NUMERICAL EVOLUTIONS OF BINARY BLACK HOLE SYSTEMS

A Dissertation Presented to the Faculty of the Graduate School of Cornell University in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Daniel Hemberger

August 2013

© 2013 Daniel Hemberger All Rights Reserved

NUMERICAL EVOLUTIONS OF BINARY BLACK HOLE SYSTEMS

Daniel Hemberger, Ph.D.

Cornell University 2013

Numerical evolutions of binary black holes with moderate spins and mass ratios are becoming more routine. However, dealing with high spins or mass ratios has provided significant challenges. This work presents enhancements to the Spectral Einstein Code that enable such evolutions. The key improvements ensure that the computational coordinates conform to the shape of the black hole horizon and that information does not enter the computational domain through this surface. With these enhancements in place, a set of high spin binary black hole coalescences are simulated. The merger remnant properties are used to improve existing fitting formulae for final spin and energy radiated in gravitational waves. Such fitting formulae can be used to improve phenomenological template waveforms for Advanced LIGO gravitational wave searches.

BIOGRAPHICAL SKETCH

Daniel Hemberger was born in 1985 in Reading, Pennsylvania. At an early age he decided that he wanted to be a scientist, though the only scientists he knew were probably Egon Spengler and Batman. Several years later, *Bill Nye the Science Guy* gave him a slightly more realistic impression, which was equally appealing. In high school, Daniel loved studying calculus because it so elegantly and exactly described the most basic physical phenomena that otherwise would remain a mystery. His decision to pursue astrophysics is likely a consequence of how much fun he and his friends had annoying their physics teacher with silly questions about paradoxes in Special Relativity.

In 2003, Daniel attended Oberlin College, where he studied pulsars with Daniel Stinebring. During his senior year, he traveled with Dr. Stinebring to the University of Leiden in the Netherlands to complete his undergraduate Honors research. After graduating in 2007 with highest honors, he began his graduate studies in numerical relativity at Cornell University with Saul Teukolsky.

Daniel's primary avocation is music, and he has been in ensembles of varying professionalism his entire life. While his knowledge of physics does not have a significant practical impact on his musical pursuits (and vice versa), he feels that they are symbiotic and that one enhances the appreciation of the other. This dissertation is dedicated to my grandparents.

ACKNOWLEDGEMENTS

Many people have played an important role in my graduate student career, and I would like to thank each of them for their invaluable contributions:

- First and foremost, my thesis advisor Saul Teukolsky for his guidance, motivation, and patience. I may never be able to answer questions as succinctly and expertly as Saul, but I will know I have succeeded when I can *ask* questions like Saul.
- Larry Kidder and Geoffrey Lovelace for their mentorship and friendship.
- My collaborators for sharing their expertise and enduring our marathon teleconferences.
- The original authors (Larry Kidder, Mark Scheel, and Harald Pfeiffer) and all subsequent contributors to the Spectral Einstein Code.
- All those with whom I have had enlightening discussions, especially my sixth-floor compatriots: Andy Bohn, Mike Boyle, François Hébert, Kris Henriksson, Curran Muhlberger, and Will Throwe.
- Jim Cordes and Éanna Flanagan for serving on my special committee and educating me at various points in my career – Jim as my REU advisor and Éanna as my instructor for two semesters of General Relativity coursework.
- Dan Stinebring, my undergraduate advisor, for going above and beyond to help set me on a course for Cornell.
- The Ithaca winters for teaching me that I can move somewhere warmer.
- My family and friends for their unwavering support, which has meant everything to me.

TABLE OF CONTENTS

				Page
	Biog	graphica	al Sketch	. iii
	Ded	ication		. iv
	Ack	nowledg	gements	. v
	Tab	le of Co	ontents	. vi
	List	of Figu	ures	. viii
	List	of Tabl	les	. xii
	Pref	ace		. xiii
1	Int	roduct	tion	1
	1.1	Gravit	cational-wave astronomy	. 1
		Ι	Historical perspective	. 1
		II	Detection techniques	. 3
	1.2	Overv	iew	. 5
		Ι	Control systems for choosing computational coordinates	. 5
		II	High-spin binary black hole evolutions	. 7
\mathbf{R}	efere	nces		9
2	Dy	namic	al Excision Boundaries in Spectral Evolutions of Binary Bla	ck
	Hol	e Spac	etimes	12
	2.1	Introd		. 13
	2.2	Contro	of Theory	. 15
	2.3	Contro	D Systems in SpEC	. 18
		l	Definition of control errors and control signals	. 20
			Averaging out noise	. 23
	0.4		Dynamic timescale adjustment	. 24
	2.4	Contro	al Systems for Mans	. 27
		l		. 30
		11	Scaling	
		TTT	Scaling	. 32
		III	Scaling	. 32 . 34
		III IV	Scaling	. 32 . 34 . 36
		III IV V	Scaling	. 32 . 34 . 36 . 41
		III IV V VI	Scaling	. 32 . 34 . 36 . 41 . 44
	2.5	III IV V VI Size co	Scaling	. 32 . 34 . 36 . 41 . 44 . 50

	II Apparent horizon tracking
	III Adaptive switching of size control
2.6	Merger and ringdown
	I Initialization
2.7	Control systems for efficiency
2.8	Summary
A1	Implementation of exponentially-weighted averaging
A2	Details on computing the CutX map weight function
A3	Estimation of zero-crossing times

References

7	1
•	•

3	\mathbf{Fin}	al spin and radiated energy in numerical simulations of binary black	
	hole	es with equal masses and equal, aligned or anti-aligned spins	75
	3.1	Introduction	76
	3.2	Simulation methods	78
	3.3	Simulations	81
		I Mass and spin	82
	3.4	Results	88
		I Final spin	89
		II Radiated energy	91
		III Extremality	96
	3.5	Conclusions	98
	A1	Method for constructing our fitting formulae	99
Re	eferei	nces	107

