tube of the event horizon \mathcal{N} , showing n^μ (the timelike normal to Σ_t), s^μ (the spatial normal of \mathcal{S}_t), and ℓ^μ (the null normal to the EH world-tube \mathcal{N}).

are indicated by the green dashed curves in Figure 3.1. The union of these three components is a smooth self-intersecting surface with the topology of a sphere (as suggested by Kip Thorne[6, 1]).

Let us first consider how to represent the surface to be evolved. Apparent horizon finders often parameterize a surface by giving the radius, relative to a fixed point, as a function of angular coordinates, i.e. $r = f(\theta, \phi)$. Such a star-shaped surface is insufficient here, because the surface will be self-intersecting for $t < t_{\text{CEH}}$ and will cease to be star-shaped even before then (see Fig. 1.1 of [9]). The axisymmetric EH finder presented in [1] parameterized the surface by $\rho = s(z, t)$, where z is a coordinate along the axis of symmetry, and ρ is the cylindrical radius. This allows some mild form of self-intersection, like, for instance, the $t = 13.5M$ snapshot in Figure 3.1. However, at earlier times, the locus of future null generators of the horizon “bulges outward” and becomes multivalued when considered as a function of z , cf. $t = 9M$ in Figure 3.1. In this case, the parameterization of [1] fails even for an axisymmetric configuration. In this paper, we use a parametric representation of \mathcal{S}_t , i.e. $r^i = r^i(t, u, v)|_t^2$. The full 3-dimensional null hypersurface \mathcal{N} being constructed is represented as a 3-parameter surface in spacetime:

$$r^\mu(t, u, v) = [t, r^i(t, u, v)]. \quad (3.8)$$

We wish to find an equation that will allow us to evolve \mathcal{S}_t in such a way as to trace out the null 3-surface \mathcal{N} . Further, we would like this equation to have the property that for fixed (u_0, v_0) , the curve $r^\mu(t, u_0, v_0)$ traces out a null geodesic. This allows us to directly compare the surface obtained by the surface method to the surface obtained for equivalent initial conditions by the geodesic method.

For the curve $r^\mu(t, u, v)|_{u,v}$ to be null, its tangent $\partial r^\mu(t, u, v)/\partial t$ must be outgoing and null, i.e.

$$\frac{\partial r^\mu}{\partial t} = \ell^\mu, \quad (3.9)$$

²Although (u, v) are typically used to represent null coordinates, here they represent arbitrary parametric coordinates

where ℓ^μ is a null normal to \mathcal{S}_t . ℓ^μ can be written as

$$\ell^\mu = c(n^\mu + s^\mu), \quad (3.10)$$

where n^μ is the timelike unit normal to Σ_t , s^μ is the spatial outward-pointing unit normal to \mathcal{S}_t (cf. Figure 3.2) and c is an overall scaling. Consistency of (3.8) and (3.9) requires that ℓ^μ is normalized such that $\ell^t = 1$. To find the value of c from the condition $\ell^t = 1$, first notice that from the 3+1 decomposition,

$$n^\mu = \frac{1}{N} [1, -\beta^i], \quad (3.11)$$

where N and β^i are the lapse and shift fields. Also, since s^μ lies within the spatial slice Σ_t , we may write $s^\mu = [0, s^i]$, so that (3.10) becomes

$$\ell^\mu = c \left[\frac{1}{N}, s^i - \frac{1}{N} \beta^i \right]. \quad (3.12)$$

Thus $\ell^t = 1$ implies $c = N$, and we can write our final evolution equation for the spatial components of r^i ,

$$\frac{\partial r^i}{\partial t} = N s^i - \beta^i. \quad (3.13)$$

In order to find the unit normal s^i to the spatial surface \mathcal{S}_t , we follow the standard procedure for a surface parameterized as $r^i(u, v)$, i.e.

$$\tilde{s}^i = \gamma^{il} \epsilon_{ljk} \frac{\partial r^j}{\partial u} \frac{\partial r^k}{\partial v}, \quad (3.14a)$$

$$\rho = \sqrt{\gamma_{ij} \tilde{s}^i \tilde{s}^j}, \quad (3.14b)$$

$$s^i = \rho^{-1} \tilde{s}^i. \quad (3.14c)$$

where ϵ_{ljk} is the antisymmetric tensor and where we have chosen the sign of the root such that s^i points outward for a right-handed choice of coordinates.

This evolution equation for the surface method (3.13) is very different from the evolution equations for the geodesic equation (3.7b)-(3.7a). The surface method does not require derivatives of the metric, but derivatives $\partial_u r^i$, $\partial_v r^i$ along the surface; the geodesic method, in contrast, requires derivatives of the metric, but treats each geodesic completely independently. Nevertheless, due to our choice of evolution equation (3.9), each point on the parameterization of the surface traces its own geodesic; see Section 3.3 for a proof.

3.2.3 Level-set Method

The level-set method [1, 10, 3, 11, 12, 2, 13] utilizes a function $f = f(t, x^i)$ defined on the full spacetime (or at least, a region of spacetime covering the vicinity of the expected location of the EH). The function f is determined such that $f = \text{const}$ contours (i.e. level-sets) represent null surfaces i.e. $g^{\alpha\beta}\partial_\alpha f \partial_\beta f = 0$. In the 3+1 decomposition, this becomes [12, 2, 13],

$$\partial_t f = \beta^i \partial_i f \mp N \sqrt{\gamma^{ij} \partial_i f \partial_j f}, \quad (3.15)$$

where the \mp accommodates both ingoing and outgoing null surfaces, with the minus sign being appropriate for outgoing null surfaces if the gradient $\partial_i f$ is outward-pointing.

Libson et al. [1] had previously made use of (3.15), but parameterized the $f = 0$ contour based on axisymmetry. The motivation of evolving (3.15) directly in the volume is to remove any assumptions of symmetries.

Unfortunately, when trying to implement the level-set method in SpEC, we encountered two fundamental problems. The first difficulty is related to the characteristic speed of the level-set method. Simply put, all $f = \text{const}$ contours approach the event horizon, therefore new contours need to be filled in at the boundaries of the region in which f is evolved (i.e. the outer boundary and possibly one or more inner boundaries if black hole excision is employed). To see this, note that the characteristic speed of (3.15) relative to a spatial direction \bar{n}_i is

$$v = N \bar{n}_i \frac{\partial^i f}{\sqrt{\gamma^{ij} \partial_i f \partial_j f}} - \bar{n}_i \beta^i, \quad (3.16)$$

where the sign of the first term depends on the gradient $\partial_i f$ being outward pointing. For most coordinate systems of interest, lapse N and shift β^i behave such that $v > 0$ at the outer boundary and at any excision boundaries (if present). When integrating (3.15) backward in time, well-posedness requires boundary conditions at these boundaries. Our preferred numerical techniques are spectral methods because of their promise to achieve exponential convergence for smooth problems. Spectral methods are very sensitive to the existence of an underlying well-posed continuum problem and therefore require boundary conditions. Unfortunately there is no particular physical reasoning to suggest a choice of boundary condition. While essentially any choice of boundary condition that results in f being continuous rendered our spectral level-set implementation stable, and convergent to at least first order, we have been unable to find a boundary condition that ensures that f remains smooth and thus leads to the desired exponential convergence, not even in the single black hole case. A full finite-difference evolution of f would be less sensitive to the lack of proper boundary conditions (see [13]), but would be much slower for finding an EH in spectral-code metric data (due to interpolations from the spectral to the finite-difference grid) and much less accurate.

The second fundamental difficulty lies in singular behaviour of the function f in certain cases. Let us consider an equal-mass head-on merger as depicted in Figure 3.1. Assume f to be smooth, and let us focus on the value of f at the point of symmetry, marked with C in Figure 3.1. We assume that $\partial_i f$ is outward-pointing near the event horizon. At late times, after the merger, f will be negative at C, because C is inside the event horizon. Throughout the whole simulation, $\partial_i f = 0$ at C by symmetry, and therefore, (3.15) implies that $\partial_t f = 0$ there, so that f at C remains fixed at a finite negative value. At merger, however, the $f = 0$ contour passes through C. Therefore, f must be singular³. Any method for solving the level-set equations that assumes a smooth and regular solution (including finite-difference methods that do not explicitly treat the singularity) will therefore produce results that differ from the exact solution at the singular point. In [13], one-sided finite-difference stencils are carefully chosen so as to not differentiate across the singularity.

Because of these two issues we have stopped development of a spectral implementation of the level-set method. These problems arise because of properties of the function f , which is merely a *tool* to represent the actual surface of interest, $f = 0$. This surface itself is well-behaved and smooth, suggesting it will be possible to evolve this surface directly. Geodesic and surface methods do precisely this, and so we focus on these two methods in the remainder of this paper.

3.3 Proof that points in the Surface Method follow geodesics

Consider a 2-dimensional family of null geodesics, $q^\mu(t, u, v)$, where u, v label different geodesics. Assume the parameter t along the geodesic coincides with the coordinate time of the underlying black hole simulation, i.e. $q^0(t, u, v) = t$. This family of geodesics traces out a three-dimensional null surface \mathcal{N} , parameterized by coordinates t, u, v : $q^\mu(t, u, v)$, where t is the parameter along each null curve, and u, v are the parameters relating each null curve to nearby null curves. In this parameterization, we can write the outgoing null normal $\ell^\mu = \partial q^\mu / \partial t|_{u,v}$, i.e. a coordinate derivative $\vec{\ell} = \partial_t$ within the (t, u, v) coordinates of \mathcal{N} . Displacement vectors that relate each null curve to its neighbours are given by $\vec{m} = \partial / \partial u$, $\vec{n} = \partial / \partial v$. Since coordinate derivatives commute, we have

$$\ell^\mu \nabla_\mu m^\nu = m^\mu \nabla_\mu \ell^\nu. \quad (3.17)$$

Let us consider the rate of change of the inner product $\ell^\mu m_\mu$ as we change the time t along a geodesic (i.e. for fixed u and v):

$$\partial_t(\ell^\mu m_\mu) = \ell^\nu \nabla_\nu(\ell^\mu m_\mu) = m_\mu \ell^\nu \nabla_\nu(\ell^\mu) + \ell_\mu \ell^\nu \nabla_\nu(m^\mu). \quad (3.18)$$

³Even with re-initializations of f , as performed in [13], the same argument applies to that time interval between re-initializations during which the topology of the EH changes.

From (3.17), the second term of (3.18) vanishes,

$$\ell_\mu \ell^\nu \nabla_\nu (m^\mu) = \ell_\mu m^\nu \nabla_\nu (\ell^\mu) = \frac{1}{2} m^\nu \nabla_\nu (\ell^\mu \ell_\mu) = 0. \quad (3.19)$$

Substituting the formula for parallel transport of ℓ^μ along the geodesics, $\ell^\nu \nabla_\nu \ell^\mu = \kappa \ell^\mu$ (with $\kappa = 0$ if t is affine), (3.18) finally becomes

$$\partial_t (\ell^\mu m_\mu) = m_\mu \ell^\nu \nabla_\nu (\ell^\mu) = \kappa m_\mu \ell^\mu. \quad (3.20)$$

A similar calculation results in $\partial_t (\ell^\mu n_\mu) = \kappa \ell^\mu n_\mu$.

So far, this appendix only discusses the geodesic method. We now use the results just obtained to show that surface and geodesic methods will construct the same null surface \mathcal{N} . Both methods start with the same two-dimensional surface at some late time t_0 , and the tangent $\dot{q}^\mu(t_0, u, v)$ to the geodesics at t_0 is chosen to be normal to the 2-surface. Therefore, at t_0 , $\ell^\mu = \dot{q}^\mu$, and the surfaces resulting from evolving both the geodesic and surface methods will coincide at times infinitesimally near t_0 . Because $\ell^\mu m_\mu = \ell^\mu n_\mu = 0$ initially, (3.20) implies that $\ell^\mu m_\mu = \ell^\mu n_\mu = 0$ at all other times. Thus, the tangent to the geodesics always remains orthogonal to the surface described by the positions of all the geodesics at a given time t . Since \dot{q}^μ is normal to that surface, null, outgoing, and has $\dot{q}^0 = 1$, it is identical at all times to ℓ^μ as constructed by the surface method. Therefore, we see that the surfaces obtained by the geodesic and surface methods agree, and both techniques trace out the same \mathcal{N} given the same initial conditions.

3.4 An Introduction to SpEC

The Spectral Einstein Code (SpEC for short) is a software project collaboration between the California Institute of Technology (Caltech), Cornell University, and the Canadian Institute for Theoretical Astrophysics (CITA). SpEC has been under development since the early 2000's, primarily by lead designers Larry Kidder (Cornell), Mark Scheel (Caltech) and Harald Pfeiffer (CITA). SpEC is very different from the majority of 3D merger-capable codes under development, in that it approaches the problem of solving Einstein's equations through the use of spectral methods for computing derivatives, rather than the usual finite-differencing methods (see [14, 15] for recent results from SpEC, and [16, 17] for recent work from other numerical collaborations).

Spectral techniques compute derivatives of a function on a domain by decomposing the function into a set of analytic basis functions. The choice of which type of basis functions to use differs based on requirements such as the topology of the domain. For example, a simple periodic domain is best modeled using a Fourier decomposition of sines and cosines, while a line segment is best treated with Chebyshev polynomials. For an S2 surface, spherical harmonics are the most effective choice

of basis functions. Once the function is decomposed into a set of basis functions, the derivative of the function can be expressed in the form of derivatives of the basis functions, which are analytic. Spectral methods have the advantage that, for C_∞ smooth solutions, the error in the spectral representation decreases exponentially with added basis functions, while errors in finite differencing decrease only as a power of the number of grid-points. Since added basis functions in the spectral method represent an equivalent increase in computation time to added grid-points in the finite difference method, this represents a significant improvement in computation time for equal result.

SpEC uses a multi-domain approach to simulate the spacetime. The spatial slice is divided into a large number of adjacent subdomains of various sizes and topologies, which together make up the overall computational domain. The multiple domains allow users to increase resolution in specific regions of space where it is most needed, without resorting to adaptive mesh refinement, which is more difficult for spectral codes as it is for finite-difference codes. Derivatives of quantities are calculated on individual subdomains using a spectral basis function decomposition of the region of space within the subdomain [18]. Time-stepping is primarily performed using the Dormand-Prince 5th order algorithm, which is a dense timestepper, in the sense that one can compute the values at all intermediate times with almost the same level of accuracy as the end-points of the time-step [19].

There are a number of different versions of Einstein's equations that can be used for numerical simulations. SpEC uses the generalized harmonic formulation⁴, in which the coordinates obey wave equations with a source function that can be freely chosen by the user. In the generalized harmonic formulation, as with all other formulations, the Einstein equations are divided into a set of hyperbolic evolution equations, which define how quantities change in time, and a set of elliptic constraint equations, which must be satisfied separately on every time-slice. One can observe the behavior of the constraints, which are quantities that denote the extent to which the constraint equations are violated. SpEC includes special boundary conditions designed to preserve constraint values at the boundary, without allowing constraint violations to enter from outside the domain [22, 23]. There are also additional terms added to the generalized harmonic equations which serve the function of exponentially damping out any constraint violations that appear [24].

One other major feature of the SpEC code is the use of excision. A characteristic analysis of Einstein's equations in the generalized harmonic form shows that within the event horizon of the black hole, all characteristic fields are falling into the black hole. As a result, if one were to excise a region of the spacetime wholly within the black hole, no boundary conditions would be necessary there, since all information flows out of the computational domain⁵. This requires ensuring that as the black hole changes shape, the excision surface remains wholly within the horizon. Now, when

⁴developed by Helmut Friedrich [20], and used in the first successful complete binary black hole merger simulation (involving inspiral, merger, and ringdown) by Frans Pretorius [21].

⁵Note that this property is not true for all gauges, such as the BSSN gauge, but rather holds true only in some cases, such as with the generalized harmonic gauge.

evolving mergers of binary black holes, some mechanism must be employed to stop the excision region moving out of the black holes. Unlike finite-differencing systems with adaptive mesh refinement, spectral methods are global methods and rely on fixed regions of space being covered by basis functions. Therefore, it is not possible to turn grid-points on and off to move the excision region with the black hole. Moving the excision region would require a complete re-grid of the spacetime, a process which is extremely expensive. Instead, SpEC employs dual coordinate frames. As the binary black holes orbit and merge, SpEC maintains information about two separate coordinate systems. In the inertial coordinate system, the black holes orbit and merge, but the evolution itself is performed in a co-moving frame, which tracks the black holes. A complex proportional-integral control system modifies the mapping between the inertial and co-moving frame, ensuring that the black holes remain centered on the excision regions, which do not move in the co-moving frame. This control system is vital for performing mergers, and is constantly under improvement.

3.5 Numerical Implementation

Compared to the implementation of the geodesic method, implementing the surface method is somewhat more complex due to the presence of derivatives along the surface in (3.14a). Apart from this, the geodesic method and surface method share rather uniform implementation details. We shall first discuss those aspects that only apply to the surface method, and then follow with aspects applicable to both methods.

We represent the surface $r^i(t, u, v)$ with spectral methods (e.g. [25]). These methods approximate a desired function $U(\mathbf{x}, t)$ as a truncated expansion in basis functions ϕ_k , for instance Chebyshev polynomials or spherical harmonics:

$$U(\mathbf{x}, t) = \sum_{k=0}^{N-1} \tilde{U}_k(t) \phi_k(\mathbf{x}), \quad (3.21)$$

where N is the order of the expansion. The fundamental advantage of spectral methods lies in their fast convergence: For smooth problems and a suitable choice of basis functions, the error of the approximation (3.21) decreases *exponentially* with the number of basis functions per dimension [25]. Derivatives of the function U are computed via the (analytically known) derivatives of the basis functions. Each set of basis functions has an associated set of collocation points \mathbf{x}_i ; a matrix multiplication translates between function values at the collocation points, $U(\mathbf{x}_i)$, and spectral coefficients \tilde{U}_k .

For the surface method, we represent each Cartesian component of $r^i(t, u, v)$ (cf. (3.8)) as an

expansion in scalar spherical harmonics,

$$r^i(t, u, v) = \sum_{\ell=0}^L \sum_{m=-\ell}^{m=+\ell} \tilde{A}_{\ell m}^i(t) Y_{\ell m}(u, v). \quad (3.22)$$

This expansion assumes that at fixed t , the surface has topology S_2 . Note that (3.22) allows the surface to intersect itself, as necessary in a binary merger for $t < t_{\text{CEH}}$ (cf. Figure 3.1). Self-intersection is possible because the coordinates u and v are *not* assumed to be standard spherical angular coordinates, i.e. relations like $\cos(u) = z/\sqrt{x^2 + y^2 + z^2}$ will in general not hold.

For spherical harmonics $Y_{\ell m}(u, v)$, the collocation points form a rectangular grid in (u, v) , with the u values chosen so that $\cos(u)$ are the roots of the Legendre polynomial of order $L + 1$, and with the v values being uniformly distributed in the interval $[0, 2\pi]$. There are in total

$$N = 2(L + 1)^2 \quad (3.23)$$

collocation points. The evolution equations (3.13) require derivatives $\partial_u r^i$ and $\partial_v r^i$, which are computed by transformation to spectral coefficients, application of recurrence relations and inverse transform (using the SpherePack library [26]). These derivatives are then substituted into (3.13)–(3.14c) to compute $\partial_t r^i$, which is evolved at the collocation points.

We represent each Cartesian component r^i as an expansion in scalar spherical harmonics (see (3.22)) in order to re-use the infrastructure already developed for our spectral evolution code, which represents tensors of arbitrary rank in this manner to simplify our spectral expansions and to simplify communication of tensor quantities across subdomains of different shapes (see, e.g., [22, 27]). An alternative approach would be to represent r^i in terms of *vector* spherical harmonics, i.e.,

$$r^i(t, u, v) = \sum_{\ell=0}^L \sum_{m=-\ell}^{m=+\ell} \tilde{A}_{\ell m}^i(t) Y_{\ell m}^i(u, v). \quad (3.24)$$

The downside of choosing a scalar spherical harmonic representation is that the equation we impose on the highest order *vector* spherical harmonics is incorrect, and this leads to an instability. This difficulty with expanding vector quantities in a scalar spherical harmonic basis is cured [22] by performing the following “filtering” operation at each timestep: first transform r^i to a vector spherical harmonic basis, then remove the $\ell = L$ and $\ell = L - 1$ coefficients, and then transform back. The removal of both the highest and second highest tensor harmonic modes is necessary, since transforming an n -th rank tensor from a tensor spherical harmonics to scalar spherical harmonics requires scalar harmonics of up to $L^{\text{scalar}} = L^{\text{tensor}} + n$. We filter two modes because we wish to correctly represent the spatial derivatives of r^i (see (3.14a)), which are effectively rank 2.

The geodesic method simply evolves the ODEs (3.7a)–(3.7b). While each geodesic is evolved inde-

pendently, we find it nevertheless convenient to represent them as a two-dimensional grid, $q^i(t, u, v)$ where parameters u and v label each geodesic. We use the same parameters u and v for geodesic and surface method, and for this paper, we choose to locate the geodesics at the *same* (u, v) values as the collocation points of the surface method. We note that this choice is based on convenience to simplify comparison between the two methods; geodesics can be placed at any location, and indeed, we plan as a future upgrade of the geodesic method an adaptive placement of geodesics to help resolve interesting features like caustics.

Let us now discuss aspects common to the implementation of the geodesic and surface methods: At some late time $t = t_{\text{end}}$ long after merger, we initialize the EH surface by choosing it to be the AH at that time. Our AH finder parameterizes the radius of the AH as a function of standard azimuthal and longitudinal angles on S_2 , $r_{\text{AH}}(t_{\text{end}}, \theta, \phi)$, i.e.

$$r_{\text{AH}}^i(t_{\text{end}}, \theta, \phi) = r_{\text{AH}}(t_{\text{end}}, \theta, \phi) \begin{pmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{pmatrix}. \quad (3.25)$$

When initializing the event horizon surface, we choose (u, v) to coincide with the standard spherical angular coordinates (θ, ϕ) , i.e. we set

$$r^i(t_{\text{end}}, u, v) = r_{\text{AH}}^i(t_{\text{end}}, u, v), \quad \text{surface method}, \quad (3.26)$$

$$q^i(t_{\text{end}}, u, v) = r_{\text{AH}}^i(t_{\text{end}}, u, v), \quad \text{geodesic method}. \quad (3.27)$$

For the geodesic method we further set $p^i(t_{\text{end}}, u, v) = s_{\text{AH}}^i$, where s_{AH}^i is the unit normal to the apparent horizon, which is computed similarly to (3.14a)-(3.14c). Time stepping is conducted using a 4th order Runge-Kutta algorithm.

Both methods require interpolation of certain quantities like the spatial metric γ_{ij} onto the grid points of the surface $r^i(t, u, v)$. For the spectral evolutions of the Caltech-Cornell group [28, 29, 27] the evolution data is represented as spectral expansions in space (for each fixed time t) and spatial interpolation is performed by evaluating the appropriate spectral expansions (3.21) at the desired spatial coordinates r^i . Evolution data is available at discrete evolution times t_n and temporal interpolation is performed with 6-th order Lagrange interpolation (i.e. utilizing 3 time slices on either side of the required time).⁶

Finally, we define an area element \sqrt{h} on the surface as the root of the determinant of the induced

⁶The spectral spatial interpolation is computationally more expensive than temporal Lagrangian interpolation. Whenever the domain decomposition for the Einstein evolution is identical for all timesteps involved in a temporal interpolation, the time interpolation is performed before the spatial interpolation. In that case, only *one* spectral spatial interpolation is necessary (on the time-interpolated data), rather than six, speeding up the computation.

metric,

$$h = \frac{1}{\sin^2 u} \det \begin{pmatrix} \gamma_{ij} \partial_u r^i \partial_u r^j & \gamma_{ij} \partial_u r^i \partial_v r^j \\ \gamma_{ij} \partial_v r^i \partial_u r^j & \gamma_{ij} \partial_v r^i \partial_v r^j \end{pmatrix}. \quad (3.28)$$

The area of the evolved surface is then given by

$$A(t) = \int dA = \int \sqrt{h(t, u, v)} \sin u \, du \, dv. \quad (3.29)$$

Explicitly pulling out the factor $\sin u$ in (3.28) and (3.29) ensures that \sqrt{h} is a constant for a coordinate sphere in Euclidean space; this will simplify Figure 5.5 below. Since all the geodesics (or surface grid points) are on a Legendre-Gauss grid, we compute the derivatives in (3.28) spectrally, and we evaluate (3.29) by Legendre-Gauss quadrature. For binary black hole mergers before merger, we sometimes evaluate h based on finite-difference derivatives $\partial_u r^i$ and $\partial_v r^i$. This is discussed in detail in Section 5.1.4.

Bibliography

- [1] J. Libson, J. Massó, , E. Seidel, W.-M. Suen, and P. Walker. Event horizons in numerical relativity: Methods and tests. *Phys. Rev.*, D53:4335–4350, 1996.
- [2] Scott A. Caveny and Richard A. Matzner. Adaptive event horizon tracking and critical phenomena in binary black hole coalescence. *Phys. Rev. D*, D68:104003, 2003.
- [3] J Massó, E Seidel, W-M Suen, and P Walker. Event horizons in numerical relativity. II. analyzing the horizon. *Phys. Rev. D*, 59:064015, 1999.
- [4] Stuart L. Shapiro and Saul A. Teukolsky. Gravitational collapse to neutron stars and black holes: Computer generation of spherical spacetimes. *Astrophys. J.*, 235:199–215, 1980.
- [5] Scott A. Hughes, Charles R. Keeton, II, Paul Walker, Kevin T. Walsh, Stuart L. Shapiro, and Saul A. Teukolsky. Finding black holes in numerical spacetimes. *Phys. Rev. D*, 49(8):4004–4015, 1994.
- [6] Peter Anninos, David Bernstein, Steven Brandt, Joseph Libson, Joan Massó, Edward Seidel, Larry Smarr, Wai-Mo Suen, and Paul Walker. Dynamics of apparent and event horizons. *Phys. Rev. Lett.*, 74(5):630, 1995.
- [7] William H. Press, Saul A. Teukolsky, William T. Wetterling, and Brian P. Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, England, second edition, 1992.

- [8] John C. Butcher. *Numerical methods for ordinary differential equations*. J. Wiley, New Jersey, 2nd edition from wiley, 1st edition from hoboken edition, 2003.
- [9] Jonathan Thornburg. Event and apparent horizon finders for 3+1 numerical relativity. *Living Rev. Rel.*, 10, June 2007. 3.
- [10] Paul Walker. *Horizons, Hyperbolic Systems, and Inner Boundary Conditions in Numerical Relativity*. PhD thesis, University of Illinois at Urbana-Champaign, 1998.
- [11] Scott A. Caveny. *Tracking Black Holes in Numerical Relativity: Foundations and Applications*. PhD thesis, University of Texas at Austin, 2002.
- [12] Scott A. Caveny, Matthew Anderson, and Richard A. Matzner. Tracking black holes in numerical relativity. *Phys. Rev. D*, D68:104009, 2003.
- [13] Peter Diener. A new general purpose event horizon finder for 3D numerical spacetimes. *Class. Quantum Grav.*, 20:4901–4917, 2003.
- [14] Bela Szilagyi, Lee Lindblom, and Mark A. Scheel. Simulations of Binary Black Hole Mergers Using Spectral Methods. *Phys. Rev. D*, 80:124010, 2009.
- [15] Tony Chu, Harald P. Pfeiffer, and Mark A. Scheel. High accuracy simulations of black hole binaries: spins anti-aligned with the orbital angular momentum. *Phys. Rev. D*, 2009.
- [16] Manuela Campanelli, Carlos Lousto, Bruno Mundim, Hiroyuki Nakano, Yosef Zlochower, and Hans-Peter Bischof. Advances in simulations of generic black-hole binaries. *Phys. Rev. D*, 27:084034, 2010.
- [17] U. Sperhake, V. Cardoso, F. Pretorius, E. Berti, T. Hinderer, and N. Yunes. Ultra-relativistic grazing collisions of black holes. arXiv:1003.0882 (gr-qc).
- [18] L. E. Kidder, M. A. Scheel, S. A. Teukolsky, Eric D. Carlson, and G. B. Cook. Black hole evolution by spectral methods. *Phys. Rev. D*, 62(8):084032, 2000.
- [19] J. R. Dormand and P. J. Prince. A family of embedded runge-kutta formulae. *J. Comp. Applied Math.*, 6:19–26, 1980.
- [20] H. Friedrich. Hyperbolic reductions for Einstein’s equations. *Class. Quantum Grav.*, 13(6):1451–1469, 1996.
- [21] Frans Pretorius. Numerical relativity using a generalized harmonic decomposition. *Class. Quantum Grav.*, 22(2):425–451, 2005.

- [22] Michael Holst, Lee Lindblom, Robert Owen, Harald P. Pfeiffer, Mark A. Scheel, and Lawrence E. Kidder. Optimal constraint projection for hyperbolic evolution systems. *Phys. Rev. D*, 70:084017, 2004.
- [23] Lee Lindblom, Mark A. Scheel, Lawrence E. Kidder, Harald P. Pfeiffer, Deirdre Shoemaker, and Saul A. Teukolsky. Controlling the growth of constraints in hyperbolic evolution systems. *Phys. Rev. D*, 69:124025, 2004.
- [24] Carsten Gundlach, Jose M. Martin-Garcia, Gioel Calabrese, and Ian Hinder. Constraint damping in the Z4 formulation and harmonic gauge. *Class. Quantum Grav.*, 22:3767–3774, 2005.
- [25] J. P. Boyd. *Chebyshev and Fourier Spectral Methods*. Springer-Verlag, Berlin, 1989.
- [26] J. C. Adams and P. N. Swarztrauber. Spherepack 3.0.
<http://www.scd.ucar.edu/css/software/spherepack>.
- [27] Mark A. Scheel, Harald P. Pfeiffer, Lee Lindblom, Lawrence E. Kidder, Oliver Rinne, and Saul A. Teukolsky. Solving Einstein’s equations with dual coordinate frames. *Phys. Rev. D*, 74:104006, 2006.
- [28] Harald P. Pfeiffer, Duncan A. Brown, Lawrence E. Kidder, Lee Lindblom, Geoffrey Lovelace, and Mark A. Scheel. Reducing orbital eccentricity in binary black hole simulations. *Class. Quantum Grav.*, 24(12):S59–S81, 2007.
- [29] Michael Boyle, Duncan A. Brown, Lawrence E. Kidder, Abdul H. Mroué, Harald P. Pfeiffer, Mark A. Scheel, Gregory B. Cook, and Saul A. Teukolsky. High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions. *Phys. Rev. D*, 76:124038, 2007.

Chapter 4 Application to Kerr Spacetime

Initial tests of the event horizon finder were conducted using the Kerr spacetime in Kerr-Schild coordinates (See §33.6 of [1]):

$$g_{\mu\nu} \equiv \eta_{\mu\nu} + 2Hl_\mu l_\nu. \quad (4.1)$$

Here H is a scalar function of the coordinates, $\eta_{\mu\nu}$ is the Minkowski metric, and l^μ is a null vector. In Cartesian coordinates (t, x, y, z) , the functions H and l^μ for a black hole of mass M and dimensionless spin parameter a/M in the z direction are

$$H = \frac{Mr_{\text{BL}}^3}{r_{\text{BL}}^4 + a^2 z^2}, \quad (4.2a)$$

$$l_\mu = \left(1, \frac{xr_{\text{BL}} + ay}{r_{\text{BL}}^2 + a^2}, \frac{yr_{\text{BL}} - ax}{r_{\text{BL}}^2 + a^2}, \frac{z}{r_{\text{BL}}} \right), \quad (4.2b)$$

where $r_{\text{BL}}(x, y, z)$ is the Boyer-Lindquist radial coordinate, defined by

$$r_{\text{BL}}^2 = \frac{1}{2} (x^2 + y^2 + z^2 - a^2) + \left(\frac{1}{4} (x^2 + y^2 + z^2 - a^2)^2 + a^2 z^2 \right)^{1/2}. \quad (4.3)$$

If we define the Kerr-Schild spherical coordinates in the straightforward way ($r = \sqrt{x^2 + y^2 + z^2}$, $\cos(\theta) = z/r$, etc), we find that the event horizon of the Kerr black hole in these coordinates is given by

$$r_{\text{Kerr}}(\theta, \phi) = \sqrt{\frac{r_+^4 + r_+^2 a^2}{r_+^2 + a^2 \cos^2 \theta}}, \quad (4.4)$$

where $r_+ \equiv M + \sqrt{M^2 - a^2}$. The surface area of the event horizon is given by

$$A_{\text{Kerr}} = 8\pi M(M + \sqrt{M^2 - a^2}). \quad (4.5)$$

For our tests on the Kerr spacetime we choose the same initial surface for both the surface and geodesic methods: a coordinate sphere of radius $r = 2.5M$, which does not coincide with the horizon. The evolution begins at $t_{\text{end}} = 0$ and proceeds backward in time towards negative t . Because we choose to place geodesics coincident with the collocation points of the surface method (see Section 3.5), we can use the highest angular index L as a measure of resolution. The total number of geodesics or grid points is given by (3.23). The choice of spin in the z direction is for convenience. We have repeated the numerical tests below for spins of several different orientations, and we find no substantial difference in either stability or accuracy.




Figure 4.1: **Geodesic method** applied to a Kerr black hole. The **top panels** show the area difference between the computed and exact solution, normalized by the area of the exact solution. The **bottom panels** show the difference between the computed and exact location of the EH, as measured by (4.6). These data are shown for two series of runs: In the **left panels** we keep the dimensionless spin of the black hole fixed at $a/M = 0.6$ and vary the resolution L of the EH finder. In the **right panels** we vary the spin parameter a/M at fixed resolution. In all cases, the EH finder starts at $t = 0$ and the geodesics are evolved backward in time.

In order to test our methods of finding an EH, we use two measures of error. The first measures the error in the coordinate location of the event horizon. We define

$$\Delta r(u, v) = r(u, v) - r_{\text{Kerr}}(\theta(u, v), \phi(u, v)). \quad (4.6)$$

where $r(u, v)$, $\theta(u, v)$, and $\phi(u, v)$ are the Kerr-Schild radial and angular coordinates of the surface, which are found from either the surface-method variables $r^i(u, v) = [x(u, v), y(u, v), z(u, v)]$ or the geodesic-method variables $q^i(u, v) = [x(u, v), y(u, v), z(u, v)]$ in the usual way, e.g., $x(u, v) = r(u, v) \sin \theta(u, v) \cos \phi(u, v)$. Specifically, we will use the root-mean-square of Δr over all grid points or geodesics, which we shall denote by $||\Delta r||$, as a global measure of the error.

Our second error measure is the deviation of the area of our surface from the Kerr value,

$$\Delta A = A(t) - A_{\text{Kerr}}, \quad (4.7)$$




Figure 4.2: Approach of the tracked null surface onto the event horizon of Kerr black holes with various spins. The symbols show the numerical data (the same data as in the lower right panel of Figure 4.1), and the solid lines are representative least-squares fits. Table 4.1 compares the numerically computed e-folding time to the surface gravity of the black hole.

where $A(t)$ is determined by Equation (3.29).

4.1 Kerr spacetime using the geodesic method

Figure 4.1 shows errors in the AH surface as computed using the geodesic method for a Kerr black hole. The error measure $\|\Delta r\|$, (4.6), does not change with L because the evolution of each geodesic is independent of the total number of geodesics. The error measure $|\Delta A|$, (4.7), does depend on L , but only because the computation of the surface area depends on all geodesics. It is clear from Figure 4.1 that the geodesic method can stably model Kerr black holes of any spin.

At $t_{\text{end}} = 0$, we start the EH finder with an initial surface that does *not* coincide with the EH of Kerr. Therefore, Figure 4.1 shows initial transients as the surface being followed by the EH finder approaches the EH of Kerr. Figure 4.2 shows an enlargement of this phase. We find that the tracked surface approaches the Kerr EH exponentially when integrating backward in time,

$$\|\Delta r\| \propto e^{t/\tau}. \quad (4.8)$$

The time scale τ depends on the spin of the Kerr background. It has been shown in a number of coordinate systems [2, 3, 4] that the e-folding time for a non-spinning black hole is $\tau = 4M$. This is not true in all coordinate systems: for example, in Schwarzschild coordinates $\tau = 2M$. In Section 4.2, we generalize this result to show that null geodesics, perturbed from the Kerr EH, diverge from the

a/M	Mg_H	M/τ	$Mg_H - M/\tau$
0.0	$1/4 = 0.25$	0.249998	$2 \cdot 10^{-6}$
0.2	0.247449	0.247440	$9 \cdot 10^{-6}$
0.4	0.239110	0.239093	$1.7 \cdot 10^{-5}$
0.6	0.222222	0.222212	$1.0 \cdot 10^{-5}$
0.8	$3/16 = 0.1875$	0.187500	$< 10^{-6}$
0.9	0.151784	0.151784	$< 10^{-6}$
0.99	0.061814	0.061814	$< 10^{-6}$

Table 4.1: Exponential approach of the null surface to the correct event horizon location. Mg_H represents the (dimensionless) surface-gravity of a Kerr black hole with spin a/M . M/τ is the numerical rate of approach as determined by fits to the data shown in Figure 4.2.

EH exponentially with an e-folding time equal to $1/g_H$, where

$$g_H = \frac{\sqrt{M^2 - a^2}}{2M(M + \sqrt{M^2 - a^2})} \quad (4.9)$$

is the surface gravity of the horizon in Kerr-Schild coordinates¹. In Table 4.1, we compare the numerically computed e-folding time τ (obtained by least-squares fits) to g_H , and find excellent agreement.

4.2 Proof of Surface Gravity conjecture

We consider a null geodesic $q^\mu(t)$ that asymptotes to a horizon generator $q_H^\mu(t)$ for $t \rightarrow -\infty$, i.e.

$$q^\mu(t) = q_H^\mu(t) + \delta q^\mu(t) \quad (4.10)$$

with $\delta q^\mu(t) \rightarrow 0$ as $t \rightarrow -\infty$. In the discussion of Figure 4.2 we have asserted that

$$\delta q^\mu(t) \propto e^{g_H t}, \quad (4.11)$$

where g_H is the surface gravity of the black hole, and where the coordinates x^μ are Kerr-Schild coordinates, cf. (4.1)–(4.3). To confirm this assertion, one can substitute (4.10) into the geodesic equation and expand to linear order in δq^μ (where we assume that δq^μ , $\delta \dot{q}^\mu$, and $\delta \ddot{q}^\mu$ are of the same order). One then needs to show that the resulting linear equation indeed has the solution (4.11).

The linearization of the geodesic equation is most easily performed in adopted coordinates. We have performed the analysis in “rotating spheroidal Kerr-Schild coordinates” $x^{\mu'} = (t, r_{\text{BL}}, \theta, \phi)$,

¹The surface gravity of a black hole is an analogous quantity to the surface gravity of a Newtonian body, which is the gravitational acceleration felt by an observer at rest on the surface of the body

related to the standard Kerr-Schild coordinates of (4.1)–(4.3) by the coordinate transformation

$$x = \sqrt{r_{\text{BL}}^2 + a^2} \sin \theta \cos(\phi + \Omega_H t), \quad (4.12)$$

$$y = \sqrt{r_{\text{BL}}^2 + a^2} \sin \theta \sin(\phi + \Omega_H t), \quad (4.13)$$

$$z = r_{\text{BL}} \cos \theta. \quad (4.14)$$

The time t is not transformed. Horizon generators have the form $q^{\mu'} = [t, r_+, \theta_0, \phi_0]$, with $r_+ = M + \sqrt{M^2 - a^2}$ and θ_0, ϕ_0 constants, i.e. $\dot{q}_H^{\mu'} \propto [1, 0, 0, 0]$. In these coordinates, we have considered the geodesic equation in affine parameterization, (3.2) and have indeed confirmed

$$\delta q^{\mu'} \propto e^{g_H t} \quad (4.15)$$

to leading order in $\delta q^{\mu'}$. Exponential divergence from a horizon generator—as in (4.15)—is a property present in a quite general class of coordinate systems. For instance, consider the coordinate transformation

$$t' = t + f(x^i), \quad x^{i'} = x^{i'}(x^i), \quad (4.16)$$

where the Jacobian $\partial x^{i'}/\partial x^i$ and its inverse are finite in a neighborhood of the horizon. In this case, $\delta q^{\mu'}$ and δq^μ are related merely by a multiplication by the Jacobian, so the exponential behavior $e^{g_H t}$ is the same in both coordinate systems. The coordinate transformation (4.12)–(4.14) falls into this class, and therefore (4.15) implies (4.11).