LIST OF FIGURES

1	A generic control circuit. The simulation outputs a measure of error, Q , which	
	is used by the control system. The control system then outputs a signal, U ,	
	which changes the behavior of the simulation.	16
2	The speed in (x, t) coordinates, v, is plotted for a family of gains K_P with	
	$K_D = 1$ fixed. The wave velocity is $c = -0.2$ and the desired velocity is	
	$v_d = 0.5$ (dashed line). The controller turns on at $t = 2$. For gains in the stable	
	region, where $K_P \geq 0.25$, v settles down to v_d . One overdamped solution with	
	$K_P = 0.01$ is plotted for comparison (dotted line).	18
3	a) Computational domain in grid coordinates; the black hole centers are at	
	rest and the excision boundaries are spherical. b) Same domain in inertial	
	coordinates near merger; the excision boundaries move and distort to track	
	the apparent horizons	19
4	Two-dimensional projection of the domain decomposition near the two black	
	holes A and B. Shown are boundaries between subdomains. Each subdomain	
	takes the shape of a spherical shell, a distorted cylindrical shell, or a distorted	
	cylinder. Additional spherical-shell subdomains (not shown) surround the	
	outer boundary of the figure and extend to large radius. This domain decom-	
	position is explained in detail in the Appendix of [5]. The red, magenta, and	
	blue surfaces are those for which the function $f_A(r_A, \theta_A, \phi_A)$ from Eq. (2.4.44)	
	has a discontinuous gradient, as described in Sec. 2.4.V.	37
5	Snapshot of the grid, viewed in the mapped frame, from a binary black	
	hole evolution with $M_A = 8M_B$ shortly before the common horizon forms.	
	Left: without $\mathcal{M}_{\text{Skew}}$. Right: with $\mathcal{M}_{\text{Skew}}$. The gray areas are excised	
	regions. In the left panel, the grid is compressed as the excision boundaries	
	track increasingly skewed apparent horizons, and the evolution is terminated	
	because of excessively large constraint violation in the compressed region.	
	This problem is resolved by the inclusion of $\mathcal{M}_{\mathrm{Skew}}$.	39

6	A diagram of the skew map in the mapped frame. The horizon surfaces are drawn in black and the excision surfaces are drawn in red. The detted red line	
	drawn in black and the excision surfaces are drawn in red. The dotted red line	
	represents the me segment connecting the excision surface centers, which is	
	parallel to the x-axis. The normal to either horizon at the intersection point	
	with this segment is represented by the straight black lines. The cutting plane	
	is represented by the (skewed) blue curve, and the normal to this plane at x_C^*	
	is represented by the central green line. The green lines near the two excision	
	surfaces are constructed parallel to the central green line, where parallelism is	
	defined assuming a Euclidean background. There are three angles involved	
	in the skew control system – the angle between the normal to the skewed	
	cutting plane and the x-axis, Θ^{j} , and the angle between the normal to either	
	horizon and the normal to the skewed cutting plane, Θ_H^j . Skew control acts	
	to minimize Θ_H^j . We measure Θ_H^j in the unmapped (distorted) frame, but	
	Θ_H^j is invariant under the skew map for $W \approx 1$.	40
7	Two-dimensional projection of the domain decomposition near the two black	
	holes A and B. The function ρ in \mathcal{M}_{CutX} is zero on the magenta boundaries,	
	one on the red boundaries, and goes linearly between zero and one in the	
	"radial" direction between these boundaries. The gradient of ρ is discontinuous	
	across the solid magenta, red, and blue boundaries. Dotted lines show the	
	subdomain boundaries under the action of \mathcal{M}_{CutX} .	45
8	Snapshot of the grid, viewed in the inertial frame, from a binary black hole	
	evolution with $M_A = 8M_B$ shortly before the common horizon forms. Top:	
	without \mathcal{M}_{CutX} . Bottom: with \mathcal{M}_{CutX} . The gray areas are excised regions.	
	In the top panel, the grid is compressed as the larger excision boundary	
	approaches the cutting plane; this is especially evident in the long narrow	
	subdomains immediately adjacent to the larger excision boundary. Using	
_	\mathcal{M}_{CutX} relieves the grid compression.	46
9	Minimum characteristic speed on one of the excision surfaces of an equal mass	
	binary with equal aligned spins of $\chi = 0.97$, shown just before merger [25].	
	Two cases are shown: one without characteristic speed control (solid) and one	
	that is restarted from the uncontrolled run at $t = 6255M$ with characteristic	
	speed control turned on and $v_T = 5 \times 10^{-5}$ (dashed). The uncontrolled run is	
	terminated because of negative characteristic speeds shortly after $t = 6290M$,	50
10	but the controlled run continues through merger and ringdown.	53
10	Schematic diagram indicating the way the $\mathcal{M}_{\text{CutX}}$ weight function is defined.	
	The value of the weight function $\rho(x^*)$ is one on the red curves and in the	
	black region enclosed by red curves, it vanishes in the white regions (inside	
	the two inner magenta curves and outside the outer magenta curve) and it	
	changes linearly from zero to one in the gray regions between the red and	c-
	magenta curves.	67

11	Schematic diagram illustrating the algorithm used to compute the $\mathcal{M}_{\text{CutX}}$ weight function inside the cut-sphere. The point x^i represents an arbitrary point in the M region. R_H is the radius of the magenta sphere, and x_M is the x^0 -component of the point on the magenta sphere that intersects the	
	ray pointing from C_H^i toward x^i . The weight function $\rho(x^i)$ is zero at the intersection of the ray with the magenta curve, it is unity at the intersection of the ray with the cutting plane, and changes linearly in between. Similarly, for a point in the <i>E</i> region, between the magenta sphere and the spherical part of the red cut-sphere, $\rho(x^i)$ changes linearly from zero (at the intersection of the ray with the inner, magenta sphere) to unity (at the intersection with	
12	the outer, red cut-sphere.) \ldots	67 68
13	Effect of spin on horizon geometry. This image shows the intrinsic Ricci scalar on the apparent horizons in simulation $S_{0.85}^{++}$. The proper separation of the horizons along the line connecting their centers is about 1.7 <i>M</i> . Both spin effects (gradients as a function of polar angle) and tidal bulges (dark red regions near the intersection with the line connecting the horizon centers) can be seen	80
14	Normalized constraint violations for $S_{0.9}^{}$. For each resolution level, k , we plot $ \mathcal{C} $, the volume-averaged L^2 -norm of the generalized harmonic constraint energy divided by the volume-averaged L^2 -norm of the dynamical field gradients. This measure is defined in Eq. (71) of Ref. [50]. As the resolution level increases, the constraints decrease. Jumps in the constraints are attributed to changes in the domain structure, and the spike around $t \sim 3500M$ corresponds	00
15	to the merger	82
16	Differences in the final masses and spins between resolution levels. For each case, we compare χ_f and M_f of the highest resolution to the two lower resolutions. Note that, except for the older $S_{0.44}^{\pm\pm}$ simulations, all differences are $\leq 10^{-4}$. Differences in the initial masses and spins behave similarly.	86
17	Convergence of the final spin for $S_{0.9}^{}$ as a function of the ℓ of the horizon finder. We plot the difference between $\chi_f(\ell)$ and $\chi_f(\ell = 20)$ for each resolution. The adaptive horizon finder for this case chose $\ell = 8$ at the final time of the simulation.	87

х

- In the top panel, we plot our preferred fitting formula (solid line), the fourthorder polynomial in Eq. (3.4.1), and a comparison with a second-order polynomial (dashed line) for χ_f as a function of χ_i . Our data points are plotted as polygons, where more sides indicates higher resolution level. In the bottom panel, we plot the residuals of the fourth-order polynomial. We indicate our fit parameter (dotted line) and total prediction (dashed line) uncertainties (defined in Appendix A1), which in this case are nearly identical. Note that the residuals for the two lower resolution runs for $S_{0.44}^{++}$ are too large to fit in this panel.
- 19 Final spin as a function of initial spin. In the left panels we plot our data (circles) along with the fitting formula (red line) with error estimates (dashed) from several other studies. The top panel is from Ref. [29], the middle is from Ref. [27], and the bottom is from Ref. [28]. In the right panels we plot the difference between our data and the corresponding fitting formula on the left. The value r quantifies the size of the systematic error compared to the fourth-order polynomial.