4.3 Kerr spacetime using the surface method

We now turn our attention to the surface method. For a Schwarzschild black hole, the surface method with the standard tensor spherical harmonic filtering is stable, as shown by the “F=0” line in the left panel of Figure 4.3. However, the method is unstable for spinning black holes and fails within about $10M$ for spin $a/M = 0.6$ (see the “F=0” line in the right panel of Figure 4.3).

Therefore, we perform additional filtering for spinning black holes. After each timestep, we compute

$$R(u, v) = \sqrt{\delta_{ij} r^i(u, v) r^j(u, v)}, \quad (4.17)$$

expand $R(u, v)$ in scalar spherical harmonics,

$$R(u, v) = \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} \tilde{R}_{\ell m} Y_{\ell m}(u, v), \quad (4.18)$$




Figure 4.3: Effect of filtering using (4.19) and (4.20) for the **surface method**. Shown are evolutions with the same angular resolution $L = 18$, but for different numbers F of truncated modes in (4.19). **Left panel:** For a Schwarzschild black hole, the surface method is stable with or without this additional filtering. **Right panel:** For a Kerr black hole with $a/M = 0.6$, $F = 7$ performs best. The EH finder starts at $t = 0$ and the surface is evolved backward in time.

and truncate the highest F modes of this expansion:

$$\tilde{R}_{\ell m} \rightarrow 0, \quad \text{for } \ell > L - F. \quad (4.19)$$

From these filtered coefficients, we reconstruct the filtered radius-function $R_F(u, v)$ and replace

$$r^i \rightarrow \frac{R_F}{R} r^i. \quad (4.20)$$

The right panel shows that with appropriate choice of F , the horizon of a Kerr black hole with spin $a/M = 0.6$ can be followed for thousands of M . Unfortunately, we do not understand the effect of F on stability, and therefore a parameter search through possible values for F is required.

With this additional filtering in place, we now examine the convergence and accuracy of the surface method. Figure 4.4 shows the convergence behaviour of the surface method. From the top plots, we can see that for a black hole of moderate spin ($a/M = 0.6$), the surface method is accurate and convergent, although long-term stability issues remain. Also, the surface area computed by the surface method appears to be more accurate than the location of the surface, cf. upper vs. lower panels of Figure 4.4. This arises, because for a small change $\delta \tilde{A}_{lm}^i$ in an expansion coefficient \tilde{A}_{lm}^i in (3.22) with $\ell \neq 0$, the change in $\|\Delta r\|$ is linear in $\delta \tilde{A}_{lm}^i$, whereas the change in area is quadratic. The high accuracy of A_{EH} is a welcome feature, since the EH area is one of the most important results of an EH finder. Unfortunately, the surface method is not capable of tracking the horizon




Figure 4.4: **Surface Method**, applied to a Kerr black hole. The **top panels** show the normalized area difference between the computed and exact solution. The **bottom panels** show the difference between the computed and exact location of the EH, as measured by (4.6). These data are shown for two series of runs: In the **left panels** we keep the dimensionless spin of the black hole fixed at $a/M = 0.6$ and vary the resolution L of the EH finder. In the **right panels** we vary the spin a at fixed resolution. The value F denotes the number of truncated modes during filtering according to (4.19). For each case, we show the value of F that provides the most accurate evolution. Also, in all cases, the EH finder starts at $t_{\text{end}} = 0$ and the surface is evolved backward in time. Compare to Figure 4.1.

for spins $a/M \gtrsim 0.8$ for a useful length of time.

While the geodesic method appears superior in these Kerr tests, there are two main benefits to implementing the surface method. Firstly, it is computationally more efficient. The bulk of processing time is spent on interpolating the metric data from the simulation, and the surface method requires the metric only (10 components) whereas the geodesic method requires the metric, as well as its spatial and time derivatives (50 components). Secondly, the surface method can be used to check the errors in the geodesic method in circumstances where the surface method performs well, i.e. lower spins.

For these tests, the initial set of geodesics (or surface) is chosen to be a sphere of radius $2.5M$. In this case it requires a time $\gtrsim 100M$ for either method to converge onto the actual event horizon. This shows that for cases in which the actual EH is unknown, it is important to have a near-stationary

situation at the end of the simulation, so that the initial guess (generally taken to be the AH) has time to converge onto the true EH. The length of this interval will depend on the desired accuracy, the quality of the initial guess and the spin of the black hole. For example, during a time $\Delta t = 10/g_H$ (i.e. $40M$ for $a/M = 0$, but $160M$ for $a/M = 0.99$) the tracked surface will have approached the EH to a fraction $e^{-10} \simeq 5 \cdot 10^{-5}$ of the distance between the initial guess and the EH.

Bibliography

- [1] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. Freeman, New York, New York, 1973.
- [2] J. Libson, J. Massó, , E. Seidel, W.-M. Suen, and P. Walker. Event horizons in numerical relativity: Methods and tests. *Phys. Rev.*, D53:4335–4350, 1996.
- [3] Scott A. Caveny and Richard A. Matzner. Adaptive event horizon tracking and critical phenomena in binary black hole coalescence. *Phys. Rev. D*, D68:104003, 2003.
- [4] Peter Diener. A new general purpose event horizon finder for 3D numerical spacetimes. *Class. Quantum Grav.*, 20:4901–4917, 2003.

Chapter 5 Head-on Mergers

5.1 Head-on Nonspinning Binary Black Hole Merger

5.1.1 Details of BBH evolution

When looking for a straightforward dynamical spacetime where tracking the event horizon is of interest, one of the standard scenarios is the head on merger of two equal-mass non-spinning black holes [1, 2, 3, 4, 5]. First, the SpEC code is utilized to evolve the solution of Einstein’s equations for the head-on merger. Initially the holes are at rest, $r \simeq 4.5M$ apart, where $M = M_A + M_B$ is the total mass at $t = 0$ (because the black holes are non-spinning, we take the irreducible mass as the black hole mass, $M_{A/B} = M_{\text{irr } A/B} = \sqrt{A_{\text{AH } A/B}/(16\pi)}$). Initial data is constructed by solving the conformal thin sandwich equations [6, 7] with the same setup as in [8], but setting the orbital frequency $\Omega_0 = 0$. This data is then evolved with the SpEC code using the dual coordinate frame technique described in [8] and with a domain decomposition with two excision spheres. A common apparent horizon appears at $t = t_{\text{CAH}} = 17.83M$. Shortly thereafter, at $t = t_{\text{regrid}} = 18.96M$, the original domain decomposition with two excision boundaries is replaced by a set of concentric spherical shells with one larger excision boundary. The new excision boundary lies somewhat inside the common apparent horizon, but outside the original excision boundaries. The region very close to the original excision boundaries, and between them, is dropped, and is no longer evolved. Data is interpolated from the highest resolution merger run onto three resolutions of this new domain decomposition. The simulation is continued up to $t = 95M$ and the final mass of the merged black hole is $M_{\text{final}} = 0.9493M$.

The simulation is performed at three progressively higher resolutions, named ‘N0’ through ‘N2.’ The SpEC code does not strictly enforce the Hamiltonian or momentum constraints, nor the artificial constraints that arise from the first-order reduction of the Generalized Harmonic formulation of Einstein’s equations [9]. As such, it is important to monitor the values of these constraints during the simulation, as shown in Fig. 5.1. We normalize the constraints by an appropriate norm of the derivatives of the evolved variables (see (71) of [9] for the precise definition) and integrate constraint violations and normalization only outside the two individual apparent horizons or the common apparent horizon for this run.




Figure 5.1: Evolution of a head-on BBH merger: **normalized constraint violations**. The left panel shows the complete evolution. The right panel enlarges the time around merger, with formation of a common apparent horizon and time of regriding indicated by ‘CAH’ and ‘regrid,’ respectively. The discontinuity at t_{CAH} arises because the constraints are computed only outside the common AH for $t > t_{\text{CAH}}$. At t_{regrid} , the constraints jump because of the different numerical truncation error of the ringdown domain decomposition.

5.1.2 EH finder behaviour

Since the EH finder follows the EH backward in time, we begin our discussion with the ringdown phase of the head-on merger. Initial data for both the geodesic and surface methods is taken from the apparent horizon at $t = 81.24M$, about $60M$ after appearance of a common AH.

We run both the geodesic and surface methods for angular resolutions $L = 7, 15, 23, \dots, 47$ and compute the area $A(t)$ of the tracked surface for these runs. We do not employ filtering as per (4.19) for the surface method.

Figure 5.2 plots the relative differences between $A(t)$ computed with different angular resolution. This plot exhibits several noteworthy features, which we discuss in the next few paragraphs:

During the ringdown phase, $t \gtrsim 20M$, both the surface and geodesic methods perform admirably: Even at low resolution $L = 7$, the area is computed to better than 10^{-6} and this error drops rapidly below 10^{-12} as L is increased. The rapid convergence with L in the ringdown regime is not too surprising, because the angular resolution of the merger simulation is $L_{\text{evolution}} = 25$. Therefore, angular modes $\ell > 25$ of the EH finder carry only information about the way in which the surface parameters (u, v) deviate from the (θ, ϕ) coordinates of the simulation. As can be seen from the excellent convergence for $t \gtrsim 20M$ in Figure 5.2, such deviations are not very important. We also note that the long-term instability exhibited by the surface method during the Kerr test is not apparent.

Close to merger and before merger, $t \lesssim 20M$, the tracked surface becomes very distorted and




Figure 5.2: Effect of changing the resolution of the EH finder when applied to the BBH evolution at fixed high resolution. Shown are relative differences in the area $A(t)$ of the tracked surface. The label “7-15” denotes the difference between simulations with $L_1 = 7$ and $L_2 = 15$, normalized by $A(t)$ of L_2 . Vertical lines on the graph denote the formation of common event and apparent horizons. Note that the time scale of both plots change at $t/M = 25$.

therefore requires much higher angular resolution. This is apparent in the comparatively larger errors in $A(t)$ for $t_{\text{CEH}} < t \lesssim 20M$. In this time interval, the errors in the surface method grow more rapidly than those of the geodesic method. We attribute this to a degradation of the convergence rate of the spectral expansion (3.22). The surface method relies on the spectral expansion in an essential way to compute the derivatives that enter into (3.14a). In contrast, evolution of geodesics is independent of the spectral expansion and the spectral series is used only to compute the surface area via (3.29).

At the point of merger, when the surface being tracked by the EH finders intersects itself for the first time, the error in the area-computation suddenly increases drastically in either method. The reasons for this are quite different for the two methods: The geodesic method evolves individual geodesics perfectly fine through t_{CEH} . The large errors in Figure 5.2 arise because of the use of spectral integration to compute the surface area: At a caustic, the surface-area element \sqrt{h} , (3.28), tends to zero, resulting in a non-smooth integrand in the area integral (3.29), destroying exponential convergence of the spectral area integration. Below, we will explain how we employ finite-difference integration instead. We shall address area calculation for $t < t_{\text{CEH}}$ in Section 5.1.4, where we also discuss how to compute the area of the EH excluding the future generators of the EH.

The surface method exhibits additional, more fundamental, problems at t_{CEH} , when the surface being tracked intersects itself in a caustic with $\sqrt{h} \rightarrow 0$. At such a point, the tangents to the surface, $\partial_u r^i$ and $\partial_v r^i$ are either no longer linearly independent, or one of them is zero, cf. (3.28). Therefore

the surface normal s^i in (3.14a) is ill-defined.

While the surface method presently cannot evolve through merger, it nevertheless yields valuable consistency checks with the geodesic method during the ringdown phase. Figure 5.3 presents such a comparison between the two methods and examines the effect of varying the resolution of the underlying binary black hole simulation. The top panels show differences between the results of the geodesic method applied to evolutions with different resolutions (labelled “G:N#–N#”). As the underlying resolution is increased, the differences become smaller. Likewise, the lines labelled “S:N#–N#” show the analogous differences when running the surface method. When the surface method works, $t \gtrsim 15M$, it is more accurate than the geodesic method. For times close to the formation of the common event horizon, $t \lesssim 15M$, errors in the surface method grow very rapidly and render our current implementation essentially useless. The bottom panels of Fig. 5.3 show differences between surface and geodesic method at the same resolution of the evolved data. This difference decreases with increasing N , as it should. During ringdown, $t \gtrsim 15M$, the difference is essentially equal to the error in the geodesic method; for $t \lesssim 15M$ it is dominated by errors in the surface method.

The right panels in Figure 5.3 examine the surface area $A(t)$. No clear convergence is apparent for $t \gtrsim 20M$, perhaps because the surface area of the event horizon can be calculated with great accuracy even at low values of N . Given the lack of clear convergence, we shall take as our error estimate for the post-merger area the square sum of the following three error measures: a) the change in $A(t)$ between the geodesic method applied to the head-on simulation at the two highest resolutions (i.e. “G:N1–N2”), b) the change in $A(t)$ between the geodesic and surface methods (i.e. “N2:S–G”) and finally, c) the change in $A(t)$ in the geodesic method at $L = 47$, $N2$ when doubling the timestep (from $0.056M$ to $0.112M$; the effect of this is small and not shown in Figure 5.3). This combined error estimate is plotted in Figure 5.4.

5.1.3 Quasinormal Modes during Ringdown

After the merger, the distorted merged black hole rings down into a stationary black hole. During this phase, the area of the event horizon, A_{EH} will approach its final value A_{Final} , and one expects that the apparent horizon approaches the event horizon. This is explored in Figure 5.4. This plot also contains the error estimates obtained from Figures 5.2 and 5.3. Figure 5.4 shows that the areas of the common AH and EH differ by about 10% when the common AH first appears, though this difference drops to 0.1% within about $3M$. After this rapid initial drop, the ringdown is clearly apparent. The area of both EH and AH approach their final area exponentially, and this approach is resolved through about five orders of magnitude. A least-squares fit of $\log[A_f - A_{\text{EH}}(t)]$ to the function $C - \lambda_{\text{obs}}t$ for $30M \lesssim t \lesssim 70M$, yields $\lambda_{\text{obs}} = 0.181M_{\text{final}}^{-1}$. There are furthermore periodic features visible in the EH and AH areas, with seven periods clearly distinguishable. The period of




Figure 5.3: Error estimates for the **surface** and **geodesic** methods, with surface resolution $L = 47$. The left panels show the root-mean-square pointwise deviation between the different runs, whereas the right panels show the differences in the surface area. The lines labelled “G:N#-N#” (“S:N#-N#”) in the upper panels show the difference between the geodesic method (surface method) when applied to merger simulations of different resolution N . The lines labelled “N#:S-G” in the lower panels show the differences between the surface and geodesic methods for a given N (where N_0 , N_1 , and N_2 are resolutions of the merger simulation). Note that the time scale of all plots change at $t = 25M$.




Figure 5.4: Surface area differences between the EH, AH, and the final area, normalized by the final area. Also plotted are error estimates for $A_{EH}(t)$ and $A_{AH}(t)$.

oscillation is found to be $\tau_{\text{osc}} = 8.00M$, therefore $\omega_{\text{obs}} = 0.745M_{\text{final}}^{-1}$.

Decay rate λ_{obs} and frequency ω_{obs} can be related to quasi-normal modes of a Schwarzschild black hole as follows: The quasinormal mode parameters of a perturbed black hole are typically defined with reference to oscillations in the metric fields, which can be written as

$$\delta g_{\mu\nu} \propto e^{-\lambda t} \sin(\omega t), \quad (5.1)$$

where λ is the decay coefficient and ω is the angular frequency of the metric oscillation. Therefore

$$\delta \dot{g}_{\mu\nu} \propto -\lambda e^{-\lambda t} \sin(\omega t) + \omega e^{-\lambda t} \cos(\omega t). \quad (5.2)$$

The energy flux through the horizon, and therefore the change of its mass is $\dot{M} \propto |\delta \dot{g}_{\mu\nu}|^2$, so we have

$$\frac{\dot{A}}{A} \propto \dot{M} \propto \frac{e^{-2\lambda t}}{2} [\lambda^2 + \omega^2 + (\omega^2 - \lambda^2) \cos(2\omega t) - \lambda \omega \sin(2\omega t)]. \quad (5.3)$$

Thus the observed values $(\lambda_{\text{obs}}, \omega_{\text{obs}})$ should be *twice* the values (λ, ω) of a quasi-normal mode. Indeed, the lowest quasinormal mode of a perturbed Schwarzschild black hole is the $\ell = 2, n = 0$ mode, with [10] $\lambda_{20} = 0.08896M_{\text{final}}^{-1}$ and $\omega_{20} = 0.37367M_{\text{final}}^{-1}$. Consistent with (5.3), we find that $\lambda_{\text{obs}} - 2\lambda_{20} = 0.003M_{\text{final}}^{-1}$, and $\omega_{\text{obs}} - 2\omega_{20} = 0.002M_{\text{final}}^{-1}$.

5.1.4 Treatment of Merger

Before examining the merger phase in detail, we must develop tools to analyse the topology change the event horizon undergoes during merger. As seen in Figure 3.1, prior to merger, the surface found by the event horizon finder is the union of the two individual event horizons and the set of future generators of the joint event horizon. The event horizon itself consists of two topological spheres. At merger, $t = t_{\text{CEH}}$, the topology of the event horizon changes to a sphere. For $t < t_{\text{CEH}}$, generators of the event horizon continuously enter the event horizon at the cusps on the event horizons of the two approaching holes. The geodesic method traces geodesics perfectly fine through merger back to the start of the head-on binary black-hole evolution, and the trajectories of the geodesics are convergent as the resolution of the underlying evolution is increased, see the top left panel of Figure 5.3. In this section, we address two questions relevant to analysing the output of the geodesic method: First, when going toward earlier times, some geodesics leave the event horizon; how does one decide whether a given geodesic is still on the event horizon, or whether it is merely a future generator of the event horizon? Second, how can one compute the area of the event horizon (i.e. not counting the area of the locus of future generators)?

Let us first consider the area element \sqrt{h} of the EH surface, with h given by (3.28), which requires derivatives ∂_u, ∂_v along the surface, thus connecting neighbouring geodesics. Because we place the geodesics at a (u, v) grid consistent with spherical harmonic basis functions, we can use spectral differentiation to compute these derivatives (and have done so, up to this point in the paper). Convergence of this spectral expansion, however, becomes increasingly slow for $t \lesssim t_{\text{CEH}}$, and therefore, we compute henceforth the derivatives $\partial_u r^i$ and $\partial_v r^i$ with second order finite difference stencils.

Figure 5.5 plots the area element \sqrt{h} as a function of time for a few representative geodesics. This figure was obtained from our highest resolution run using 20,000 geodesics. To reduce CPU cost, these geodesics were initialized at $t = 19.8M$ from the $L = 47$ run of the surface method. For some geodesics in Figure 5.5, \sqrt{h} approaches zero at a certain time. This feature can be used to determine whether a given geodesic is still on the horizon: We first note that the change of area element along a given null geodesic (i.e. for fixed u, v) is proportional to the expansion of this particular geodesic:

$$\partial_t \log \sqrt{h} = \frac{\partial_t(\sqrt{h})}{\sqrt{h}} \propto \theta. \quad (5.4)$$

The constant of proportionality depends on the parameterization of the null geodesic. Note that by Raychaudhuri's equation, the expansion of a generator of the event horizon must be non-negative, $\theta \geq 0$. Figure 5.5 shows the area element as a function of time for a few representative geodesics.

At late time $t = t_{\text{end}}$ where the final black hole has settled down, we start with geodesics on the apparent horizon, which will be very close to the event horizon. Therefore, we assume that at t_{end} all




Figure 5.5: **Left panel:** Area element \sqrt{h} along a few representative geodesics during the head-on merger. Each geodesic is labelled by the angle α between the initial location of the geodesic (at t_{end}) and the axis of symmetry. Three types of behaviour are apparent: Geodesics entering the horizon from \mathcal{I}^- ($\alpha = 85^\circ$); geodesics entering the horizon from an area in the vicinity of the individual event horizons before merger ($\alpha = 72^\circ$ or 76°), and geodesics remaining on the horizon throughout. The right panels show the time derivative of \sqrt{h} , highlighting the clear signature when a geodesic enters the horizon.

tracked geodesics are generators of the event horizon. Consistently with this assumption, Figure 5.5 shows that $\partial_t \log \sqrt{h}$ starts out very close to zero, and increases as we approach the dynamical time region around merger. If a generator remains on the event horizon, $\partial_t \log \sqrt{h}$ will eventually decrease again and approach zero at very early times before the merger. Generators leaving the event horizon must do so at points where generators cross, according to a theorem by Penrose [11, 12]. For the head-on merger, at such a point nearby geodesics cross and pass through each other. Just after the geodesic enters the horizon, the horizon generators diverge from each other and their expansion is *positive* (and so is $\partial_t \log \sqrt{h}$). Just before the caustic points, nearby future generators of the event horizon converge toward the caustic point with *negative* expansion. In fact, at the caustic, $\partial_t \sqrt{h}$ changes sign discontinuously, as can be seen in Figure 5.5.

Therefore, the largest time at which the expansion of a geodesic passes through zero will be the time it joins the event horizon,

$$\partial_t \log \sqrt{h} \begin{cases} \leq 0, & t = t_{\text{join}}, \\ > 0, & t > t_{\text{join}}. \end{cases} \quad (5.5)$$

In practice, we keep track of (5.5) with a mask function $f_M(u, v)$, which is initially identical to unity. As we evolve backward in time, we evaluate $\partial_t \log \sqrt{h}$ at each time step, and if it drops below some tolerance `-tol` for a point (u_o, v_o) we set $f_M(u_o, v_o) = 0$ for that geodesic. The tolerance `tol` is necessary to avoid misidentifications due to numerical truncation error at very early or late times, where $\partial_t \log \sqrt{h} \rightarrow 0$ for event horizon generators. Because $\partial_t \log \sqrt{h}$ changes so rapidly at a caustic, the precise value for `tol` is not very important; we use `tol = 10-3`.

For generic situations, generators can also leave the EH at points where finitely separated generators cross (a “cross-over point” in the language of Husa & Winicour [13]). At such points, \sqrt{h} remains positive, and criterion (5.5) reduces to a necessary but not sufficient condition that a generator has left the horizon, i.e. t_{join} from (5.5) will be a lower bound for the actual time when a particular geodesic leaves the horizon. Cross-over points can be found by constructing the surface of the event horizon as a set of triangles, and checking every timestep to see if any geodesics have passed through any of the surface triangles. This technique is discussed in more detail in Chapter 6.

The area of the event horizon (consisting of the two disjoint components for $t < t_{\text{CEH}}$) is found by multiplying \sqrt{h} by the mask function f_M , and integrating:

$$A_{\text{EH}} = \int f_M(u, v) \sqrt{h(u, v)} \sin u \, du \, dv. \quad (5.6)$$

For $t < t_{\text{CEH}}$, there are two major sources of error in this integral: First, each geodesic can either be on or off the horizon. When f_M changes discontinuously from 1 to 0 for a geodesic, the area of the event horizon will change discontinuously. Note that this will occur at different times for




Figure 5.6: Convergence of the surface area of the event horizon during merger. The lower plot shows results for placement of the geodesic pole parallel to the axis of symmetry (i.e. consistent with axisymmetry), the upper plot has a geodesic axis perpendicular to the axis of symmetry of the merger. In both cases, geodesics are tracked using the geodesic method; derivatives for \sqrt{h} (cf. (3.28)) are computed with finite-differences; geodesics are removed from the event horizon based on (5.5). Lines are the difference between each resolution and the next highest.

different resolutions. The severity of this effect will depend on how many geodesics enter the horizon simultaneously, as illustrated by Figure 5.6. This figure shows the convergence of the event horizon area with increasing number of geodesics, and for two distinct orientations of the geodesics. In either case, the geodesics are initialized at $t = 19.86M$ from the $L = 47$ surface method determining the event horizon during ringdown, and in either case the geodesics are placed on a rectangular (u, v) grid as detailed in Section 3.5. In the lower panel of Figure 5.6, the geodesics are oriented *respecting* the axisymmetry (i.e. the $u = 0$ polar axis is aligned with the axis of symmetry), whereas in the upper panel the $u = 0$ axis is perpendicular to the axis of symmetry. The lower panel of Figure 5.6, with geodesics respecting the symmetry, shows much larger variations in the area as the resolution is increased. This arises because due to the symmetry, a full *ring* of geodesics leaves simultaneously, thus amplifying the discontinuity of $A_{\text{EH}}(t)$. For perpendicular orientation of the geodesics, individual geodesics leave the horizon, resulting in smaller jumps; this is the configuration we will use in the next section to examine the physics of the black hole merger.

The second source of error in the evaluation of (5.6) arises because the integrand is not smooth once geodesics have left the horizon. For fixed $t < t_{\text{CEH}}$, \sqrt{h} approaches zero linearly toward the caustic; off the horizon, $f_M \sqrt{h} \equiv 0$ by virtue of the mask function, so overall, the integrand is only continuous, and we cannot expect exponential convergence of the integral, despite using a Gauss-quadrature formula to evaluate (5.6).¹

5.1.5 Analysis of Merger Phase

When evolving geodesics backward, we find that the first geodesic leaves the horizon at $t_{\text{CEH}} = 14.58M$, the time of merger. However, it should be noted that the point at which an observer sees the EH change topology is not invariant because the curve traced by the cusps of the two black holes is spacelike [14]. Figure 5.7 shows the surface area of the EH and the common and individual AHs during the merger phase. The common apparent horizon forms at $t_{\text{CAH}} = 17.8M$, and we track the individual apparent horizons up to $t = 18.8M$. The area of the individual apparent horizons is remarkably constant. Up to formation of the common event horizon, its fractional increase is less than 10^{-5} ; up to common apparent horizon, its fractional increase is $5 \cdot 10^{-5}$, and even when we stop tracking the inner horizons, their area has increased by only $1.6 \cdot 10^{-4}$. In contrast, A_{EH} varies significantly more and at significantly earlier times, as can be seen from the inset.

To examine the relation between individual apparent horizons and event horizons, we plot in Figure 5.8 the difference $\Delta A \equiv A_{\text{EH}} - (A_{\text{AH,A}} + A_{\text{AH,B}})$. For times $1 \lesssim t/M \lesssim 12.5$, ΔA grows exponentially with an e-folding time of $2.12M$. This e-folding time is within a few percent of the surface gravity of a black hole with the *initial mass* of the black holes in the head-on simulation.

¹For $t > t_{\text{CEH}}$, A_{EH} in Figure 5.6 is limited by the finite-difference derivatives used to compute \sqrt{h} . Better accuracy can be obtained using spectral derivatives, as can be seen from the right panels of Figure 5.3. For the analysis of the merger below, this difference is invisible.




Figure 5.7: Area of the event horizon and of the apparent horizons before merger and during merger. The vertical dotted lines indicate formation of a common event horizon and appearance of a common apparent horizon; the inset shows an enlargement for early time.




Figure 5.8: Difference between event-horizon area and the sum of the individual apparent horizon areas. The vertical dotted lines indicate formation of common event horizon and appearance of a common apparent horizon. This difference is computed using finite-difference derivatives on the event horizon, and we see a clear improvement in the exponential change in the area at higher order finite-differencing.




Figure 5.9: Spacetime diagram of the head-on merger. The pale lines denote geodesics that will join the event horizon. Some of these geodesics come from past null infinity, but others come from a region close to the individual event horizons (cf. the arrow and the circled geodesics on the far black hole).

This confirms that as geodesics are integrated backwards in time, the individual components of the event horizon approach the individual apparent horizons with the expected rate. If our code were free from all numerical errors, the curve in Figure 5.8 would continue to decrease exponentially as one proceeds backwards in time. Indeed, when we use 2^{nd} order finite-differencing, this curve saturates at $\Delta A/A \approx 0.1\%$ at $t = 0$, and in addition, a feature in ΔA appears at $t \approx 5M$ because the EH area falls below $A_{AH,A} + A_{AH,B}$ and therefore ΔA changes sign. These effects disappear at higher-order finite-differencing. As seen in the 2^{nd} order finite-differencing case, there is a limit to the accuracy that can be obtained in this model. Although that accuracy limit was not reached in the 6^{th} order case, it still exists. In order to achieve better accuracy for the event horizon surface area at very early times when the two holes are widely separated, the EH must be split into two individual surfaces to be evolved separately.

Additionally, perhaps surprisingly, for the head-on binary black hole merger only *some* of the future null generators of the horizon start at past null infinity. A significant fraction of the generators rather start close to the individual event horizons of the black holes before merger. This can be seen in the spacetime diagram in Figure 5.9, most clearly for the geodesic pointed to with an arrow. These




Figure 5.10: Initial configuration of the head-on BBH spinning merger from [20]. The holes move primarily along the x axis, but they also accelerate in the $-y$ (downward) direction due to frame dragging. For the particular simulation on which the event horizon was evaluated, $d = 7.804/M_{\text{ADM}}$. This Figure appears in [20] as Figure 1, and in Appendix A as Figure A.1.

geodesics begin to diverge from the individual event horizon as the second black hole approaches. The increased gravity of both black holes causes such geodesics then to “turn around” and join the event horizon at the seam of the pair of pants.

5.2 Head-on Spinning BBH Merger

5.2.1 Configuration

We now turn to a more complex binary black hole simulation. In recent years, numerical simulations of binary black hole mergers have shown that in some circumstances, the resultant black hole can have a linear velocity [15, 16, 17, 18, 19]. This velocity, known as the black hole recoil, or “kick”, is a result of the fact that in these circumstances, gravitational radiation is emitted anisotropically, with energy emission in some preferred direction, and the black hole gaining momentum in the opposite direction. In order to investigate such scenarios in SpEC, a set of simulations was performed in a configuration intended to generate a straightforward kick. The merger configuration was of a head-on merger in the x -direction of two black holes of spin $0.5M$ and $-0.5M$ respectively in the z -direction (see Figure 5.10). These simulations are detailed in [20]. Work on this simulation was performed by multiple researchers, headed by Geoffrey Lovelace at Cornell University. The paper [20] is in the process of review at Physical Review D at the time of writing, and is reproduced with minor typographical changes as Appendix A.

The goal of this set of simulations is to develop physical insight into the behavior of momentum flow in highly dynamical spacetimes such as the strong-field region near the black-hole horizons in

a merging binary, and to understand how the flow of momentum contributes to the black hole kick. Many numerical simulations of black holes measure the total linear ADM momentum, but in this work, the Landau-Lifshitz momentum flow formalism is applied to the spacetime [21].

In this formalism, a mapping between the curved spacetime and an auxiliary flat spacetime (AFS) is chosen, and general relativity is reinterpreted as a field theory defined on this flat spacetime. The AFS has a set of translational Killing vectors which we use to define a localized, conserved linear momentum. In particular, we calculate i) a momentum density, ii) the momentum enclosed by horizons, and iii) the momentum enclosed by distant coordinate spheres. In the asymptotically flat region around a source, there is a preferred way to choose the mapping between the curved spacetime and the AFS; consequently, in this limit item iii) is gauge-invariant. In general, though, the choice of mapping is arbitrary, and it follows that items i) and ii) are necessarily gauge-dependent.

By examining the linear momentum flow in a dynamical spacetime—and living with the inevitable gauge dependence—we hope to develop strong intuition for the behavior of BBHs. Since some of the results presented are gauge-dependent, it is envisioned that different numerical relativity groups will choose “preferred” gauges based on the coordinates of their numerical simulations. While there is no reason, *a priori*, why simulations in different gauges should agree, one result of [20] is that there is surprisingly good agreement (for the current configuration) between the horizon-enclosed momenta, calculated using spectral and moving-puncture evolutions of similar initial data, even though the simulations use very different gauge conditions for the spectral and puncture simulations, two of the most commonly used gauge conditions in numerical relativity.

5.2.2 Overview of 4-Momentum Conservation in the Landau-Lifshitz Formalism

In this section², we briefly review the Landau-Lifshitz formulation of gravity and the statement of 4-momentum conservation within this theory. Landau and Lifshitz, in their *Classical Theory of Fields* (hereafter referred to as LL), reformulated general relativity as a nonlinear field theory in flat spacetime [22]. (Chap. 20 of MTW [12] and a paper by Babak and Grishchuk [23] are also helpful sources that describe the formalism.) Landau and Lifshitz develop their formalism by first laying down arbitrary asymptotically Lorentz coordinates on a given curved (but asymptotically-flat) spacetime. They use these coordinates to map the curved (i.e. physical) spacetime onto an auxiliary flat spacetime (AFS) by enforcing that the coordinates on the AFS are globally Lorentz. The auxiliary flat metric takes the Minkowski form, $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$.

In this formulation, gravity is described by the physical metric density

$$\mathbf{g}^{\mu\nu} := \sqrt{-g}g^{\mu\nu} , \tag{5.7}$$

²This section taken from Section II of [20], Section A.II of this thesis.

where g is the determinant of the covariant components of the physical metric, and $g^{\mu\nu}$ are the contravariant components of the physical metric. When one defines the superpotential

$$H^{\mu\alpha\nu\beta} := \mathfrak{g}^{\mu\nu} \mathfrak{g}^{\alpha\beta} - \mathfrak{g}^{\mu\alpha} \mathfrak{g}^{\nu\beta} , \quad (5.8)$$

the Einstein field equations take the field-theory-in-flat-spacetime form

$$H^{\mu\alpha\nu\beta}{}_{,\alpha\beta} = 16\pi\tau^{\mu\nu} . \quad (5.9)$$

Here $\tau^{\mu\nu} := (-g)(T^{\mu\nu} + t_{\text{LL}}^{\mu\nu})$ is the total effective stress-energy tensor, indices after the comma denote partial derivatives or, equivalently, covariant derivatives with respect to the flat auxiliary metric), and the Landau-Lifshitz pseudotensor $t_{\text{LL}}^{\mu\nu}$ (a real tensor in the auxiliary flat spacetime) is given by Eq. (100.7) of LL [22] or equivalently Eq. (20.22) of MTW [12]:

$$\begin{aligned} 16\pi(-g)t_{\text{LL}}^{\alpha\beta} &= \mathfrak{g}^{\alpha\beta}{}_{,\lambda} \mathfrak{g}^{\lambda\mu}{}_{,\mu} - \mathfrak{g}^{\alpha\lambda}{}_{,\lambda} \mathfrak{g}^{\beta\mu}{}_{,\mu} \\ &+ \frac{1}{2} g^{\alpha\beta} g_{\lambda\mu} \mathfrak{g}^{\lambda\nu}{}_{,\rho} \mathfrak{g}^{\rho\mu}{}_{,\nu} \\ &- g^{\alpha\lambda} g_{\mu\nu} \mathfrak{g}^{\beta\nu}{}_{,\rho} \mathfrak{g}^{\mu\rho}{}_{,\lambda} - g^{\beta\lambda} g_{\mu\nu} \mathfrak{g}^{\alpha\nu}{}_{,\rho} \mathfrak{g}^{\mu\rho}{}_{,\lambda} \\ &+ g_{\lambda\mu} g^{\nu\rho} \mathfrak{g}^{\alpha\lambda}{}_{,\nu} \mathfrak{g}^{\beta\mu}{}_{,\rho} \\ &+ \frac{1}{8} (2g^{\alpha\lambda} g^{\beta\mu} - g^{\alpha\beta} g^{\lambda\mu}) \\ &\times (2g_{\nu\rho} g_{\sigma\tau} - g_{\rho\sigma} g_{\nu\tau}) \mathfrak{g}^{\nu\tau}{}_{,\lambda} \mathfrak{g}^{\rho\sigma}{}_{,\mu} \end{aligned} \quad (5.10)$$

Due to the symmetries of the superpotential—they are the same as those of the Riemann tensor—the field equations (5.9) imply the differential conservation law for 4-momentum

$$\tau^{\mu\nu}{}_{,\nu} = 0 . \quad (5.11)$$

Eq. (5.11) is equivalent to $T^{\mu\nu}{}_{;\nu} = 0$, where the semicolon denotes a covariant derivative with respect to the physical metric.

In both LL and MTW, it is shown that the total 4-momentum of any isolated system (measured in the asymptotically flat region far from the system) is

$$p_{\text{tot}}^{\mu} = \frac{1}{16\pi} \oint_{\mathcal{S}} H^{\mu\alpha 0j}{}_{,\alpha} d\Sigma_j , \quad (5.12)$$

where $d\Sigma_j$ is the surface-area element of the flat auxiliary metric, and \mathcal{S} is an arbitrarily large surface surrounding the system. This total 4-momentum satisfies the usual conservation law

$$\frac{dp_{\text{tot}}^{\mu}}{dt} = - \oint_{\mathcal{S}} \tau^{\mu j} d\Sigma_j . \quad (5.13)$$




Figure 5.11: The regions of space around and inside a binary-black-hole system. This Figure appears in [20] as Figure 4, and in Appendix A as Figure A.4.

See the end of Section III of [21] for a brief proof of why this holds for black holes.

Because this paper focuses on BBHs, we will make a few further definitions that will be used frequently in our study. First, we label the two³ black holes in the binary (and the regions of space within their horizons) by A and B , and denote their surfaces (sometimes the hole's event horizon and other times the apparent horizon) by ∂A and ∂B , as shown in Fig. 5.11. We let \mathcal{E} stand for the region outside both bodies but inside the arbitrarily large surface S where the system's total momentum is computed (in our case, this is taken to be a fixed coordinate sphere inside the outer boundary of the numerical-relativity computational grid).

With the aid of Gauss's theorem and the Einstein field equations (5.9), one can reexpress Eq. (5.12) for the binary's total 4-momentum as a sum over contributions from each of the bodies and from the gravitational field in the region \mathcal{E} outside them:

$$p_{\text{tot}}^\mu = p_A^\mu + p_B^\mu + p_{\text{field}}^\mu. \quad (5.14a)$$

Here

$$p_A^\mu := \frac{1}{16\pi} \oint_{\partial A} H^{\mu\alpha 0j}{}_{,\alpha} d\Sigma_j \quad (5.14b)$$

³After the holes merge, there is only one horizon, which we label ∂C . Equations (5.13)–(5.15) hold after removing terms with subscript B and then substituting $A \rightarrow C$.

is the 4-momentum of body A (an equivalent expression holds for body B), and

$$p_{\text{field}}^\mu := \int_{\mathcal{E}} \tau^{0\mu} d^3x \quad (5.14c)$$

is the gravitational field's 4-momentum in the exterior of the black holes. We define an effective velocity of black hole A (with similar expressions holding for hole B) by

$$v_{\text{LL}}^j := \frac{p_A^j}{p_A^0}. \quad (5.15)$$

In analogy to Eq. (5.13) for the rate of change of the binary's total 4-momentum, one can write the corresponding equation for the rate of change of the 4-momentum of body A :

$$\frac{dp_A^\mu}{dt} = - \oint_{\partial A} (\tau^{\mu k} - \tau^{\mu 0} v_A^k) d\Sigma_k. \quad (5.16)$$

Equation (5.16) describes the flow of field 4-momentum into and out of body A (the second term comes from the motion of the boundary of body A with local coordinate velocity v_A^k).⁴

In [20], we use Eqs. (5.13)–(5.15) as the basis for our study of momentum flow in black-hole binaries. The actual values of the body and field 4-momenta, computed in the above ways, will depend on the arbitrary mapping between the physical spacetime and the AFS; this is the gauge-dependence that will be discussed in Sec. IV B of [20].