90

92

- In the top panel, we plot our preferred fitting formula (solid line), the hyperbolic function in Eq. (3.4.4), and a comparison with a second-order polynomial (dashed line) for E_{rad} as a function of χ_i. Our data points are plotted as polygons, where more sides indicates higher resolution level. In the bottom panel, we plot the residuals of the hyperbolic function. We indicate our fit parameter (dotted line) and total prediction (dashed line) uncertainties (defined in Appendix A1).
 E_{rad} as a function of initial spin. In the left panels we plot our data (circles)
- 21 E_{rad} as a function of initial spin. In the left panels we plot our data (circles) along with the fitting formula (red line) with error estimates (dashed) from several other studies. The top panel is from Ref. [29], the middle is from Ref. [30], and the bottom is from Ref. [31]. In the right panels we plot the difference between our data and the corresponding fitting formula on the left. The value r quantifies the size of the systematic error compared to the 3-parameter hyperbola.
 22 Directed acyclic graph displaying the conditional dependence structure of the Dependence structure of the structure of the displaying the conditional dependence structure of the structure of the displaying the structure of the structure for the structure of the displaying the structure of the structure of the displaying the structure of the structure of the structure of the displaying the conditional dependence structure of the displaying the structure of the structure for the structure of the

LIST OF TABLES

1	Initial data parameters (radial velocity \dot{d}_0 and angular velocity Ω_0) at sepa-	
	ration $d_0 = 15.366M$ for the ten new simulations with target spin, χ_0 . Also	
	included is the approximate number of orbits until merger. Here M is the	
	sum of the Christodoulou masses at time $t = 0$	81
2	Dimensionless spins measurements. For each case, we provide the initial	
	spin magnitude of each hole and the final spin magnitude of the remnant at	
	the highest resolution. Adding the number in parentheses to the last two	
	significant digits gives the value at the next highest resolution.	84
3	Christodoulou mass measurements. For each case, we provide the total initial	
	mass of the black holes, the final mass of the remnant, and the radiated energy	
	computed from Eq. $(3.3.1)$ at the highest resolution. Adding the number	
	in parentheses to the last two significant digits gives the value at the next	
	highest resolution.	85
	0	

Preface

Citations to Previously Published Work

This dissertation includes material from work that has been previously published or submitted as follows.

- Daniel A. Hemberger, Mark A. Scheel, Lawrence E. Kidder, Béla Szilágyi, Geoffrey Lovelace, Nicholas W. Taylor, and Saul A. Teukolsky. Dynamical Excision Boundaries in Spectral Evolutions of Binary Black Hole Spacetimes. *Class. Quantum Grav.*, 30(11):115001, 2013
- D. A. Hemberger, G. Lovelace, T. J. Loredo, L. E. Kidder, M. A. Scheel, B. Szilágyi, N. W. Taylor, and S. A. Teukolsky. Final spin and radiated energy in numerical simulations of binary black holes with equal masses and equal, aligned or anti-aligned spins. *Submitted to Phys. Rev. D, arXiv:1305.5991*, 2013

Chapter 1

INTRODUCTION

Contents

1.1	Gravitational-wave astronomy	1
1.2	Overview	5

1.1 Gravitational-wave astronomy

I Historical perspective

In 1916, Albert Einstein published his general theory of relativity, which describes gravity as a geometric property of spacetime. In this theory, matter and radiation influence the curvature of spacetime, and the curvature of spacetime in turn influences the dynamics of the matter and radiation. In the same year, Karl Schwarzschild discovered a solution to Einstein's equations that later was found to imply the existence of black holes.

For almost all terrestrial applications, Newton's law of universal gravitation from the 1600's produces the same equations of motion as general relativity. Deviations between the two theories become dramatic for large masses moving at significant fractions of the speed of light. General relativity has great astrophysical relevance because one will not find such extreme systems on Earth; instead we must search the rest of the universe, which is rich with exploding stars, black holes, and colliding galaxies.

An immediate success of general relativity was its ability to explain the anomalous perihelion precession of Mercury. Even taking into account subtle effects, such as the gravitational tugs on Mercury from the other planets, Mercury's observed perihelion precession disagreed with Newton's theory by 43 arcseconds per century, or about 7% (according to an 1895 calculation). General relativity predicted a correction of exactly this amount.

Another deviation from Newton's law of gravitation was observed during the 1919 solar eclipse, when Arthur Eddington confirmed Einstein's prediction for the amount of light bending caused by the sun. It was this success that made Einstein a household name.

These early tests of general relativity probed the weak-field limit, where gravitational fields are weak and Einstein's equations can be linearized. Observing nonlinear effects in the strong-field regime is the next great test of Einstein's theory.

The first moderately strong-field test of general relativity was the indirect detection of gravitational waves from the Hulse-Taylor binary neutron star system in 1974. The prediction for orbital decay due to the emission of gravitational waves [1] was confirmed by using pulsar timing to monitor the orbital period of this binary [2]. It is an indirect detection because the gravitational waves are only inferred from electromagnetic observations.

Almost 40 years later, gravitational waves have still not been directly detected [3], but considerable progress has been made toward this goal. Since 1992, several facilities have been built around the world to detect gravitational waves, such as LIGO, Virgo, and GEO. These facilities use laser interferometry to measure the change in separation between a set of test masses as gravitational waves pass by. Additional strategies, such as resonant bar mass detectors and pulsar timing arrays, have also been employed.

Advanced LIGO, an upgrade to the LIGO facilities taking place right now and scheduled for completion in 2014, is expected to increase the theoretical detection rates by orders of magnitude [4]. It is likely that gravitational waves from black hole and neutron star events will be detected, allowing Einstein's theory to be tested for the first time in the fully nonlinear, strong-field regime. With possibly daily detections, we may only be a couple years away from an era of multi-messenger astronomy, where gravitational and electromagnetic radiation observations are used in conjunction to probe the current and future mysteries of the cosmos.

II Detection techniques

Different gravitational wave detectors are sensitive to different sources of gravitational waves. Pulsar timing arrays observe in the nHz frequency band, which corresponds to the stochastic background arising from the inspiral of supermassive black hole binaries at the centers of galaxies. LIGO observes in the 10–10⁴ Hz band, which corresponds to the inspiral and merger of binary systems containing neutron stars and stellar-mass black holes, as well as other burst and continuous sources. A space-based mission, such as LISA, would observe in the frequency range between pulsar timing arrays and LIGO, which targets supermassive black hole merger events, extreme mass-ratio inspirals, and white dwarf binaries.