5.2.3 Landau-Lifshitz Velocities of Event and Apparent Horizons

The effective velocities of the two individual apparent horizons (pre-merger), the common apparent horizon (post-merger) are calculated and compared to the effective velocity of the event horizon. Both the holes' co-ordinate velocities and Landau-Lifshitz velocities are computed. The results in Figure 5.12 show that at times before and after merger, the Landau-Lifshitz velocity as computed on the event and apparent horizons agree quite well. It is also clear that computing the LL-velocity on the event horizon is vital in order to fully understand the transition between plunge, and a recoiling final black hole.

We would like to compare our quantitative results of the effective velocity v_{LL}^y calculated using the event horizon surface (Fig. 5.13) with qualitative observations of the event horizon's dynamics (Fig. 5.14). We find that the greatest variation in *both* the event horizon geometry *and* the value of v_{LL}^y occurs over a period of about $\Delta t = 13M_{\text{ADM}}$ from $t = 28M_{\text{ADM}}$ to $t = 41M_{\text{ADM}}$. At time $t = 27.7M_{\text{ADM}}$, the cusps of the event horizon just begin to become noticeable (Figs. 5.14 a & b). One can see in Fig. 5.13 that this is the time at which v_{LL}^y changes from decreasing to increasing.

⁴In the case that the body's event horizon is stationary (i.e. sufficiently far from merger), $v_A^k = dx_A^k_{\text{cm}}/dt$, the center of mass velocity of body A . However, if the body's event horizon is dynamical (i.e. during the merger phase), then v_A^k is the local coordinate velocity of the event horizon surface, $v_A^k = dx_{\partial A}^k/dt$.




Figure 5.12: The velocity of the individual and merged black holes. The Landau-Lifshitz velocity $v_{LL}^y := p_{LL}^y / p_{LL}^t$, where p_{LL}^μ is the Landau-Lifshitz 4-momentum enclosed, is measured on the individual and common apparent horizons (labeled AH and AHC, respectively) and also on the event horizon (labeled EH). For comparison, the coordinate velocities v_{coord}^y of the apparent horizons are also shown. This Figure appears in [20] as Figure 8, and in Appendix A as Figure A.8.




Figure 5.13: The effective velocity v_{LL}^y calculated on the event horizon surface, with the specified snapshots in Fig. 5.14 of the event horizon surface marked: a,b, $t = 27.7M_{\text{ADM}}$; c, $t = 30.8M_{\text{ADM}}$; d, $t = 31.6M_{\text{ADM}}$; e, $t = 35.5M_{\text{ADM}}$; f, $t = 40.8M_{\text{ADM}}$. This Figure appears in [20] as Figure 13, and in Appendix A as Figure A.13.

Shortly after⁵, at $t = 31.1M_{\text{ADM}}$, the two separate event horizons coalesce into a common event horizon, and the common event horizon rapidly expands to form a convex shape by $t = 35.5M_{\text{ADM}}$ (Figs. 5.14 d & e). At this time, we note that v_{LL}^y is rapidly increasing (Fig. 5.13, arrow e); this rapid increase corresponds to the quickly expanding event horizon surface.

We interpret this process as the merging black holes “swallowing” the gravitational field momentum between the holes. The resulting change in v_{LL}^y can be divided into two distinct portions: i) one that results from the changing event horizon surface in space, i.e. the field momentum swallowed by the black holes [mathematically, the *second* term, in Eq. (5.16)] and ii) a second that results from the change of field momentum at the black holes’ surface, i.e. the field momentum *flowing* into the black holes [mathematically, the *first* term, in Eq. (5.16)]. While this distinction is clearly coordinate dependent, it could, after further investigation, nevertheless provide an intriguing and intuitive picture of the near-zone dynamics of merging black hole binaries.

Bibliography

- [1] Peter Anninos, David Bernstein, Steven Brandt, Joseph Libson, Joan Massó, Edward Seidel, Larry Smarr, Wai-Mo Suen, and Paul Walker. Dynamics of apparent and event horizons. *Phys. Rev. Lett.*, 74(5):630, 1995.
- [2] J. Libson, J. Massó, , E. Seidel, W.-M. Suen, and P. Walker. Event horizons in numerical relativity: Methods and tests. *Phys. Rev.*, D53:4335–4350, 1996.
- [3] R. A. Matzner, E. Seidel, S. L. Shapiro, L. L. Smarr, W.-M. Suen, S. A. Teukolsky, and J. Winicour. Geometry of a black hole collision. *Science*, 270(5238):941–947, 1995.
- [4] J Massó, E Seidel, W-M Suen, and P Walker. Event horizons in numerical relativity. II. analyzing the horizon. *Phys. Rev. D*, 59:064015, 1999.
- [5] Scott A. Caveny and Richard A. Matzner. Adaptive event horizon tracking and critical phenomena in binary black hole coalescence. *Phys. Rev. D*, D68:104003, 2003.
- [6] James W. York. Conformal “thin-sandwich” data for the initial-value problem of general relativity. *Phys. Rev. Lett.*, 82(7):1350–1353, Feb 1999.
- [7] Harald P. Pfeiffer and James W. York. Extrinsic curvature and the Einstein constraints. *Phys. Rev. D*, 67(4):044022, Feb 2003.

⁵Note that at $t = 31.1M_{\text{ADM}}$, we (smoothly) modify our gauge condition (see Eq. B11 of [20]) and the surrounding discussion]. The separate event horizons coalesce at time $t = 31.1M_{\text{ADM}}$ as well; this is a coincidence.




Figure 5.14: Snapshots of the event horizons at the times indicated in Fig. 5.13: a,b, $t = 27.7M_{\text{ADM}}$; c, $t = 30.8M_{\text{ADM}}$; d, $t = 31.6M_{\text{ADM}}$; e, $t = 35.5M_{\text{ADM}}$; f, $t = 40.8M_{\text{ADM}}$. All snapshots are looking down the z-axis to the x-y plane, except for shot a, which is slightly skewed (slightly rotated about the y axis) to better see the geodesic structure. In shot a, the future generators of the horizon are visible as small blue dots. Note how the future generators map out a surface that meets the event horizon at the event horizon's cusps; this is where the future generators join the horizon. The data shown are from the high-resolution evolution N2.C. This Figure appears in [20] as Figure 14, and in Appendix A as Figure A.14.

- [8] Mark A. Scheel, Harald P. Pfeiffer, Lee Lindblom, Lawrence E. Kidder, Oliver Rinne, and Saul A. Teukolsky. Solving Einstein's equations with dual coordinate frames. *Phys. Rev. D*, 74:104006, 2006.
- [9] Lee Lindblom, Mark A. Scheel, Lawrence E. Kidder, Robert Owen, and Oliver Rinne. A new generalized harmonic evolution system. *Class. Quantum Grav.*, 23:S447–S462, 2006.
- [10] Kostas D. Kokkotas and Bernd G. Schmidt. Quasi-normal modes of stars and black holes. *Living Rev. Rel.*, 2, August 1999. 2.
- [11] C.M. DeWitt and J.A. Wheeler. *Battelle Rencontres - 1967 Lectures in Mathematics and Physics*. W.A. Benjamin, Inc., New York, New York, first edition, 1968.
- [12] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. Freeman, New York, New York, 1973.
- [13] Sascha Husa and Jeffrey Winicour. The asymmetric merger of black holes. *Phys. Rev. D*, 60(8):084019, Sep 1999.
- [14] S. L. Shapiro, S. A. Teukolsky, and J. Winicour. Toroidal black holes and topological censorship. *Phys. Rev. D*, 52:6982, 1995.
- [15] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, James R. van Meter, and M. Coleman Miller. Getting a kick out of numerical relativity. *Astrophys. J.*, 653:L93–L96, 2006.
- [16] J. A. Gonzalez, M. D. Hannam, U. Sperhake, B. Brügmann, and S. Husa. Supermassive recoil velocities for binary black-hole mergers with antialigned spins. *Phys. Rev. Lett.*, 98:231101, 2007.
- [17] F. Herrmann, Ian Hinder, D. Shoemaker, P. Laguna, and Richard A. Matzner. Gravitational recoil from spinning binary black hole mergers. *Astrophys. J.*, 661:430–436, 2007.
- [18] Michael Koppitz, Denis Pollney, Christian Reisswig, Luciano Rezzolla, Jonathan Thornburg, Peter Diener, and Erik Schnetter. Recoil velocities from equal-mass binary-black-hole mergers. *Phys. Rev. Lett.*, 99:041102, 2007.
- [19] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, and David Merritt. Maximum gravitational recoil. *Phys. Rev. Lett.*, 98:231102, 2007.
- [20] Geoffrey Lovelace, Yanbei Chen, Michael Cohen, Jeffrey D. Kaplan, Drew Keppel, Keith D. Matthews, David A. Nichols, Mark A. Scheel, and Ulrich Sperhake. Momentum flow in black-hole binaries: II. Numerical simulations of equal-mass, head-on mergers with antiparallel spins. arXiv:0907.0869 (gr-qc), 2009.

- [21] Drew Keppel, David A. Nichols, Yanbei Chen, and Kip S. Thorne. Momentum flow in black hole binaries: I. Post-Newtonian analysis of the inspiral and spin-induced bobbing. *Phys. Rev. D*, 80:124015, 2009.
- [22] L. D. Landau and E. M. Lifshitz. *Classical Theory of Fields*. Addison Wesley, Reading, MA, second edition, 1962.
- [23] S. V. Babak and L. P. Grishchuk. Energy-momentum tensor for the gravitational field. *Phys. Rev. D*, 61:024038, 1999.

Chapter 6 Binary Black Hole Inspirals and Mergers

6.1 Introduction to Binary Black Hole Inspirals

Binary black hole (BBH) inspirals and mergers are considered by many the key problem in the field of numerical relativity. The collisions of black holes are among the most energetic events in the universe, and, given their expected prevalence, a prime potential source for current and future gravitational wave detectors. Advanced LIGO expects to see between 2 and 4000 events per year, given an average black hole mass of $10M_{\odot}$ [1], while LISA, a planned space-based detector, expects to see approximately 10 mergers per year for black holes with mass range $10^5 M_{\odot} \lesssim M \lesssim 10^6 M_{\odot}$ [2].

Progress towards successful simulations of the inspiral-merger-ringdown process was greatly advanced by Pretorius' development of the generalized harmonic evolution system [3], which made use of finite-differencing and adaptive mesh refinement. The SpEC collaboration in 2008 succeeded in evolving a full equal-mass non-spinning BBH inspiral through 16 orbits, merger and ringdown [4]. The primary difficulty a spectral code faces for inspirals and mergers is the difficulty of coping with the black holes' movement. In order to track the black holes without needing to regularly regrid the computational domain, SpEC makes use of a dual coordinate system [5]. The evolution is performed in a special coordinate frame, which relates to the inertial frame through a complex set of coordinate maps that control the size, shape and position of the black hole. In this way, the black hole stays relatively stationary in the coordinate frame, while in the inertial frame it is free to move according to the physics of the problem. Additionally, the generalized harmonic equations provide a certain freedom in the way the gauge fields are specified. The keys to successful inspiral-merger simulations are improved distortion mappings, and the use of a different gauge during the merger, which gradually overtakes the gauge used for the inspiral phase.

Subsequently, another inspiral simulation from SpEC became available for event horizon finding. This simulation, published in [6], breaks all the remaining symmetries of the equal-mass non-spinning inspiral, with a 2:1 mass ratio, and low ($a/M \simeq 0.4$) spins in random directions. In this chapter we discuss insights arising from simulating event horizons for these two systems.

6.2 Collision Detection

A key challenge when evolving event horizons is to be able to accurately determine when geodesics merge onto the horizon. In the event horizon finder, the geodesics being tracked are not all on the event horizon for all time. Some of these geodesics have merged onto the black hole during

the time period covered by the simulation. As such, distinguishing these geodesics from others is not automatic. The set of such merger points has been studied extensively [7, 8], where it is argued that this set of points forms a 2-dimensional surface segment of the full 3-dimensional null hypersurface, and that these points can be classified into two types. Caustics points, discussed previously in Section 5.1.4, occur when neighboring rays converge. These points are not generic, but rather serve as a boundary to the set of generic meeting points, termed cross-over points, in which non-neighboring points converge. In the head-on axisymmetric case (Section 5.1), the symmetry condenses the set of merger points in such a way that all merger points are on the boundary of the set of merger points, and are therefore caustic points. Interestingly, we find that in the spinning head-on case (Section 5.2), despite the lack of pure axisymmetry, the set of merger points is also composed entirely of caustic points. In these cases, it was sufficient to perform a collision detection analysis designed to catch only caustic points (see Section 5.1.5). However, this is not the case for mergers that result from inspirals. Therefore, a new technique for detecting crossover points is required.

In order to detect crossover points, we resort to a straightforward point-surface collision detection algorithm. We model the event horizon as a set of triangles. These triangles are easily defined, since the geodesics are initially placed on the collocation points of an S2 surface, which is a rectangular grid, and the property “neighbor-ness” (i.e. knowing which geodesics are to the left/right/above/below any given geodesic) is maintained throughout the simulation. Since represent our geodesic set as an expansion in vector spherical harmonics, the number of geodesics in a surface of resolution L is $2(L + 1)^2$, and the number of triangles in the surface is $4(L + 1)^2$. The polar regions of the S2 parameter-surface not covered by the rectangular grid have triangles defined by placing an artificial pole point using the mean x, y, z coordinates of the top row of geodesics on the S2 rectangular grid. Triangles formed above (below) the rectangular grid each use two neighboring geodesics on the top (bottom) row, and the north (south) pole point of the S2 configuration. The algorithm compares every triangle with every geodesic point, to determine whether the geodesic has passed through that triangle between the current and previous time-step. If the number of geodesics on the horizon in N , the number of triangles is $2N$, and the algorithm scales as $\mathcal{O}(N^2)$. Unfortunately, we have not been able to develop an algorithm where the number of operations scales with a better than $\mathcal{O}(N^2)$ behavior.

Determining whether the point has passed through the triangle proceeds as follows (see Figure 6.1 for a diagram): Suppose that the positions of the three geodesics that comprise the vertices of the triangle at time t_0 are p_0, q_0, r_0 , and the position of the potentially intersecting geodesic is a_0 . At time t_1 these positions are p_1, q_1, r_1 and a_1 . We assume that the geodesics move linearly in space between time t_0 and t_1 . Thus $p(t) = p_0 + t(p_1 - p_0) = p_0 + t\bar{p}$, and similarly for q, r and a . We now




Figure 6.1: Diagram of collision detection between a moving triangle and a moving point.

define the normal of the triangle at time t_0

$$n_0 = (q_0 - p_0) \times (r_0 - p_0), \quad (6.1)$$

where we have assumed that the orientation of the triangle points is anti-clockwise. As a function of time, the normal is

$$\begin{aligned} n(t) &= (q(t) - p(t)) \times (r(t) - p(t)) \\ &= (q_0 - p_0 + t(\bar{q} - \bar{p})) \times (r_0 - p_0 + t(\bar{r} - \bar{p})) \\ &= (q_0 - p_0) \times (r_0 - p_0) + t[(\bar{q} - \bar{p}) \times (r_0 - p_0) + (q_0 - p_0) \times (\bar{r} - \bar{p})] \\ &\quad + t^2(\bar{q} - \bar{p}) \times (\bar{r} - \bar{p}). \end{aligned} \quad (6.2)$$

Since $p_0, q_0, r_0, \bar{p}, \bar{q}, \bar{r}$ are known quantities, we can write Equation 6.2 as

$$n(t) = n_0 + \alpha t + \beta t^2. \quad (6.3)$$

Now, any given plane P has the property that

$$\forall i \in P, \quad i \cdot n_P = D, \quad (6.4)$$

where D is a constant, and n_P is the normal of the plane. Now, $D(t) = p(t) \cdot n(t)$, a cubic equation,

so our geodesic $a(t)$ and the triangle $\{p, q, r\}(t)$ are coplanar at times t which satisfy the equation

$$p(t) \cdot n(t) - a(t) \cdot n(t) = n(t) \cdot (p(t) - a(t)) = 0. \quad (6.5)$$

Equation 6.5 is a cubic with algebraic roots, which can be solved for analytically. For every root found between $t_0 < t \leq t_1$, it is a simple matter to check whether $a(t_{\text{root}})$ is within the triangle $\{p, q, r\}(t_{\text{root}})$, rather than merely being co-planar.

There are of course a few important boundary conditions to be checked, such as ensuring that the geodesic being tested for intersection is not one of the geodesics that make up the triangle, or including special cases for when the cubic equation is degenerate, but the algorithm itself is quite robust and effective. Although the algorithm is, as mentioned above, of order $\mathcal{O}(N^2)$, the expense of the algorithm is mitigated by two factors. Firstly, since the algorithm involves analytically solving an at most cubic equation, the run time of each individual instance is very small, on the order of microseconds. Secondly, the looping condition is sufficiently simple that it can be parallelized over multiple cores without any significant overhead. In practice, with the current maximum resolution of just under 30,000 geodesics, the run-time is not prohibitive. Additionally, since the caustic point-finding algorithm from Section 5.1.4 is run prior to the collision detection algorithm we can determine the caustic vs. crossover structure of the event horizon.

6.3 Topological structure of the Event Horizon for an Inspiral-Merger

As with the head-on merger from Chapter 5, the start point for event horizon finding is late in the simulation, after the final black hole achieved a mostly stationary state. We continue to use the geodesic method (see Section 3.2.1), with resolution $L = 47$ for the ringdown portion. The merger section of the simulation is of most interest to us however, specifically the topological structure of merger.

Husa and Winicour [7] posit that mergers of binary black holes in a non-axisymmetric configuration generically result in an intermediate toroidal state of the event horizon. Having found merger at a single point occurring in not only the axisymmetric head-on merger, but also the head-on spinning merger (where axisymmetry is broken), we were strongly motivated to determine the topological behaviour of the event horizon for inspiral mergers, where axisymmetry is broken in no uncertain terms. In both the equal-mass and generic inspirals, our results show that the event horizons merge at a point, with no intermediate toroidal phase (see Figure 6.2).

In order to understand why no toroidal intermediate stage is found, we need to further understand the topological structure of the event horizon null hypersurface in the case of a binary inspiral and





Figure 6.2: Slices through the event horizon hypersurface at the exact point of merger (to within numerical accuracy). **Upper panel:** Equal mass non-spinning 16-orbit inspiral at $t/M = 3932.779$, the precise point of merger to within $\pm 0.006t/M$. Here M is the sum of the ADM masses. **Lower panel:** Generic 2:1 mass ratio random spin-direction 1.5 orbit inspiral at $t/M = 117.147$, the precise point of merger to within to within $\pm 0.005t/M$. Here M is the sum of the Christodolou masses, since the black holes have spin. At prior times the two black hole horizons are disjoint. At later times, the two horizons have merged into a single S^2 surface. No toroids are evident in the limit of our accuracy.






Figure 6.3: Diagrams of the event horizon null hypersurface in axisymmetric and non-axisymmetric mergers. The merger is along the z -axis. In both panels, the regions $\mathcal{C} \cup \mathcal{X}$ are spacelike. **Left panel:** In the axisymmetric case, the caustic/crossover set is reduced to a single line of caustic points, the “inseam” of the “pair of pants,” labeled \mathcal{C} . The x direction is suppressed but, since the x and y directions are identical for axisymmetry, the identical diagram would obtain if we were to suppress y in favour of x . **Right panel:** In the non-axisymmetric case, such as an inspiral (where we have “unwound” the legs of the “pair of pants”), the set of crossover points \mathcal{X} is evident, bounded on both sides by “inseams” \mathcal{C} . Unlike the axisymmetric case, here the x and y directions are not identical. Since the caustic/crossover set of points is a 2-surface, the diagram we would obtain by suppressing y in favour of x would look identical to the left panel case, except that the single caustic “inseam” would be composed of crossover points.

merger. In [7], Husa and Winicour consider two sets of points. One set, labeled \mathcal{C} , is the set of all caustic points in the spacetime, where neighbouring event horizon geodesics cross. The other set of points, \mathcal{X} , is the set of all crossover points in the spacetime, where non-neighbouring event horizon geodesics cross. They show that the set of points \mathcal{X} is an open 2-surface on the event horizon null hypersurface, and that this set is bounded by the caustic set \mathcal{C} . They further show that the behavior of this 2-surface of caustic/crossover points is governed by the topology of the merger. In an axisymmetric prolate merger (such as our headon case), the 2-surface is reduced by the symmetry, resulting in the single boundary line of caustic points we see as being the “inseam” of the “pair of pants,” as shown in the left panel of Figure 6.3. In the non-axisymmetric case, the set of caustic and crossover points is a 2-surface on the event horizon, as shown in the case of a binary black hole inspiral in the right panel of Figure 6.3 (where we show the merger in a corotating frame).

The question of whether toroidal horizons can be found in the intermediate stages of binary black hole merger can be answered by considering the various ways in which these “pair of pants” diagrams can be sliced. The fact that the set caustic/crossover points $\mathcal{C} \cup \mathcal{X}$ is a spacelike 2-surface on a non-axisymmetric event horizon hypersurface (and, for an axisymmetric case, the line of points \mathcal{C} is a spacelike line) provides some freedom in the allowed spacelike slicings of this surface.

Let us first consider whether a non-trivial topology might be obtained in the axisymmetric case.




Figure 6.4: This figure shows a 2-dimensional slice through the event horizon null hypersurface in an axisymmetric merger. The horizontal direction in the right panel could be either x or y . We attempt to construct a slice \mathcal{S}_1 in x (or y) from point \mathcal{P} that intersects the black hole. This slice is clearly not spacelike. Since \mathcal{N} is spacelike only at \mathcal{C} , only a non-intersecting slice such as \mathcal{S}_0 can be spacelike.

In order to do so, we need to consider how such a slice may be constructed. Clearly, if we were to construct flat slices of the null hypersurface in the left panel of Figure 6.3, we would produce a slicing in which the merger occurred at a point. However, we can attempt to construct slices in which the lapse is somewhat retarded near the “crotch.” In Figure 6.4 we examine a 2-dimensional slice in $\{t, y\}$ through the center of the hypersurface. It is clear that if we choose a central point for the slice before the merger of the black holes, we cannot extend a spacelike slice from this central point in either the x or y directions in such a way as to encounter the black holes. Only in the z direction can we encounter the black holes.

This changes however, when we consider the non-axisymmetric case. In this case, the x and y directions are different, due to the presence of the 2-surface \mathcal{X} . Taking a $\{t, y\}$ 2-slice of the event horizon in Figure 6.5, we now have a hypersurface slice that is spacelike both at \mathcal{C} , and along the line \mathcal{X} . Thus, given a point \mathcal{P} below the “crotch” of the event horizon, we can construct three distinct slices, each with different behaviour. Slice \mathcal{S}_0 does not encounter the event horizon at all. Slice \mathcal{S}_1 encounters the event horizon four times, twice in the null region, and twice in the spacelike region. Finally, slice \mathcal{S}_2 encounters the event horizon four times in the spacelike region. Note that in the x direction, the cut through the event horizon is identical to Figure 6.4. Therefore, if we slice our spacetime using slices \mathcal{S}_1 or \mathcal{S}_2 , our slice encounters the event horizon four times in the z and y directions, and not at all in the x direction. This is precisely a toroidal intermediate stage. Such a slice can be seen in three dimensions in Figure 6.6. Additionally, it is important to note that there is a distinction between the behavior of slices \mathcal{S}_1 and \mathcal{S}_2 in Figure 6.5. When a slice intersects the event horizon at a point that is a member of $\mathcal{C} \cup \mathcal{X}$, that point is the point where two generators of the event horizon pass through each other as they merge onto the event horizon. Consequently, that




Figure 6.5: This figure shows a 2-dimensional slice through the event horizon null hypersurface in a non-axisymmetric merger. Unlike the previous figure, the horizontal direction in the right panel is not interchangeable between x and y . We construct three slices $\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2$ from the starting point \mathcal{P} which all intersect the event horizon in different ways. Since $\mathcal{C} \cup \mathcal{X}$ is spacelike, all these slices are spacelike.

point is not a smooth part of the event horizon, whereas when the slice intersects the event horizon at a point in $\mathcal{N} \setminus (\mathcal{C} \cup \mathcal{X})$ that point is a smooth part of the event horizon. Therefore, \mathcal{S}_1 corresponds to a toroidal intermediate stage where the torus has a non-smooth (i.e. sharp) inner edge, and \mathcal{S}_2 corresponds to a stage where a line segment on the outside of the torus is also sharp-edged.

6.4 Topological Structure of Simulated Event Horizons

Having shown that toroidal event horizons are possible with the appropriate choice of slicing, and having found no intermediate toroidal phase in either the equal-mass non-spinning inspiral or the generic 2:1 mass ratio spinning inspiral, three important questions arise:

1. What is the structure of caustic and crossover points for the simulations we have performed, and how do those results relate to the structure discussed in the previous section?
2. Can a fully spacelike reslicing of an existing simulation of a binary black hole merger be performed in such a way that the event horizon has an intermediate toroidal shape?
3. Can a Generalized Harmonic evolution, from identical initial data but with a different choice of gauge, generate a different slicing of the spacetime such that an intermediate toroidal event horizon phase exists.

According to Figure 6.6, for a non-axisymmetric merger, an early slice through the event horizon before merger should show each black hole with a linear cusp. Unlike the axisymmetric case, where all geodesics merged onto the event horizon at a point, the non-axisymmetric merger should show each black hole with a linear scar on its surface, through which geodesics merge onto the horizon.




Figure 6.6: A 3-dimensional representation of slice \mathcal{S}_1 from Figure 6.5 Here we clearly see the continuation of the slice in the z direction, which creates a toroidal intermediate black hole. The toroidal region of the slice is that part of the slice that has dipped through the crossover region \mathcal{X} , with the center of the torus at \mathcal{P} .




Figure 6.7: A snapshot of the geodesics being followed by the event horizon finder at time $t/M = t_{\text{merger}}/M - 0.067$, for the equal-mass inspiral. The small dots are geodesics currently on the event horizon. The larger points represent geodesics in the process of merging onto the event horizon. Crosses represent points merging through caustic points, while circles represent points merging through crossovers. In this slice, it is clear that the cusp on the black hole is linear, composed of crossover points with caustics at the end-points.

The cusp should be composed of crossover points, except at the boundaries, which are caustic points. In Figure 6.7, we see that the black hole does indeed have a linear cusp on its surface. The caustic points occur at the edges of the cusp.

At the point of merger our slicing remains consistent with a flat slicing through Figure 6.5. In this figure, the crossover region dips slightly below the caustic lines, and so the flat slice at the point of merger is expected to have a rough “X”-shape of crossover points, meeting at the merger point, and anchored at the edges of the black hole cusps by caustic points. In Figure 6.8, we see that this is indeed the case.

After the point of merger between the two event horizons, our slicing continues to be consistent with a flat slicing through Figure 6.5. Such a slice would show a “bridge” between the black holes, with merger points along each side. These merger points would be a line of crossover points between each black hole, anchored at each end by a caustic point. The “X”-shape of the merger has disconnected, resulting in two line segments of merger points. This is clearly visible in Figure 6.9.

It is clear from these results that our simulation is consistent with the topological structure discussed by Husa and Winicour in [7], and outlined in Section 6.3 above. Our slicing represents a flat slicing (with respect to Figures 6.4-6.6) through the structure of the event horizon, but this does not preclude the possibility of other spacelike slicings producing toroidal intermediate stages




Figure 6.8: A snapshot of the geodesics being followed by the event horizon finder at time $t/M = t_{\text{merger}}/M$, the exact point of merger in the equal-mass inspiral simulation. The small dots are geodesics currently on the event horizon. The larger points represent geodesics in the process of merging onto the event horizon. Crosses represent points merging through caustics, while circles represent points merging through crossovers. Although finding the exact point of merger is difficult given limited numerical time-accuracy, we can extrapolate the “X”-shape of the cusps to see that the merger point is clearly a crossover point.




Figure 6.9: A snapshot of the geodesics being followed by the event horizon finder at time $t/M = t_{\text{merger}}/M + 0.039$, shortly after merger, for the equal-mass inspiral. The small dots are geodesics currently on the event horizon. The larger points represent geodesics in the process of merging onto the event horizon. Crosses represent points merging through caustics, while circles represent points merging through crossovers. The “bridge” between the two black holes has two lines of merger points running on either side of it, with the majority being crossover points anchored by caustics at either end.

during merger. Unfortunately, the two remaining questions remain unsolved at this time. There has not been sufficient time within the framework of this thesis to develop software that would reslice existing simulations with new spacelike slicings, in order to confirm that a reslicing of the same simulation could yield toroidal intermediate stages during merger. Additionally, efforts to modify the gauge conditions of the generalized harmonic evolution code in order to produce a slicing such as in Figure 6.6 have not as yet been successful. These questions remain unanswered as of this writing.

Bibliography

- [1] B. Abbott et al. LIGO: The laser interferometer gravitational-wave observatory. *Rep. Prog. Phys.*, 72:076901, 2009.
- [2] S. Hughes. A brief survey of LISA sources and science. *Proceedings of the 6th International LISA Symposium*, pages 13–20, 2007.
- [3] Frans Pretorius. Evolution of binary black-hole spacetimes. *Phys. Rev. Lett.*, 95(12):121101, 2005.
- [4] M. Scheel, M. Boyle, T. Chu, L. Kidder, K. Matthews and H. Pfeiffer. High-accuracy waveforms for binary black hole inspiral, merger, and ringdown. *Phys. Rev. D*, 79:024003, 2009.
- [5] Mark A. Scheel, Harald P. Pfeiffer, Lee Lindblom, Lawrence E. Kidder, Oliver Rinne, and Saul A. Teukolsky. Solving Einstein’s equations with dual coordinate frames. *Phys. Rev. D*, 74:104006, 2006.
- [6] Bela Szilagyi, Lee Lindblom, and Mark A. Scheel. Simulations of Binary Black Hole Mergers Using Spectral Methods. *Phys. Rev. D*, 80:124010, 2009.
- [7] Sascha Husa and Jeffrey Winicour. The asymmetric merger of black holes. *Phys. Rev. D*, 60(8):084019, Sep 1999.
- [8] L. Lehner, N Bishop, R. Gomez, B Szilagyi, and J Winicour. Exact solutions for the intrinsic geometry of black hole coalescence. *Phys. Rev. D*, 60:044005, 1999.

Chapter 7 Conclusion

This thesis examines three different methods for locating event horizons in dynamical black hole spacetimes: the geodesic method, the surface method and the level-set method. All three methods rely on the principle that outgoing null geodesics exponentially approach the event horizon when followed backward in time. We implement both the geodesic and surface methods, the latter implemented without the assumption of axisymmetry as done in earlier work [1]. Overall, we find that the geodesic method is more robust, with the capability to accurately follow highly spinning black holes (tested up to $a/M = 0.99$), as well as the merger of two black holes. For the head-on merger, we find that the surface-area element \sqrt{h} of the geodesic congruence is an excellent diagnostic of whether and when a geodesic joins the event horizon at a caustic, cf. (5.5) of Chapter 5.

Errors due to tangential drift of the geodesics—as explained in [1]—are not apparent in our simulations. The observed good properties of the geodesic method might be related to the improvements in accuracy of the spacetime metric since the early tests [1], as well as the ability to interpolate the metric spectrally to the geodesic locations. Because each geodesic is evolved independently, the geodesic method parallelizes trivially. Tracking of the cusp of the disjoint components of the event horizon before merger, as well as computation of A_{EH} is currently not as highly accurate as we would like, since comparatively few geodesics cover the region close to the cusps. Our current scheme calls for significantly increasing the number of geodesics at some time $t > t_{\text{merger}}$, where the spherical harmonic representation of the event horizon is still very accurate. In the future, we intend to improve accuracy at the cusps by ceasing to rely on spherical harmonic representations of the event horizon, and placing geodesics iteratively or adaptively, based on knowledge of where the cusps are occurring.

The surface method is less robust and exhibits a long-term instability when applied to Kerr black holes with spins $a/M \lesssim 0.6$, and rapid blow-up for larger spins. Nevertheless during the ringdown phase $t > t_{\text{CEH}}$ of the head-on axisymmetric merger, the surface method locates the event horizon with comparable accuracy to the geodesic method and provides an important independent test of the geodesic method. However, when the surface being tracked self-intersects in a caustic point, our current method for defining the normal breaks down because $\partial r^i / \partial v = 0$ in equations (3.14a)-(3.14c) of Chapter 3, and thus our current implementation of the surface method fails.

The level-set method, finally, is not implemented in this paper. It requires boundary conditions for the level-set function f ; furthermore f can become singular during a black hole merger. Both reasons made it unduly difficult to implement this method in our spectral code. In conclusion, we find that the geodesic method, the oldest of the three methods considered, to be the most accurate

and useful in our tests.

In improving our code to handle more generic situations such as non-axisymmetric mergers, where crossover points are expected (see [2]), we have developed a collision detection algorithm to detect mergers of geodesics onto the event horizon. This algorithm, unfortunately, requires a number of operations that increases with the square of the number of geodesics. We find that this algorithm leaves us with a clear picture of the structure of the merger points that make up the cusp of the black hole event horizon. We have further developed theory by Husa & Winicour [2] and demonstrated numerically that our simulations comply with this improved picture of the topological structure of the event horizon. While we have not found any toroidal intermediate stages of the event horizon in any of our simulations, we find that they are clearly possible within the topological structure of a non-axisymmetric merger.

Open questions remain regarding the topological structure of the event horizon. Due to time constraints, we have not been able to develop software that reslices existing simulations with a new spacelike slicing. It would be valuable to know whether such a reslicing could be performed on existing data in such a way as to create a toroidal intermediate phase in the event horizon. We believe that there should be no fundamental difficulties with implementing such a reslicing, which would provide additional evidence that toroidal horizons are purely slicing-dependent in non-axisymmetric binary black hole mergers. Additionally, there is the question of whether the generalized harmonic evolution scheme implemented in SpEC could produce a toroidal event horizon phase from initial data used in current simulations, simply by modifying gauge conditions in such a way as to retard the lapse function near the merger point of the black holes. Attempts so far have not proven to be successful.

Turning our attention to applications of the event horizon finders, Figure 4.2 presents a new quantitative test of event horizon finders: When finding the EH of a Kerr black hole starting away from the true horizon, does the tracked null surface approach the true event horizon with the correct rate, namely the surface gravity g_H ? Table 4.1 confirms this for the geodesic method. For the head-on merger, both geodesic and surface method perform admirably during the ringdown phase, where we are able to clearly observe the quasinormal ringing of the single merged black hole. For both the event and apparent horizons, the frequency and damping time of the ringing matches the ($\ell = 2, n = 0$) mode of the Schwarzschild quasinormal ringing spectrum to within 2% for the decay rate and 0.3% for the frequency.

In Chapter 5, we show further applications of event horizon finding. We are able to compute the Landau-Lifshitz pseudotensor on the event horizon, which serves as a useful complement to computations performed on the apparent horizon. In the head-on maximum kick simulation (see Section 5.2), we see that computing the Landau-Lifshitz velocity on the event horizon provides crucial information about a stage of the merger not adequately covered by the apparent horizon

analysis (see Figures 5.12 and 5.13). Although not discussed in this thesis, we are also able to compute the scalar curvature and the Hájíček one-form (see [3, 4]) on the event horizon.

Additionally, we find that the apparent horizons provide an excellent approximation to the event horizon for the head-on merger very early before the merger, and very late after the ringdown. Thus, while in principle the apparent horizon is slice-dependent and there is no guarantee that it should coincide with the event horizon, in practice no such behaviour is found.

Further development of event horizon finding algorithms would be greatly beneficial both to the SpEC community and the greater numerical relativity community. As an accurate and robust algorithm, evolving geodesics to track the event horizon can provide great insight into the structure of black hole mergers. In the future, it would be very useful to develop improved techniques for geodesic placement, whereby a simulation of the event horizon can show the need for increased resolution in certain regions of space, which can be provided by either adaptively placing geodesics where needed, or iteratively by returning to the start of the simulation and modifying the initial geodesic placements. These improved techniques will aid us in imaging the cusps of black holes in greater detail. Another improvement of note would be development of a way of “splitting” the event horizon pre-merger. At a certain point once the horizon has been followed through merger (backwards in time), it becomes impractical to follow all the geodesics which are not on the event horizon. At this point, splitting the system into two individual horizons would be beneficial, if done in such a way as to preserve the correct dynamical behavior of the system.

In conclusion, event horizons offer an exciting new tool for understanding the structure, dynamics and topology of binary black hole mergers. In some cases, they allow us to observe dynamics that apparent horizons simply are not capable of reproducing. We believe that the event horizon finding algorithms described in this thesis represent the most accurate tracking of event horizons in complex merger simulations that have been developed to date. It is unfortunate that event horizons have been in general neglected as a tool in the arsenal of numerical relativity, but it is our hope that the results presented in this thesis will provide some impetus towards an increase in the attention paid to event horizons by numerical relativists, and an increase in their use for the purpose of better understanding black hole mergers.

Bibliography

- [1] J. Libson, J. Massó, , E. Seidel, W.-M. Suen, and P. Walker. Event horizons in numerical relativity: Methods and tests. *Phys. Rev.*, D53:4335–4350, 1996.
- [2] Sascha Husa and Jeffrey Winicour. The asymmetric merger of black holes. *Phys. Rev. D*, 60(8):084019, Sep 1999.

- [3] T. Damour. Surface effects in black hole physics. *Proc. 2nd Marcel Grossmann Mtng of Gen. Rel.*, pages 587–608.
- [4] E.ourgoulhon and J. L. Jaramillo. A 3+1 perspective on null hypersurfaces and isolated horizons. *Phys. Rep.*, 423:159–294, 2006.

Part II

Searches for Cosmic-String Gravitational-Wave Bursts in Mock LISA Data

Chapter 8 Introduction

8.1 LISA and the Mock LISA Data Challenges

LISA (Laser Interferometer Space Antenna) is a space-based gravitational wave detector mission, currently in the formulation stages as a joint project of the European Space Agency (ESA) and the National Aeronautics and Space Administration (NASA). LISA is designed to complement the sensitivity range of earth-based detectors such as LIGO and VIRGO, which are most sensitive between 10 Hz and 1 kHz (see Figure 8.1 for current estimates about LIGO and LISA sensitivities). LISA is composed of three spacecraft in an equilateral triangle formation, placed in orbit around the sun along the Earth's orbital path. With an arm length of approximately 5 million km, LISA is expected to be most sensitive to signals in the range of 10^{-4} to 1 Hz.

Although the preliminary LISA Pathfinder mission, designed to test vital aspects of LISA technology in space, has been green-lighted, the full LISA science mission has not yet been approved for funding. Despite this, and in order to increase the likelihood that the full mission will be approved, LISA continues to be in active development all around the world, under the auspices of the LISA International Science Team (LIST).

As part of a comprehensive effort to improve our ability to extract science from LISA data, the LIST has undertaken a series of Mock LISA Data Challenges (MLDCs) [1, 2]. The MLDCs consist of data sets, created using software models of expected LISA sources, overlaid with noise realizations obtained from current LISA noise models. These data sets are published openly, and members of the LISA community are invited to develop data analysis software to accurately determine the presence, number, and parameters of signals with undisclosed parameters, embedded within the noise. The aim of the MLDCs is to increase our understanding of LISA, by developing and improving algorithms to extract signals from expected LISA noise, and determining which set of parameter values in the source model most closely corresponds with the extracted signal. Additionally, the MLDC will help us understand what potential sources exist or may exist in the universe that would emit a gravitational wave signature detectable by LISA, which will hopefully lead to improved modeling of the gravitational waves emitted by these sources. For these reasons, the MLDC is an important step towards developing prototypes of data-analysis software for LISA.