The most likely sources detectable by LIGO are binary compact object coalescences. The expected signal at the detector from such a source is very small, and for this reason special techniques are necessary to extract the signal from the noise. Because these are modeled sources, a technique called optimal matched filtering can be used [5, 6]. This method employs an algorithmically chosen template waveform bank (chosen to minimize the number of templates while preserving SNR¹ to within 3%). The signal is cross-correlated with the template bank, and if the SNR reaches a certain threshold, it is flagged as a potential detection candidate [7].

The construction of every template waveform ultimately relies on numerical simulations. This is because the merger phase of the inspiral-merger-ringdown characterization is fully nonlinear and is not calculable analytically. The inspiral can be modeled with post-Newtonian

 $^{^{1}}$ SNR = signal-to-noise ratio. The threshold of 3% means that no more than 10% of events should be missed because of gaps between the waveform templates.

expansions [8], and the ringdown to a quasi-equilibrium black hole can be modeled as a perturbation to the Kerr metric [9], but the behavior during the merger can currently only be computed with numerical techniques.

Numerical simulations are very computationally expensive. A highly accurate binary black hole simulation at best costs about 10^3 – 10^4 CPU hours per orbit (depending on mass ratio and spin) using 10 GFLOP processors. This will certainly improve as numerical codes become more sophisticated and as the performance of the computing environments increases, but it is still a prohibitive expense when compared to the requirements in accuracy, number of orbits, and parameter space coverage for LIGO template banks [10, 11].

Since post-Newtonian approximations are accurate very early in the inspiral, numerical waveforms can be extended by prepending a post-Newtonian waveform in a process called *hybridization*. This procedure is successful if the overlap between hybridized waveforms matched over different regions is very high [12]. To interpolate to points in the mass-spin parameter space that have not been simulated numerically, phenomenological waveform models are used. These models are calibrated using numerical relativity results and are compared with hybridized waveforms to assess their quality [12, 13]. Because of the limited number of numerical relativity waveforms, these phenomenological waveforms comprise the majority of template banks.

Once a detection is made, detailed waveform models are necessary for *parameter estima*tion [14] (i.e. determining the masses, spins, sky location, and distance of a gravitational wave source). In addition to providing the most stringent tests of general relativity in the strong-field regime, here is where the most interesting astrophysical questions will be answered. Mass measurements will help infer mass limits and distributions of black holes and neutron stars, spin measurements will provide insight into black hole formation models and supernova processes, and localization of sources on the sky will enable coincident electromagnetic observations. Unsurprisingly, the template waveform length and accuracy requirements are much greater for parameter estimation than for detection.

1.2 Overview

Numerical simulations used in this work were generated with the Spectral Einstein Code (SpEC) [15], which is a pseudo-spectral PDE solver embedded into an infrastructure suited for solving Einstein's equations in isolated or binary compact object spacetimes. Spectral methods benefit from exponential convergence, meaning that a linear increase in resolution decreases errors exponentially, so long as the solution is smooth. Discontinuities in the fields or their derivatives will degrade the rate of convergence.

I Control systems for choosing computational coordinates

A dominant numerical challenge in evolving black hole spacetimes is the treatment of the physical singularity at the center of the black hole. One solution ("moving puncture") is to choose coordinates that effectively push the singularity off the computational domain [16], which can be achieved with the BSSN (Baumgarte-Shapiro-Shibata-Nakamura) formulation of Einstein's equations [17]. This is the most popular solution for finite difference codes [18], but there are technical problems in using spectral methods with the BSSN formulation [19].

Another solution is to excise the singularity from the computational domain. To remain well-posed, the equations must take a (symmetric) hyperbolic form (e.g. the Generalized Harmonic formulation employed by SpEC) [20]. Then, provided the excision boundary is inside the event horizon, no boundary conditions are needed because all information is flowing into the black hole along characteristics. However, excision in spectral methods is more complicated than in finite difference methods.

In finite difference excision techniques, grid points in the differencing stencil can be deactivated to avoid computations that occur near the singularity. In spectral methods, one cannot simply deactivate a collocation point because all collocation points are required to compute derivatives. To avoid computationally prohibitive alternatives, SpEC uses a reference frame where the black holes are static and stationary (called the *grid* frame) to keep the collocation points fixed from timestep to timestep. A time-dependent mapping connects this frame to the asymptotically inertial frame, and a control system monitors criteria for each component of the map to ensure that the mapping remains stable and that the grid frame remains stationary. This is called the *dual-frames method*.

The original dual-frames method [21] tracked only the black hole centers, and was sufficient for the evolution of equal mass, non-spinning binaries during the inspiral. Near merger, characteristic speeds at the excision surfaces would quickly become negative, causing the evolution equations to become ill-posed because we do not have a boundary condition to impose there. If a common horizon around the two black holes could be found sufficiently early, this problem was avoided by excising the region between the common horizon and the individual excision surfaces. The inclusion of black hole spin or unequal mass ratios caused characteristic speeds to become negative much earlier than the formation of the common horizon because the horizons were more distorted. For these configurations, relying on early common horizon excision was no longer viable.

In Chapter 2, I describe the implementation of feedback control systems in SpEC and related enhancements to our excision techniques in the dual-frames method that enabled the successful evolution of high mass ratio and nearly extremal binary systems. By controlling the shape and characteristic speeds of the horizons in addition to their centers, the excision surfaces can be made to remain pure outflow boundaries. However, this causes the computational grid to become highly distorted. To avoid the computational expense of resolving large Jacobians, additional maps and associated control systems are implemented to make the grid more flexible and conform to the shape of the excision surfaces as they conform to the shape of the horizons (we call these the "Skew" and "CutX" maps).

Enhancements to the original dual-frames method also include control systems with better stability properties. The feedback signal is a function of the difference between some quantity and its target value. This function can include derivatives, which introduce noise into the feedback loop. By using fewer numerical derivatives and including an integral term (and smoothing the feedback signal when necessary), we prevent rapid, destabilizing oscillations in the control systems.

II High-spin binary black hole evolutions

In Chapter 3, I report on a set of equal-mass binary black hole simulations with equal spin aligned (or anti-aligned) with the orbital angular momentum. Apparent horizon properties (Christodoulou mass and quasi-local spin) are used to improve phenomenological fitting formulae that predict the final mass and spin of the merger remnant as a function of initial spin. The improvements can be attributed primarily to two factors. First, spectral convergence of the code facilitates longer and more accurate evolutions for the same computational cost as finite difference codes. Second, the ability to create and evolve initial data above the "Bowen-York limit" makes high-spin data points uniquely available.

The practical impact of these improved fitting formulae is their use in approximate analytic waveform models. In particular, the quasi-normal modes in the ringdown expression will depend on the final mass and spin predictions, and even small differences in these quantities can produce considerable phase errors between the analytic and numerical waveforms [13].