The work discussed in this part of the thesis was performed within the context of MLDC 3. MLDC 3 introduced for the first time a new potential source of gravitational waves for LISA - cosmic strings. Cosmic strings are cosmological phenomena, posited to have formed during phase transitions of fields occurring in the early universe. These phase transitions are hypothesized to give




Figure 8.1: Expected LIGO and LISA sensitivity ranges, and the sources that may be found within those ranges. Sourced from the LISA Science Case (www.lisascience.org).

rise to so-called “topological defects” in quantities such as the stress-energy tensor, and can be said to evolve on their own, much as one could treat a crack that forms in melting ice as a standalone phenomenon, possessing its own equations of motion.

8.2 Cosmic Strings

There are several mechanisms by which an observable network of cosmic (super)strings could have formed in the early Universe. Basically, string formation arises from the breaking of some $U(1)$ symmetry (either global or local) as the Universe expands and cools. In the 1980s and 1990s, interest was primarily in cosmic strings arising from grand unified theories [3], but in recent years several string-theory-inspired inflationary models have also been shown to populate the Universe with a network of cosmic-scale strings [4, 5]. For instance, brane-inflation models can naturally lead to the breaking of $U(1)$ symmetries at the end of inflation, leading to the formation of both long fundamental strings and $D(k+1)$ -branes that wrap around k compact dimensions and extend in one of Nature’s three large spatial dimensions. These long strings can be stable on cosmological timescales (depending on the exact model) and could reasonably have string tensions in the range $10^{-12} \lesssim \mu \lesssim 10^{-6}$. We refer the reader to [6] for a nice review of the main physical ideas.

Simulations have shown that string networks rapidly approach an attractor: the distribution of straight strings and loops rapidly becomes independent of its initial conditions. The network properties *do* depend on two basic parameters of the strings, the string tension μ and the string reconnection probability p . The distribution of loop sizes at their birth should in principle be

derivable from μ and p , but the huge range of scales makes this a very difficult problem to solve via simulations, and today the typical loop size at birth (as a fraction of the Hubble scale) is still uncertain by many orders of magnitude. We refer the reader to Allen [7] for a brief, pedagogical introduction to string networks, and to Vilenkin and Shellard [3] for a more comprehensive review.

Once formed, string loops oscillate and therefore lose energy and shrink due to gravitational-wave (GW) emission. The spectrum of this GW background radiation is calculated to be roughly flat over many orders of magnitude in frequency, including the frequency bands where current ground-based GW interferometers (like LIGO and Virgo) and planned space-based GW interferometers (like LISA) are sensitive. It is conventional to express the energy density ρ_{GW} of GWs in terms of

$$\Omega_{GW}(f) \equiv \frac{1}{\rho_c} \frac{d\rho_{GW}}{d \ln f}, \quad (8.1)$$

where ρ_c is the Universe's closure density. The current limit on $\Omega_{GW}(f)$ from pulsar timing is $\Omega_{GW}(f \sim 2.5 \times 10^{-7} \text{ Hz}) \lesssim 4 \times 10^{-8}$ [8], and the limit from first-generation ground-based interferometers is $\Omega_{GW}(f \sim 100 \text{ Hz}) < 6.9 \times 10^{-6}$ [9]. For comparison, the Advanced LIGO detectors should be capable of detecting a stochastic background with $\Omega_{GW}(f \sim 40 \text{ Hz}) \gtrsim 10^{-9}$ [9], while LISA should be capable of detecting a string-generated background $\Omega_{GW}(f \sim 10^{-4} - 10^{-1.5} \text{ Hz}) \gtrsim 10^{-10}$ [10]. (For LISA, this threshold is set *not* by detector noise, but instead by the background from short-period Galactic binaries.)

In addition to this broadband stochastic background, Damour and Vilenkin [11, 12] pointed out that the *kinks* and *cusps* that form on cosmic strings produce short GW bursts that could also be detectable for a large range of string parameters μ and p . Kinks are discontinuities in the string's tangent direction, which arise when strings overlap and interconnect, while cusps are points on the string that become instantaneously accelerated to the speed of light. The portion of string near the cusp beams a burst of linearly polarized GWs in a narrow cone around the cusp's direction of motion. Damour and Vilenkin showed that, for current and planned GW interferometers, cusp bursts should be significantly more detectable than kink bursts, so for the rest of this work we focus on the former. GW bursts from string cusps have a universal shape $h(t) \propto |t - t_C|^{1/3}$, or equivalently $\tilde{h}(f) = \mathcal{A}|f|^{-4/3} e^{2\pi i f t_C}$. (More precisely, for observers that are not exactly at the center of the radiation cone, $\tilde{h}(f)$ carries a cut-off frequency f_{\max} which also smooths out $h(t)$ at $t = t_C$; see Chapter 9 below.)

Searches for cosmic-string bursts in LIGO–Virgo data are already being carried out, though to date there have been no detections [13]. However it is easy to see that the planned space-based GW detector LISA should be far more sensitive to string bursts than any current or planned ground-based instrument, due to two factors. To understand the first, recall that the matched-filtering

signal-to-noise ratio (SNR) for any burst is given by

$$\text{SNR}^2 \sim \int_0^{f_{\max}} \frac{f^2 |\tilde{h}(f)|^2 d(\log f)}{f S_h(f)} \quad (8.2)$$

for any single detector with noise spectral density $S_h(f)$, up to geometrical factors ~ 1 . Thus, for bursts with $|\tilde{h}(f)| \propto f^{-4/3}$, we have (roughly) $\text{SNR} \propto f_b^{-1/3}/[f_b S_h(f_b)]^{1/2}$, where f_b is the frequency where the detector has its best sensitivity. The value of $f_b^{-1/3}/[f_b S_h(f_b)]^{1/2}$ is ~ 10 times higher for LISA than Advanced LIGO, largely due to LISA's much lower sensitive frequency band. The second factor arises from the fact, discussed in Chapter 9, that a burst's cut-off frequency f_{\max} scales as $\alpha^{-1/3}$, where α is the angular separation between the beam direction (which is along the instantaneous direction of the cusp's motion) and the observer's line of sight. From this, we will show in Section 9.4 that the rate of bursts arriving at the detector, and satisfying $f_{\max} > f_b$, scales as $f_b^{-2/3}$. Hence, based on a uniform Euclidean distribution of sources, we can estimate that the distance to the *closest* burst that enters a detector's band scales as $f_b^{-2/9}$. This is also a factor ~ 10 higher for LISA than Advanced LIGO. So we conclude that in any given year, the strongest burst detected by LISA will have an SNR a factor ~ 100 larger than the strongest burst detected by Advanced LIGO. Clearly, LISA's much lower frequency range is a major advantage for string-burst searches.

While individual bursts are relatively featureless, as Polchinski [6] emphasizes, many burst detections would give us an approximate spectrum $dN/d\rho = \alpha\rho^\beta$ (where N is the number of detections and ρ is their SNR), and the two measured parameters α and β in principle determine the fundamental string parameters μ and p , at least for networks that are dominated by a single type of string. (However we note that in the large region of parameter space for which the strongest observed bursts would be much closer than the Hubble distance, the exponent β must be very close to -4 , and so measuring β may not be very constraining on the underlying string parameters; see Section 9.3.) Also, there are large regions of parameter space for which LISA would detect *both* individual string bursts from cusps *and* the broadband stochastic background from loop oscillations [14]. Clearly the measured energy density of the background in the LISA band would place one additional constraint on the string model.

Since the gravitational waveforms from cusps are both very simple and rather precisely known, it is natural to search for them using matched filtering. As we explain in more detail in Chapter 9, for any set of string parameters, one can easily compute the SNR^2 , which is essentially a measure of how well the model waveform (i.e., template) matches the data. Then, roughly speaking, finding the best-fit parameters is a matter of maximizing the SNR^2 over the six-dimensional source-parameter space. For three of the parameters (the signal's amplitude \mathcal{A} , polarization ψ , and arrival time t_C), this maximization can be performed almost trivially, using a combination of the F -statistic [15]

and the FFT. For the remaining three parameters (the two angles giving the source’s sky position, and the cut-off frequency f_{\max}), we made use of two publicly available optimization codes: PyMC [16], a Python implementation of Markov Chain Monte Carlo integration, and MultiNest [17, 18], a Fortran 90 implementation of a multimodal nested-sampling algorithm [19]. Employing two different optimization algorithms allowed us to carry out useful cross-checks. For high-SNR cases, we were able to recover Fisher-matrix error estimates, as expected.

We tested our searches using data sets from the recent third Mock LISA Data Challenge (MLDC) [1, 2]. Both our PyMC and MultiNest searches performed well in locating the global SNR maxima in parameter space, and our best-fit SNRs were within 1% of the true SNRs for all MLDC3 cases. The sources proved difficult to localize correctly on the sky, but, as we show in Chapter 10, that was due to near-degeneracies intrinsic to the problem, rather than a failure of our searches.

Two other reports on LISA string-burst searches, also developed and tested in the context of MLDC 3, have appeared recently [20, 21]. Our work differs from those in several ways: First, we use the F -statistic and FFT to improve search efficiency. Second, we present an in-depth analysis of waveform overlap (maximized over \mathcal{A} , ψ , and t_c) as a function of sky position. This analysis clarifies why, for most LISA cusp-burst detections, the source’s sky location is likely to be very poorly constrained by the data. Third, we analyze in detail some aspects of the problem that heretofore have not been carefully explored, including a suite of nearly exact symmetries (most of which were not previously noted), and the expected distribution of the maximum frequency in *observed* cusp-bursts.

Other authors have recently focused on other possible kinds of GW signatures from cosmic strings: DePies and Hogan [22] pointed out that for very small string tensions ($10^{-19} \lesssim \mu \lesssim 10^{-11}$), GWs might be detected from the oscillations of individual nearby strings, thanks to the nearly periodic nature of loop oscillations, and to the gravitational clustering of string loops near our Galaxy. Leblond and colleagues [23] showed how the breaking of metastable cosmic strings could result in detectable GW signals. In this work, however, we restrict attention to searches for cusp-bursts.

The plan for the rest of Part II is as follows: In Chapter 3 we briefly review the general form of a GW burst emitted by a cosmic-string cusp, as well as the associated signal registered by LISA. We also review how to maximize SNR cheaply over the extrinsic parameters \mathcal{A} , ψ , and t_C , using the F -statistic and the FFT (both standard tricks), and we introduce an approximate Bayesian version of the F -statistic, which is only slightly harder to compute than the standard variety. Finally, we digress slightly to discuss the expected distribution of f_{\max} for observable sources. In Chapter 10 we discuss the near-degeneracies in the space of burst signals (and therefore in source parameter space), which significantly impact one’s ability to infer the true source parameters from a measurement: to wit, there is a discrete near-symmetry between sky locations that are reflections of each other across the plane of the LISA detector; in addition, a typical signal from a generic sky location can be

mimicked to surprising accuracy by templates corresponding to a broad swath of very distant points on the sky, if the amplitude, polarization and arrival time of the templates are adjusted suitably. In Chapter 11 we give brief reviews of the MCMC and nested-sampling search concepts, and we describe the particular tunings of these methods that we found to be efficient for our GW burst searches. In Chapter 12 we describe the efficacy and accuracy of our searches in the MLDC data sets. We summarize our results and conclusions in Chapter 13. Throughout this part of the thesis we use units where $G = c = 1$; all quantities are expressed in units of seconds (to some power).

Bibliography

- [1] S. Babak et al. The Mock LISA Data Challenges: from Challenge 1b to Challenge 3. *Class. Quantum Grav.*, 25:184026, 2008.
- [2] S. Babak et al. The Mock LISA Data Challenges: from Challenge 3 to Challenge 4. *Class. Quantum Grav.*, 27:084009, December 2010.
- [3] A. Vilenkin and E. P. S. Shellard. *Cosmic Strings and Other Topological Defects*. Cambridge University Press, Cambridge, United Kingdom, 2000.
- [4] S. Sarangi and S.-H. H. Tye. Cosmic string production towards the end of brane inflation. *Physics Letters B*, 536:185–192, June 2002.
- [5] E. J. Copeland, R. C. Myers, and J. Polchinski. Cosmic superstrings II. *Comptes Rendus Physique*, 5:1021–1029, November 2004.
- [6] J. Polchinski. Introduction to Cosmic F- and D-Strings. arXiv:0412244 (hep-th), December 2004.
- [7] B. Allen. The Stochastic Gravity-Wave Background: Sources and Detection. In J. A. Miralles, J. A. Morales, and D. Saez, editors, *Some Topics on General Relativity and Gravitational Radiation*, page 3, Paris, 1997. Frontiers.
- [8] F. A. Jenet, G. B. Hobbs, W. van Straten, R. N. Manchester, M. Bailes, J. P. W. Verbiest, R. T. Edwards, A. W. Hotan, J. M. Sarkissian, and S. M. Ord. Upper Bounds on the Low-Frequency Stochastic Gravitational Wave Background from Pulsar Timing Observations: Current Limits and Future Prospects. *Astrophys. J.*, 653:1571–1576, December 2006.
- [9] LIGO Scientific Collaboration and Virgo Collaboration. An upper limit on the stochastic gravitational-wave background of cosmological origin. *Nature*, 460:990, 2009.
- [10] C. J. Hogan and P. L. Bender. Estimating stochastic gravitational wave backgrounds with the Sagnac calibration. *Phys. Rev. D*, 64(6):062002, September 2001.

- [11] T. Damour and A. Vilenkin. Gravitational Wave Bursts from Cosmic Strings. *Phys. Rev. Lett.*, 85:3761–3764, October 2000.
- [12] T. Damour and A. Vilenkin. Gravitational wave bursts from cusps and kinks on cosmic strings. *Phys. Rev. D*, 64(6):064008, September 2001.
- [13] B. P. Abbott et al. First LIGO search for gravitational wave bursts from cosmic (super)strings. *Phys. Rev. D*, 80(6):062002, September 2009.
- [14] X. Siemens, J. S. Key, and N. J. Cornish. *to be published*, 2010.
- [15] C. Cutler and B. F. Schutz. Generalized F-statistic: Multiple detectors and multiple gravitational wave pulsars. *Phys. Rev. D*, 72(6):063006, September 2005.
- [16] A. Patil, D. Huard, and C. J. Fonnesbeck. PyMC: Bayesian Stochastic Modelling in Python. *Journal of Statistical Software*, Vol 35, Issue 4, July, 2010.
- [17] Farhan Feroz and M. P. Hobson. Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analysis. *Mon. Not. Roy. Astr. Soc.*, 384:449–463, Feb 2008.
- [18] Farhan Feroz, M. P. Hobson, and M. Bridges. Multinest: an efficient and robust Bayesian inference tool for cosmology and particle physics. *Mon. Not. Roy. Astron. Soc.*, 398(4):1601–1614, 2009.
- [19] J. Skilling. Nested Sampling. In R. Fischer, R. Preuss, & U. V. Toussaint, editor, *American Institute of Physics Conference Series*, volume 735 of *American Institute of Physics Conference Series*, pages 395–405, November 2004.
- [20] Joey Shapiro Key and Neil J. Cornish. Characterizing the gravitational wave signature from cosmic string cusps. *Phys. Rev. D*, 79:043014, Feb 2009.
- [21] F. Feroz, J. R. Gair, P. Graff, M. P. Hobson, and A. Lasenby. Classifying LISA gravitational wave burst signals using Bayesian evidence. *Class. Quantum Grav.*, 27:075010, November 2010.
- [22] M. R. DePies and C. J. Hogan. Harmonic Gravitational Wave Spectra of Cosmic String Loops in the Galaxy. arXiv:0904.1052 (astro-ph), April 2009.
- [23] L. Leblond, B. Shlaer, and X. Siemens. Gravitational waves from broken cosmic strings: The bursts and the beads. *Phys. Rev. D*, 79(12):123519, June 2009.

Chapter 9 Theoretical Background

9.1 The gravitational waveform from cosmic-string bursts

The GWs arriving at the detector from string-cusp bursts are fully characterized by six parameters: the source's sky location (given in the MLDCs as the ecliptic latitude β and longitude λ), the burst's overall amplitude (at the detector) \mathcal{A} , the polarization angle ψ , the burst's time of arrival t_C , and the upper cut-off frequency f_{\max} .

If we fix the direction \hat{k} of GW propagation (i.e., we fix β and λ) and we let e_{ij}^+ and e_{ij}^\times be a pair of orthogonal polarization basis tensors for waves traveling along \hat{k} , the general burst waveform is expressed most simply in the Fourier domain as

$$\tilde{h}_{ij}(f) = [A^1 e_{ij}^+ + A^2 e_{ij}^\times] \Lambda(f) e^{2\pi i f t_C}, \quad (9.1)$$

where we adopt the MLDC approximation for $\Lambda(f)$,

$$\Lambda(f) \equiv \begin{cases} f^{-\frac{4}{3}} & f < f_{\max}, \\ f^{-\frac{4}{3}} e^{1-f/f_{\max}} & f > f_{\max}. \end{cases} \quad (9.2)$$

In terms of these variables, \mathcal{A} and ψ are given by

$$\mathcal{A} = \sqrt{(A^1)^2 + (A^2)^2}, \quad \psi = \arctan(A^2/A^1), \quad (9.3)$$

and in order of magnitude,

$$\mathcal{A} \sim \frac{\mu L^{2/3}}{D_L}, \quad f_{\max} \sim 2/(\alpha^3 L), \quad (9.4)$$

where μ is the string tension, L is the characteristic length of the cosmic string, D_L is the luminosity distance to the cusp, and α is the angle between the observer and the center of the beam, which points along the cusp's instantaneous velocity.¹

¹What Damour and Vilenkin actually show is that $|\tilde{h}(f)| \propto f^{-4/3}$ for $f \ll f_{\max}$, and that $|\tilde{h}(f)|$ falls to zero exponentially for $f \gg f_{\max}$. Equation 9.2 follows the signal model implemented in the LIGO Algorithm Library (LAL) to generate burst injections. This model is more precise than Damour and Vilenkin's description, though not necessarily very accurate. For consistency, the MLDCs adopted the LAL model.

9.2 Maximization over the extrinsic parameters

The SNR can be maximized analytically over the parameters \mathcal{A} and ψ using a version of the F -statistic, while the FFT provides a highly efficient method to maximize SNR over t_C . Let us work out the details, beginning with the F -statistic. Consider the space of cusp-burst waveforms, and fix the parameters $\Theta \equiv (\beta, \lambda, t_C, f_{\max})$. We shall build a statistic that is equal to the log-likelihood maximized over the vector space of all (A^1, A^2) . This statistic is a straightforward adaptation of the method employed in the (more complicated) cases of circular-orbit binaries [1] and GW pulsars [2, 3].

The LISA science data will consist of the time series of laser-noise-canceling TDI (Time Delay Interferometry) observables ([4], and references therein); all the available information about GWs can be recovered from a basis of three such observables, such as A , E , and T [5, 6] (these three are especially expedient since they have uncorrelated noises). Thus we represent the detector output as the vector $\mathbf{s} \equiv (s_A(t), s_E(t), s_T(t))$, and we define the natural inner product on the vector space of all possible LISA signals (see, e.g., [7]),

$$\langle \mathbf{u} | \mathbf{v} \rangle \equiv 2 \int_{-\infty}^{\infty} \frac{\tilde{u}_A(f) v_A^*(f) df}{S_A(f)} + (\text{integrals for } E \text{ and } T), \quad (9.5)$$

where $S_A(f)$ is the single-sided noise spectral density for the observable A (and similarly for $S_E(f)$ and $S_T(f)$). Assuming Gaussian noise, the log probability density of any noise realization \mathbf{n} is then just $(-1/2)\langle \mathbf{n} | \mathbf{n} \rangle$, and therefore the log-likelihood of the data \mathbf{s} given the signal model \mathbf{h} is $(-1/2)\langle \mathbf{s} - \mathbf{h} | \mathbf{s} - \mathbf{h} \rangle$.

Now, both polarization components of the burst produce a linear response in the three TDI observables,

$$\begin{aligned} A^1 \Lambda(f) e^{2\pi i f t_C} e_{ij}^+ &\rightarrow A^1 \left(F_A^+, F_E^+, F_T^+ \right) \Lambda(f) e^{2\pi i f t_C} \equiv A^1 \mathbf{h}_1(t_C), \\ A^2 \Lambda(f) e^{2\pi i f t_C} e_{ij}^\times &\rightarrow A^2 \left(F_A^\times, F_E^\times, F_T^\times \right) \Lambda(f) e^{2\pi i f t_C} \equiv A^2 \mathbf{h}_2(t_C); \end{aligned} \quad (9.6)$$

here the $F_{A,E,T}^{+,\times}$ are linear time-delay operators that encode the LISA response to plane GWs (see [4, 8], as well as the discussion in Section 10.1). The time delays change continuously as the LISA constellation orbits the Sun, but in the limit of short-lived GWs, LISA can be considered stationary, and the delays fixed. Thus, the operators can be represented as frequency-dependent complex factors $F_{A,E,T}^{+,\times}(t_C, f)$, which are the analogs of antenna patterns for ground-based interferometers. For cosmic-string bursts, this approximation is justified by the fact that most of the SNR is accumulated over several thousand seconds, to be compared with the one-year timescale of the LISA motion. In our searches, however, we always compute the full LISA response in the time domain, using *Synthetic*

LISA [4] (a software package that simulates the response function of LISA).

The best-fit values of A^1 and A^2 are those that minimize

$$\langle \mathbf{s} - A^1 \mathbf{h}_1(t_C) - A^2 \mathbf{h}_2(t_C) | \mathbf{s} - A^1 \mathbf{h}_1(t_C) - A^2 \mathbf{h}_2(t_C) \rangle. \quad (9.7)$$

It is easy to show that the optimized A^i and the log-likelihood $\log L$ are given by

$$A^i = (\Gamma^{-1})^{ij} \langle \mathbf{h}_j(t_C) | \mathbf{s} \rangle, \quad (9.8)$$

$$\log L = -\frac{1}{2} [\langle \mathbf{s} | \mathbf{s} \rangle - (\Gamma^{-1})^{ij} \langle \mathbf{h}_i(t_C) | \mathbf{s} \rangle \langle \mathbf{h}_j(t_C) | \mathbf{s} \rangle] + \text{const.}, \quad (9.9)$$

where the constant in Eq. 9.9 is just the logarithm of a volume factor, and where

$$\Gamma_{ij}(t_C) = \langle \mathbf{h}_i(t_C) | \mathbf{h}_j(t_C) \rangle. \quad (9.10)$$

For any given data \mathbf{s} , the term $\langle \mathbf{s} | \mathbf{s} \rangle$ is also a constant; the remaining piece of $\log L$, which depends on h , is known as the F -statistic, and it is given by

$$F \equiv \frac{1}{2} (\Gamma^{-1})^{ij} \langle \mathbf{h}_i(t_C) | \mathbf{s} \rangle \langle \mathbf{h}_j(t_C) | \mathbf{s} \rangle. \quad (9.11)$$

In the limit of high SNR, $F \approx \text{SNR}^2/2$, while in the absence of GWs the expectation value of F is 1. (It is 2 for GW pulsars, but in that case the F -statistic is maximized analytically over twice as many parameters.)

Using the FFT to maximize SNR over the time of arrival is also a standard technique [9]. Here we merely review the implementation details for our case. We arrive at the best-fit t_C [for a given $(\beta, \lambda, f_{\max})$] by a simple, iterative scheme. We make an initial estimate $t_C^{(0)}$ (e.g., by an initial search step in which the source is assumed to be at the ecliptic North pole), and we compute $\tilde{h}_1^{(0)}(f)$ and $\tilde{h}_2^{(0)}(f)$ using the time-delay operators evaluated for that time. Next, we calculate the overlap integrals $\langle \mathbf{h}_i(t_C) | \mathbf{s} \rangle$ at times $t_C = t_C^{(0)} + \Delta t$ by taking the inverse Fourier transform,

$$\langle \mathbf{h}_i(t_C) | \mathbf{s} \rangle = 2 \int_{-\infty}^{\infty} \left[\frac{\tilde{s}_A(f) \tilde{h}_i^{(0)}(f)^*}{S_A(f)} + (A \leftrightarrow E) + (A \leftrightarrow T) \right] e^{-2\pi i f \Delta t} df. \quad (9.12)$$

Approximating Γ_{ij} as the constant $\Gamma_{ij}(t_C^{(0)})$, we have

$$F(t_C^{(0)} + \Delta t) = \frac{1}{2} (\Gamma^{-1}(t_C^{(0)}))^{ij} \langle \mathbf{h}_i(t_C^{(0)} + \Delta t) | \mathbf{s} \rangle \langle \mathbf{h}_j(t_C^{(0)} + \Delta t) | \mathbf{s} \rangle. \quad (9.13)$$

Of course, the advantage of this approach is that we can use the FFT to obtain $F(t_C^{(0)} + n\delta t)$ cheaply for all integers n , where δt is the sampling time. We can now find the value $n = n_b$ that maximizes

F , fit a parabola to the values of F at the points $n_b - 1$, n_b , and $n_b + 1$, and locate Δt_b at the maximum of the parabola. We then set $t_C^{(1)} \rightarrow t_C^{(0)} + \Delta t_b$, replace $(\Gamma^{-1}(t_C^{(0)}))^{ij}$ by $(\Gamma^{-1}(t_C^{(1)}))^{ij}$, and iterate. The reason we are iterating is that we need to account for the change in the time-delay operators over the time Δt ; in practice, we always find that the original estimate $t_C^{(0)}$ is within ~ 500 s of the true t_C (see Section 11.1), and that a single iteration determines the best-fit t_C to ~ 0.01 s. (That is, further iterations change t_C by $\lesssim 0.01$ s.)

This completes our account of the maximization of log-likelihood over the parameters (\mathcal{A}, ψ, t_C) . The search over the remaining parameters $(\beta, \lambda, f_{\max})$, is discussed in Section 11.1.

9.3 Bayesian version of the F -statistic

As emphasized above, the F -statistic maximizes the log-likelihood over the parameters \mathcal{A} and ψ . However, since we have *prior* information on their distribution, it makes sense to use it to improve their estimation, as well as detection performance. As shown by Prix and Krishnan [10], it is straightforward to construct a Bayesian version of F (which we shall call F_B) that incorporates the prior knowledge. The exact form of F_B is somewhat unwieldy, but in this work we show how to construct an approximate version that is only slightly harder to compute than the standard F -statistic, and that is quite accurate for reasonably high SNR (i.e., for the cases of greatest interest).

Given the LISA data \mathbf{s} , let $P(\Theta, \mathcal{A}, \psi | \mathbf{s})$ be the posterior probability of the source parameters [with $\Theta \equiv (\beta, \lambda, t_C, f_{\text{rm}})$]. As per Bayes' theorem,

$$P(\Theta, \mathcal{A}, \psi | \mathbf{s}) \propto P(\mathbf{s} | \Theta, \mathcal{A}, \psi) P(\Theta, \mathcal{A}, \psi), \quad (9.14)$$

where the first factor on the right is the likelihood of measuring \mathbf{s} given the parameters, and the second is the prior parameter distribution. Given rotational invariance (no preferred source direction, no preferred polarization, and no preferred angle between our line of sight and the cusp velocity vector), and given the scaling of f_{\max} with the observing angle α given in Eq. 9.4 (which implies that the solid angle $\alpha d\alpha$ is $\propto f_{\max}^{-5/3} df_{\max}$), the prior *must* have the general form

$$\begin{aligned} P(\Theta, \mathcal{A}, \psi) d\Theta d\mathcal{A} d\psi &= (\sin \beta d\beta) d\lambda dt_C (f_{\max}^{-5/3} df_{\max}) \\ &\times (w(\mathcal{A}) d\mathcal{A}) d\psi, \end{aligned} \quad (9.15)$$

where $w(\mathcal{A})$ is a function of \mathcal{A} that encodes cosmological information. For simplicity, in the rest of this work we shall set $w(\mathcal{A}) = \mathcal{A}^{-4}$, as appropriate for a uniform distribution of strings in Euclidean space ($\mathcal{A} \propto r^{-1}$, where r is the distance to the source, implies $r^2 dr \propto \mathcal{A}^{-4} d\mathcal{A}$). This is a reasonable approximation for light strings ($\mu \lesssim 10^{-8}$), for which the strongest bursts that LISA observes would occur at $z < 1$. It is straightforward to modify the calculation below to treat any other form of $w(\mathcal{A})$.

The Bayesian version of the F -statistic corresponds to *integrating the posterior* $P(\Theta, \mathcal{A}, \psi | \mathbf{s})$ over \mathcal{A} and ψ , as opposed to *maximizing the likelihood* for the regular F -statistic. Fixing the data \mathbf{s} and the parameters Θ , let \mathbf{h}_b be the best-fit waveform with the \mathcal{A}_b and ψ_b that minimize $\langle \mathbf{s} - \mathbf{h} | \mathbf{s} - \mathbf{h} \rangle$. Defining $\Delta \mathbf{h} \equiv \mathbf{h}(\Theta, \mathcal{A}_b, \psi_b) - \mathbf{h}(\Theta, \mathcal{A}, \psi) \equiv \mathbf{h}_b - \mathbf{h}$, we have

$$\begin{aligned} \langle \mathbf{s} - \mathbf{h} | \mathbf{s} - \mathbf{h} \rangle &\equiv \langle \mathbf{s} - \mathbf{h}_b + \Delta \mathbf{h} | \mathbf{s} - \mathbf{h}_b + \Delta \mathbf{h} \rangle \\ &= \langle \mathbf{s} - \mathbf{h}_b | \mathbf{s} - \mathbf{h}_b \rangle + \langle \Delta \mathbf{h} | \Delta \mathbf{h} \rangle \end{aligned} \quad (9.16)$$

$$= \langle \mathbf{s} | \mathbf{s} \rangle - 2F + \langle \Delta \mathbf{h} | \Delta \mathbf{h} \rangle; \quad (9.17)$$

here Eq. 9.16 holds because $\Delta \mathbf{h}$ lies in the (A^1, A^2) vector subspace, to which $\mathbf{s} - \mathbf{h}_b$ is orthogonal thanks to the best-fit condition, and Eq. 9.17 follows from the very definition of F . Thus, the Bayesian F_B is defined by

$$e^{F_B(\Theta)} = e^{F(\Theta)} \int e^{-\langle \Delta \mathbf{h} | \Delta \mathbf{h} \rangle / 2} \mathcal{A}^{-4} d\mathcal{A} d\psi, \quad (9.18)$$

or equivalently

$$F_B(\Theta) = F(\Theta) - \log \left[\int e^{-\Gamma_{ij} \delta A^i \delta A^j / 2} \mathcal{A}^{-5} dA^1 dA^2 \right], \quad (9.19)$$

where we have changed variables from (\mathcal{A}, ψ) to (A^1, A^2) , defined (A_b^1, A_b^2) to be the best-fit values of the amplitude parameters and $\delta A^i \equiv A^i - A_b^i$, used the definition of Γ^{ij} , and transformed volume elements using the standard identity $dA^1 dA^2 = \mathcal{A} d\mathcal{A} d\psi$. We shall now introduce an approximation that is appropriate in the limit of high SNRs, for which the exponential $e^{-\Gamma_{ij} \delta A^i \delta A^j / 2}$ becomes ever more peaked around $\delta A^i = 0$. We therefore expand \mathcal{A}^{-5} around \mathcal{A}_b , discarding all terms higher than quadratic:

$$\mathcal{A}^{-5} \rightarrow \mathcal{A}_b^{-5} + \delta A^i \partial_i (\mathcal{A}^{-5})|_{\mathcal{A}_b} + \frac{1}{2} \delta A^i \delta A^j \partial_i \partial_j (\mathcal{A}^{-5})|_{\mathcal{A}_b}. \quad (9.20)$$

Note that this approximation effectively regularizes the divergence of $P(\Theta, \mathcal{A}, \psi | \mathbf{s})$ as $\mathcal{A} \rightarrow 0$, which arises from the \mathcal{A}^{-4} factor in the integrand. This divergence is unphysical anyway; it originates in the assumption of an infinite Euclidean universe, and so it is basically another version of Olbers' paradox. If we had used a cosmologically sensible prior, such as one based on an FRW universe, there would have been no divergence in the first place.

Because of symmetry, the linear term (and indeed all odd terms) of Eq. 9.20 brings no contribution to the Gaussian integral. Compared to the zeroth-order term, the contribution of the quadratic term is suppressed by $O(\text{SNR})^{-2}$, and the contribution of the quartic piece by $O(\text{SNR})^{-4}$, which justifies neglecting the latter. The remaining integral is trivial: defining

$$\lambda_{ij} \equiv \frac{1}{2} \mathcal{A}_b^5 \partial_i \partial_j (\mathcal{A}^{-5})|_{\mathcal{A}_b} = \frac{35}{2} \mathcal{A}_b^{-4} (A_b)_i (A_b)_j - \frac{5}{2} \mathcal{A}_b^{-2} \delta_{ij}, \quad (9.21)$$

we have

$$\mathcal{A}_b^{-5} \int e^{-\Gamma_{ij} \delta A^i \delta A^j / 2} [1 + \lambda_{ij}] d(\delta A^1) d(\delta A^2) = 2\pi \mathcal{A}_b^{-5} (\det \Gamma)^{-1/2} [1 + \lambda_{ij} (\Gamma^{-1})^{ij}], \quad (9.22)$$

and therefore

$$F_B = F - 5 \log \mathcal{A}_b - \frac{1}{2} \log \det \Gamma + \log [1 + \lambda_{ij} (\Gamma^{-1})^{ij}]. \quad (9.23)$$

where we have ignored the constant $\log \pi$ term, which is irrelevant to searches. Aesthetically, the reader may prefer to multiply the integral by a constant scale factor s^3 , where s is typical size for \mathcal{A} and the A^i (e.g., 10^{-21}), and then work with rescaled versions of \mathcal{A} , A^i , Γ_{ij} , and λ_{ij} , so that these are all within a few orders of magnitude of unity: $\bar{\mathcal{A}} \equiv \mathcal{A}/s$, $\bar{A}^i \equiv A^i/s$, $\bar{\Gamma}_{ij} \equiv s^2 \Gamma_{ij}$, and $\bar{\lambda}_{ij} \equiv s^2 \lambda_{ij}$. This leads to an equivalent representation of F_B , given by Eq. 9.23 after replacing all variables with their barred version.

The effect of the ‘‘Bayesian correction’’ terms in F_B is to penalize fits that have relatively larger amplitude parameters A^i . This is precisely what we should expect: since the amplitudes scale as $1/r$, larger A^i must come from strings that inhabit smaller volumes around the detector, which is a priori less likely. Note also that the terms involving Γ_{ij} (or its inverse or determinant) incorporate the effects of the detector response, and therefore depend on sky location; for the same A^i , they penalize sky-locations for which the LISA response is relatively poorer.

Ironically, our Bayesian correction is not quite appropriate for the sources in MLDC data sets, which have SNRs drawn from a uniform distribution, so that farther sources are *not* more likely than nearby ones, and sources from sky locations with a poor LISA response are equally likely to be detected. Thus, while our F_B (or its analog with a better cosmological model) would be optimal for a real search, it does not minimize the expected parameter-estimation error for our MLDC entries.

9.4 Distribution of f_{\max} for detected bursts

As an enlightening application of the distribution of burst parameters given in Eq. 9.15, we estimate the distribution of the cut-off frequency f_{\max} for the cosmic-string bursts that LISA would actually detect; i.e., for the bursts whose SNR is above some detection threshold ρ_{th} . We shall see that for most detections f_{\max} is in-band and is < 50 mHz. Since this section is something of a digression from the main flow of this work, we are content with providing a sketch of the derivation.

The first step is to change variables from \mathcal{A} to ρ , where ρ is the SNR of the observation (the other five parameters remain the same). Clearly $\rho \propto \mathcal{A}$. For simplicity, we estimate ρ in the low-frequency approximation to the LISA response [7]. In this approximation, the response functions factorize into




Figure 9.1: Expected distribution $dN/d(\log(f_{\max}))$ of the maximum burst frequency f_{\max} for the string bursts detectable by LISA.

a frequency-dependent term times an angle-dependent term, so we can write

$$\rho = \mathcal{A} \eta(f_{\max}) \kappa(\beta, \lambda, \psi) \quad (9.24)$$

where κ is a known function of the angles (β, λ, ψ) whose precise form is irrelevant, and

$$\eta(f_{\max}) \equiv \left[\int_0^{f_{\max}} \frac{\Lambda^2(f) df}{S_h(f)} \right]^{1/2}, \quad (9.25)$$

where $\Lambda(f)$ was defined in Eq. 9.2, and $S_h(f)$ [unlike the $S_{A,E,T}(f)$ of Eq. (9.5)] includes the frequency-dependent LISA response. The Jacobian of the transformation is just $(\eta\kappa)^{-1}$. Integrating the prior over all the angles, over the observation time, and over ρ from the detection threshold ρ_{th} up to ∞ , we are left with the probability distribution of detectable bursts

$$dN/df_{\max} \propto f_{\max}^{-5/3} \eta^3(f_{\max}). \quad (9.26)$$

In Fig. 9.1 we plot the function $dN/d(\log f_{\max})$. To evaluate η , we used the $S_h(f)$ fit given in Eqs. (26)–(31) of [11], which includes confusion noise from unresolved white-dwarf binaries, and for simplicity we approximated $\Lambda(f)$ as $f^{-4/3} \Theta(f_{\max} - f)$, with $\Theta(f_{\max} - f)$ the Heaviside function. As f_{\max} increases above ~ 10 mHz, η remains nearly constant, so at these higher frequencies $dN/d(\log f_{\max})$ scales as $f_{\max}^{-2/3}$. We find that the median value of f_{\max} is 12 mHz, and that $\sim 2/3$ of detected string bursts will have $f_{\max} \in [5, 50]$ mHz.

For this calculation we have assumed the “uniform, Euclidean” prior on the amplitude, $w(\mathcal{A}) \propto \mathcal{A}^{-4}$; however it should be clear that the qualitative conclusion would remain the same even if most

detected bursts were at cosmological distances. Of course, the results for the case of ground-based detectors like LIGO and Virgo would be completely analogous: the median f_{\max} for detected string bursts should be a factor $\sim 2\text{--}3$ higher than the frequency where $S_h(f)$ is at a minimum. Since for both ground-based and space-based GW detectors f_{\max} will be in-band for most observed bursts, it seems worthwhile to devote more effort to determining the precise shape of $\tilde{h}(f)$ around f_{\max} (instead of just patching together a power law with an exponential, as is currently done).

Bibliography

- [1] C. Cutler and É. E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral waveform? *Phys. Rev. D*, 49:2658–2697, March 1994.
- [2] P. Jaranowski, A. Królak, and B. F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Phys. Rev. D*, 58(6):063001, September 1998.
- [3] C. Cutler and B. F. Schutz. Generalized F-statistic: Multiple detectors and multiple gravitational wave pulsars. *Phys. Rev. D*, 72(6):063006, September 2005.
- [4] Michele Vallisneri. Synthetic lisa: Simulating time delay interferometry in a model lisa. *Phys. Rev. D*, 71:022001, Jan 2005.
- [5] T. A. Prince, M. Tinto, S. L. Larson, and J. W. Armstrong. The lisa optimal sensitivity. *Phys. Rev. D*, 66:122002, 2002.
- [6] M. Vallisneri, J. Crowder, and M. Tinto. Sensitivity and parameter-estimation precision for alternate lisa configurations. *Class. Quant. Gravity*, 25:065005, 2008.
- [7] C. Cutler. Angular resolution of the LISA gravitational wave detector. *Phys. Rev. D*, 57:7089–7102, June 1998.
- [8] M. Vallisneri. Geometric time delay interferometry. *Phys. Rev. D*, 72:042003, 2005.
- [9] S. Smith. Algorithm to search for gravitational radiation from coalescing binaries. *Phys. Rev. D*, 36:2901, 1987.
- [10] R. Prix and B. Krishnan. Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics. *Classical and Quantum Gravity*, 26(20):204013, October 2009.
- [11] L. Barack and C. Cutler. Confusion noise from LISA capture sources. *Phys. Rev. D*, 70(12):122002, December 2004.