There are two natural directions to expand this work: analysis of additional data from these simulations, and expansion of the parameter space. For the former, one may look at the gravitational waveforms, comparing them to post-Newtonian approximants to assess the accuracy of hybridized waveforms. For the latter, evolutions with even higher spin are being performed and will provide useful tests of the fitting formulae and conclusions for the equal mass, equal aligned spin subspace. Additionally, a catalog of simulations with more generic parameters has just been constructed [22], and the methods used here could be extended to generate more generic fitting formulae.

REFERENCES

- P. C. Peters and J. Mathews. Gravitational radiation from point masses in a Keplerian orbit. *Phys. Rev.*, 131(1):435–440, Jul 1963.
- [2] J.M. Weisberg, D.J. Nice, and J.H. Taylor. Timing Measurements of the Relativistic Binary Pulsar PSR B1913+16. Astrophys. J., 722:1030–1034, 2010.
- [3] J. Abadie et al. Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO's Sixth Science Run and Virgo's Science Runs 2 and 3. *Phys. Rev.* D, 85:082002, 2012.
- [4] J. Abadie et al. Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors. *Class. Quant. Grav.*, 27:173001, 2010.
- [5] Lee S. Finn. Detection, measurement, and gravitational radiation. *Phys. Rev. D*, 46:5236–5249, 1992.
- [6] Benjamin J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60(2):022002, Jun 1999.
- [7] B. Abbott et al. LIGO: The Laser interferometer gravitational-wave observatory. *Rept. Prog. Phys.*, 72:076901, 2009.
- [8] 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(12):124038, 2007.

- [9] S. Teukolsky. Perturbations of a rotating black hole. fundamental equations for gravitational, electromagnetic, and neutrino-field perturbations. Astrophys. J., 185:635, 1973.
- [10] Ilana MacDonald, Samaya Nissanke, and Harald P. Pfeiffer. Suitability of post-Newtonian/numerical-relativity hybrid waveforms for gravitational wave detectors. *Class. Quantum Grav.*, 28(13):134002, July 2011.
- [11] Frank Ohme. Analytical meets numerical relativity status of complete gravitational waveform models for binary black holes. *Class. Quant. Grav.*, 29:124002, 2012.
- [12] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, J. T. Whelan, B. Brügmann, P. Diener, J. Gonzalez, M. Hannam S. Husa, M. Koppitz, D. Pollney, L. Rezzolla, L. Santamaría, A. M. Sintes, U. Sperhake, and J. Thornburg. A phenomenological template family for black-hole coalescence waveforms. *Class. Quantum Grav.*, 24(19):S689– S699, 2007.
- [13] Andrea Taracchini, Yi Pan, Alessandra Buonanno, Enrico Barausse, Michael Boyle, et al. Prototype effective-one-body model for nonprecessing spinning inspiral-merger-ringdown waveforms. *Phys. Rev. D*, 86:024011, 2012.
- [14] J. Aasi et al. Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network. 2013.
- [15] http://www.black-holes.org/SpEC.html.
- [16] Manuela Campanelli, C.O. Lousto, P. Marronetti, and Y. Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Phys. Rev. Lett.*, 96:111101, 2006.
- [17] T. W. Baumgarte and S. L. Shapiro. On the numerical integration of Einstein's field equations. *Phys. Rev. D*, 59:024007, 1999. gr-qc/9810065.
- [18] Frank Loffler, Joshua Faber, Eloisa Bentivegna, Tanja Bode, Peter Diener, et al. The Einstein Toolkit: A Community Computational Infrastructure for Relativistic Astrophysics. *Class.Quant.Grav.*, 29:115001, 2012.

- [19] Wolfgang Tichy. Long term black hole evolution with the BSSN system by pseudo-spectral methods. *Phys. Rev.*, D80:104034, 2009.
- [20] 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.
- [21] 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.
- [22] Abdul H. Mroue, Mark A. Scheel, Bela Szilagyi, Harald P. Pfeiffer, Michael Boyle, Daniel A. Hemberger, Lawrence E. Kidder, Geoffrey Lovelace, Sergei Ossokine, Nicholas W. Taylor, Anil Zenginoglu, Luisa T. Buchman, Tony Chu, Evan Foley, Matthew Giesler, Robert Owen, and Saul A. Teukolsky. A catalog of 171 high-quality binary black-hole simulations for gravitational-wave astronomy. 2013. arXiv:1304.6077.

Chapter 2

Dynamical Excision Boundaries in Spectral Evolutions of Binary Black Hole Spacetimes

Contents

2.1	Introduction	13
2.2	Control Theory	15
2.3	Control Systems in SpEC	18
2.4	Control Systems for Maps	27
2.5	Size control	50
2.6	Merger and ringdown	58
2.7	Control systems for efficiency	62
2.8	Summary	63
$\mathbf{A1}$	Implementation of exponentially-weighted averaging	64
$\mathbf{A2}$	Details on computing the CutX map weight function	66
A3	Estimation of zero-crossing times	69

Simulations of binary black hole systems using the Spectral Einstein Code (SpEC) are done on a computational domain that excises the regions inside the black holes. It is imperative that the excision boundaries are outflow boundaries with respect to the hyperbolic evolution equations used in the simulation. We employ a time-dependent mapping between the fixed computational frame and the inertial frame through which the black holes move. The timedependent parameters of the mapping are adjusted throughout the simulation by a feedback control system in order to follow the motion of the black holes, to adjust the shape and size of the excision surfaces so that they remain outflow boundaries, and to prevent large distortions of the grid. We describe in detail the mappings and control systems that we use. We show how these techniques have been essential in the evolution of binary black hole systems with extreme configurations, such as large spin magnitudes and high mass ratios, especially during the merger, when apparent horizons are highly distorted and the computational domain becomes compressed. The techniques introduced here may be useful in other applications of partial differential equations that involve time-dependent mappings.

2.1 Introduction

Feedback control systems are ubiquitous in technological applications. They are found, for example, in thermostats, autopilots, chemical plants, and cruise control in automobiles. The purpose of a control system is to keep some measured output (such as the temperature in a room) at some desired value by adjusting some input (such as the power to a furnace).

In the last few years, feedback control systems have also found applications in the field of numerical relativity, particularly in simulations of binary black hole systems that employ spectral methods and excision techniques [1, 2, 3, 4, 5].

Black hole excision is a means of avoiding the physical singularities that lurk inside black holes. The idea is to solve Einstein's equations only in the region outside apparent horizons, cutting out the region inside the horizons. The boundaries of the excised regions are called excision boundaries. Causality ensures that the excision boundaries and the excised interiors cannot affect the physics of the exterior solution, and an appropriate hyperbolic formulation of Einstein's equations [6, 7] can ensure that gauge and constraint-violating degrees of freedom also do not propagate out of the excised region.

Excision is straightforward for black holes that remain stationary in the coordinates that are used in the simulation, but excision becomes more complicated when the black holes move or change shape. For numerical methods based on finite differencing, the excision boundaries can be changed at every time step by activating or deactivating appropriate grid points and adjusting differencing stencils [8, 9, 10, 11, 12, 13, 14, 15]. However, for spectral numerical methods, there is no equivalent of deactivating individual grid points; instead, spectral methods are defined in finite extended spatial regions with smooth boundaries. The nearest equivalent to the finite-difference excision approach would be interpolating all variables to a new slightly offset grid at every time step, which would be computationally expensive. Therefore, spectral numerical methods need a different approach to reconcile the need for moving black holes with the need for a fixed excision boundary inside of each black hole.

The solution [1] to this problem adopted by our group makes use of multiple coordinate systems. We call "inertial coordinates" those coordinates that asymptotically correspond to an inertial observer; in these coordinates the black holes orbit each other, have distorted shapes, and approach each other as energy is lost to gravitational radiation. Spectral methods are applied in another coordinate system, "grid coordinates", in which the excision boundaries are spherical and stationary. We connect grid coordinates with inertial coordinates by means of an analytic mapping function \mathcal{M} that depends on some set of time-dependent parameters $\lambda(t)$. These parameters must be continually adjusted so that each spherical, stationary grid-frame excision boundary is mapped to a surface in the inertial frame that follows the motion and the shape of the corresponding black hole as the system evolves. It is this adjustment of each parameter $\lambda(t)$ that is accomplished by means of feedback control systems, one control system per parameter.

In this paper, we describe in detail the mapping functions and the corresponding feedback control systems that we use to handle black hole excision with spectral methods. Some earlier implementations have been described previously [5, 1, 2, 4, 16], but there have been many improvements that now allow spectral excision methods to produce robust simulations of binary black hole systems, including those with unequal masses, high spins, and precession.

We describe these improvements here. In Sec. 2.2 we review control theory and present simple examples of control systems. In Sec. 2.3 we discuss the implementation of control systems in the SpEC [17] code. Sections 2.4 and 2.5 detail the coordinate mappings used in SpEC and the feedback control systems used to control them. In Sec. 2.6 we describe the transition to the post-merger domain with a single excision boundary. In Sec. 2.7 we describe applications of control systems in SpEC besides the ones used to adjust map parameters. We summarize in Sec. 2.8.

2.2 Control Theory

To motivate control theory, we begin with a simple example: cruise control in an automobile. Suppose we wish to control the speed v of a car that is driving up an incline of angle θ . The equation of motion for this system is

$$\frac{dv}{dt} = -\frac{\eta}{m}v + \frac{F}{m} - g\sin\theta, \qquad (2.2.1)$$

where *m* is the mass of the car, *g* is the gravitational acceleration, η is a drag coefficient, and *F* is a force supplied by the car's engine. We wish to determine *F* so as to cause the car to maintain a speed of $v = v_0$, even if the angle θ changes as the car climbs the incline. To do this, we choose the force at time *t* to be:

$$\frac{F(t)}{m} = K_P Q(t) + K_I \int_0^t Q(\tau) \, d\tau, \qquad (2.2.2)$$

where the control system is turned on at time t = 0, where K_I and K_P are constants, and where $Q(t) = v_0 - v(t)$ is the quantity that we wish to drive to zero. We call Q the control error. Substituting Eq. (2.2.2) into Eq. (2.2.1) and differentiating with respect to time yields

$$\frac{d^2Q}{dt^2} + \left(K_P + \frac{\eta}{m}\right)\frac{dQ}{dt} + K_I Q = g\cos\theta\frac{d\theta}{dt},\tag{2.2.3}$$

which is the equation for a damped, forced harmonic oscillator. By choosing K_P to produce critical damping, and choosing K_I to set a timescale, this choice will drive v toward v_0 as desired. The basic structure of a control system is easily understood as a feedback loop (see Fig. 1). The simulation (e.g., a binary black hole simulation, or a car driving up an incline) produces some measure of error, Q(t), that defines the deviation from some desired target value. This error acts as the input for the control system, which then computes a control signal U(t)(e.g., the derivative of a map parameter, or the force supplied by a car's engine) that will minimize the error Q(t) when fed back into the simulation.



Figure 1: A generic control circuit. The simulation outputs a measure of error, Q, which is used by the control system. The control system then outputs a signal, U, which changes the behavior of the simulation.

A simple and effective way to compute U(t) is to make it a linear combination of the error, Q(t), and integrals and/or derivatives of the error. The term proportional to the error acts to reduce the deviation from the desired value, the terms proportional to derivatives of the error act to oppose rapid deviations, and the terms proportional to the integrals act to reduce any persistent deviation or offset that accumulates over time. In the cruise control example, we used a proportional and an integral term only in Eq. (2.2.2).

We now turn to another example of a control system that is more closely related to the way we use control systems in binary black hole simulations. Consider two coordinate systems, (x, t) and (\bar{x}, \bar{t}) , that are related by the map

$$x = \bar{x} - V(t)\,\bar{t},\tag{2.2.4}$$

$$t = \bar{t}.\tag{2.2.5}$$

We wish to control the parameter V(t) so that a wave $f(\bar{x}, \bar{t}) = f(\bar{x} - c\bar{t})$ that propagates at speed $\bar{v} = c$ in the (\bar{x}, \bar{t}) coordinates will propagate at some arbitrary desired speed v_d in the (x, t) coordinates. According to Eq. (2.2.4), the speed of the wave in the (x, t) coordinates is v = c - V(t). Therefore we define a control error Q to be

$$Q = c - V(t) - v_d, (2.2.6)$$

and we construct a control system that drives this control error to zero. If we choose the control signal U(t) to be d^2V/dt^2 , then the simplest feedback loop that can be constructed uses only a term proportional to the error and amplified by a "gain" K_P . Then the evolution of V(t) is given by

$$\frac{d^2 V}{dt^2} = K_P Q = K_P \left[c - V(t) - v_d \right].$$
(2.2.7)

The solution to this equation is of the form

$$V(t) = c - v_d + A_1 \sin(\alpha t) + A_2 \cos(\alpha t), \qquad (2.2.8)$$

where $\alpha := \sqrt{K_P}$, which is oscillatory for $K_P > 0$ and divergent for $K_P < 0$. Thus we see that adding only a proportional term to the control signal U(t) is insufficient, since it does not reduce the control error Q.

However, if we add a derivative term to the feedback equation,

$$\frac{d^2V}{dt^2} = K_P Q + K_D \frac{dQ}{dt},\tag{2.2.9}$$

then the solution is of the form

$$V(t) = c - v_d + e^{-\frac{1}{2}K_D t} \left[B_1 \sin(\beta t) + B_2 \cos(\beta t) \right], \qquad (2.2.10)$$

where $\beta := \frac{1}{2}\sqrt{4K_P - K_D^2}$. This solution is stable with an exponentially damped envelope when $4K_P \ge K_D^2$, which will cause $v \to v_d$ as $t \to \infty$.