Chapter 10 Near-symmetries and overlap maps

10.1 Sky-position reflection across the LISA plane

There is a degeneracy in the LISA response to short-duration, linearly polarized GW sources that are located at sky positions related by a reflection across the LISA plane, as first noted in [1]. This degeneracy becomes exact in the limit of infinitely short (and linearly polarized) GW signals. To understand how this degeneracy arises, we recall that the GW response of the laser-noise-canceling TDI observables can be written as [2]

$$\text{TDI}(t) = \sum_A c_A y_{(slr)_A}(t - \Delta_A), \quad (10.1)$$

where the $y_{slr}(t)$ denote the one-way phase measurements along the six LISA laser links; the slr triplet (a permutation of 123) indexes the laser-sending spacecraft, the link, and the receiving spacecraft (see Fig. 3 of [2]); the Δ_A are time delays (sums of the inter-spacecraft times of flight), and $c_A = \pm 1$. Each phase measurement y_{slr} registers plane GWs according to

$$y_{slr}(t) = \frac{\hat{n}_l(t) \cdot [h(t_s - \hat{k} \cdot p_s(t_s)) - h(t - \hat{k} \cdot p_r(t))] \cdot \hat{n}_l(t)}{2(1 - \hat{k} \cdot \hat{n}_l(t))}. \quad (10.2)$$

To parse this equation, it is useful to think about the effect of GWs on a single laser pulse received at spacecraft r at time t : the unit vector \hat{k} points along the direction of GW propagation; h is the GW strain tensor at the solar system barycenter (SSB), which is transverse to \hat{k} ; the $p_{s,r}(t)$ are the positions of the sending and receiving LISA spacecraft; the $\hat{n}_l(t) \propto p_r(t) - p_s(t_s)$ are the photon-propagation unit vectors; and the retarded time t_s is determined by the light-propagation equation $t_s = t - |p_r(t) - p_s(t_s)|$. Thus, the GW strain tensor h is projected onto \hat{n}_l at the events $(t, p_r(t))$ and $(t_s, p_s(t_s))$ [the reception and emission of the pulse]. For plane GWs, the value of h at those events is obtained by giving h the appropriate retarded-time arguments $t - \hat{k} \cdot p_r(t)$ and $t_s - \hat{k} \cdot p_s(t_s)$.

Because the $p_i(t)$ evolve on the LISA orbital timescale of a year, LISA can be considered stationary with respect to signals of much shorter duration. In that case, the three $p_{s,r}$, evaluated at the time when the signal impinges on LISA, define a plane that contains the six \hat{n}_l . Without loss of generality, let us then express all geometric quantities in an (x, y, z) coordinate system where the LISA plane lies along x and y . We reflect the source position across the LISA plane by setting $\hat{k}_z \rightarrow -\hat{k}_z$, and multiplying h on both sides by $\text{diag}(1, 1, -1)$; this has the side-effect of rotating

the polarization angle ψ of the source.¹ Because the \hat{n}_l have no z component, all the dot products that appear in Eq. (10.2) are unchanged, except for the retarded h times: but since the spacecraft positions $p_{r,s}$ can be written as a vector in the (x, y) plane plus the position vector of the LISA center, $R = (p_1 + p_2 + p_3)/3$, the overall effect is that $\text{TDI}(t)$ acquires an additional delay of $-2\hat{k} \cdot R$.

To summarize, a linearly polarized burst from some given direction is almost perfectly mimicked, in the LISA data, by a burst whose incidence direction is reflected across the LISA plane (as determined at the time when the GWs impinge on LISA), and whose polarization and arrival-time at the SSB are suitably rotated and time-translated, respectively. This degeneracy is immediately evident as the reflection symmetry across the equator in all the plots in Fig. 10.4, which examines the F -statistic structure for the strongest source in the noiseless training data set. Even for the full LISA response (without any assumptions of stationarity), the reflection symmetry is accurate to better than one part in 10^6 (in FF), which means that SNRs $\sim 1,000$ would be required to discriminate between the two sky positions.

10.2 Broad F -statistic quasi-degeneracy across the sky

Our searches revealed an additional, approximate degeneracy in the (\mathcal{A}, ψ, t_C) -maximized overlap (i.e., the F -statistic) between linearly polarized burst signals incoming from an arbitrary sky position, and templates spread in broad patterns across the sky. This approximate degeneracy appears even if we use all three noise-uncorrelated TDI observables A , E , and T (see e.g. [3]), and it is worse (i.e., more nearly degenerate) for bursts with lower f_{\max} .


While the reflection degeneracy discussed in the last section has a clear counterpart in the analytical expression of the LISA response to polarized, plane GW waves, this broad degeneracy seems harder to understand analytically. To explore it, in Fig. 10.1 we present a representative set of *fitting-factor* (FF) sky maps: each map corresponds to a target signal with the sky position and polarization indicated by the dot (and with unit amplitude and arbitrary arrival time); the contours in each map represent the overlap between the target signal and templates across the sky, maximized over the amplitude, polarization, and arrival time of the templates. By definition, $-1 \leq \text{FF} \leq 1$, but for our signals FF is very close to one across much of the sky, so we actually graph $-\log_{10}(1 - \text{FF})$ (e.g., contour “4” corresponds to $\text{FF} = 0.9999$). In all maps (and to label each map) we use latitude and longitude coordinates defined with respect to the instantaneous LISA plane. To compute the FFs, we work with the frequency-domain representation of burst waveforms *and* of the LISA response, modeling the LISA formation as a stationary, equilateral triangle; this is the same

¹For a suitable definition of the polarization angle (as given in Appendix A of [2]), the rotation is just $\psi \rightarrow -\psi$. Now, a generic non-linearly polarized signal can be described by the linear combination of two orthogonally polarized signals; the effect of the reflection considered here is then not just an overall rotation, but also a relative sign change between the two polarizations. This destroys the reflection degeneracy for generic sources, unless yet another source parameter can be adjusted to reverse the sign change.

approximation was used in [3] to compute LISA sensitivity curves. (Unequal armlengths will change the FFs somewhat, but our maps are roughly consistent with the probability distributions found in our searches, which used a full model of the LISA orbits.)

Looking at Fig. 10.1, and specifically at the large square multiple plot at the top (corresponding to a target source with latitude $\beta = \pi/3$), we observe a high-FF cell around the true position of the target source (the dot), with a mirror cell reflected across the LISA plane, at $\beta' = -\beta$. The two cells sit on a “circle in the sky” of higher FF; unlike the case of two ground-based interferometric detectors, this pattern cannot be explained by simple timing considerations, but originates from a more complicated matching of geometric elements. One side of the circle crosses the equator with higher FF, and indeed our searches often yield broken-circle distributions. In the limit of the target source moving to the equator, the two cells coalesce into one; for a target source at the pole, the maps exhibit symmetries that oscillate between two- and four-fold as a function of polarization. The bottom panel shows that FFs are considerably closer to one for bursts with lower-frequency cut offs, although the structure of the maps is qualitatively the same. The appearance of double linked circles in some maps is due to the fact that the highest displayed FF contour is set at 0.9999 (indexed by “4”); single circles would be seen to form at even higher FF.

We note that Figure 10.1 presents sky maps for reduced ranges of the target source’s λ and ψ , which are however representative of the full ranges. Because of a number of symmetries, the map for any β and λ can be obtained by appropriately shifting and reflecting one of the maps in the figure. To wit (and as exemplified in Fig. 10.2):

1. Rotating the source’s sky position by $2\pi/3$ around an axis perpendicular to the LISA plane is equivalent to relabeling the three LISA spacecraft (and the TDI observables), so the available geometric information about incoming GW signals must remain the same. Therefore $\text{map}[\beta, \lambda + 2\pi/3, \psi]$ can be obtained by shifting $\text{map}[\beta, \lambda, \psi]$ circularly by $2\pi/3$ along λ' . This degeneracy was first mentioned in [1].
2. Furthermore, there is symmetry in the geometric relation between the LISA spacecraft and sources on either side of a LISA triangle bisector: . With the definition of polarization given in [2], this results in $\text{map}[\beta, \lambda, \psi]$ reproducing $\text{map}[\beta, 2\pi/3 - \lambda, -\psi]$, modulo a λ' reflection and circular shift by $2\pi/3$.
3. Moving on to polarization, letting $\psi \rightarrow \psi + \pi/2$ amounts to reversing the sign of the polarization tensor, a change that is absorbed by the F -statistic. It follows that $\text{map}[\beta, \lambda, \psi + \pi/2] = \text{map}[\beta, \lambda, \psi]$.
4. Last, there is a non-obvious symmetry corresponding to reversing the sign of k and ψ for both target source and templates (i.e., to considering signals incoming from the antipodal sky




Figure 10.1: FF maps for high- (top) and low-frequency (bottom) bursts: $-\log_{10}(1 - \text{FF})$ contours are computed between (β, λ, ψ) target sources (with $\beta = 0, \pi/3, \pi/2$, $\lambda \in [0, \pi/3]$, $\psi \in [0, \pi/4]$) and (β', λ') templates across the sky ($\beta' \in [-\pi/2, \pi/2]$, $\lambda' \in [-\pi, \pi]$, each small square). Because of the symmetries discussed in Section 10.2, these λ and ψ ranges exhaust the variety of maps seen across their entire ranges. The target-source latitude $\beta = \pi/3$ is also representative of latitudes intermediate between the equator $\beta = 0$ and the pole $\beta = \pi/2$. At the equator, ψ has no effect on the maps (except for $\psi = \pi/4$, where there is no LISA response); at the pole, λ is degenerate, and ψ is defined consistently with the $\lambda = 0$ meridian.




Figure 10.2: Symmetries between FF maps, as explained in the main text, exemplified for the case of $\beta = \pi/6$, $\lambda = \pi/9$, $\psi = \pi/6$.

position). Because the burst GWs are invariant w.r.t. time inversion about t_C , it turns out that the LISA response to $(-k, -\psi)$ signals equals the time-inverted and time-shifted response to the original (k, ψ) signals (see subsection 10.2.1). Now, the inner product (9.5) is manifestly invariant w.r.t. the time inversion and translation of both u and v ! Thus, this results in $\text{map}[\beta, \lambda + \pi/3, \psi]$ reproducing $\text{map}[\beta, \lambda, \psi]$, modulo a circular shift by $\pi/3$.

Perhaps the most concise way to characterize the breadth of the degeneracy pattern is to plot, for each map, the fraction of the sky with FF below a given level. We do this in Fig. 10.3, where each of the superimposed lines corresponds to a choice of λ and ψ across their entire ranges; the target source latitude is kept fixed to the representative value of $\pi/3$. We can see that for high-frequency bursts (left plot), roughly half of the sky has $\text{FF} > 0.995$, and 2% (about 800 square degrees) has $\text{FF} > 0.9999$. The plot is even more dramatic for low-frequency bursts, where around 25% has $\text{FF} > 0.9999$. The significance of high FFs with respect to the determination of the source's sky position is roughly as follows: for the likelihood of any sky position to decrease by a factor e , FF must descend below $1 - 1/\text{SNR}_{\text{opt}}^2$, where SNR_{opt} is the optimal SNR for a given source. Thus $\text{FF} > 0.9999$ contains the relevant uncertainty region for $\text{SNR} \sim 100$.




Figure 10.3: Fraction of the sky with $\text{FF}(A + E + T) > 1 - 10^{-x}$, for target-source $\beta = \pi/3$, and uniformly distributed (λ, ψ) , where each pair corresponds to one of the superimposed curves. The curves were obtained by generating 40×40 maps as for Fig. 10.1, assigning a weight to each pixel corresponding to its area in the sky, sorting the resulting sequence by increasing FF, and computing normalized cumulative weights.

10.2.1 Proof of the fourth FF-map symmetry

A simple way to see this is to consider a one “arm” or a simple-Michelson TDI response (this entails no loss of generality, since Michelson TDI variables are a basis for all possible observables [3], and the derivation would proceed very similarly for first- and second-generation TDI Michelson variables). For instance, using the notation of [2] and of Eq. (10.2), consider

$$\begin{aligned}
 \text{arm}_{12}(k, \psi; t) &= y_{231}(t) + y_{13'2}(t - L) \\
 &= \frac{1}{2} \frac{n_3 \cdot [h(t - k \cdot p_1) - h(t - L - k \cdot p_2)] \cdot n_3}{1 - k \cdot n_3} \\
 &\quad + \frac{1}{2} \frac{n_{3'} \cdot [h(t - L - k \cdot p_2) - h(t - 2L - k \cdot p_1)] \cdot n_{3'}}{1 - k \cdot n_{3'}};
 \end{aligned} \tag{10.3}$$

Now $n_3 = -n_{3'}$, and the dot product of $n_3 \otimes n_3$ with the polarization tensor for a linearly polarized plane GW with (k, ψ) and $(-k, -\psi)$ can be seen to be the same using the formulas of [2, Appendix A]. Let us then drop those products, and concentrate on the time arguments of the h , as well as the geometric projection factors $1 - k \cdot n_j$. Now we let $k \rightarrow -k$, exchange n_3 with $-n_{3'}$ in the denominator, and time-advance the whole expression by $2L$:

$$\frac{h(t + 2L + k \cdot p_1) - h(t + L + k \cdot p_2)}{1 - k \cdot n_{3'}} + \frac{h(t + L + k \cdot p_2) - h(t + k \cdot p_1)}{1 - k \cdot n_3}; \tag{10.4}$$




Figure 10.4: Posterior-probability structure for the strongest source (#3) in the noiseless training data set from MLDC 3.4. **Left:** in this sky map, the density of dots (MultiNest equal-weight “resamples”) is proportional to the posterior probability, maximized over \mathcal{A} , ψ and t_C , and marginalized over f_{\max} . Crosses mark the true location of the source, and its LISA-plane-reflected counterpart. The map is plotted in the area-preserving Mollweide projection, which we adopt throughout the rest of this paper. **Right:** F -statistic as a function of ecliptic latitude and longitude, for the same sky locations as in the left panel. Here F is offset by a constant $\simeq 3,029$, and it is only slightly higher for the neighborhoods of the true and reflected sky locations than for the arcs connecting them.

after time-inverting the argument of the h (without loss of generality, let $t_C = 0$), we can match the terms one by one with the original expression, yielding, Q.E.D.,

$$\text{arm}_{12}(-k, -\psi; t + 2L) = -\text{arm}_{12}(k, \psi; -t). \quad (10.5)$$

10.3 Effects of degeneracies on searches

The broad quasi-degeneracy pattern is observed clearly in the posterior probability plots produced by our MultiNest runs (see Section 11.2). Figure 10.4 was obtained for the strongest source (with an SNR $\simeq 78$) in the noiseless² MLDC 3.4 training data set. In the left-panel sky map, the density of the dots is proportional to the posterior, maximized over \mathcal{A} , ψ and t_C , and marginalized over f_{\max} . As expected, the dots cluster around the true and reflected locations, but they extend around a thick circle that cuts through the instantaneous LISA plane at the time of the burst. In the right panel, we see that the F -statistic decreases only slightly across the circle.

Of course, detector noise will somewhat modify the noiseless posterior distribution. Figure 10.5 shows the posteriors computed for the noisy MLDC 3.4 training data set, and for five more data sets with the same source and different noise realizations, created using `lisatools` [4] with the MLDC 3.4 noise priors, but different pseudorandom-number seeds. Because FFs are consistently

²In a truly noiseless data set, the source SNR would be infinite, and it would be possible to determine its source parameters exactly. Figure 10.4 is instead produced with the usual statistical characterization of noise, for a noise realization that just *happens* to be identically zero.




Figure 10.5: Effect of different noise realizations on the posterior-probability structure for the strongest source (#3) in the noisy MLDC 3.4 training data set, and in five more data sets with the same source and different noise realizations.

high across the circle, it is possible for detector noise to displace the best-fit sky location by large angular distances, while significantly altering the structure of the circle.

In three of the plots of Fig. 10.5, the best-fit point ends up very close to the instantaneous LISA plane. Now, sources from those locations elicit a strongly suppressed response in the TDI observables, because they come close to being cross-polarized with respect to the LISA arms. However, by construction the F -statistic will raise the template amplitude correspondingly to achieve a good fit to the signal, as shown in the left panel of Fig. 10.6 for the strongest source (#3) in the (noisy) MLDC 3.4 training data set. Thus, a “straight” maximum-likelihood search can easily lead to a best-fit \mathcal{A} that is orders of magnitude larger than its true value. We have dubbed this phenomenon a *mirage*, because it makes sources appear much stronger and closer than they truly are.

It seems that mirages were not noticed by the other research groups who participated in the MLDC 3 searches for string-cusp bursts [1, 5]. We conjecture that the reason is as follows. While the F -statistic provides the best-fit \mathcal{A} and ψ for any sky location and f_{\max} , the other groups used stochastic algorithms that treat all parameters alike. Since the mirage occurs in regions of parameter space that are far removed from the true parameters, and in a subspace in which the \mathcal{A} and ψ parameters are rather precisely correlated, it is difficult for these searches to end up in these regions. (Given sufficient time, they *would* arrive there, but if one did not know that the mirages existed, one could easily be fooled into thinking that the search had converged before it actually had.)

Such mirages motivated our development of the Bayesian F_B -statistic (Section 9.3), which penalizes the large-amplitude, nearby-source fits that are a priori very unlikely. Best-fit sky locations are




Figure 10.6: **Left:** the best-fit value for the template amplitude, as computed by the F -statistic, increases dramatically for sky positions close to the instantaneous LISA plane, as shown here for source #3 in the noisy MLDC 3.4 training data set. **Right:** the Bayesian F_B -statistic shifts the best-fit sky locations away from the instantaneous LISA plane, as seen here for the six data sets of Fig. 10.5. In some cases, the best-fit location moves to the other side of the sky; this is not significant, given that reflected points have essentially the same posterior probability against the same source.

correspondingly pushed away from the instantaneous LISA plane, as illustrated in the right panel of Fig. 10.6 for the six signal-cum-noise realizations of Fig. 10.5. Unfortunately, while F_B does tend to disfavor mirage-like fits, it does not necessarily lead to best fits that are any closer to the true locations. The broad quasi-degeneracy described in Section 10.2 implies that good fits exist over much of sky, even when Bayesian priors are called into play.

Bibliography

- [1] Joey Shapiro Key and Neil J. Cornish. Characterizing the gravitational wave signature from cosmic string cusps. *Phys. Rev. D*, 79:043014, Feb 2009.
- [2] Michele Vallisneri. Synthetic lisa: Simulating time delay interferometry in a model lisa. *Phys. Rev. D*, 71:022001, Jan 2005.
- [3] M. Vallisneri, J. Crowder, and M. Tinto. Sensitivity and parameter-estimation precision for alternate lisa configurations. *Class. Quant. Gravity*, 25:065005, 2008.
- [4] lisatools.googlecode.com.
- [5] F. Feroz, J. R. Gair, P. Graff, M. P Hobson, and A. Lasenby. Classifying LISA gravitational wave burst signals using Bayesian evidence. *Class. Quantum Grav.*, 27:075010, November 2010.

Chapter 11 Search Methods

11.1 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) methods are used to efficiently integrate (and by extension, explore) arbitrary functions f defined over moderate-to-large-dimensional spaces with complex or computationally expensive integration measures P [1], when neither analytic techniques nor simple gridding techniques are feasible. MCMC methods work by creating a *Markov chain* of points that are asymptotically distributed according to P . Each next point in the chain is chosen by proposing a new candidate randomly as a function of the current point, and by choosing either the current point or the candidate on the basis of an appropriate criterion that involves their P . For any function f with finite expectation value with respect to P and for sufficiently long chains, the average value of f on the chain approaches the P -weighted average of f on the full space.

In applications of MCMC methods to Bayesian inference in signal analysis [2], P is typically the posterior probability. In this work, P is either e^F or e^{F_B} , evaluated on the 3-dimensional parameter space $(\beta, \lambda, f_{max})$, or sometimes a subspace. Our Metropolis–Hastings MCMC searches were performed using the *PyMC* software package [3] for the Python programming language. We computed F and F_B as described in Section 3, using *Synthetic LISA* [4] to obtain the GW polarizations $\mathbf{h}_{1,2}(t_C)$. *Synthetic LISA* was designed to perform highly accurate calculations of LISA’s TDI responses for any gravitational waveform impinging on LISA (e.g., for burst waveforms it does not use the approximation that LISA is stationary over the timescale of the burst), but this generality and accuracy come at some cost in speed; we find that each computation of $F(t_C)$ or $F_B(t_C)$ takes 2–3 seconds on a $\simeq 3$ GHz processor. Since single MCMC chains cannot be easily parallelized, we typically compute multiple chains, with each chain beginning in a different location in the parameter space.

Given a data set, we find it useful to initially localize the bursts in time, at least roughly. To do this, we create a waveform template with arbitrary values for the sky position (β, λ) and f_{max} , and compute $F(t)$ for all possible times t using the standard inverse Fourier transform trick described in Section 9.2. The peaks of $F(t)$ correspond to the best matches for the template in the data set. In a search on actual LISA data, we would need to carefully choose a detection threshold, to separate true GW bursts from random noise peaks. However, because MLDC 3.4 was the first challenge involving a search for cosmic strings in Mock LISA Data, the SNRs of the injected bursts were sufficiently high that the peaks from the bursts could be found in $F(t)$ by eye. Because the sky-position for our template was arbitrary, the true values of t_C (the arrival times of the signal at the SSB, not at LISA)

could differ from the times t_{\max} that maximize $F(t)$ by up to $\sim 10^3$ s. In practice, we narrowed the search to time windows $t_C \in [t_{\max} - 2000 \text{ s}, t_{\max} + 2000 \text{ s}]$, using a longer-than-necessary window for additional safety. We use each t_{\max} as the starting point for a three-stage search:

1. For the first stage, we use the fact that the best-fit value of f_{\max} has only very weak dependence on the sky position (β, λ) , so we choose a random sky position and perform a 1-D search over f_{\max} . Now, not all the MLDC 3.4 sources have a well-defined f_{\max} , which is chosen randomly (with uniformly distributed logarithm) between 10^{-3} and 10 Hz. (We noted in Section 9.4 that the true prior must scale as $f_{\max}^{-5/3}$, but rigorous verisimilitude was not a goal of this Challenge.) Thus, f_{\max} can be above the 0.5 Hz Nyquist frequency of the data set, in which case f_{\max} cannot be determined, other than to say that is > 0.5 Hz. For those signals with f_{\max} below Nyquist, we find that $\sim 1,000$ iterations are sufficient to obtain a very good estimate.
2. We now fix f_{\max} to this value, and search over the sky position (β, λ) . For this second stage, we use eight chains of $\sim 1,000$ iterations each, starting from different sky locations. Because of the reflection symmetry across the LISA plane for burst sources (see Section 10.1), two nearly equal local modes are found at this stage. For each mode, the point of highest probability among all chains is chosen as the starting point for the third stage of the search.
3. In this final stage, we search over all three $(\beta, \lambda, f_{\max})$, restricting the MCMC proposal distribution to a very narrow Gaussian in order to explore only the immediate vicinity of the starting points. We generate one chain for each of the two modes, and define our best fit as the highest-probability point of both chains.

We note that because of the computational limitations discussed above, none of our MCMC runs performed enough iterations to enter the regime of convergence. Therefore, we regard the chains as searches (maximizations) rather than explorations (integrations), and use the maxima attained by the chains as estimates of the true mode of the distributions.

11.2 MultiNest

MultiNest [5, 6] is a publicly available implementation of the *nested-sampling* algorithm for computing the Bayesian evidence of a model given a set of data. Nested sampling works by picking a set of N “live” points (typically 1,000) at random from parameter space and then systematically replacing the point with the least P with a randomly chosen point¹ of higher P . In this way the set of live points is gradually attracted toward the modes of the distribution. As the algorithm proceeds, the number of random draws required to find a suitable replacement for the worst point tends to increase

¹This random choice must take into account the prior distributions of the parameters. Indeed, MultiNest requires that the n -dimensional parameter space first be mapped into the n -dimensional unit hypercube, from which MultiNest draws samples assuming a uniform distribution. Any non-uniform priors must be taken into account in this mapping.

sharply. In order to alleviate this problem, MultiNest groups live points into ellipses, using the k - and x -means point-clustering algorithms [7]. The ellipses are designed to identify and encompass the regions of parameter space that will attract a high concentration of live points. The proposed replacements are then drawn randomly not from the entire space, but from these ellipses.

Nested sampling, like MCMC, provides a way to converge efficiently onto the (local) modes of a distribution. While this method was designed primarily to calculate the Bayesian evidence (an important concern to determine detection confidence for weak sources), we find that it also performs well at locating local maxima. Indeed, we found it relatively simple to implement a MultiNest-based search for cosmic-string bursts. Again, since we use the F -statistic and the FFT trick to maximize the likelihood over (\mathcal{A}, ψ, t_C) , we define P as e^F or e^{F_B} , and search on the remaining three parameters $(\beta, \lambda, f_{\max})$. With 1,000 live points, we find that the code converges well after approximately 10,000 point replacements, or 10 replacements per live point.

Since the probability function is identical to that used for our PyMC searches, the results from the two methods should be in good agreement. We found that this was indeed the case for both the training and challenge data. However, we prefer our MultiNest-based search, for several reasons. First, it is easily parallelized. While multiple CPUs can be used for multiple chains in MCMC, the long computation time for the log-likelihood results in none of our chains reaching the convergent regime in a reasonable run time. Although techniques such as parallel tempering and chain mixing increase the utility of a multi-chain approach, they require significantly longer chains than we were able to achieve given our choice to use exact templates (as computed with *Synthetic LISA*) rather than their static-LISA approximation. By comparison, we can easily leverage multiple CPUs for significant speed gains in MultiNest, where multiple candidate replacement points can be prepared in parallel, and unexamined candidates saved for later use. Second, since our MCMC chains do not reach the convergent regime (as discussed in Section 11.1), we are more confident in the results provided by the MultiNest algorithm, which does converge according to a well-defined criterion (a tolerance on the computed evidence). Finally, MultiNest performs well even without the somewhat elaborate three-stage procedure we use with PyMC.

11.3 High-SNR limit and the Fisher-Matrix formalism

For signals with sufficiently high SNR, the Fisher-matrix formalism provides a useful test of how accurately our codes are calculating the posterior probability. Consider a single burst immersed in noise, and imagine dialing up the burst's amplitude. As the SNR increases, the contour of constant likelihood that encloses a given fraction of the total probability (say, 68% for the $1\text{-}\sigma$ contour) shrinks to encompass an ever smaller region of parameter space. (Actually, because of the discrete symmetry described in Section 10, in our case two disjoint contours shrink onto two distinct regions: one region

that is close to the true parameter values, and another that is related to it by reflection across the LISA plane.) The smaller the region, the better the log-likelihood function within the contour is described by a constant (the maximum value) plus the second partial derivative term (the Hessian) in a Taylor expansion. The matrix of partial second derivatives of the log-likelihood is given by

$$-\frac{1}{2}\partial_\mu\partial_\nu\langle\mathbf{s}-\mathbf{h}|\mathbf{s}-\mathbf{h}\rangle=\langle\partial_\mu\partial_\nu\mathbf{h}|\mathbf{s}-\mathbf{h}\rangle-\Gamma_{\mu\nu}, \quad (11.1)$$

where $\Gamma_{\mu\nu}$ is the Fisher matrix [8], defined by

$$\Gamma_{\mu\nu}\equiv\left\langle\frac{\partial}{\partial x^\mu}\mathbf{h}\left|\frac{\partial}{\partial x^\nu}\mathbf{h}\right.\right\rangle. \quad (11.2)$$

Here $\mathbf{h}(x^\mu)$ is the waveform (a function of all the parameters x^μ), $\langle\cdots|\cdots\rangle$ is the inner product defined in Eq. (9.5), and the partial derivatives are evaluated at the local maximum. [In a slight abuse of notation, we are using Greek indices to distinguish the Gamma matrix $\Gamma_{\mu\nu}$ on the full parameter space from its restriction to the two-dimensional subspace (A^1, A^2) , which we defined as Γ_{ij} in Section 9.2.] In the high-SNR limit, the posterior distribution function near a local mode approaches a Gaussian, and the second term on the right-hand side of Eq. (11.1) dominates, so by integration of a Gaussian exponential the covariance matrix of the parameters (restricted to parameter values near the given mode) approaches the inverse of the Fisher matrix. To wit: let x_b^μ be the local best-fit parameter values, let $\Delta x^\mu\equiv x^\mu-x_b^\mu$, and let $\overline{\Delta x^\mu\Delta x^\nu}$ be the posterior-weighted average of $\Delta x^\mu\Delta x^\nu$ (where the averaging is restricted to a neighborhood of the given mode); then

$$\overline{\Delta x^\mu\Delta x^\nu}\rightarrow(\Gamma^{-1})^{\mu\nu}\quad\text{as}\quad\text{SNR}\rightarrow\infty. \quad (11.3)$$

Thus, an especially simple test of the posterior distribution generated by our MultiNest runs is just to check that, for high SNR, the ‘‘variance factor’’ $\overline{\Delta x^\mu\Delta x^\mu}/(\Gamma^{-1})^{\mu\mu}$ approaches one for all μ . As our test case, we choose the strong source (#3) from the MLDC 3.4 noiseless training data set. As shown in the top plot in Fig. 11.1, near both modes the posterior distribution is more ‘‘banana-shaped’’ than ellipsoidal, so we would not expect the Fisher-matrix approximation to be very accurate. The bottom six plots in Fig. 11.1 show the posterior distribution for each parameter separately, and compare these with Gaussian distributions based on the inverse Fisher matrix. We see that in this case, for which the SNR is ≈ 78 , the marginalized posteriors do not have Gaussian shapes, and the Fisher matrix provides only a rough estimate of the actual variances; the variance factor ranges between 0.6 and 8.8. In Fig. 11.2 we show the posterior distribution for the same source, with an increased SNR $\approx 1,000$. The agreement is much better.

We regard Fig. 11.2 as additional confirmation that our search codes are working as expected. By contrast, we regard Fig. 11.1 as a warning that for LISA detections of string-bursts, even at SNR

~ 80 , the Fisher-matrix approximation cannot be relied on to predict parameter-estimation errors accurately.

Bibliography

- [1] Jun S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer, New York, 2001.
- [2] Nelson Christensen and Renate Meyer. Markov chain monte carlo methods for bayesian gravitational radiation data analysis. *Phys. Rev. D*, 58:082001, Sep 1998.
- [3] A. Patil, D. Huard, and C. J. Fonnesbeck. PyMC: Bayesian Stochastic Modelling in Python. *Journal of Statistical Software*, Vol 35, Issue 4, July, 2010.
- [4] Michele Vallisneri. Synthetic lisa: Simulating time delay interferometry in a model lisa. *Phys. Rev. D*, 71:022001, Jan 2005.
- [5] Farhan Feroz and M. P. Hobson. Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analysis. *Mon. Not. Roy. Astr. Soc.*, 384:449–463, Feb 2008.
- [6] Farhan Feroz, M. P. Hobson, and M. Bridges. Multinest: an efficient and robust Bayesian inference tool for cosmology and particle physics. *Mon. Not. Roy. Astron. Soc.*, 398(4):1601–1614, 2009.
- [7] John A. Hartigan. *Clustering Algorithms*. John Wiley & Sons, New York-London-Sydney, 1975.
- [8] M. Vallisneri. Use and abuse of the fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects. *Phys. Rev. D*, 77:042001, 2008.




Figure 11.1: Comparison of MultiNest posterior distributions with Fisher-matrix estimates, in the case of the strongest source (#3) of the MLDC 3.4 noiseless training data. The top plot shows that the posterior distribution on the sky is more “banana-shaped” than ellipsoidal. The next six plots compare the true posterior distribution (restricted to the neighborhood of the “true” mode) with Gaussian distributions of variance $\sigma_{\mu}^2 = (\Gamma^{-1})^{\mu\mu}$. The variance factor, defined as $\sigma_{\text{fit}}^2 / \sigma_{\text{Fisher}}^2$, ranges between 0.6 and 8.8.




Figure 11.2: Same as Fig. 11.1, except that the source's SNR is now 1,000. In this case, the posterior is fit very well by the Fisher-matrix prediction. Even at this high SNR, a secondary maximum is present around the reflected location, but it is not shown in this plot.

Chapter 12 Results from the Mock LISA Data Challenges

The purpose of the MLDCs is to stimulate the development and evaluate the performance of LISA data-analysis tools and methods. In each challenge, data sets containing simulated noise plus GW signals of undisclosed source parameters are made publicly available and all interested research groups are invited to test their algorithms on these *blind* challenge data. Each challenge includes also training data sets with published source parameters, to help groups develop and calibrate their codes. The MLDCs are becoming more realistic with each new challenge, encompassing a larger number and variety of sources.

The third MLDC was the first to include a search for bursts from cosmic strings, MLDC 3.4. This data set consisted of 2^{21} samples with a cadence of 1 s (for a total of ~ 1 month), and it included a few randomly chosen string-burst signals injected into purely instrumental noise (i.e., the data set did *not* include signals from other types of sources, or the confusion noise from unresolvable Galactic binaries). The sky positions of the injected sources were chosen randomly from a uniform sky distribution; the polarizations ψ were drawn uniformly from $[0, \pi]$; and f_{\max} was drawn $[10^{-3}, 10]$ Hz with a uniformly distributed logarithm.

MLDC 3.4 called for a random number (a Poisson deviate of mean 5) of injected bursts, with SNRs drawn uniformly from $[10, 100]$. As discussed in Section 9.3, these priors for f_{\max} and SNR are not astrophysically realistic, but the intent for this challenge was less to maintain astrophysical realism than to test search algorithms for a wide range of source parameters (i.e, a wider range than one would obtain from a handful of detections with realistic parameters). As it turned out, the MLDC 3.4 data set contained exactly three string bursts, all with SNRs in the range 36–45. Of course, the realistic expectation is that most detections will have SNRs within 50% of the detection threshold, which is likely to be ~ 6 . Thus, all the MLDC 3.4 bursts had SNRs a factor 4–5 higher than will be typical.

In this challenge, the exact spectral densities of instrumental noise were randomized and undisclosed, but they were guaranteed to lie within fairly narrow ranges. In our searches, we ignored this feature, to little apparent damage, by taking the TDI observables to have the standard MLDC noise spectral densities as assumed in the other MLDC challenges. Explicit expressions for these $S_A(f)$, $S_E(f)$ and $S_T(f)$ are given in [1].

In our entries to MLDC 3.4 and in this thesis,¹ we report the best-fit parameters found by our

¹The values shown in this thesis are somewhat different from the values we submitted for MLDC 3.4, which can be viewed at www.tapir.caltech.edu/mldc. Our algorithms have improved since the conclusion of MLDC 3, and to keep this thesis current with our research effort, here we have chosen to report our newer results. In some cases, our newer best-fit parameters are actually further from the true parameter values than our original entries. Nevertheless, the values reported here arise from a more correct analysis of the data.

searches (i.e., the maxima of F or F_B). In fact, because there are always two parameter sets that fit the data almost equally well, due to the reflection symmetry described in Section 10.1, for each burst we report the best-fit parameters of both modes. Table 12.1 lists the true and best-fit parameters, and Table 12.2 the corresponding estimation errors; Figure 12.1 shows sky plots of the posterior distributions derived from our MultiNest searches.

Certain aspects of the results presented in Table 1 and Figure 10 require clarification. For Source #0, the true sky location is ruled out by parameter estimation. This should not be surprising: in the high-SNR regime, the variance of SNR^2 over the ensemble of noise realizations is of the order of the number of source parameters; thus the likelihood at the best-fit parameters can exceed the likelihood at the true parameters by large exponential factors. For Source #1, we find that the maximum of F_B lies outside the two regions of the sky where the posterior probability is concentrated. In the Table we report instead on the maxima that lie *within* the large, high-probability clusters. The outlying maximum lies close to the LISA plane, and so it resembles the mirages discussed in Section 10.3. In this case, however, the best-fit amplitude is only a factor of two higher than the true value, so the Bayesian correction term implicit in F_B does not strongly disfavor it. For Source #0, MultiNest converged to values of f_{\max} above the Nyquist frequency, although one of the MCMC chains managed to lock onto a better value.

In summary, we find that both the PyMC and MultiNest searches perform well at locating the peaks of the posterior, and that the best fits found by the two methods are mostly consistent. In this sense, both techniques are successful. However, because of the broad degeneracy of the posterior across the sky (described in Section 10.2), we find that instrument noise will generally shift the best-fit parameters rather far from their true values. Because the LISA response introduces strong correlations between sky position and the parameters (\mathcal{A}, ψ, t_C) , these come to have large errors as well. Thus, we should not hope for accurate sky locations in LISA detections of string bursts with $\text{SNR} \sim 40$, and the situation will only be worse for typical LISA detections with $\text{SNR} \lesssim 10$.