Notice that this control system allows us to choose a V(t) such that v is the opposite sign of c and the wave is left-going in the (x, t) frame instead of right-going. This behavior can be seen in the example in Fig. 2.

The overdamped solution (dotted line in Fig. 2) has a persistent offset that can be ameliorated by adding an integral term to the feedback equation, as was done in the cruise



Figure 2: The speed in (x, t) coordinates, v, is plotted for a family of gains K_P with $K_D = 1$ fixed. The wave velocity is c = -0.2 and the desired velocity is $v_d = 0.5$ (dashed line). The controller turns on at t = 2. For gains in the stable region, where $K_P \ge 0.25$, v settles down to v_d . One overdamped solution with $K_P = 0.01$ is plotted for comparison (dotted line).

control example. In principle we could continue to add more terms, but in practice, it is usually sufficient to use a PID (proportional-integral-derivative) controller, which has terms proportional to the control error, its integral, and its derivative. When the underlying system is unknown, this is the best controller to use [18].

2.3 Control Systems in SpEC

Control systems are used in SpEC for several purposes. The most important is their role in handling moving, excised black holes in a spectral evolution method. We use a dual-frame method [1] in which the grid is fixed in some coordinates (t, x^i) but the components of dynamical fields are expressed in a different coordinate system (\bar{t}, \bar{x}^i) . We call (t, x^i) the grid coordinates and (\bar{t}, \bar{x}^i) the inertial coordinates. Figure 3 shows an example domain decomposition in both coordinate systems. The two coordinate systems are connected by a map \mathcal{M} that depends on time-dependent parameters $\lambda(t)$. The excision boundaries are exactly spherical and stationary in grid coordinates. In inertial coordinates, the apparent horizons move and distort as determined by the solution of Einstein's equations supplemented by our gauge conditions. The parameters $\lambda(t)$ need to be controlled so that the excision boundaries in the inertial frame follow the motion and shapes of the apparent horizons. We use control systems to accomplish this.



Figure 3: a) Computational domain in grid coordinates; the black hole centers are at rest and the excision boundaries are spherical. b) Same domain in inertial coordinates near merger; the excision boundaries move and distort to track the apparent horizons.

The particular maps that we use will be discussed in Sec. 2.4. In this section we will describe how we construct the control system for a general parameter $\lambda(t)$, including how we define the relationship between the control error Q(t) and the control signal U(t), how we smooth out noise in the control system, and how we dynamically adjust the feedback parameters and timescales.

I Definition of control errors and control signals

We represent a general time-dependent map parameter $\lambda(t)$ as a polynomial in time with a piecewise constant Nth derivative:

$$\lambda(t) = \sum_{n=0}^{N} \frac{1}{n!} (t - t_i)^n \lambda_i^n, \quad \text{for} \quad t_i \le t < t_{i+1}, \quad (2.3.1)$$

where for each time interval $t_i \leq t < t_{i+1}$ the quantities λ_i^n are constants.

At the beginning of each new time interval t_i , we set the constants λ_i^n in Eq. (2.3.1) as follows. First for n = N we set $\lambda_i^N = U(t_i)$, where U(t) is the control signal defined in detail below. For n < N we set $\lambda_i^n = d^n \lambda / dt^n |_{t=t_i}$, where the derivative is evaluated at the end of the previous time interval. In this way all the derivatives of $\lambda(t)$ except the Nth derivative are continuous across intervals. The goal will be to compute the control signal U(t) so as to drive the map parameter $\lambda(t)$ to some desired behavior. Before we describe how to compute the control signal U(t), we first discuss the control error Q, which will be used in the computation of U(t).

To appropriately define the control error Q, one must answer the question of how a small change in the map parameter corresponds to a change in the observed variables. If the control error is defined to be too large, then the controller will consistently overshoot its target, potentially leading to unstable behavior; conversely, if the control error is defined to be too small, then the controller may never be able to reach its target value.

If there exists a target value of $\lambda(t)$, call it λ_{target} , that does not depend on the map but may depend on other observable quantities in the system (call them A, B, \ldots), then we define

$$Q = \lambda_{\text{target}}(A, B, \ldots) - \lambda. \tag{2.3.2}$$

The goal is to drive Q to zero and thereby drive λ to λ_{target} .

If instead, as often happens for nonlinear systems, the target value of λ depends on λ itself, even indirectly, then we define Q differently using a generalization of Eq. (2.3.2): we

require that λ attains its desired value when $Q \to 0$, and we require that

$$\frac{\partial Q}{\partial \lambda} = -1 + \mathcal{O}(Q). \tag{2.3.3}$$

The primary motivation for this condition is its anticipated use in relating time derivatives of Q to those of λ (see Eq. (2.3.8), below). Note that in either Eqs. (2.3.2) or (2.3.3), Q could in principle be multiplied by an arbitrary factor; however, if this were done, that factor would need to be taken into account in the computation of the control signal U(t) below. Without loss of generality we assume no additional scaling.

In the case of several map parameters $\lambda_a(t)$ with corresponding Q_a , where *a* is an index that labels the map parameters, the desired value of some λ_a may depend on a different map parameter λ_b . In this case, we generalize Eq. (2.3.3) and require that each Q_a satisfy

$$\frac{\partial Q_a}{\partial \lambda_b} = -\delta_{ab} + \mathcal{O}(Q), \qquad (2.3.4)$$

where δ_{ab} is a Kronecker delta. This criterion ensures that we can treat each λ_a independently when all control errors are small. A way of understanding Eq. (2.3.4) is to consider a set of Q'_a that are obtained via Eq. (2.3.3) without regard to coupling between different λ_a . Then a set of Q_a satisfying Eq. (2.3.4) can be obtained by diagonalizing the matrix $\partial Q'_a / \partial \lambda_b$. In the remainder of this section we assume that if there are multiple map parameters λ_a , the corresponding control errors Q_a satisfy Eq. (2.3.4). We therefore drop the *a* indices and write equations for U(t) in terms of a single Q(t) satisfying Eq. (2.3.3) that represents the control error for a single $\lambda(t)$.

The control error Q(t) is a function of several observables. In the case of the mapping functions that are designed to move and distort the excision boundaries to follow the motion and shapes of the apparent horizons, Q(t) is some function of the current position or shape of one or more apparent horizons. The precise definition of Q(t) is different for each map parameter, and depends on the details of how each map parameter couples to the observables. We will discuss the control error Q for each of the map parameters in Sec. 2.4. But it is not necessary to know the exact form of the control error in order to compute the control signal U(t); it suffices to know only that $U(t_i)$ is equal to λ_i^N in Eq. (2.3.1), and that the control error Q obeys Eq. (2.3.3).

We now turn to the computation of the control signal U(t). There is some flexibility in the control law determining U(t), so long as key feedback mechanisms are in place (as shown in Sec. 2.2). In SpEC, we use either a standard PID controller,

$$U(t) = a_0 \int Q(t) dt + a_1 Q(t) + a_2 \frac{dQ}{dt},$$
(2.3.5)

or a special PD (proportional-derivative) controller,

$$U(t) = \sum_{k=0}^{K} a_k \frac{d^k Q}{dt^k},$$
(2.3.6)

where typically K = 2.