We emphasize that we believe that these large parameter-estimation errors are *not* a result of bugs or lack of convergence in our search methods, but are simply the consequence of the broad parameter-space degeneracy of cusp-burst signals. Besides the consistency between our PyMC and MultiNest results, we performed an additional test by verifying that parameter-estimation accuracy improves when we boost the SNR to $\sim 1,000$, as shown in Table 12.3 for source #3 in the noisy MLDC 3.4 training data set. For such high SNR, the MultiNest best-fit parameters are reassuringly close to the true values.

parameter	true value	MCMC #1	MCMC #2	MN #1	MN #2
Source 0					
β [rad]	0.556	0.551	0.119	0.543	0.933
λ [rad]	3.711	5.843	0.005	5.858	5.295
f_{\max} [Hz]	0.030	> 0.5	0.044	> 0.5	> 0.5
ψ [rad]	3.319	2.936	2.776	2.926	1.914
\mathcal{A} [10^{-21}]	0.86636	3.0368	1.1394	2.903	3.142
t_C [10^6 s]	1.60216	1.60288	1.60305	1.60289	1.60265
SNR	44.610	44.985	44.842	44.987	44.993
Source 1					
β [rad]	-0.444	-0.753	0.256	-0.658	0.221
λ [rad]	3.167	0.015	3.486	0.076	3.502
f_{\max} [Hz]	0.0010842	0.0010927	0.0010932	0.001087	0.001085
ψ [rad]	5.116	4.233	5.023	4.275	5.019
\mathcal{A} [10^{-21}]	2.7936	1.6528	1.6585	1.621	1.688
t_C [10^6 s]	1.07269	1.07349	1.07266	1.07352	1.07265
SNR	36.691	36.704	36.702	36.703	36.704
Source 2					
β [rad]	-0.800	0.179	1.154	0.141	1.176
λ [rad]	0.217	0.271	2.746	0.259	2.876
f_{\max} [Hz]	6.1495	0.030	0.025	0.026	0.030
ψ [rad]	4.661	4.631	5.225	4.630	5.129
\mathcal{A} [10^{-21}]	0.85403	1.0319	1.0285	1.007	1.016
t_C [10^6 s]	0.60001	0.60015	0.59949	0.60015	0.59949
SNR	41.378	41.497	41.496	41.495	41.496

Table 12.1: True source parameter values and MCMC and MultiNest best fits for the MLDC 3.4 challenge data set. When the estimated f_{\max} is larger than the 0.5 Hz Nyquist frequency.

parameter		MCMC #1	MCMC #2	MN #1	MN #2
Source 0					
Δsky [rad]		1.680	2.278	1.695	1.140
$\Delta \log_{10} f_{\max}$		> 1.222	0.169	> 1.222	> 1.222
$\Delta\psi$ [rad]		0.383	0.543	0.394	1.405
$\Delta \log \mathcal{A}$		1.254	0.274	1.209	1.288
Δt_C [s]		716.38	881.18	722.40	485.39
ΔSNR		0.375	0.232	0.378	0.383
Source 1					
Δsky [rad]		1.944	0.766	2.039	0.742
$\Delta \log_{10} f_{\max}$		3.37×10^{-3}	3.59×10^{-3}	1.270×10^{-3}	4.083×10^{-4}
$\Delta\psi$ [rad]		0.884	0.093	0.842	9.758×10^{-2}
$\Delta \log \mathcal{A}$		0.525	0.521	0.544	0.504
Δt_C [s]		794.28	41.06	828.39	43.95
ΔSNR		0.014	0.011	0.012	0.013
Source 2					
Δsky [rad]		0.980	2.662	0.942	2.690
$\Delta \log_{10} f_{\max}$		2.316	2.396	2.377	2.318
$\Delta\psi$ [rad]		0.030	-0.564	0.031	0.467
$\Delta \log \mathcal{A}$		0.189	0.186	0.165	0.174
Δt_C [s]		141.40	519.79	145.06	522.02
ΔSNR		0.119	0.118	0.117	0.118

Table 12.2: Differences between true source parameter values and MCMC and MultiNest best fits, for the MLDC 3.4 challenge data set. The Δsky error is measured in radians along the geodesic arc between the true and best-fit sky positions.

parameter	true value	boosted	best fit	best fit (boosted)
β [rad]	0.239		-0.036	0.247
λ [rad]	1.090		1.204	1.092
f_{\max} [Hz]	1.152×10^{-2}		1.161×10^{-2}	1.151×10^{-2}
ψ [rad]	0.399		0.571	0.394
\mathcal{A} [10^{-21}]	2.647	37.26	2.204	37.26
t_C [10^6 s]	2.060273		2.060245	2.060272
SNR	78.122	1082.9278	78.137	1082.9291

parameter	error	error (boosted)
Δ_{sky} [rad]	0.297	8.617×10^{-3}
$\Delta \log_{10} f_{\max}$	3.6×10^{-3}	1.2×10^{-4}
$\Delta \psi$ [rad]	0.171	4.9×10^{-3}
$\Delta \log \mathcal{A}$	0.183	0.0125
Δt_C [s]	27.89	1.45
ΔSNR	0.015	1.3×10^{-3}

Table 12.3: Parameter accuracy achieved by MultiNest for source #3 in the MLDC 3.4 training data set, with the original and boosted SNR.




Figure 12.1: MultiNest sky-location posteriors for sources 0–2 in the MLDC 3.4 challenge data set. The density of the dots is proportional to the posterior probability (including the F_B prior correction described in Section 9.3), maximized over t_C , and marginalized over \mathcal{A} , ψ , and f_{\max} . Crosses and circles indicate the true and best-fit locations, respectively. For source 1, the stars indicate the location of *mirage* best fits discarded by F_B .

Bibliography

- [1] M. Vallisneri, J. Crowder, and M. Tinto. Sensitivity and parameter-estimation precision for alternate lisa configurations. *Class. Quant. Gravity*, 25:065005, 2008.

Chapter 13 Conclusions

In this thesis I have reported work by my collaborators and I on development of string-burst search pipelines, which rely on the F -statistic and the FFT to efficiently maximize the likelihood over (\mathcal{A}, ψ) and t_C , respectively, and which are based on the publicly available PyMC and MultiNest libraries to maximize over the remaining parameters $(\beta, \lambda, f_{max})$. Both of our pipelines proved reasonably efficient (MultiNest more so, due to greater gains from parallelization). We tested our searches by checking that they yielded mutually consistent best fits, and that posteriors results agreed with Fisher-matrix estimates for sufficiently large SNR. Given the relative simplicity of string-burst signals, we expected that off-the-shelf optimization codes like PyMC and MultiNest be would sufficiently powerful for this search, which our work has verified.

Although the few string-burst injections in MLDC 3.4 all had SNR ~ 40 , it did not prove possible to localize them on the sky to better than \sim one radian. We showed that this result is just what should be expected, on the basis of the broad degeneracy illustrated by the fitting-factor maps of Section 10.2. Determinations of \mathcal{A} and ψ are correspondingly poor—again to be expected, since these parameters are strongly correlated with the sky location in the signal measured by LISA. While so far we have analyzed only a handful of bursts in detail, there is every reason to presume that LISA will suffer from poor parameter-estimation accuracy for string-burst detections. In future work, we intend to verify or disprove this presumption by analyzing a much larger sample of bursts drawn from an astrophysically sensible distribution. This thesis also included:

1. the proof of the near-degeneracy between (linearly polarized) burst signals from directions that are reflections of each other across the LISA plane (which had been noted elsewhere, but heretofore not explained analytically);
2. the first detailed look at string-burst fitting factors (FF) as a function of sky position, revealing very high FF over a large fraction of the sky;
3. the analysis of four discrete symmetries (three of which were not previously discussed) *between* different fitting function maps;
4. the derivation of an approximate, easily computed Bayesian version of the F -statistic, based on realistic priors;
5. a calculation of the expected distribution of f_{max} for detected bursts.

We envisage two broad directions for future work. First, so far we have concentrated on finding the physical parameters of a single string-burst. Using these sorts of results as input, the next step

will be to determine how well LISA can answer questions about the string network (e.g., are there different types of strings? What are μ and p for each class?) based on an observed population of string-bursts, plus any information from a cosmic-string stochastic background. Second, so far our searches have been designed for single bursts in Gaussian noise of known spectral density. We need to generalize our methods to the cases where the noise level and shape are not precisely known (and so must be determined from the data), and where the burst signals are superimposed on a realistic LISA data set containing confusion noise from millions of individually unresolvable sources (mostly white-dwarf binaries) plus tens of thousands of resolvable signals from a variety of sources (especially white-dwarf binaries, EMRIs, and merging massive black hole binaries).

Appendix A Momentum flow in black-hole binaries: II. Numerical simulations of equal-mass, head-on mergers with antiparallel spins.

This paper has been submitted for publication to Physical Review D, and is awaiting approval and publication. It can be found on arXiv with the reference gr-qc/0907.0869.

Authors of this paper are:

Geoffrey Lovelace ¹

Yanbei Chen, Michael Cohen, Jeffrey D. Kaplan, Drew Keppel, Keith D. Matthews, David A. Nichols, Mark A. Scheel, Ulrich Sperhake ²

A.1 Abstract

Research on extracting science from binary-black-hole (BBH) simulations has often adopted a “scattering matrix” perspective: given the binary’s initial parameters, what are the final hole’s parameters and the emitted gravitational waveform? In contrast, we are using BBH simulations to explore the nonlinear dynamics of curved spacetime. Focusing on the head-on plunge, merger, and ringdown of a BBH with transverse, antiparallel spins, we explore numerically the momentum flow between the holes and the surrounding spacetime. We use the Landau-Lifshitz field-theory-in-flat-spacetime formulation of general relativity to define and compute the density of field energy and field momentum outside horizons and the energy and momentum contained within horizons, and we define the effective velocity of each apparent and event horizon as the ratio of its enclosed momentum to its enclosed mass-energy. We find surprisingly good agreement between the horizons’ effective and coordinate velocities. During the plunge, the holes experience a frame-dragging-induced acceleration orthogonal to the plane of their spins and their infall (“downward”), and they reach downward speeds of order 1000 km/s. When the common apparent horizon forms (and when the event horizons merge and their merged neck expands), the horizon swallows upward field momentum that resided between the holes, causing the merged hole to accelerate in the opposite (“upward”) direction. As the merged hole and the field energy and momentum settle down, a pulsational burst of gravitational waves is emitted, and the merged hole has a final effective velocity of about 20 km/s upward, which agrees with the recoil velocity obtained by measuring the linear momentum carried to infinity by the

¹Center for Radiophysics and Space Research, Cornell University, Ithaca, New York, 14853

²Theoretical Astrophysics 350-17, California Institute of Technology, Pasadena, CA 91125

emitted gravitational radiation. To investigate the gauge dependence of our results, we compare generalized harmonic and BSSN-moving-puncture evolutions of physically similar initial data; although the generalized harmonic and BSSN-moving-puncture simulations use different gauge conditions, we find remarkably good agreement for our results in these two cases. We also compare our simulations with the post-Newtonian trajectories and near-field energy-momentum.

A.2 Introduction

A.2.1 Motivation

Following Pretorius’s 2005 breakthrough [1], several research groups have developed codes to solve Einstein’s equations numerically for the inspiral, merger, and ringdown of colliding binary black holes (BBHs). Most simulations of BBH mergers to date have adopted the moving-puncture method [2, 3], and spectral methods [4] have also successfully simulated BBH mergers.

A major goal of current research is to successfully extract the physical content of these simulations. Typically, efforts toward this goal adopt a “scattering matrix” approach. Information obtained from numerical simulations on a finite set of islands in the seven-dimensional³ parameter space is being extrapolated, by various research groups, to design complicated functions that give the final parameters of the merged hole and the emitted gravitational waveforms as functions of the binary’s initial parameters.

In this paper, however, we take a different perspective: we focus our attention on the *nonlinear dynamics of curved spacetime* during the holes’ merger and ringdown. Following Ref. [5] (paper I in this series), our goal is to develop physical insight into the behavior of highly dynamical spacetimes such as the strong-field region near the black-hole horizons in a merging binary. As in paper I, we focus this study on the distribution and flow of linear momentum in BBH spacetimes. In contrast to paper I’s description of the pre-merger motion of the holes in the post-Newtonian approximation, in this paper we study the momentum flow during the plunge, merger, and ringdown of merging black holes in fully relativistic simulations.

A.2.2 Linear momentum flow in BBHs and gauge dependence

Typically, numerical simulations calculate only the *total* linear momentum of a BBH system and ignore the (gauge-dependent) linear momenta of the individual black holes. However, linear momentum has been considered by Krishnan, Lousto and Zlochower [6]. Inspired by the success of quasilocal angular momentum (see, e.g., [7] for a review) as a tool for measuring the spin of an individual black hole, Krishnan and colleagues proposed an analogous (but gauge-dependent) formula

³One parameter for the mass ratio and six for the individual spins; additional parameters might arise from eccentric orbits and the apparent dependence, in at least some configurations, of the recoil on the initial phase of the binary.

for the quasilocal linear momentum, and they calculate this quasilocal linear momentum for, e.g., the highly-spinning, unequal-mass BBH simulations in Ref. [8]. This quasilocal linear momentum is also used to define an orbital angular momentum in Ref. [9].

In this paper, we adopt a different, complementary method for measuring the holes’ linear momenta: for the first time, we apply the Landau-Lifshitz momentum-flow formalism (described in paper I and summarized in Sec. A.3) to numerical simulations of merging black holes. In this formalism, a mapping between the curved spacetime and an auxiliary flat spacetime (AFS) is chosen, and general relativity is reinterpreted as a field theory defined on this flat spacetime. The AFS has a set of translational Killing vectors which we use to define a localized, conserved linear momentum. In particular, we calculate i) a momentum density, ii) the momentum enclosed by horizons, and iii) the momentum enclosed by distant coordinate spheres. In the asymptotically flat region around a source, there is a preferred way to choose the mapping between the curved spacetime and the AFS; consequently, in this limit item iii) is gauge-invariant. In general, though, the choice of mapping is arbitrary, and it follows that items i) and ii) are necessarily gauge-dependent.

By examining the linear momentum flow in a dynamical spacetime—and living with the inevitable gauge dependence—we hope to develop strong intuition for the behavior of BBHs. As discussed in Sec. I C of paper I, we envision different numerical relativity groups choosing “preferred” gauges based on the coordinates of their numerical simulations. While there is no reason, *a priori*, why simulations in different gauges should agree, one of our hopes from paper I is realized *for the cases we consider*; namely, in this paper, we calculate the horizon-enclosed momentum using generalized harmonic and BSSN-moving-puncture evolutions of similar initial data, and we do find surprisingly good agreement (cf. Figs. A.8 and A.15), even though the simulations use manifestly different gauge conditions [Eqs. (A.14) for the generalized harmonic simulations and Eqs. (A.51)–(A.52) for the BSSN-moving-puncture simulations]. These are two of the most commonly used gauge conditions in numerical relativity.

Therefore, we continue to hope that in general—for the gauges commonly used in numerical simulations—the momentum distributions for evolutions of physically similar initial data will turn out to be at least qualitatively similar. If further investigation reveals this to be the case, then different research groups can simply use the coordinates used in their simulations as the “preferred coordinates” for constructing the mapping to the AFS. Otherwise, we would advocate (as in Sec. I C of paper I) that different numerical-relativity groups construct the mapping to the AFS by first agreeing on a choice of “preferred” coordinates (e.g., a particular harmonic gauge) and then transforming the results of their simulations to those coordinates.

A.2.3 BBH mergers with recoil

A particularly important application of this approach is an exploration of the momentum flow in BBH mergers with recoil. The gravitational recoil or kick effect arising in a BBH coalescence has attracted a great deal of attention in recent years in the context of a variety of astrophysical scenarios including the structure of galaxies [10, 11, 12], the reionization history of the universe [13], the assembly of supermassive black holes [14, 15, 16, 17, 18] and direct observational signatures [19, 20, 21]. For a long time, estimates of the recoil magnitude were based on approximative techniques [22, 23, 24, 25]; accurate calculations in the framework of fully nonlinear general relativity have only become possible in the aftermath of important breakthroughs in the field of numerical relativity [1, 2, 3].

Several groups have used numerical simulations to study the kick resulting from the merger of non-spinning and spinning binaries (see, e.g., [26, 27, 28, 29, 30, 31]). Most remarkably, recoil velocities of several thousand km/s have been found for binaries with equal and opposite spins in the orbital plane [30, 32, 33], and variants thereof with hyperbolic orbits even reach recoil velocities of 10^4 km/s [34]. Given the enormous astrophysical repercussions of such large recoil velocities, the community is now using various approaches to obtain a better understanding of the kick as a function of the initial BBH parameters [35, 36, 37, 38, 39, 40] resulting in phenomenological fitting formulas; see [41, 9, 38, 42, 43, 8] and references therein.

On the other hand, our understanding of the local dynamics in these extraordinarily violent events is still rather limited. Some insight into the origin of the holes' kick velocity has been obtained by examining the individual multipole moments of the emitted gravitational waves [44, 45] and by approximating the recoil analytically using post-Newtonian [24, 46], effective-one-body [25], and black-hole-perturbation theory [47]. Pretorius has presented an intuitive picture which describes aspects of the so-called superkick configurations (which generate velocities in the thousands of km/s) in terms of the frame-dragging effect (cf. Fig. 5 of Ref. [48]).

Investigating the momentum distribution and flow in recoiling BBH mergers could help to build further intuition into the nonlinear dynamics of the spacetime and their influence on the formation of kicks. Paper I made some headway into the former issue but could not address the latter. Specifically, paper I examined the distribution and the flow of linear momentum in BBH spacetimes using the Landau-Lifshitz formalism in the post-Newtonian approximation. It then specialized this approach to the extreme-kick configuration [30, 32, 33], which is a system of inspiraling, BBHs with equal and anti-parallel spins in the orbital plane. During inspiral, the two black holes simultaneously and sinusoidally bob perpendicularly to the orbital plane; in paper I, this motion was first recognized as arising from the combined effect of frame dragging and spin-curvature coupling and then was found to arise from the exchange of momentum between the near-zone gravitational field and the black holes.

Because paper I analyzed the system at a post-Newtonian level, its analysis could not be extended




Figure A.1: Initial configuration of the head-on BBH considered in this paper. The holes move primarily along the x axis, but they also accelerate in the $-y$ (downward) direction due to frame dragging. See Table A.1 for the value of $d = 2x_0$.

to merger and beyond. Consequently, it was not possible to address how the nonlinear dynamics in the pre-merger near zone transitions into the final behavior of the merged black hole. This paper (paper II) lets us begin to address this transition as we study momentum flow during the plunge, merger, and ringdown of BBHs in *full numerical relativity*. Our study allows us, for example, to examine how accurately Pretorius’s intuitive picture applies during the merger and ringdown of a recoiling BBH merger.

A.2.4 Overview and summary

As a first step toward analyzing the momentum flow in superkicks, in this paper we apply the Landau-Lifshitz momentum-flow formalism to a much simpler case: the head-on plunge, merger, and ringdown of an equal-mass BBH. The holes initially have antiparallel spins of equal magnitude that are transverse to the holes’ head-on motion (Fig. A.1). Primarily, the holes simply fall toward each other in the $\pm x$ direction. However, each hole’s spin drags the space around itself, causing the other hole to accelerate in the downward, $-y$ direction.

How does this frame dragging relate to the final kick velocity of the merged hole? To address this question, we compute the 4-momentum p^μ inside each apparent horizon using the Landau-Lifshitz formalism; we then define an *effective velocity* as

$$v_{\text{LL}}^i := \frac{p^i}{p^0}. \quad (\text{A.1})$$

In Sec. A.5, we find that this effective velocity behaves similarly to the apparent horizons’ coordinate




Figure A.2: The effective velocity v_{LL}^y for the individual (red dotted line) and common (green dashed line) apparent horizons and for the event horizon (black solid line). The inset shows the velocity of the common apparent horizon at late times.

velocities⁴

The effective y velocity for the generalized harmonic simulation described in Sec. A.4.1 is shown in Fig. A.2. Before the merger, the individual apparent horizons do indeed accelerate in the $-y$ (“down”) direction, eventually reaching velocities of order 10^3 km/s. However, when the common apparent horizon forms, it pulsates; during the first half-pulsation, the horizon expands and accelerates to $\sim 10^3$ km/s in the up ($+y$) direction. This happens because as the common horizon forms and expands, it swallows not only the downward linear momentum inside each individual horizon but also a large amount of upward momentum in the gravitational field between the holes (Fig. A.3). During the next half-pulsation, as the horizon shape changes from oblate to prolate (cf. Fig. A.11), the horizon swallows a net downward momentum, thereby losing most of its upward velocity. Eventually, after strong damping of the pulsations, the common horizon settles down to a very small velocity of about 23 km/s in the $+y$ direction (inset of Fig. A.2), which (Sec. A.5) is consistent with the kick velocity inferred from the emitted gravitational radiation.

This momentum flow between field and holes is also described quite beautifully in the language of the holes’ *event horizon*. Unlike apparent horizons, the event horizon evolves and expands continuously in time, rather than discontinuously. As the event horizon expands, it continuously swallows surrounding field momentum, and that swallowing produces a continuous evolution of the event

⁴By coordinate velocity, we mean the velocity of the center of the apparent horizon, as measured in our asymptotically inertial coordinates.




Figure A.3: A contour plot of the y component of the momentum density at the moment when the common apparent horizon forms. The common horizon encloses the momentum inside the individual horizons and also the momentum in the gravitational field. The grey-shaded region and solid, red contours indicate positive momentum density, while the white-shaded region and blue, dashed contours indicate negative momentum density. The individual apparent horizons are shaded black, and the common apparent horizon is shown as a thick black line.

horizon's velocity, an evolution that is nearly the same as for the apparent-horizon velocity. Figure A.2 shows how the effective velocity of the event horizon smoothly transitions from matching the individual apparent horizons' velocities to matching the common apparent horizon's velocity. For further details, see Sec. A.5.1 and especially Figs. A.13 and A.14.

In the remainder of this paper, we discuss our results and the simulations that are used to obtain them. In Sec. A.3, we briefly review the Landau-Lifshitz formalism and momentum conservation. The simulations themselves are presented in Sec. A.4. We analyze the simulations' momentum flow in Sec. A.5 and conclude in Sec. A.6. In the appendices, we describe in greater depth the numerical methods used for the simulations presented in this paper.

A.3 4-Momentum Conservation in the Landau-Lifshitz Formalism

In this section, we briefly review the Landau-Lifshitz formulation of gravity and the statement of 4-momentum conservation within this theory. Landau and Lifshitz, in their *Classical Theory of Fields* (hereafter referred to as LL), reformulated general relativity as a nonlinear field theory in flat spacetime [49]. (Chap. 20 of MTW [50] and a paper by Babak and Grishchuk [51] are also helpful sources that describe the formalism.) Landau and Lifshitz develop their formalism by first laying down arbitrary asymptotically Lorentz coordinates on a given curved (but asymptotically-flat) spacetime. They use these coordinates to map the curved (i.e. physical) spacetime onto an auxiliary flat spacetime (AFS) by enforcing that the coordinates on the AFS are globally Lorentz. The auxiliary flat metric takes the Minkowski form, $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$.

In this formulation, gravity is described by the physical metric density

$$\mathbf{g}^{\mu\nu} := \sqrt{-g}g^{\mu\nu} , \quad (\text{A.2})$$

where g is the determinant of the covariant components of the physical metric, and $g^{\mu\nu}$ are the contravariant components of the physical metric. When one defines the superpotential

$$H^{\mu\alpha\nu\beta} := \mathbf{g}^{\mu\nu}\mathbf{g}^{\alpha\beta} - \mathbf{g}^{\mu\alpha}\mathbf{g}^{\nu\beta} , \quad (\text{A.3})$$

the Einstein field equations take the field-theory-in-flat-spacetime form

$$H^{\mu\alpha\nu\beta}{}_{,\alpha\beta} = 16\pi\tau^{\mu\nu} . \quad (\text{A.4})$$

Here $\tau^{\mu\nu} := (-g)(T^{\mu\nu} + t_{\text{LL}}^{\mu\nu})$ is the total effective stress-energy tensor, indices after the comma denote partial derivatives or, equivalently, covariant derivatives with respect to the flat auxiliary

metric), and the Landau-Lifshitz pseudotensor $t_{\text{LL}}^{\mu\nu}$ (a real tensor in the auxiliary flat spacetime) is given by Eq. (100.7) of LL [49] or equivalently Eq. (20.22) of MTW [50]:

$$\begin{aligned}
16\pi(-g)t_{\text{LL}}^{\alpha\beta} &= \mathfrak{g}^{\alpha\beta}{}_{,\lambda}\mathfrak{g}^{\lambda\mu}{}_{,\mu} - \mathfrak{g}^{\alpha\lambda}{}_{,\lambda}\mathfrak{g}^{\beta\mu}{}_{,\mu} \\
&+ \frac{1}{2}g^{\alpha\beta}g_{\lambda\mu}\mathfrak{g}^{\lambda\nu}{}_{,\rho}\mathfrak{g}^{\rho\mu}{}_{,\nu} \\
&- g^{\alpha\lambda}g_{\mu\nu}\mathfrak{g}^{\beta\nu}{}_{,\rho}\mathfrak{g}^{\mu\rho}{}_{,\lambda} - g^{\beta\lambda}g_{\mu\nu}\mathfrak{g}^{\alpha\nu}{}_{,\rho}\mathfrak{g}^{\mu\rho}{}_{,\lambda} \\
&+ g_{\lambda\mu}g^{\nu\rho}\mathfrak{g}^{\alpha\lambda}{}_{,\nu}\mathfrak{g}^{\beta\mu}{}_{,\rho} \\
&+ \frac{1}{8}(2g^{\alpha\lambda}g^{\beta\mu} - g^{\alpha\beta}g^{\lambda\mu}) \\
&\times (2g_{\nu\rho}g_{\sigma\tau} - g_{\rho\sigma}g_{\nu\tau})\mathfrak{g}^{\nu\tau}{}_{,\lambda}\mathfrak{g}^{\rho\sigma}{}_{,\mu}
\end{aligned} \tag{A.5}$$

Due to the symmetries of the superpotential—they are the same as those of the Riemann tensor—the field equations (A.4) imply the differential conservation law for 4-momentum

$$\tau^{\mu\nu}{}_{,\nu} = 0. \tag{A.6}$$

Eq. (A.6) is equivalent to $T^{\mu\nu}{}_{;\nu} = 0$, where the semicolon denotes a covariant derivative with respect to the physical metric.

In both LL and MTW, it is shown that the total 4-momentum of any isolated system (measured in the asymptotically flat region far from the system) is

$$p_{\text{tot}}^\mu = \frac{1}{16\pi} \oint_{\mathcal{S}} H^{\mu\alpha 0j}{}_{,\alpha} d\Sigma_j, \tag{A.7}$$

where $d\Sigma_j$ is the surface-area element of the flat auxiliary metric, and \mathcal{S} is an arbitrarily large surface surrounding the system. This total 4-momentum satisfies the usual conservation law

$$\frac{dp_{\text{tot}}^\mu}{dt} = - \oint_{\mathcal{S}} \tau^{\mu j} d\Sigma_j. \tag{A.8}$$

See the end of Section III of [5] for a brief proof of why this holds for black holes.

Because this paper focuses on BBHs, we will make a few further definitions that will be used frequently in our study. First, we label the two⁵ black holes in the binary (and the regions of space within their horizons) by A and B , and denote their surfaces (sometimes the hole's event horizon and other times the apparent horizon) by ∂A and ∂B , as shown in Fig. A.4. We let \mathcal{E} stand for the region outside both bodies but inside the arbitrarily large surface \mathcal{S} where the system's total momentum is computed (in our case, this is taken to be a fixed coordinate sphere inside the outer boundary of the numerical-relativity computational grid).

⁵After the holes merge, there is only one horizon, which we label ∂C . Equations (A.8)–(A.10) hold after removing terms with subscript B and then substituting $A \rightarrow C$.




Figure A.4: The regions of space around and inside a binary-black-hole system.

With the aid of Gauss's theorem and the Einstein field equations (A.4), one can reexpress Eq. (A.7) for the binary's total 4-momentum as a sum over contributions from each of the bodies and from the gravitational field in the region \mathcal{E} outside them:

$$p_{\text{tot}}^\mu = p_A^\mu + p_B^\mu + p_{\text{field}}^\mu. \quad (\text{A.9a})$$

Here

$$p_A^\mu := \frac{1}{16\pi} \oint_{\partial A} H^{\mu\alpha 0j}{}_{,\alpha} d\Sigma_j \quad (\text{A.9b})$$

is the 4-momentum of body A (an equivalent expression holds for body B), and

$$p_{\text{field}}^\mu := \int_{\mathcal{E}} \tau^{0\mu} d^3x \quad (\text{A.9c})$$

is the gravitational field's 4-momentum in the exterior of the black holes. We define an effective velocity of black hole A (with similar expressions holding for hole B) by

$$v_{\text{LL}}^j := \frac{p_A^j}{p_A^0}. \quad (\text{A.10})$$

In analogy to Eq. (A.8) for the rate of change of the binary's total 4-momentum, one can write the corresponding equation for the rate of change of the 4-momentum of body A :

$$\frac{dp_A^\mu}{dt} = - \oint_{\partial A} (\tau^{\mu k} - \tau^{\mu 0} v_A^k) d\Sigma_k. \quad (\text{A.11})$$

Equation (A.11) describes the flow of field 4-momentum into and out of body A (the second term comes from the motion of the boundary of body A with local coordinate velocity v_A^k).⁶

We will use Eqs. (A.8)–(A.10) as the basis for our study of momentum flow in black-hole binaries. The actual values of the body and field 4-momenta, computed in the above ways, will depend on the arbitrary mapping between the physical spacetime and the AFS; this is the gauge-dependence that will be discussed in Sec. A.5.2.

A.4 Simulations of head-on BBH collisions with anti-aligned spins

In order to investigate the gauge dependence of our results, we compare simulations of the same physical system using two separate methods that employ different choices of coordinates. One method is a pseudospectral excision scheme based on generalized harmonic coordinates; the other is a finite-difference moving-puncture scheme that uses the Baumgarte-Shapiro-Shibata-Nakamura (BSSN) [52, 53] formulation, 1+log slicing, and a gamma-driver shift condition (henceforth referred to as “BSSN-moving-puncture gauge”; for details see Appendix A.9.2). The coordinates used in the two methods differ both for the initial data and during the evolution. In this section we summarize the construction of initial data and the evolution scheme for both methods, and we present convergence tests and estimate numerical uncertainties. Further details about our numerical methods are given in Appendices A.8 and A.9.

A.4.1 Generalized harmonic

Quasiequilibrium excision data

The evolutions described in Sec. A.4.1 begin with quasiequilibrium excision data constructed using the method of Ref. [54]. This method requires the arbitrary choice of a conformal three-metric; we choose this metric to be flat almost everywhere but curved (such that the metric is nearly that of a single Kerr-Schild hole) near the horizons.

Our initial data method also requires us to choose an outer boundary condition on a shift vector β^i ; for a general binary that is orbiting and inspiraling, we use⁷

$$\beta^i = (\mathbf{\Omega}_0 \times r)^i + \dot{a}_0 r^i + V_0^i, \quad r \rightarrow \infty, \quad (\text{A.12})$$

⁶In the case that the body’s event horizon is stationary (i.e. sufficiently far from merger), $v_A^k = dx_{A \text{ cm}}^k/dt$, the center of mass velocity of body A . However, if the body’s event horizon is dynamical (i.e. during the merger phase), then v_A^k is the local coordinate velocity of the event horizon surface, $v_A^k = dx_{\partial A}^k/dt$. See Sec. A.5.1 for a discussion of the dynamics of the event horizon.

⁷The shift vector β^i used here and in Appendix A.8 for the construction of initial data is not the same as the shift vector used during our evolutions. Except for Sec. A.4.1 and Appendix A.8, we always use β^i to refer to the shift *during the evolution*.

Set	x_o/M_{ADM}	$M_{\text{irr}}/M_{\text{ADM}}$	$M_{\text{Chr}}/M_{\text{ADM}}$	S_z/M_{Chr}^2
S1	3.902	0.4986	0.5162	± 0.5000
P1	4.211	0.4970	0.5146	± 0.5000
P2	8.368	0.4802	0.5072	± 0.5091
H1	14.864	0.4870	0.5042	± 0.5000

Table A.1: Parameters of the initial data configurations studied in this work. Model S1 (see Sec. A.4.1) gives the parameters used to construct a set of Superposed-Kerr-Schild quasiequilibrium excision initial data. Model H1 (see Appendix A.8.2) gives the parameters for the larger separation Superposed-Harmonic-Kerr initial data set. Both S1 and H1 were used in generalized harmonic, pseudospectral evolutions. P1 and P2 provide the Bowen-York parameters for the two systems evolved with the BSSN-moving-puncture method. The holes are initially separated by a coordinate distance $d = 2x_0$ and are located at coordinates $(x, y, z) = (\pm x_0, 0, 0)$. For clarity, only 4 significant figures are shown.

where Ω_0 is the angular velocity, $\dot{a}_0 r^i$ is the initial radial velocity, and V_0^i is a translational velocity. Note that Eq (A.12) is different from the choice made in Ref. [54]. In this paper we confine our focus to collisions that are head-on, which we define as $\Omega_0 = \dot{a}_0 = 0$. However, V_0^i must be nonzero to make the total linear momentum of the initial data vanish.

Table A.1 summarizes the initial data used in this paper. The Arnowitt-Deser-Misner (ADM) mass M_{ADM} (Eq. (11.2.14) in Ref. [55]; see also [56, 57]), the irreducible mass M_{irr} and Christodoulou mass M_{Chr} of one of the holes are listed, where M_{Chr} is related to M_{irr} and the spin of the hole S_z by

$$M_{\text{Chr}}^2 = M_{\text{irr}}^2 + \frac{S_z^2}{4M_{\text{irr}}^2}. \quad (\text{A.13})$$

Table A.1 also shows the dimensionless spin S_z/M_{Chr}^2 ; by definition, this measure of the spin lies in the interval $-1 \leq S_z/M_{\text{Chr}}^2 \leq 1$.

For set S1 listed in Table A.1, V_0^i is adjusted so that the initial effective velocity of the entire spacetime $v_{\text{tot}}^i := p_{\text{tot}}^i/p_{\text{tot}}^0$ is smaller than 0.1 km/s, which is approximately the size of our numerical truncation error (cf. Fig. A.9): $(|v_{\text{tot}}^x|, |v_{\text{tot}}^y|, |v_{\text{tot}}^z|) = (4 \times 10^{-4}, 5 \times 10^{-2}, 2 \times 10^{-3})$ km/s at time $t = 0$.

The construction of initial data is described in more detail in Appendix A.8.

Generalized harmonic evolutions

We evolve the quasiequilibrium excision data described in Sec. A.4.1 pseudospectrally, using generalized harmonic gauge [58, 59, 60, 61], for which the coordinates x^μ satisfy the gauge condition

$$g_{\mu\nu} \nabla^\rho \nabla_\rho x^\mu = H_\nu(x^\rho, g_{\sigma\tau}) \quad (\text{A.14})$$

where H_ν is a function of the coordinates and the spacetime metric. In this subsection, we summarize the computational grid used for our generalized harmonic evolutions, and we briefly discuss our numerical accuracy. Details of our pseudospectral evolution methods are given in Appendix A.9.1.

Our computational grid covers only the exterior regions of the black holes (“black hole excision”): there is an artificial inner boundary just inside each apparent horizon. No boundary conditions are needed at these boundaries because of causality; note that the formulation of Einstein’s equations we use [61] admits only causal characteristic speeds, even for gauge modes and constraint-violating modes. The grid extends to a large radius $r_{\max} \sim 400M_{\text{ADM}}$. A set of overlapping subdomains of different shapes (spherical shells near each hole and far away; cylinders elsewhere) covers the entire space between the excision boundaries and $r = r_{\max}$.

Because different subdomains have different shapes and the grid points are not distributed uniformly, we describe the resolution of our grid in terms of the total number of grid points summed over all subdomains. We label our resolutions N0, N1, and N2, corresponding to approximately 55^3 , 67^3 , and 79^3 grid points, respectively. After merger, we regrid onto a new computational domain that has only a single excised region (just inside the newly-formed apparent horizon that encompasses both holes). This new grid has a different resolution (and a different decomposition into subdomains) from the old grid. We label the resolution of the post-merger grid by A , B , and C , corresponding to approximately 63^3 , 75^3 , and 87^3 gridpoints, respectively. We label the entire run using the notation ‘ $Nx.y$ ’, where the characters before and after the decimal point denote the pre-merger and post-merger resolution for that run. Thus, for example, ‘N2. B ’ denotes a run with approximately 67^3 grid points before merger, and 75^3 grid points afterward. On the outermost portion of the grid (farther than $\sim 200M_{\text{ADM}}$), we use a coarser numerical resolution than we do elsewhere. (We only measure the gravitational wave flux, linear momentum, etc., at radii of $r \leq 160M_{\text{ADM}}$.)

To demonstrate the convergence of our evolutions, we plot the constraint violation in Fig. A.5 for several resolutions. The quantity plotted is the L^2 norm of all the constraints of the generalized harmonic system, normalized by the L^2 norm of the spatial gradients of all the dynamical fields, as defined by Eq. (71) of Ref. [61]. The left portion of the plot depicts the constraint violation during the plunge, the right third of the plot shows the constraint violation during the ringdown, and the middle panel shows the constraints shortly before and shortly after the common apparent horizon forms. Throughout the evolution, we generally observe exponential convergence, although the convergence rate is smaller near merger. After merger, there are two sources of constraint violations: those generated by numerical truncation error after merger (these depend on the resolution of the post-merger grid) and those generated by numerical truncation error before merger and are still present in the solution (these depend on the resolution of the pre-merger grid). We see from Fig. A.5 that the constraint violations after merger are dominated by the former source. Also, at about $t = 200M_{\text{ADM}}$, the constraint violation increases noticeably (but is still convergent); at this




Figure A.5: Constraint violation at different numerical resolutions for the generalized harmonic evolutions S1. The common apparent horizon forms at time $t = 34.73M_{\text{ADM}}$. Labels of the form $Nx.y$ indicate the grid resolution, where the pre-merger resolution is labeled (from coarse to fine) by $x = 0, 1, 2$ and the post-merger resolution is labeled by $y = A, B, C$. The constraints decrease exponentially with higher resolution; the convergence rate is smaller near merger.

time, the outgoing gravitational waves have reached the coarser, outermost region of the grid.

Finally, in Fig. A.6, we demonstrate the accuracy of the recoil velocity $v_{\text{kick}} = 22$ km/s inferred from the gravitational wave signal Ψ_4 , which asymptotically is related to the gravitational wave amplitudes h_+ and h_\times by

$$\Psi_4 = \frac{d^2}{dt^2} h_+ - i \frac{d^2}{dt^2} h_\times. \quad (\text{A.15})$$

We extract the spin-weighted spherical harmonic coefficients of $\Psi_4(t)$ from the simulation as described in Ref. [4], and we integrate these coefficients over time to obtain $\dot{h}^{\ell m}(t)$, which are the spin-weighted spherical harmonic coefficients of $\dot{h} = \dot{h}_+ - i\dot{h}_\times$. For each (ℓ, m) , the integration constant is chosen so that the average value of $\dot{h}^{\ell m}(t)$ is zero. The $\dot{h}^{\ell m}(t)$ are then used to compute the 4-momentum flux of the gravitational waves from Eqs. (3.14)–(3.19) of Ref. [62]. Integrating this flux over time yields the total radiated energy-momentum, p_{rad}^μ . The recoil velocity can then be computed from energy-momentum conservation: $v_{\text{kick}}^i = -p_{\text{rad}}^i / M_{\text{final}}$, where $M_{\text{final}} := M_{\text{ADM}} - E_{\text{rad}}$ and E_{rad} is the energy radiated to infinity. For set S1, we obtain a radiated energy of $E_{\text{rad}}/M_{\text{ADM}} = (5.6840 \pm 0.0008) \times 10^{-4}$, where the quoted error includes truncation error




Figure A.6: Recoil velocity for initial data set S1 inferred from the gravitational wave signal Ψ_4 extracted at $r_{\text{extr}} = 160M_{\text{ADM}}$ at the highest resolution (upper panel). Differences between several coarser resolutions and the highest resolution result are plotted in the lower panel.

and uncertainty from extrapolation to infinite radius (as discussed below). The top panel of Fig. A.6 shows the recoil velocity as a function of time for our highest resolution simulation, while the lower panel shows differences between the highest resolution (N2.C) and lower resolutions. From these differences, we estimate a numerical uncertainty for the final recoil velocity of 5×10^{-3} km/s for N1.B and 2×10^{-5} km/s for N2.B.

This numerical uncertainty includes only the effects of numerical truncation error; however, there are other potential sources of uncertainty in the simulations that must also be considered. The first is the spurious “junk” gravitational radiation that arises because the initial data do not describe a perfect equilibrium situation. This radiation is not astrophysically realistic, but by carrying a small amount of energy-momentum that contributes to the measured p_{rad}^μ at large distances, the spurious radiation does affect our determination of the final recoil velocity. In our investigation of momentum flow (Sec. A.5), we do not correct for the initial data’s failure to be in equilibrium; here we estimate the contribution of the resulting spurious radiation to the final recoil velocity. First, we note that for head-on collisions, the physical gravitational waves are emitted predominantly after merger. Therefore, we estimate the influence of the spurious radiation by examining the accumulated recoil velocity at time $t = \Delta t + r$, where r is the radius of the extraction surface and Δt is a cutoff time. Because the holes merge so quickly (because they begin at so small an initial separation), the spurious and physical contributions to the recoil are not clearly distinguishable in

Fig. A.6. Varying Δt between $31.1M_{\text{ADM}}$ and $38.3M_{\text{ADM}}$ (the common event and apparent horizons form at $t = 31.1M_{\text{ADM}}$ and $t = 34.7M_{\text{ADM}}$, respectively), we estimate that the spurious radiation contributes approximately 1 km/s (about 5%) to the recoil velocity—a much larger uncertainty than the truncation error. The same variation of Δt implies that the spurious radiation contributes about 10% of the total radiated energy E_{rad} .)

Another potential source of uncertainty in v_{kick}^i arises from where on the grid we measure the gravitational radiation. In particular, the quantity Ψ_4 in Eq. (A.15) should ideally be measured at future null infinity. Instead, we measure Ψ_4 on a set of coordinate spheres at fixed radii, compute v_{kick}^i on each of these spheres, and extrapolate the final equilibrium value of v_{kick}^i to infinite radius (cf. Fig. A.12). We estimate our uncertainty in the extrapolated value by comparing polynomial extrapolation of orders 1, 2, and 3; we find an uncertainty of 3×10^{-3} km/s for the quadratic fit. Note that if we had not extrapolated to infinity, but had instead simply used the value of v_{kick}^y at our largest extraction sphere ($r = 160M_{\text{ADM}}$), we would have made an error of 0.85 km/s, which is much larger than the uncertainty from numerical truncation error. Finally, we mention that our computation of Ψ_4 is not strictly gauge invariant unless Ψ_4 is evaluated at future null infinity. As long as gauge effects in Ψ_4 fall off faster than $1/r$ as expected, extrapolation of v_{kick}^y to infinity should eliminate this source of uncertainty.

A.4.2 BSSN-moving-puncture

Bowen-York puncture data

In order to address the importance of gauge dependence for our calculations using the Landau-Lifshitz formalism, we also simulate BBH mergers using the so-called moving puncture method, which employs the covariant form of “1+log” slicing [2, 63] for the lapse function α and a “Gamma-driver” condition (based on the original “Gamma-freezing” condition introduced in [64]) for the shift vector. The precise evolution equations for the gauge variables as well as further technical details of our puncture simulations are given in Appendix A.9.2.

Our simulations start with puncture initial data [65] provided in our case by the spectral solver of Ref. [66]. The initial data are fully specified in terms of the initial spin $\vec{S}_{1,2}$, linear momentum $\vec{P}_{1,2}$ and initial coordinate position $\vec{x}_{1,2}$ as well as the bare mass parameters $m_{1,2}$ of either hole [67]. In order to assess the impact of the initial binary separation, we consider two models as specified in Table A.1. There we also list the total black-hole mass M_{Chr} and normalize all quantities using the total ADM mass M_{ADM} . The main difference between the two configurations is the initial separation of the holes. The lapse and shift are initialized as $\alpha = \gamma^{-1/6}$ and $\beta^i = 0$, where γ is the determinant of the physical three-metric.

BSSN-moving-puncture evolutions

The evolution of the puncture initial data is performed using sixth order spatial discretization of the BSSN equations combined with a fourth order Runge-Kutta time integration. Mesh refinement of Berger-Oliger [68] type is implemented using Schnetter’s CARPET package [69, 70]. The prolongation operator is of fifth order in space and quadratic in time. Outgoing radiation boundary conditions are implemented using second-order accurate advection derivatives (see, for example, Sec. VI in Ref. [71]).

Using the notation of Sec. II E of Ref. [72] the grid setup in units of M_{ADM} for these evolutions is given by (rounded to 3 significant digits)

$$\begin{aligned} &\{(202, 101, 58.8, 25.2, 12.6) \times (3.15, 1.58, 0.788), h\}, \\ &\{(201, 100, 58.5, 25.1) \times (6.27, 3.13, 1.57, 0.784), h\}, \end{aligned}$$

respectively. Here h denotes the resolution on the innermost refinement level. For model P1 we perform a convergence analysis by setting h to $h_c = M_{\text{ADM}}/49.5$, $h_m = M_{\text{ADM}}/57.1$ and $h_f = M_{\text{ADM}}/64.7$, respectively, for coarse, medium and fine resolution. Model P2 is evolved using $h = M_{\text{ADM}}/49.8$.

Before we discuss the physical results from the BSSN-moving-puncture simulations, we estimate the numerical errors due to discretization, finite extraction radius and the presence of unphysical gravitational radiation in the initial data.

In order to study the dependence of the results on resolution, we have evolved model P1 of Table A.1 using different resolutions h_c , h_m and h_f on the finest level and correspondingly larger grid spacings by a factor of two on each consecutive level. Numerical simulations based on finite differencing techniques incur a numerical error of polynomial dependence on the grid resolution h because derivatives in the differential equations are discretized via Taylor expansion. A numerical result f_h will therefore differ from the continuum limit f by a discretization error $e(h) := f_h - f = \text{const} \times h^n + \dots$ where n is the order of convergence and the dots denote higher order terms. In our case, the lowest order ingredient in the code arises in the prolongation in time which is second order accurate. The consistency of the code can then be tested by calculating the order of convergence according to

$$Q_n := \frac{f_{h_c} - f_{h_m}}{f_{h_m} - f_{h_f}}, \tag{A.16}$$

where $f(h_c)$, $f(h_m)$ and $f(h_f)$ denote the numerical solution at coarse, medium and fine resolution, respectively. Inserting the above mentioned error function $e(h)$ and ignoring higher order terms, the




Figure A.7: Gravitational recoil for model P1 as estimated from the gravitational wave signal Ψ_4 extracted at $r_{\text{ex}} = 73.5 M_{\text{ADM}}$ using the highest resolution (upper panel). Differences in the recoil obtained at coarse, medium and fine resolution rescaled for second order convergence (lower panel).

expected convergence factor for a scheme of n^{th} order accuracy is

$$Q_n = \frac{h_c^n - h_m^n}{h_m^n - h_f^n}. \quad (\text{A.17})$$

The kick velocity from the high resolution simulation, as inferred from the gravitational radiation flux at $r_{\text{ex}} = 73.5 M_{\text{ADM}}$, is shown in the upper panel of Fig. A.7. The bottom panel shows the differences between the velocities obtained at the different resolutions scaled for second order convergence using a factor $Q_2 = 1.49$. By using Richardson extrapolation we estimate the error in the final kick for the fine resolution run to be 1 km/s or 5 %. We similarly find overall second order convergence for the velocity derived from the components of the Landau-Lifshitz tensor as integrated over the apparent horizon. The error in that quantity barely varies throughout the entire simulation and stays at a level just below $\Delta v_{\text{LL}} \approx 50$ km/s and 60 km/s for fine and coarse resolution respectively.

The gravitational wave signal is further affected by the use of finite extraction radius and linear momentum contained in the spurious initial radiation. We estimate the uncertainty due to the finite extraction radius by fitting the final kick velocity obtained for the medium resolution simulation of model P1 at radii $r_{\text{ex}} = 31.5 \dots 94.5 M_{\text{ADM}}$ in steps of $10.5 M_{\text{ADM}}$. The resulting final kick velocities

are well approximated by a polynomial of the form $a_0 + a_1/r_{\text{ex}} + a_2/r_{\text{ex}}^2$. For $r_{\text{ex}} = 73.5 M$ we thus obtain an uncertainty of 0.4 km/s corresponding to a relative error of 2.2 %.

Finally we take into account contributions from the spurious initial radiation by discarding the wave signal up to $t - r_{\text{ex}} = \Delta t$. For model P1 it is not entirely clear where exactly the spurious wave signal stops and the physical signal starts. By varying Δt from 30 to 45 M_{ADM} we obtain an additional error of about ± 1 km/s. For model P2 no such problem arises because of the smaller amplitude of the spurious radiation and because the longer pre-merger time enables the junk radiation to escape the system long before the merger happens. We estimate the resulting total uncertainty by summing the squares of the individual errors and obtain 7.5 % and 5.5 % for models P1 and P2, respectively.

Using these uncertainties, the gravitational wave emission for model P1 results in a total radiated energy of $E_{\text{rad}}/M_{\text{ADM}} = (0.042 \pm 0.008)$ % and a recoil velocity $v_{\text{kick}} = (20.3 \pm 1.5)$ km/s. For model P2 the result is $E_{\text{rad}}/M_{\text{ADM}} = (0.0555 \pm 0.0023)$ % and $v_{\text{kick}} = (19.7 \pm 1.1)$ km/s.

A.5 Momentum flow

In this section, we turn to the momentum flow during the evolutions described in Sec. A.4. First, in Sec. A.5.1 we measure the momentum of the holes during plunge, merger, and ringdown during a generalized harmonic evolution of initial data set S1 (Table A.1), focusing on the momentum density and the inferred Landau-Lifshitz velocity v_{LL}^y along and opposite the frame-dragging direction (which in this paper are chosen to be the $\mp y$ direction, respectively). In Sec. A.5.2, we look at the momentum flow in a BSSN-moving-puncture simulation with similar initial data, and by comparing the BSSN-moving-puncture and generalized harmonic simulations, we investigate the influence of the choice of gauge on our results. Then, in Sec. A.5.3 we compare the momentum density and velocity of the holes with post-Newtonian predictions.

A.5.1 Generalized harmonic results

Throughout the generalized harmonic evolutions summarized in Sec. A.4.1, we measure the 4-momentum density by explicitly computing the Landau-Lifshitz pseudotensor [Eq. (A.5)]. Because our evolution variables are essentially the spacetime metric $g_{\mu\nu}$ and its first derivative $g_{\mu\nu,\rho}$, we are able to compute the momentum density without taking any additional numerical derivatives. Besides measuring the momentum density, we also measure the 4-momentum p_A^μ [Eq. (A.9b)] enclosed by i) the apparent horizons, ii) the event horizon, and iii) several spheres of large radius. From the enclosed momentum, we evaluate the effective velocity v_{LL}^j [Eq. (A.10)].




Figure A.8: The velocity of the individual and merged black holes. The Landau-Lifshitz velocity $v_{LL}^y := p_{LL}^y / p_{LL}^t$, where p_{LL}^μ is the Landau-Lifshitz 4-momentum enclosed, is measured on the individual and common apparent horizons (labeled AH and AHC, respectively) and also on the event horizon (labeled EH). For comparison, the coordinate velocities v_{coord}^y of the apparent horizons are also shown. The data shown are from the high-resolution evolution N2.C.




Figure A.9: Convergence of v_{LL}^y with resolution. Specifically, differences between v_{LL}^y at the highest resolution N2.C and at various lower resolutions are shown. Labels of the form Nx.y indicate the grid resolution, where the pre-merger resolution is labeled (from coarse to fine) by $x = 0, 1, 2$ and the post-merger resolution is labeled by $y = A, B, C$. The difference between the second-highest and highest resolution is below 0.1 km/s except near merger, when it grows as large as 1 km/s.

Apparent horizons

The effective velocities of the apparent horizons are shown in Fig. A.8 (dashed curves). To demonstrate convergence, Fig. A.9 shows the differences between apparent-horizon effective velocities computed at different resolutions. During the plunge, the difference between the medium and fine resolution is less than 0.1 km/s until shortly before merger, when it reaches a few tenths of a km/s. Shortly after merger, the difference between the highest and medium continuation resolutions between N2.B and N2.C falls from about 1 km/s to about 0.1 km/s.

For comparison, Fig. A.8 also shows the apparent horizons' coordinate velocities (dotted curves); the coordinate and effective velocities agree qualitatively during the plunge and quantitatively during the merger. There is no reason to expect this observed agreement *a priori*; this is one sense in which our gauge choice appears to be “reasonable.” Also, Fig. A.8 shows that the effective velocities of individual apparent horizons and the event horizon agree well until shortly before merger, when the event horizon's velocity smoothly transitions to agree with the common apparent horizon's (cf. Sec. A.5.1 below).

Because of frame-dragging, during the plunge the individual apparent horizons accelerate in the downward ($-y$) direction, eventually reaching velocities of thousands of km/s. But when the common apparent horizon appears, its velocity is much closer to zero and quickly changes sign, eventually reaching speeds of about 1000 km/s in the $+y$ direction (i.e., in the direction *opposite* the frame-dragging direction). Then, as the common horizon rings down, the velocity relaxes to a final kick velocity of about 20 km/s in the $+y$ direction.

After merger, why have the horizon velocities suddenly changed from thousands of km/s in the frame-dragging direction to over a thousand km/s in the opposite direction? The answer can be seen in Fig. A.10, which plots contours of constant y -momentum density at several times. At $t = 0$, the momentum density has an irregular shape, because the initial data is initially not in equilibrium. By time $t = 26.92M_{\text{ADM}}$, the momentum density has relaxed. When the common apparent horizon forms (at time $t = 34.73M_{\text{ADM}}$), it encloses not only the momentum of the individual apparent horizons but also the momentum in the gravitational field between the holes.

It turns out that the net momentum outside the individual horizon but inside the common horizon points in the $+y$ direction; as the common horizon expands, it absorbs more and more of this upward momentum. Fig. A.11 compares the common apparent horizon's effective velocity to its area and shape; the latter is indicated by the pointwise maximum and minimum of the horizon's intrinsic scalar curvature. During the first half-period of oscillation (to the left of the leftmost dashed vertical line), the common horizon expands (as seen by its increasing area); as it expands, the upward-pointing linear momentum it encloses causes v_{LL}^y to increase. After the first half-period, the horizon shape is maximally oblate (cf. panel B on the right side of Fig. A.11), and v_{LL}^y is at its maximum value of about 1000 km/s.




Figure A.10: Contour plots of the y (up-down) component of the momentum density, which points along or opposite of the holes' motion due to frame dragging. Adjacent contours correspond to a factor of 10 difference in the magnitude of the momentum density. Contours of positive y momentum density are shown as solid red lines, while contours of negative y momentum density are shown as dashed blue lines. The region containing positive y momentum density is shaded grey. The regions inside the apparent horizons are shaded black, except for the upper right panel, where the region inside the *individual* horizons is shaded black, while the common apparent horizon is indicated by a thick black line. The data shown are from the high-resolution evolution N2.C.




Figure A.11: *Left:* A comparison of the common apparent horizon's effective velocity and the horizon's shape and area. The top panel shows the horizon's effective velocity v_{LL}^y . The middle panel shows the pointwise minimum and maximum of the horizon's dimensionless intrinsic scalar curvature; both $M_{\text{Chr}}^2 R_{\text{min}}(R)$ and $M_{\text{Chr}}^2 R_{\text{max}}(R)$ relax to the Schwarzschild value of $1/2$ as the horizon rings down. (The first four local minima of $M_{\text{Chr}}^2 R_{\text{min}}(R)$ are indicated by vertical dashed lines.) The bottom panel shows the area A of the common apparent horizon normalized by the total area of the individual horizons at $t = 0$. The data shown are from the high-resolution evolution N2.C. *Right:* The dimensionless intrinsic scalar curvature $M_{\text{Chr}}^2 R$ of the common apparent horizon at the times labeled A–F in the left panel. The horizon begins peanut-shaped, then rings down, eventually settling down to a sphere with a constant curvature $M_{\text{Chr}}^2 R = 0.5$.

After another half-period of oscillation, the apparent horizon becomes prolate and encloses enough downward-pointing momentum that v_{LL}^y has decreased to only about +200 km/s. After one additional full period, the effective velocity has fallen to nearly zero. As the horizon is ringing down, the momentum density in the surrounding gravitational field also oscillates: the final four panels in Fig. A.10 show how the momentum density relaxes to a final state as the horizon relaxes to that of a boosted Schwarzschild black hole.

As the horizon rings down, gravitational waves are emitted, and these waves carry off a small amount of linear momentum. The net radiated momentum is only a small fraction of the momenta of the individual holes at the time of merger: the final effective velocity of the merged hole is about 20 km/s in the upward-pointing direction, or about 1% of the individual holes' downward velocity just before merger.

Various measures of the final velocity of the merged hole are shown in Fig. A.12. The kick velocity v_{kick}^y , which is inferred from the outgoing gravitational waves, is measured on four coordinate spheres (with radii R of $100M_{ADM}$, $120M_{ADM}$, $140M_{ADM}$, and $160M_{ADM}$); the effective velocity is




Figure A.12: A comparison of various measures of the final velocity of the merged hole in simulation S1. The kick velocity inferred from the gravitational-wave flux (thin dashed lines) and the Landau-Lifshitz effective velocities v_{LL}^y (thin solid lines) are measured on spheres of radius $100M_{ADM}$, $120M_{ADM}$, $140M_{ADM}$, and $160M_{ADM}$. The value of the kick velocity at the final time is extrapolated to $r = \infty$ (black cross). The effective velocity measured on the common apparent horizon (thick solid line) and the coordinate velocity (thick dashed line) are also shown. The data shown are from the high-resolution evolution N2.C.




Figure A.13: The effective velocity v_{LL}^y calculated on the event horizon surface, with the specified snapshots in Fig. A.14 of the event horizon surface marked: a,b, $t = 27.7M_{\text{ADM}}$; c, $t = 30.8M_{\text{ADM}}$; d, $t = 31.6M_{\text{ADM}}$; e, $t = 35.5M_{\text{ADM}}$; f, $t = 40.8M_{\text{ADM}}$.

measured on the same coordinate spheres. At late times, we find that the effective velocity v_{LL}^y has no significant dependence on the radius of the extraction surface: it simply approaches the coordinate velocity v_{coord}^y of the common apparent horizon. The dependence of v_{kick}^y on the extraction radius is expected, since our method of extracting Ψ_4 at finite radius has gauge-dependent contributions that vanish as $R \rightarrow \infty$. When v_{kick}^y is extrapolated to infinite radius⁸, however, it does agree well (within 0.2 km/s) with v_{LL}^y . Also, the effective velocity v_{LL}^y calculated on the horizon also agrees fairly well (within about 0.5 km/s) with v_{LL}^y measured on distant spheres.

Event horizon

We would like to compare our quantitative results of the effective velocity v_{LL}^y calculated using the event horizon surface (Fig. A.13) with qualitative observations of the event horizon's dynamics (Fig. A.14). We find that the greatest variation in *both* the event horizon geometry *and* the value of v_{LL}^y occurs over a period of about $\Delta t = 13M_{\text{ADM}}$ from $t = 28M_{\text{ADM}}$ to $t = 41M_{\text{ADM}}$. At time $t = 27.7M_{\text{ADM}}$, the cusps of the event horizon just begin to become noticeable (Figs. A.14 a & b). One can see in Fig. A.13 that this is the time at which v_{LL}^y changes from decreasing to increasing. Shortly after⁹, at $t = 31.1M_{\text{ADM}}$, the two separate event horizons coalesce into a common event horizon, and the common event horizon rapidly expands to form a convex shape by $t = 35.5M_{\text{ADM}}$

⁸To extrapolate, we fit the velocities v_{kick}^y at the final time to a function of radius R of the form $a_0 + a_1/R + a_2/R^2$.

⁹Note that at $t = 31.1M_{\text{ADM}}$, we (smoothly) modify our gauge condition [Eq. (A.43) and the surrounding discussion]. The separate event horizons coalesce at time $t = 31.1M_{\text{ADM}}$ as well; this is a coincidence.




Figure A.14: Snapshots of the event horizons at the times indicated in Fig. A.13: a,b, $t = 27.7M_{\text{ADM}}$; c, $t = 30.8M_{\text{ADM}}$; d, $t = 31.6M_{\text{ADM}}$; e, $t = 35.5M_{\text{ADM}}$; f, $t = 40.8M_{\text{ADM}}$. All snapshots are looking down the z-axis to the x-y plane, except for shot a, which is slightly skewed (slightly rotated about the y axis) to better see the geodesic structure. In shot a, the future generators of the horizon are visible as small blue dots. Note how the future generators map out a surface that meets the event horizon at the event horizon's cusps; this is where the future generators join the horizon. The data shown are from the high-resolution evolution N2.C.

(Figs. A.14 d & e). At this time, we note that v_{LL}^y is rapidly increasing (Fig. A.13, arrow e); this rapid increase corresponds to the quickly expanding event horizon surface.

We interpret this process as the merging black holes “swallowing” the gravitational field momentum between the holes. The resulting change in v_{LL}^y can be divided into two distinct portions: i) one that results from the changing event horizon surface in space, i.e. the field momentum swallowed by the black holes [mathematically, the *second* term, in Eq. (A.11)] and ii) a second that results from the change of field momentum at the black holes’ surface, i.e. the field momentum *flowing* into the black holes [mathematically, the *first* term, in Eq. (A.11)]. While this distinction is clearly coordinate dependent, it could, after further investigation, nevertheless provide an intriguing and intuitive picture of the near-zone dynamics of merging black hole binaries.

A.5.2 BSSN-moving-puncture results and gauge

As summarized in Sec. A.3, the Landau-Lifshitz formalism that we have applied to our numerical simulations is based on a mapping between the curved spacetime of the simulation and an auxiliary flat spacetime. In the asymptotically-flat region far from the holes, there is a preferred way to construct this mapping. Consequently, when the surface of integration is a sphere approaching infinite radius, Eq. (A.9b) gives a *gauge-invariant* measure of the system’s total 4-momentum (see, e.g., Sec. 20.3 of Ref. [50]). However, when the surface of integration is in the strong-field region of the spacetime (e.g., when the surface is a horizon), the 4-momentum enclosed is gauge dependent. The momentum density, being given by a pseudotensor, is always gauge dependent.

The gauge-dependence of the effective velocity can be investigated at late times—when the spacetime has relaxed to its final, stationary configuration—by comparing the velocity obtained on the horizon with gauge-invariant measures of the kick velocity (Fig. A.12). At the final time in our generalized harmonic simulation, the effective velocities of the apparent and event horizons agree within tenths of a km/s with the (extrapolated) kick velocity inferred from the gravitational-wave flux; at late times, the horizon effective velocities also agree with the effective velocity measured on coordinate spheres of large radius. At least at late times, then, the effective velocity v_{LL}^y is not significantly affected by our choice of gauge.

But how strong is the influence of gauge on our results in the highly-dynamical portion of the evolution, when we have no gauge-invariant measure of momentum or velocity? To investigate this, we have evolved initial data that are physically similar using two manifestly different gauge conditions: i) the generalized harmonic condition used in our spectral evolutions, and ii) the “1+log” slicing and “Gamma-driver” shift conditions used in our BSSN-moving-puncture evolutions.

Figs. A.15 and A.16 display the velocity obtained from the horizon integral of the components of the Landau-Lifshitz tensor in the BSSN-moving-puncture evolutions described in Sec. A.4.2. For comparison we also plot the velocity obtained for model S1 in the left panel (dashed curve). The most




Figure A.15: Velocity perpendicular to the line of sight associated with the horizon integrals of the the Landau-Lifshitz tensor obtained for models P1 and S1. The shaded area represents the numerical uncertainty for P1. During the pre-merger phase, in each simulation the velocities of the individual horizons are identical.

remarkable feature in these plots is a large temporary acceleration of the black holes in the frame-dragging direction. The magnitude of the velocity reaches about 4500 km/s, which is of the order of the superkicks first reported in Refs. [32, 30]. In contrast to those inspiraling configurations, however, the black hole motion reverses during the merger and settles down to a small value of -30 ± 50 km/s.

In order to examine to what extent this behavior is dependent on specific properties of the puncture evolution (such as the particular form of the spurious radiation, which differs in our generalized harmonic and BSSN-moving-puncture evolutions), we have performed the following additional simulations. First, we have changed the gauge parameter η in Eq. (A.52) to 0.75 and 1.25. We do not observe a significant change in the behavior of the effective velocity for this modification.

Second, in order to gain further insight into the dependence of the effective velocity on the initial separation of the black holes, we have increased the initial separation of the holes to allow for a longer pre-merger interaction phase; We study the evolution of the second model P2 in Table A.1. This simulation has been performed with the LEAN code as summarized in Sec. A.4.2 using a resolution $h_c = M_{ADM}/49.8$. The resulting velocity is shown in Fig. A.16 and represents numerical uncertainties as gray shading. The remarkable similarity between the figure and its counterpart Fig. A.15 for model P1 demonstrates that the numerical results are essentially independent of the initial separation.




Figure A.16: Same as Fig. A.15 for model P2 of Table A.1.

Comparing the effective velocities for simulations S1 and P1 in Fig. A.15, the qualitative behavior of the apparent horizons' effective velocities agrees. In both the generalized harmonic and BSSN-moving-puncture simulations:

1. during the plunge, the individual apparent horizons accelerate to speeds larger than 1000 km/s in the frame dragging direction,
2. when the common horizon forms, its velocity is much smaller in magnitude, because the common horizon has enclosed momentum pointing opposite the frame-dragging direction, and
3. the velocity relaxes to a value of only tens of km/s that (within numerical uncertainty) agrees with the kick velocity measured using the gravitational-wave flux.

These results are particularly encouraging because two popular gauge choices used in the NR community give remarkable overall agreement. While this qualitative agreement certainly does not constitute a proof of a gauge independence of our findings, we feel encouraged in our hope that the gauge dependence in practice is not too severe, at least for the set of gauges actually used in numerical simulations. Most importantly from a practical point of view, it appears possible that such local descriptions can be derived from the current generation of BBH codes without the different numerical relativity groups having to agree upon one and the same gauge choice for comparing their momentum densities and effective velocities. Future investigations using a wider class of coordinate conditions should further clarify the significance of gauge choices in this context.

A.5.3 Comparison with post-Newtonian predictions

In this section we compare our results to post-Newtonian predictions. For each comparison, first the S1 data set (Table A.1) is presented along with post-Newtonian predictions of a corresponding initial configuration, then the H1 data set (Table A.1) is presented along with its post-Newtonian predictions. The post-Newtonian trajectories for spinning point particles were generated by evolving the post-Newtonian equations of motion [73, 74]. The difference between the two data sets are: i) set H1 begins with a larger initial separation than set S1, and ii) set H1 is evolved in a nearly harmonic gauge. Comparing evolutions of data sets S1 and H1 illustrates how these two effects improve the comparisons one can make with post-Newtonian predictions.

The left panels of Figs. A.17–A.19 show the comparison between the highest-resolution evolution (N2.C) of initial data set S1 and several orders of post-Newtonian predictions. The right panels of Figs. A.17–A.19 show analogous comparisons with an evolution of initial data set H1.

Figure A.17 shows that the bulk, longitudinal motions (i.e., motion in the x direction) agree both qualitatively and quantitatively with post-Newtonian predictions through most of the plunge (i.e., a few M_{ADM} before the formation of the common apparent horizon) for both data sets. In the left panel of Fig. A.17, we have added another 2.5 PN curve that is offset vertically such that the 2.5 PN coordinate velocity agrees exactly with the numerical effective velocity at $t \approx 18.34M_{\text{ADM}}$; this is done in order to account for the period of initial relaxation in the S1 data set. Quantitative agreement is then found between 2.5 PN predictions and both the effective and coordinate velocities from $t \approx 5M_{\text{ADM}}$ through $t \approx 20M_{\text{ADM}}$. The right panel of Fig. A.17, which has less of an initial relaxation due to the increased separation, shows excellent agreement between both the effective and coordinate velocities and the 2.0 PN and 2.5 PN predictions.

For the minor (yet more interesting) transverse motion (i.e., the motion along the y direction), we find only qualitative agreement between the numerical data and post-Newtonian predictions—spin-orbit coupling [more specifically, frame-dragging plus spin-curvature coupling, see Eq. (5.11) of paper I and discussions thereafter] cause the holes to move in the $-y$ direction during the plunge, reaching speeds of order 1000 km/s before the holes merge. The post-Newtonian expansion scheme we adopt (paper I and Refs. [73, 74]) uses a harmonic gauge and a *physical* spin supplementary condition (SSC) of $S^{\alpha\beta}u_\beta = 0$, where $S^{\alpha\beta}$ is the spin angular momentum tensor of the black hole and u^β its four velocity (see e.g., Sec. II B of paper I).

Specifically, in paper I, the authors found that for an equal-mass binary with anti-aligned spins at leading 1.5PN order, the black holes' effective velocity v_{LL}^y is not equal to the post-Newtonian coordinate velocity of a point particle; rather, the coordinate velocity is 3/2 times the effective velocity. Roughly speaking, this difference arises from the fact that in the Landau-Lifshitz description one defines the momentum in terms of a surface integral over a body of finite size. This introduces effects due to the field momentum within the body that are not present in a point-particle description.




Figure A.17: A comparison of numerical and post-Newtonian longitudinal velocities (i.e., v^x/c) versus time. The predicted coordinate velocities at several post-Newtonian orders are shown as broken curves. *Left:* A comparison of S1 numerical data and post-Newtonian predictions. The numerical and post-Newtonian curves agree qualitatively. When the 2.5 PN curve is offset by a certain amount, it agrees quantitatively with the coordinate velocity v_{coord}^x and the effective velocity v_{LL}^x . *Right:* A comparison of H1 numerical data and PN predictions. The effective velocity v_{LL}^x (thick black line) closely tracks the coordinate velocity v_{coord}^x ; both numerical curves also agree well with the 2.0 PN and 2.5 PN curves.




Figure A.18: A comparison of numerical and post-Newtonian transverse velocities (i.e., v^y in km/s) versus time. The left panel shows numerical results from simulation S1, while the right panel shows numerical results from simulation H1. The predicted coordinate velocity at several post-Newtonian orders are shown as broken curves. The effective velocity is shown in black; it has been rescaled by a factor of $3/2$ in order to aid comparison with the post-Newtonian point-particle velocities, as discussed in Sec. A.5.3. The turn around in the 2.5 PN curves is due to a 2.5 order spin-orbit term becoming quite large at a separation of roughly $2M_{\text{ADM}}$. One can argue this is due to the post-Newtonian approximation breaking down at this small separation.




Figure A.19: A comparison of numerical and post-Newtonian velocities. In the figure, v^y in km/s is plotted against v^x/c . The effective velocity v_{LL}^y of the highest-resolution (N2.C) evolution of initial data S1 (Table A.1) on the left and of the evolution of initial data H1 (Table A.1) on the right are shown as a thick black line. The predicted coordinate velocity at several post-Newtonian orders are shown as broken curves. The transverse effective velocities only agree qualitatively with post-Newtonian predictions; however, the coordinate velocity agrees very well with post-Newtonian predictions. In the left panel, the coordinate velocity has been artificially truncated shortly before merger because at that point we do not have a good measure of the coordinate velocity. The effective velocity has been rescaled by a factor of $3/2$ to aid comparison with the post-Newtonian point-particle velocities, as discussed in Sec. A.5.3.

Paper I’s Secs. II B, II C, and V C, as well as its Table I explain this fact in greater detail. Because the majority of the comparison between post-Newtonian and numerical-relativity results takes place at separations and speeds during which the leading, 1.5PN-order terms contribute most strongly, we continue to use the factor of 3/2 to convert between coordinate and effective velocities for higher post-Newtonian terms.

In Figs. A.18 and A.19, we compare the post-Newtonian point-particle y velocity with the numerical coordinate y velocity and 3/2 of the numerical effective y velocity v_{LL}^y . For the comparison to the S1 data set, we find qualitative agreement with both the effective and coordinate velocities and the post-Newtonian predictions. We think this agreement is not better because of the large initial relaxations present in the S1 data set related to small initial separation. The small separation of the black holes also poses problems for the post-Newtonian approximation. As one can see, in Fig. A.18 the 2.5 PN curve decelerates and the velocity changes sign. This happens because a next-to-leading-order, spin-orbit term becomes significantly larger at this point (a post-Newtonian separation of roughly $2M_{\text{ADM}}$). This suggests that the post-Newtonian approximation is moving out of its domain of convergence. However, in the H1 comparison, we find excellent agreement between the coordinate velocity and the 2.5 PN prediction but only qualitative agreement between the effective velocity and post-Newtonian predictions. In these figures, offsets of -433 km/s (for S1 data) and -38 km/s (for H1 data) have been used to make 2.5 PN coordinate velocity agree better with numerical results. Such offsets can be motivated as follows. Our numerical initial data were chosen such that the initial total momentum of the entire spacetime vanishes. This, in our post-Newtonian scheme, corresponds to nonvanishing initial y velocities of (see Table I of paper I)

$$v_{\text{coord}}^y = \frac{\chi}{4(r_0/M_{\text{ADM}})^2}, \quad (\text{A.18})$$

where χ is the spin parameter of each hole, and r_0 their initial separation. This corresponds to -616 km/s for the S1 data, and -42 km/s for H1 data. Again, the agreement is qualitative for S1 data, and quantitative for H1 data.

One final comparison we make between the H1 data set and post-Newtonian predictions is the near-field momentum density, shown in Fig. A.20. The numerical data comes from the harmonic evolution H1, while the 1.5 PN momentum density is computed from Eqs. (A2a)-(A2c) in paper I using the numerical hole trajectories. The left panels, comparing the initial data to the predicted post-Newtonian momentum density, show differences which are presumably due to differences in the post-Newtonian and numerical initial data, such as the numerical initial data being out of equilibrium. The center panels show the momentum densities agree very well once enough time has elapsed for the spacetime to relax and for the spurious radiation to be emitted but before the holes have fallen too close together. The right panels make a final comparison just before the holes get




Figure A.20: Comparison of numerical (top row) and post-Newtonian (bottom row) y momentum density. The numerical data comes from the harmonic evolution H1 described in Appendix A.8.2. The 1.5 PN momentum density is computed from Eqs. (A2a)–(A2c) in paper I using the numerical hole trajectories. As in Fig. A.10, contours represent powers of 10 in y momentum density. The positive y momentum density contours are shown in red, negative in blue. The region of positive y momentum density is shaded grey. In the numerical plots the apparent horizons are shown in black.

close enough to merge and shows differences appearing between the numerical data and the post-Newtonian predictions very near the holes—which could be an indication of the breakdown of the post-Newtonian approximation.

These comparisons with post-Newtonian predictions have yielded several interesting results. The primary result of these comparisons is the surprisingly good agreement found between post-Newtonian predictions and the coordinate velocities, especially from the harmonic gauge evolution. Also, the longitudinal effective and coordinate velocities track each other; consequently, the longitudinal effective velocity agrees with post-Newtonian predictions. The transverse effective velocities agree qualitatively with the post-Newtonian predictions in the sense that they both indicate that the holes accelerate in the expected frame-dragging direction to speeds of order 1000 km/s. Finally, we have also found the qualitative agreement between harmonic gauge numerical data and post-Newtonian extends to the near-zone momentum density after the initial data relaxes but before the holes have fallen too close together.

A.6 Conclusion

With the goal of building up greater physical intuition, we have used the Landau-Lifshitz momentum-flow formalism to explore the nonlinear dynamics of fully relativistic simulations of a head-on BBH plunge, merger, and ringdown. We have defined and computed an effective velocity of the black holes in terms of the momentum and mass-energy enclosed by their horizons, and we have interpreted the holes' transverse motion—which reaches speeds of order 1000 km/s—as a result of momentum flow between the holes and the gravitational field of the surrounding spacetime. We have found that the merged hole's final effective velocity—about 20 km/s—agrees with the recoil velocity implied by the momentum carried off by the emitted gravitational waves.

Our measures of linear momentum and effective velocity are gauge dependent. Nonetheless, after comparing simulations of comparable initial data in generalized harmonic and BSSN-moving-puncture gauges, we have observed remarkably weak gauge dependence for the generalized harmonic and BSSN-moving-puncture evolutions discussed in this paper. Additionally, we have found surprisingly good agreement between the holes' effective and coordinate velocities, and at late times, the holes' final effective velocities and gauge-invariant measures of the kick velocity agree.

These results motivate future explorations of momentum flow in fully-relativistic numerical simulations that are more astrophysically realistic. We are particularly eager to investigate simulations of superkick BBH mergers (the inspiral of a superkick configuration was considered using the post-Newtonian approximation in paper I). Other future work includes studies of the linear and angular momentum flow in inspiraling (rather than head-on) mergers as well as mergers with larger spins.

A.7 Acknowledgments

We are pleased to acknowledge Michael Boyle, Jeandrew Brink, Lawrence Kidder, Robert Owen, Harald Pfeiffer, Saul Teukolsky, and Kip Thorne for helpful discussions. This work was supported in part by the Sherman Fairchild Foundation, the Brinson Foundation, the David and Barbara Groce Fund at Caltech, NSF grants PHY-0652952, DMS-0553677, PHY-0652929, PHY-0601459, PHY-0652995, PHY-0653653, DMS-0553302 and NASA grants NNX09AF96G and NNX09AF97G. Some calculations were done on the Ranger cluster under NSF TeraGrid grant PHY-090003.

A.8 Appendix A: Excision initial data

A.8.1 Superposed-Kerr-Schild (SKS) initial data

The initial data for the generalized harmonic simulations presented in this paper was constructed using the methods described in Ref. [54]. In this appendix, we describe in more detail these initial data (which we summarize in Sec. A.4.1).

The usual 3+1 decomposition splits the spacetime metric $g_{\mu\nu}$ into a spatial metric γ_{ij} , lapse α , and shift β^i , i.e.

$$\begin{aligned} ds^2 &= g_{\mu\nu} dx^\mu dx^\nu = -\alpha^2 dt^2 \\ &+ \gamma_{ij} (dx^i + \beta^i dt)(dx^j + \beta^j dt). \end{aligned} \quad (\text{A.19})$$

On the initial spatial slice (at time $t = 0$), the initial data must specify the spatial metric γ_{ij} and the extrinsic curvature K_{ij} , which is related to the time derivative of the spatial metric by

$$\partial_t \gamma_{ij} = -2\alpha K_{ij} + 2\nabla_{(i} \beta_{j)}. \quad (\text{A.20})$$

We use the quasiequilibrium formalism [75, 76, 77, 78, 79], in which γ_{ij} and K_{ij} are expanded as

$$\begin{aligned} \gamma_{ij} &= \psi^4 \tilde{\gamma}_{ij}, \\ K_{ij} &= A_{ij} + \frac{1}{3} \gamma_{ij} K. \end{aligned} \quad (\text{A.21})$$

The conformal metric $\tilde{\gamma}_{ij}$, the trace of the extrinsic curvature K , and their time derivatives can be chosen freely. We adopt the quasiequilibrium choices

$$\begin{aligned} \tilde{u}_{ij} &:= \partial_t \tilde{\gamma}_{ij} = 0, \\ \partial_t K &= 0. \end{aligned} \quad (\text{A.22})$$

The remaining free data are based on a weighted superposition of two boosted, spinning Kerr-Schild black holes (Eqs. (45)–(46) of Ref. [54]):

$$\tilde{\gamma}_{ij} := f_{ij} + \sum_{a=1}^2 e^{-r_a^2/w_a^2} (\gamma_{ij}^a - f_{ij}), \quad (\text{A.23})$$

$$K := \sum_{a=1}^2 e^{-r_a^2/w_a^2} K_{a..} \quad (\text{A.24})$$

Here f_{ij} is the metric of flat space, r_a is the Euclidean distance from the center of the apparent horizon of hole a , and γ_{ij}^a and K_a are the spatial metric and mean curvature of a boosted (with velocity \tilde{v}^i), spinning (with spin \tilde{S}/\tilde{M}^2) Kerr-Schild black hole centered at the initial position of hole a . In this paper we choose $\tilde{v}^i = 0$ (since we seek data describing holes falling head-on from rest), $\tilde{M} = 0.39M_{\text{ADM}}$, and $\tilde{S}/\tilde{M}^2 = 0.5$. The Gaussian weighting parameter is chosen to be $w_a = d/3$, where d is the initial coordinate separation between the two holes; note that this choice causes the conformal metric to be flat everywhere except near each hole. The holes are located at coordinates $(x, y, z) = (x_0 \equiv \pm d/2, 0, 0)$.

These free data are then inserted into the extended conformal thin sandwich (XCTS) equations (e.g., Eqs. (13)–(15) of Ref. [76])¹⁰, which are then solved for the conformal factor ψ , the lapse α , and the shift β^i . The XCTS equations are solved using a spectral elliptic solver [80] on a computational domain with i) a very large outer boundary (which is chosen to be a coordinate sphere with radius $10^9\tilde{M}$), and ii) with the region inside the holes' apparent horizons excised. The excision surfaces \mathcal{S} are surfaces of constant Kerr radius r_{Kerr} , where

$$\frac{x^2 + y^2}{r_{\text{Kerr}}^2 + \tilde{S}_a^2/\tilde{M}_a^2} + \frac{z^2}{r_{\text{Kerr}}^2} = 1. \quad (\text{A.25})$$

The excision surfaces are the apparent horizons of the holes; this is enforced by the boundary condition given by Eq. (48) of Ref. [76]. On the apparent horizon, the lapse satisfies the boundary condition

$$\alpha\psi = 1 + \sum_{a=1}^2 e^{-r_a^2/w_a^2} (\alpha_a - 1) \text{ on } \mathcal{S}, \quad (\text{A.26})$$

where α_a is the lapse of the Kerr-Schild metric corresponding to hole a . The shift satisfies

$$\beta^i = \alpha s^i - \Omega_r \xi^i \text{ on } \mathcal{S}. \quad (\text{A.27})$$

The first term in Eq. (A.27) implies that the holes are initially at rest, and the second term determines

¹⁰The XCTS equations are also given by Eqs. (37a)–(37d) of Ref. [54], aside from the following typographical error: the second term in square brackets on the right-hand-side of Eq. (37c) should read $(5/12)K^2\psi^4$ (not $(5/12)K^4\psi^4$).

the spin of the hole; to make the spin point in the $\pm z$ direction with magnitude $S/M_{\text{Chr}}^2 = 0.5$ (measured using the method described in Appendix A of Ref. [54]), we choose $M_{\text{ADM}}\Omega_r = \mp 0.244146$ and $\xi^i = \partial_\phi$, where ∂_ϕ is the rotation vector on the apparent horizon corresponding to rotation about the $+z$ axis.

On the outer boundary \mathcal{B} , the spacetime metric is flat:

$$\psi = 1 \text{ on } \mathcal{B}, \quad (\text{A.28})$$

$$\alpha\psi = 1 \text{ on } \mathcal{B}. \quad (\text{A.29})$$

Our initial data are constructed [Eq. (A.27)] in a frame comoving with the black holes. Thus, an asymptotic rotation, expansion, and translation in the comoving shift β^i cause the holes to initially have radial, angular, or translational velocity in the inertial frame; this corresponds to the boundary condition

$$\beta^i = (\mathbf{\Omega}_0 \times \mathbf{r})^i + \dot{a}_0 r^i + V_0^i \text{ on } \mathcal{B}. \quad (\text{A.30})$$

We choose $\dot{a}_0 = 0$ and $\mathbf{\Omega}_0 = \mathbf{0}$. To make the total momentum of the initial data vanish, we choose $V^y = -0.001444$ and $V^x = V^z = 0$. This choice gives the holes an initial coordinate velocity of $0.001444 = 433 \text{ km/s}$ in the $-y$ direction (cf. Fig A.8). Note that the initial data are evolved in inertial, not comoving, coordinates, so that the shift during the evolution is different from the comoving shift β^i obtained from the XCTS equations: the former asymptotically approaches zero, not a constant vector V_0^i .

A.8.2 Superposed-Harmonic-Kerr (SHK) initial data

We also present a simulation, H1 in Table A.1, that is similar to S1 except that the initial separation between the holes is larger and the gauge is nearly harmonic. The construction of this Superposed-Harmonic-Kerr initial data for this run follows that of the Superposed-Kerr-Schild (S1) initial data described in Appendix A.8. The differences are as follows.

The first difference is our choice of coordinates. In Appendix A.8, the quantities γ_{ij}^a , K_a , and α_a that appear in Eqs. (A.23), (A.24), and (A.26) refer to the three-metric, the trace of the extrinsic curvature, and the lapse function of the Kerr metric in Kerr-Schild coordinates. Here we still use Eqs. (A.23), (A.24), and (A.26), but γ_{ij}^a , K_a , and α_a now refer to the three-metric, the trace of the extrinsic curvature, and the lapse function of the Kerr metric in fully harmonic coordinates, Eqs. (22)-(31), (41) and (43) of Ref. [81]. Furthermore, the computational domain is excised on surfaces

of constant Boyer-Lindquist radius, r_{BL} , where

$$\frac{x^2 + y^2}{(r_{\text{BL}} - \tilde{M}_a)^2 + \tilde{S}_a^2 / \tilde{M}_a^2} + \frac{z^2}{(r_{\text{BL}} - \tilde{M}_a)^2} = 1. \quad (\text{A.31})$$

The initial coordinate separation was chosen to be $d = 29.73M_{\text{ADM}}$ and the Gaussian weighting parameter that appears in Eqs. (A.23), (A.24), and (A.26) is $w_a = d/9$. To obtain $S/M_{\text{Chr}}^2 = \{0, 0, \pm 0.5\}$ we choose $\Omega_r = \mp 0.261332/M_{\text{ADM}}$ in Eq. (A.27), and to make the total momentum vanish we choose $V_0^y = -0.0000582185$ in Eq. (A.30).

Solving the XCTS equations results in initial data that is approximately harmonic. Harmonic coordinates satisfy $\nabla^c \nabla_c x^a = 0$, or equivalently, $\Gamma_a := \Gamma_{ab}{}^b = 0$. We can evaluate the degree to which the harmonic gauge condition is satisfied in our initial data by examining the normalized magnitude of Γ_a :

$$f := \left(\frac{\sum_a |\Gamma_a|^2}{\frac{1}{4} \sum_a \sum_b |\Gamma_{ab}{}^b|^2} \right)^{1/2}. \quad (\text{A.32})$$

The denominator consists of the sum of squares of terms that must cancel to produce $\Gamma_a = 0$, so that $f = 1$ corresponds to complete violation of the harmonic coordinate condition. On the apparent horizons $f < 0.049$, while in the asymptotically flat region far from the holes $f < 0.0083$. In the regions where the Gaussians in Eqs. (A.23), (A.24) and (A.26) transition the XCTS free data from harmonic Kerr to conformally flat we cannot expect the data to be strongly harmonic, and we find that $f < 0.12$.

The techniques employed in the spectral evolution from this SHK initial data follow those used for the SKS initial data as described in Appendix A.9.1. In particular, the generalized harmonic gauge source function, H_a (Eq. A.14), is constructed by demanding that $\tilde{H}_{a'}$ remains frozen to its value in the initial data. The evolution proceeds in nearly harmonic gauge because of the way the initial data is constructed.

Three of these H1 evolutions were performed at resolutions of approximately 61^3 , 67^3 and 72^3 grid points. The constraints were found to be convergent. The data presented in this paper is taken from the highest resolution run.

These simulations are specifically constructed to provide data for comparison with PN approximations, so we are restricted to remain in our approximately harmonic gauge. However, currently this gauge choice prevents us from continuing our H1 evolutions beyond the plunge phase; we have not observed the formation of a common horizon.

A.9 Appendix B: Numerical methods for evolutions

A.9.1 Generalized harmonic evolutions

We evolve the initial data summarized in Sec. A.4.1 using the Caltech-Cornell pseudospectral code SPEC. This code and the methods it employs are described in detail in Refs. [82, 83, 4]. Some of these methods have been simplified for the head-on problem discussed here, and others have been modified to account for a nonzero center-of-mass velocity, so we will describe them here.

We evolve a first-order representation [61] of the generalized harmonic system [58, 59, 60]. We handle the singularities by excising the black hole interiors from the computational domain. Our outer boundary conditions [61, 84, 85] are designed to prevent the influx of unphysical constraint violations [86, 87, 88, 89, 90, 91, 92] and undesired incoming gravitational radiation [93, 94] while allowing outgoing gravitational radiation to pass freely through the boundary.

We find the event horizon using the techniques of Ref. [95], except that for calculating derivatives of quantities on the event horizon surface, we use a 6th order finite differencing stencil, which is an improvement on the 2nd order stencil used in [95]. (The formation of cusps on the event horizon prevents us from taking spectral derivatives there.)

We employ the dual-frame method described in Ref. [82]: we solve the equations in an “inertial frame” that is asymptotically Minkowski, but our domain decomposition is fixed in a “comoving frame” that is allowed to shrink, translate and distort relative to the inertial frame. The positions of the centers of the black holes are fixed in the comoving frame; we account for the motion of the holes by dynamically adjusting the coordinate mapping between the two frames. Note that the comoving frame is referenced only internally in the code as a means of treating moving holes with a fixed domain. Therefore all coordinate quantities (e.g. black hole trajectories) mentioned in this paper are inertial-frame values unless explicitly stated otherwise.

The mapping from comoving to inertial coordinates is changed several times during the run. During the plunge phase, we denote the mapping by $\mathcal{M}_p(x^i, x'^i)$, where primed coordinates denote the comoving frame and unprimed coordinates denote the inertial frame. Explicitly, $\mathcal{M}_p(x^i, x'^i)$ is the mapping

$$x = F(r', t) \sin \theta' \cos \phi', \quad (\text{A.33})$$

$$y = F(r', t) \sin \theta' \sin \phi' + e^{-r'^2/r_T'^2} Y(t), \quad (\text{A.34})$$

$$z = F(r', t) \cos \theta' \cos \phi', \quad (\text{A.35})$$

where

$$F(r', t) := r' \left[a(t) + (1 - a(t)) \frac{r'^2}{R_0'^2} \right]. \quad (\text{A.36})$$

Here $a(t)$ and $Y(t)$ are functions of time, (r', θ', ϕ') are spherical polar coordinates in the comoving frame centered at the origin, and R'_0 and r'_T are constants. For the choice $R'_0 = \infty$ and $r'_T = \infty$, the mapping is simply an overall contraction by $a(t) \leq 1$ plus a translation $Y(t)$ in the y direction. Choosing R'_0 equal to the outer boundary radius R'_{\max} and choosing $r'_T \sim R'_{\max}/6$ causes the map to approach the identity near the outer boundary; this prevents the outer boundary from falling close to the strong-field region during merger, and makes it easier to keep the outer boundary motion smooth through the merger/ringdown transition. The functions $a(t)$ and $Y(t)$ are determined by dynamical control systems as described in Ref. [82]. These control systems adjust $a(t)$ and $Y(t)$ so that the centers of the apparent horizons remain stationary in the comoving frame. For the evolutions presented here, we use $R'_0 = 532.2M_{\text{ADM}} = 1.1R'_{\max}$ and $r'_T = 31.21M_{\text{ADM}} = 4d_o$, where d_o is the initial separation of the holes.

The gauge freedom in the generalized harmonic system is fixed via a freely specifiable gauge source function H_a that satisfies the constraint

$$0 = \mathcal{C}_a := \Gamma_{ab}{}^b + H_a, \quad (\text{A.37})$$

where $\Gamma^a{}_{bc}$ are the spacetime Christoffel symbols. To choose this gauge source function, we define a new quantity \tilde{H}_a that transforms like a tensor and agrees with H_a in inertial coordinates (i.e. $\tilde{H}_a = H_a$). Then we choose \tilde{H}_a so that the constraint (A.37) is satisfied initially, and we demand that $\tilde{H}_{a'}$ is constant in the moving frame.

Shortly before merger (at time $t_1 = 31.1M_{\text{ADM}}$), we make two modifications to our algorithm to reduce numerical errors and gauge dynamics during merger. First, we begin controlling the size of the individual apparent horizons so that they remain constant in the comoving frame, and therefore they remain close to their respective excision boundaries. This is accomplished by changing the map between comoving and inertial coordinates as follows. We define the map $\mathcal{M}_{\text{AH}_1}(\tilde{x}^i, x'^i)$ for black hole 1 as

$$\tilde{x} = x'_{\text{AH}_1} + \bar{r} \sin \theta' \cos \phi', \quad (\text{A.38})$$

$$\tilde{y} = y'_{\text{AH}_1} + \bar{r} \sin \theta' \sin \phi', \quad (\text{A.39})$$

$$\tilde{z} = z'_{\text{AH}_1} + \bar{r} \cos \theta', \quad (\text{A.40})$$

$$\bar{r} := r' - e^{-(r'-r'_0)^3/\sigma_1^3} \lambda_1(t), \quad (\text{A.41})$$

where (r', θ', ϕ') are spherical polar coordinates centered at the (fixed) comoving-coordinate location of black hole 1, which we denote as $(x'_{\text{AH}_1}, y'_{\text{AH}_1}, z'_{\text{AH}_1})$. The constant R'_{AH_1} is the desired average radius (in comoving coordinates) of black hole 1. Similarly, we define the map $\mathcal{M}_{\text{AH}_2}(\tilde{x}^i, x'^i)$ for black hole 2. Then the full map from the comoving coordinates x'^i to the inertial coordinates x^i is

given by

$$\mathcal{M}_m(x^i, x'^i) := \mathcal{M}_p(x^i, \bar{x}^i) \mathcal{M}_{\text{AH}_2}(\bar{x}^i, \tilde{x}^i) \mathcal{M}_{\text{AH}_1}(\tilde{x}^i, x'^i). \quad (\text{A.42})$$

The constants σ_1 , σ_2 , and r'_0 are chosen to be $0.780M_{\text{ADM}}$, $0.780M_{\text{ADM}}$, and $1.01M_{\text{ADM}}$, respectively. The functions $\lambda_1(t)$ and $\lambda_2(t)$ are determined by dynamical control systems that drive the comoving-coordinate radius of the apparent horizons towards their desired values $R'_{\text{AH}_1} = R'_{\text{AH}_2} = 1.56M_{\text{ADM}}$. Note that in comoving coordinates, the shape of the horizons is not necessarily spherical; only the average radius of the horizons is controlled.

The second change we make at time $t_1 = 31.1M_{\text{ADM}}$ is to smoothly roll gauge source function H_a to zero by adjusting $\tilde{H}_{a'}(t)$ according to

$$\tilde{H}_{a'}(t) = \tilde{H}_{a'}(t_1) e^{-(t-t_1)^2/\tau^2}, \quad (\text{A.43})$$

where $\tau = 0.5853M_{\text{ADM}}$. This choice makes it easier for us to continue the evolution after the common horizon has formed, and it also reduces gauge dynamics that otherwise cause oscillations in the observed Landau-Lifshitz velocity v_{LL}^y during the ringdown.

When the two black holes are sufficiently close to one another, a new apparent horizon suddenly appears, encompassing both black holes. At time $t_m = 34.73M_{\text{ADM}}$ (which is shortly after the common horizon forms), we interpolate all variables onto a new computational domain that contains only a single excised region, and we choose a new comoving coordinate system so that the merged (distorted, pulsating) apparent horizon remains spherical in the new comoving frame. This is accomplished in the same way as described in Section II.D. of [4], except that here the map from the new comoving coordinates to the inertial coordinates contains an additional translation in the y direction that handles the nonzero velocity of the merged black hole. In [4] a third change, namely a change of gauge, was necessary to continue the simulation after merger. But in the simulations discussed here, Eq. (A.43) has caused H_a to fall to zero by the time of merger, and we find it suffices to simply allow H_a to remain zero after merger.

For completeness, we now explicitly describe the map from the new comoving coordinates x''^i to

the inertial coordinates x^i . This map is given by

$$x = r \sin \theta'' \cos \phi'', \quad (\text{A.44})$$

$$y = r \sin \theta'' \sin \phi'' + e^{-r''^2/r_T''^2} Y(t), \quad (\text{A.45})$$

$$z = r \cos \theta'', \quad (\text{A.46})$$

$$r = \tilde{r} \left[1 + \sin^2(\pi \tilde{r} / 2R''_{\max}) \times \left(A(t) \frac{R''_{\max}}{R''_{\max}} + (1 - A(t)) \frac{R''_{\max}^3}{R''_{\max} R_0''^2} - 1 \right) \right], \quad (\text{A.47})$$

$$\tilde{r} = r'' - q(r'') \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \lambda_{\ell m}(t) Y_{\ell m}(\theta'', \phi''), \quad (\text{A.48})$$

(r'', θ'', ϕ'') are spherical polar coordinates in the new comoving coordinate system, R''_{\max} is the value of r'' at the outer boundary, and r_T'' is a constant chosen to be $31.21M_{\text{ADM}}$. The function $q(r'')$ is given by

$$q(r'') = e^{-(r'' - R''_{\text{AH}})^3 / \sigma_q^3}, \quad (\text{A.49})$$

where R''_{AH} is the desired radius of the common apparent horizon in comoving coordinates. The function $A(t)$ is

$$A(t) = A_0 + (A_1 + A_2(t - t_m)) e^{-(t - t_m) / \tau_A}, \quad (\text{A.50})$$

where the constants A_0 , A_1 , and A_2 are chosen so that $A(t)$ matches smoothly onto $a(t)$ from Eq. (A.36): $A(t_m) = a(t_m)$, $\dot{A}(t_m) = \dot{a}(t_m)$, and $\ddot{A}(t_m) = \ddot{a}(t_m)$. The constant τ_A is chosen to be on the order of $5M$. The functions $Y(t)$ and $\lambda_{\ell m}(t)$ are determined by dynamical control systems that keep the apparent horizon spherical and centered at the origin in comoving coordinates; see [4] for details.

A.9.2 BSSN-moving-puncture evolutions

In addition to the generalized evolutions, we have performed a second set of simulations using the so-called *moving puncture* technique [3, 2] using the LEAN code [72, 96]. This code is based on the CACTUS computational toolkit [97] and uses mesh refinement provided by the CARPET package [69, 70]. Initial data are provided in the form of the TWOPUNCTURES thorn by Ansorg's spectral solver [66] and apparent horizons are calculated with Thornburg's AHFINDERDIRECT [98, 99].

The most important ingredient in this method for the present discussion is the choice of coordinate conditions. A detailed study of alternative gauge conditions in the context of moving puncture type black-hole evolutions is given in Ref. [100]. In particular, they demonstrate how the common choice of a second order in time evolution equation for the shift vector β^i can be integrated in time analytically and thus reduced to a first order equation. Various test simulations performed with the LEAN code

confirm their Eq. (26) as the most efficient method to evolve the shift vector. In contrast to the shift, moving puncture codes show little variation in the evolution of the lapse function. Here we follow the most common choice so that our gauge conditions are given by

$$\partial_t \alpha = \beta^i \partial_i \alpha - 2\alpha K, \quad (\text{A.51})$$

$$\partial_t \beta^i = \beta^m \partial_m \beta^i + \frac{3}{4} \tilde{\Gamma}^i - \eta \beta^i. \quad (\text{A.52})$$

$\tilde{\Gamma}^i$ is the contracted Christoffel symbol of the conformal 3-metric, K the trace of the extrinsic curvature [see for example Eq. (1) of [72]] and η a free parameter set to 1 unless specified otherwise. For further details about the moving puncture method and the specific implementation in the LEAN CODE code we refer to Sec. II of Ref. [72]. Except for the use of sixth instead of fourth order spatial discretization [101], we did not find it necessary to apply any modifications relative to the simulations presented in that work.

The calculation of the 4-momentum in the LEAN code is performed in accordance with the relations listed in Sec. A.3. The only difference is that in a BSSN code the four metric and its derivatives are not directly available but need to be expressed in terms of the 3-metric γ_{ij} , the extrinsic curvature K_{ij} as well as the gauge variables lapse α and shift β^i . The key quantity for the calculation of the 4-momentum is the integrand in Eq. (A.7). A straightforward calculation gives it in terms of the canonical ADM variables

$$\partial_\alpha H^{0\alpha 0j} = \frac{1}{\chi^3} \left[\frac{3}{\chi} \gamma^{jm} \partial_m \chi + \gamma^{km} \gamma^{jn} \partial_k \gamma_{mn} \right], \quad (\text{A.53})$$

$$\begin{aligned} \partial_\alpha H^{i\alpha 0j} &= \frac{1}{\chi^3} \left[2\alpha (K^{ij} - \gamma^{ij} K) + \gamma^{ij} \partial_m \beta^m - \gamma^{im} \gamma_m \beta^j \right] \\ &\quad - \beta^i \partial_\alpha H^{0\alpha 0j}, \end{aligned} \quad (\text{A.54})$$

where $K := K^i{}_i$ and $\chi := \det \gamma^{-1/3}$ have been used for convenience because they are fundamental variables in our BSSN implementation.

Bibliography

- [1] Frans Pretorius. Evolution of binary black-hole spacetimes. *Phys. Rev. Lett.*, 95(12):121101, 2005.
- [2] Manuela Campanelli, Carlos O. Lousto, Pedro Marronetti, and Yosef Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Phys. Rev. Lett.*, 96(11):111101, 2006.

- [3] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter. Gravitational-wave extraction from an inspiraling configuration of merging black holes. *Phys. Rev. Lett.*, 96(11):111102, 2006.
- [4] Mark A. Scheel, Michael Boyle, Tony Chu, Lawrence E. Kidder, Keith D. Matthews, and Harald P. Pfeiffer. High-accuracy waveforms for binary black-hole inspiral, merger, and ringdown. *Phys. Rev. D*, 79:024003, 2009.
- [5] Drew Keppel, David A. Nichols, Yanbei Chen, and Kip S. Thorne. Momentum flow in black hole binaries: I. Post-Newtonian analysis of the inspiral and spin-induced bobbing. *Phys. Rev. D*, 80:124015, 2009.
- [6] Badri Krishnan, Carlos O. Lousto, and Yosef Zlochower. Quasi-Local Linear Momentum in Black-Hole Binaries. *Phys. Rev.*, D76:081501(R), 2007.
- [7] László B. Szabados. Quasi-local energy-momentum and angular momentum in General Relativity: A review article. *Living Rev. Rel.*, 12(4), 2009.
- [8] Carlos O. Lousto and Yosef Zlochower. Modeling gravitational recoil from precessing highly-spinning unequal-mass black-hole binaries. *Phys. Rev. D*, 79:064018, 2009.
- [9] Carlos O. Lousto and Yosef Zlochower. Further insight into gravitational recoil. *Phys. Rev. D*, 77:044028, 2008.
- [10] M. Boylan-Kolchin, C.-P. Ma, and E. Quataert. Core formation in galactic nuclei due to recoiling black holes. *Astrophys. J*, 613:L37–L40, 2004. astro-ph/0407488.
- [11] A. Gualandris and D. Merritt. Ejection of Supermassive Black Holes from Galaxy Cores. *Astrophys. J*, 678:780–797, 2008. arXiv:0708.0771 [astro-ph].
- [12] S. Komossa and D. Merritt. Gravitational Wave Recoil Oscillations of Black Holes: Implications for Unified Models of Active Galactic Nuclei. *Astrophys. J. Lett.*, 689:l89, 2008. arXiv:0811.1037 [astro-ph].
- [13] P. Madau, M. J. Rees, M. Volonteri, F. Haardt, and S. P. Oh. Early reionization by mini-quasars. *Astrophys. J.*, 604:484–494, 2004.
- [14] Z. Haiman. Constraints from gravitational recoil on the growth of supermassive black holes at high redshift. *Astrophys. J*, 613:36–40, 2004. astro-ph/0404196.
- [15] P. Madau and E. Quataert. The effect of gravitational-wave recoil on the demography of massive black holes. *Astrophys. J.*, 606:L17–L20, 2004.

- [16] David Merritt, Milos Milosavljevic, Marc Favata, Scott A. Hughes, and Daniel E. Holz. Consequences of gravitational radiation recoil. *Astrophys. J.*, 607:L9–L12, 2004.
- [17] M. Volonteri. Gravitational recoil: Signatures on the massive black hole population. *Astrophys. J.*, 663:L5–L8, 2007. astro-ph/0703180.
- [18] L. Blecha and A. Loeb. Effects of gravitational-wave recoil on the dynamics and growth of supermassive black holes. *Mon. Not. R. Astron. Soc.*, 390:1311, 2008. arXiv:0805.1420 [astro-ph].
- [19] A. Loeb. Observable Signatures of a Black Hole Ejected by Gravitational Radiation Recoil in a Galaxy Merger. *Phys. Rev. Lett.*, 99:041103, 2007. astro-ph/0703722.
- [20] S. Komossa, H. Zhou, and H. Lu. A recoiling supermassive black hole in the quasar SDSSJ092712.65+294344.0? *Astrophys. J.*, 678:L81, 2008. arXiv:0804.4585 [astro-ph].
- [21] K. Menou, Z. Haiman, and B. Kocsis. Cosmological Physics with Black Holes (and Possibly White Dwarfs). *New Astron. Rev.*, 51:884–890, 2008. arXiv:0803.3627 [astro-ph].
- [22] M. J. Fitchett. The influence of gravitational wave momentum losses on the centre of mass motion of a newtonian binary system. *Mon. Not. Roy. Astr. Soc.*, 203:1049–1062, 1983.
- [23] Marc Favata, Scott A. Hughes, and Daniel E. Holz. How black holes get their kicks: Gravitational radiation recoil revisited. *Astrophys. J.*, 607:L5–L8, 2004.
- [24] Luc Blanchet, Moh'd S. S. Qusailah, and Clifford M. Will. Gravitational recoil of inspiralling black-hole binaries to second post-newtonian order. *Astrophys. J.*, 635:508, 2005.
- [25] Thibault Damour and Achamveedu Gopakumar. Gravitational recoil during binary black hole coalescence using the effective one body approach. *Phys. Rev. D*, 73(12):124006, 2006.
- [26] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, James R. van Meter, and M. Coleman Miller. Getting a kick out of numerical relativity. *Astrophys. J.*, 653:L93–L96, 2006.
- [27] Jose A. Gonzalez, Ulrich Sperhake, Bernd Brügmann, Mark Hannam, and Sascha Husa. Maximum kick from nonspinning black-hole binary inspiral. *Phys. Rev. Lett.*, 98:091101, 2007.
- [28] F. Herrmann, Ian Hinder, D. Shoemaker, P. Laguna, and Richard A. Matzner. Gravitational recoil from spinning binary black hole mergers. *Astrophys. J.*, 661:430–436, 2007.
- [29] Michael Koppitz, Denis Pollney, Christian Reisswig, Luciano Rezzolla, Jonathan Thornburg, Peter Diener, and Erik Schnetter. Recoil velocities from equal-mass binary-black-hole mergers. *Phys. Rev. Lett.*, 99:041102, 2007.

- [30] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, and David Merritt. Large merger recoils and spin flips from generic black hole binaries. *Astrophys. J. Lett.*, 659(1):L5–L8, 2007.
- [31] Wolfgang Tichy and Pedro Marronetti. Binary black hole mergers: Large kicks for generic spin orientations. *Phys. Rev.*, D76:061502(R), 2007.
- [32] J. A. Gonzalez, M. D. Hannam, U. Sperhake, B. Brügmann, and S. Husa. Supermassive recoil velocities for binary black-hole mergers with antialigned spins. *Phys. Rev. Lett.*, 98:231101, 2007.
- [33] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, and David Merritt. Maximum gravitational recoil. *Phys. Rev. Lett.*, 98:231102, 2007.
- [34] James Healy, Frank Herrmann, Ian Hinder, Deirdre M. Shoemaker, Pablo Laguna, and Richard A. Matzner. Superkicks in hyperbolic encounters of binary black holes. *Phys. Rev. Lett.*, 102:041101, 2009.
- [35] L. Boyle, M. Kesden, and S. Nissanke. Binary black hole merger: symmetry and the spin expansion. *Phys. Rev. Lett.*, 100:151101, 2008. arXiv:0709.0299 [gr-qc].
- [36] L. Boyle and M. Kesden. The spin expansion for binary black hole merger: new predictions and future directions. *Phys. Rev. D*, 78:024017, 2008. arXiv:0712.2819 [astro-ph].
- [37] J. D. Schnittman and A. Buonanno. The distribution of recoil velocities from merging black holes. *Astrophys. J.*, 662:L63, 2007. astro-ph/0702641.
- [38] John G. Baker, William D. Boggs, Joan Centrella, Bernard J. Kelly, Sean T. McWilliams, M. Coleman Miller, and James R. van Meter. Modeling kicks from the merger of generic black-hole binaries. *Astrophys. J.*, 682:L29, 2008.
- [39] W. Tichy and P. Marronetti. The final mass and spin of black hole mergers. *Phys. Rev. D*, 78:081501(R), 2008. arXiv:0807.2985 [gr-qc].
- [40] C. O. Lousto, M. Campanelli, and Y. Zlochower. Remnant masses, spins and recoils from the merger of generic black-hole binaries. 2009. arXiv:0904.3541 [gr-qc].
- [41] John G. Baker, William D. Boggs, Joan Centrella, Bernard J. Kelly, Sean T. McWilliams, M. Coleman Miller, and James R. van Meter. Modeling kicks from the merger of non-precessing black-hole binaries. *Astrophys. J.*, 668:1140, 2007.
- [42] J. A. Gonzalez, U. Sperhake, and B. Brügmann. Black-hole binary simulations: the mass ratio 10:1. *Phys. Rev. D*, 79:124006, 2009. arXiv:0811.3952 [gr-qc].

- [43] L. Rezzolla. Modelling the final state from binary black-hole coalescences. *Class. Quant. Grav.*, 26:094023, 2009. arXiv:0812.2325 [gr-qc].
- [44] J. D. Schnittman, A. Buonanno, J. R. van Meter, J. G. Baker, W. D. Boggs, J. Centrella, B. J. Kelly, and S. T. McWilliams. Anatomy of the binary black hole recoil: A multipolar analysis. *Phys. Rev. D*, 77:044031, 2008.
- [45] Sarah H. Miller and R.A. Matzner. Multipole analysis of kicks in collision of binary black holes. *Gen. Rel. Grav.*, 41:525, 2009.
- [46] Etienne Racine, Alessandra Buonanno, and Lawrence E. Kidder. Recoil velocity at 2PN order for spinning black hole binaries. *Phys. Rev. D*, 80:044010, 2009.
- [47] Yasushi Mino and Jeandrew Brink. Gravitational radiation from plunging orbits: Perturbative study. *Phys. Rev. D*, 78:124015, 2008.
- [48] Frans Pretorius. Binary Black Hole Coalescence. 2007. arXiv:0710.1338 [gr-qc].
- [49] L. D. Landau and E. M. Lifshitz. *Classical Theory of Fields*. Addison Wesley, Reading, MA, second edition, 1962.
- [50] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. Freeman, New York, New York, 1973.
- [51] S. V. Babak and L. P. Grishchuk. Energy-momentum tensor for the gravitational field. *Phys. Rev. D*, 61:024038, 1999.
- [52] M. Shibata and T. Nakamura. Evolution of three-dimensional gravitational waves: Harmonic slicing case. *Phys. Rev. D*, 52:5428–5444, 1995.
- [53] T. W. Baumgarte and S. L. Shapiro. On the Numerical integration of Einstein’s field equations. *Phys. Rev. D*, 59:024007, 1998. gr-qc/9810065.
- [54] Geoffrey Lovelace, Robert Owen, Harald P. Pfeiffer, and Tony Chu. Binary-black-hole initial data with nearly-extremal spins. *Phys. Rev. D*, 78:084017, 2008.
- [55] Robert M. Wald. *General Relativity*. University of Chicago Press, Chicago and London, 1984.
- [56] R. Arnowitt, S. Deser, and Charles W. Misner. The dynamics of general relativity. In L. Witten, editor, *Gravitation: An Introduction to Current Research*. Wiley, New York, 1962.
- [57] James W. York, Jr. Kinematics and dynamics of general relativity. In Larry L. Smarr, editor, *Sources of Gravitational Radiation*, pages 83–126. Cambridge University Press, Cambridge, England, 1979.

- [58] Helmut Friedrich. On the hyperbolicity of Einstein's and other gauge field equations. *Commun. Math. Phys.*, 100(4):525–543, 1985.
- [59] David Garfinkle. Harmonic coordinate method for simulating generic singularities. *Phys. Rev. D*, 65(4):044029, 2002.
- [60] Frans Pretorius. Numerical relativity using a generalized harmonic decomposition. *Class. Quantum Grav.*, 22(2):425–451, 2005.
- [61] Lee Lindblom, Mark A. Scheel, Lawrence E. Kidder, Robert Owen, and Oliver Rinne. A new generalized harmonic evolution system. *Class. Quantum Grav.*, 23:S447–S462, 2006.
- [62] M. Ruiz, R. Takahashi, M. Alcubierre, and D. Núñez. Multipole expansions for energy and momenta carried by gravitational waves. *Gen. Relativ. Gravit.*, 40:1705–1729, 2008.
- [63] Carles Bona, Joan Massó, Edward Seidel, and Joan Stela. First order hyperbolic formalism for numerical relativity. *Phys. Rev. D*, 56:3405–3415, 1997.
- [64] Miguel Alcubierre, Bernd Brügmann, Peter Diener, Michael Koppitz, Denis Pollney, Edward Seidel, and Ryoji Takahashi. Gauge conditions for long-term numerical black hole evolutions without excision. *Phys. Rev. D*, 67:084023, 2003.
- [65] Steve Brandt and Bernd Brügmann. A simple construction of initial data for multiple black holes. *Phys. Rev. Lett.*, 78(19):3606–3609, 1997.
- [66] Marcus Ansorg, Bernd Brügmann, and Wolfgang Tichy. A single-domain spectral method for black hole puncture data. *Phys. Rev. D*, 70:064011, 2004.
- [67] Jeffrey M. Bowen and James W. York, Jr. Time-asymmetric initial data for black holes and black-hole collisions. *Phys. Rev. D*, 21(8):2047–2056, 1980.
- [68] M. J. Berger and J. Olinger. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.*, 53:484–512, 1984.
- [69] Erik Schnetter, Scott H. Hawley, and Ian Hawke. Evolutions in 3d numerical relativity using fixed mesh refinement. *Class. Quantum Grav.*, 21:1465, 2004.
- [70] Carpet – Adaptive Mesh Refinement for the Cactus Framework. <http://www.carpetcode.org>.
- [71] M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi. Gauge conditions for long-term numerical black hole evolutions without excision. *Phys. Rev. D*, 67:084023–084042, 2003.
- [72] Ulrich Sperhake. Binary black-hole evolutions of excision and puncture data. *Phys. Rev. D*, 76:104015, 2007.

- [73] Guillaume Faye, Luc Blanchet, and Alessandra Buonanno. Higher-order spin effects in the dynamics of compact binaries. I. Equations of motion. *Phys. Rev. D*, 74(10):104033, 2006.
- [74] H. Tagoshi, A. Ohashi, and B. J. Owen. Gravitational field and equations of motion of spinning compact binaries to 2.5 post-Newtonian order. *Phys. Rev. D*, 63:044006, 2001.
- [75] Gregory B. Cook. Corotating and irrotational binary black holes in quasicircular orbits. *Phys. Rev. D*, 65(8):084003, Mar 2002.
- [76] Gregory B. Cook and Harald P. Pfeiffer. Excision boundary conditions for black-hole initial data. *Phys. Rev. D*, 70(10):104016, Nov 2004.
- [77] Matthew Caudill, Greg B. Cook, Jason D. Grigsby, and Harald P. Pfeiffer. Circular orbits and spin in black-hole initial data. *Phys. Rev. D*, 74(6):064011, 2006.
- [78] Eric Gourgoulhon, Philippe Grandclément, and Silvano Bonazzola. Binary black holes in circular orbits. I. A global spacetime approach. *Phys. Rev. D*, 65:044020, 2002.
- [79] Philippe Grandclément, Eric Gourgoulhon, and Silvano Bonazzola. Binary black holes in circular orbits. II. numerical methods and first results. *Phys. Rev. D*, 65:044021, 2002.
- [80] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky. A multidomain spectral method for solving elliptic equations. *Comput. Phys. Commun.*, 152:253–273, 2003.
- [81] Gregory B. Cook and Mark A. Scheel. Well-behaved harmonic time slices of a charged, rotating, boosted black hole. *Phys. Rev. D*, 56(8):4775–4781, 1997.
- [82] Mark A. Scheel, Harald P. Pfeiffer, Lee Lindblom, Lawrence E. Kidder, Oliver Rinne, and Saul A. Teukolsky. Solving Einstein’s equations with dual coordinate frames. *Phys. Rev. D*, 74:104006, 2006.
- [83] Michael Boyle, Duncan A. Brown, Lawrence E. Kidder, Abdul H. Mroué, Harald P. Pfeiffer, Mark A. Scheel, Gregory B. Cook, and Saul A. Teukolsky. High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions. *Phys. Rev. D*, 76:124038, 2007.
- [84] Oliver Rinne. Stable radiation-controlling boundary conditions for the generalized harmonic Einstein equations. *Class. Quantum Grav.*, 23:6275–6300, 2006.
- [85] Oliver Rinne, Lee Lindblom, and Mark A. Scheel. Testing outer boundary treatments for the Einstein equations. *Class. Quantum Grav.*, 24:4053–4078, 2007.
- [86] J. M. Stewart. The Cauchy problem and the initial boundary value problem in numerical relativity. *Class. Quantum Grav.*, 15(9):2865–2889, 1998.

- [87] Helmut Friedrich and Gabriel Nagy. The initial boundary value problem for Einstein's vacuum field equation. *Commun. Math. Phys.*, 201(3):619–655, 1999.
- [88] James M. Bardeen and Luisa T. Buchman. Numerical tests of evolution systems, gauge conditions, and boundary conditions for 1D colliding gravitational plane waves. *Phys. Rev. D*, 65:064037, Mar 2002.
- [89] Béla Szilágyi, Bernd Schmidt, and Jeffrey Winicour. Boundary conditions in linearized harmonic gravity. *Phys. Rev. D*, 65(6):064015, Feb 2002.
- [90] Gioel Calabrese, Jorge Pullin, Oscar Reula, Olivier Sarbach, and Manuel Tiglio. Well posed constraint-preserving boundary conditions for the linearized Einstein equations. *Commun. Math. Phys.*, 240:377–395, 2003.
- [91] Béla Szilágyi and Jeffrey Winicour. Well-posed initial-boundary evolution in general relativity. *Phys. Rev. D*, 68(4):041501(R), Aug 2003.
- [92] Lawrence E. Kidder, Lee Lindblom, Mark A. Scheel, Luisa T. Buchman, and Harald P. Pfeiffer. Boundary conditions for the Einstein evolution system. *Phys. Rev. D*, 71:064020, 2005.
- [93] Luisa T. Buchman and Olivier C. A. Sarbach. Towards absorbing outer boundaries in general relativity. *Class. Quantum Grav.*, 23:6709–6744, 2006.
- [94] L. T. Buchman and O. C. A. Sarbach. Improved outer boundary conditions for Einstein's field equations. *Class. Quantum Grav.*, 24:S307–S326, 2007.
- [95] Michael Cohen, Harald P. Pfeiffer, and Mark A. Scheel. Revisiting event horizon finders. *Class. Quantum Grav.*, 26:035005, 2009.
<http://iopscience.iop.org/0264-9381/26/3/035005/>, ©2009 IOP Publishing Ltd.
- [96] Ulrich Sperhake, Bernd Brügmann, Jose Gonzalez, Mark Hannam, and Sascha Husa. Head-on collisions of different initial data. In *Proceedings of the eleventh Marcel Grossmann Meeting*, 2007.
- [97] The Cactus Computational Toolkit. <http://www.cactuscode.org>.
- [98] J. Thornburg. Finding apparent horizons in numerical relativity. *Phys. Rev. D*, 54:4899–4918, 1996. gr-qc/9508014.
- [99] J. Thornburg. A fast apparent-horizon finder for 3-dimensional Cartesian grids in numerical relativity. *Class. Quantum Grav.*, 21:743–766, 21 January 2004. gr-qc/0306056.

- [100] James R. van Meter, John G. Baker, Michael Koppitz, and Dae-Il Choi. How to move a black hole without excision: gauge conditions for the numerical evolution of a moving puncture. *Phys. Rev. D*, 73:124011, 2006.
- [101] Sascha Husa, José A. González, Mark Hannam, Bernd Brügmann, and Ulrich Sperhake. Reducing phase error in long numerical binary black hole evolutions with sixth order finite differencing. *Class. Quantum Grav.*, 25:105006, 2008.