We set the constants a_k so that the system damps Q to zero on some timescale τ_d that we choose. We assume that τ_d is longer than the interval $t_{i+1} - t_i$ defined in Eq. (2.3.1), so that we can approximate Q(t) and $\lambda(t)$ as smooth functions rather than as functions with piecewise constant Nth derivatives. Under this assumption, we write

$$U(t) = d^N \lambda / dt^N. \tag{2.3.7}$$

We also write

$$dQ/dt = (\partial Q/\partial \lambda)(d\lambda/dt),$$

= $-d\lambda/dt.$ (2.3.8)

In the first line we have neglected the time dependence of other parameters besides λ that enter into Q under the assumption that the control system timescale is shorter than the timescales of the quantities that we want to control. In the second line we have used Eq. (2.3.3)
and we have assumed that Q is small. Similarly we write

$$d^{2}Q/dt^{2} = \frac{\partial Q}{\partial \lambda} \frac{d^{2}\lambda}{dt^{2}} + \frac{\partial^{2}Q}{\partial \lambda^{2}} \left(\frac{d\lambda}{dt}\right)^{2},$$

$$= \frac{\partial Q}{\partial \lambda} \frac{d^{2}\lambda}{dt^{2}},$$

$$= -\frac{d^{2}\lambda}{dt^{2}},$$
 (2.3.9)

where in the second line we have retained only terms of linear order in dQ/dt (and therefore in $d\lambda/dt$). For higher derivatives we continue to retain only terms linear in Q and its derivatives, so from Eq. (2.3.7) we obtain

$$U(t) = d^N \lambda / dt^N,$$

= $-d^N Q / dt^N.$ (2.3.10)

For the PID controller and N = 2, combining Eqs. (2.3.5) and (2.3.10) yields

$$-\frac{d^2Q}{dt^2} = a_0 \int Q(t) \, dt + a_1 \, Q(t) + a_2 \frac{dQ}{dt}.$$
 (2.3.11)

If we choose $a_0 = 1/\tau_d^3$, $a_1 = 3/\tau_d^2$, and $a_2 = 3/\tau_d$, then the solution to Eq. (2.3.11) will be exponentially damped on the timescale τ_d ,

$$Q \propto e^{-t/\tau_d}.\tag{2.3.12}$$

The same exponential damping holds for the PD controller, Eq. (2.3.6), for appropriate choices of the parameters a_k . For K = 2 and N = 3, the parameters a_k are identical to those in the PID case above.

II Averaging out noise

The PID controller, Eq. (2.3.5), is computed by measuring the control error Q, its time integral, and its time derivative. The PD controller, Eq. (2.3.6), may require multiple derivatives of Q(t) depending on the order K. Generally only Q, and not its derivatives or integrals, is available in the code at any given time step. The simplest way to compute the derivatives of Q is by finite differencing in time, and the simplest way to compute the integral is by a numerical quadrature.

We measure the control error Q at time intervals of length τ_m , where we choose $\tau_m < t_{i+1}-t_i$. The measuring time interval τ_m is then used as the time step for finite difference stencils and quadratures.

The measured Q is typically a function of apparent horizon locations or shapes, and this measured Q includes noise caused by the finite resolution of the evolution, and the finite residual and finite number of iterations of the apparent horizon finder. Taking a numerical derivative of Q amplifies the noise, and then this noise is transferred to the control signal via Eqs. (2.3.5) or (2.3.6), and then to the map. If the noise amplitude is too large, the control system will become unstable. The PID controller generally handles noise better than the PD controller for two reasons. First, each successive numerical derivative amplifies noise even further, so using only one derivative instead of two (or more) results in a more accurate control signal. Second, the inclusion of an integral term acts to further smooth the control signal.

In some cases, however, the noise in Q can be problematic even for the PID controller. In these cases we implement direct averaging of the control error in one of two ways: 1) we perform a polynomial fit of order N to the previous M measurements of the control error, where M > N, or 2) we perform an exponentially-weighted average, with timescale τ_{avg} , of all previous control error measurements and their derivatives and integrals. The latter is our preferred method, which we describe in detail in Appendix A1.

III Dynamic timescale adjustment

In the previous section we have identified four timescales relevant for each control system. The first is the damping timescale τ_d ; this describes how quickly the control error Q(t) falls to zero, and therefore how quickly the map parameter $\lambda(t)$ approaches its desired value. The second timescale is the control interval $\Delta t_i := t_{i+1} - t_i$, which represents how often the Nth derivative of the function $\lambda(t)$ in Eq. (2.3.1) is updated. The third is the measurement timescale τ_m , which indicates how often the control error Q is measured. The fourth is the averaging timescale τ_{avg} , which is used to smooth the control error Q (and its derivatives and integrals) for use in computing the control signal U(t). These timescales are not all independent; for example we have assumed $\Delta t_i < \tau_d$ in deriving Eq. (2.3.11), and we have assumed $\tau_m < \Delta t_i$ so that we can obtain smooth measurements of the derivatives of Q.

Because binary black hole evolutions are nonlinear dynamic systems, we adjust the damping timescale, τ_d , throughout the simulation. We then set the three other timescales, Δt_i , τ_m , and τ_{avg} , based on the current value of the damping timescale τ_d , as we now describe.

The timescale τ_d should be shorter than the timescale on which the physical system changes; otherwise, the control system cannot adjust the map parameters quickly enough to respond to changes in the system. But if the timescale τ_d is too small, then the measurement timescale τ_m on which we compute the apparent horizon must also be small, meaning that frequent apparent horizon computations are needed; this is undesirable because computing apparent horizons is computationally expensive. We would like to adjust τ_d in an automatic way so that it is relatively large during the binary black hole inspiral, decreases during the plunge and merger, and increases again as the remnant black hole rings down. In the canonical language of control theory, we would like to implement "gain scheduling" [19], tuning the behavior of the control system for different operating regimes.

We do this as follows: For all map parameters (except the $\ell = 0$ component of the horizon shape map, which is treated differently; see Sec. 2.5.III), the damping timescale is a generic function of Q and \dot{Q} , i.e., the error in its associated map parameter and its derivative. Whenever we adjust the control signal U(t) at interval t_i , we also tune the timescale in the following way:

$$\tau_d^{i+1} = \beta \tau_d^i, \tag{2.3.13}$$

where typically

$$\beta = \begin{cases} 0.99, & \text{if } \dot{Q}/Q > -1/2\tau_d \text{ and } |Q| \text{ or } |\dot{Q}\tau_d| > Q_t^{\text{Max}} \\ 1.01, & \text{if } |Q| < Q_t^{\text{Min}} \text{ and } |\dot{Q}\tau_d| < Q_t^{\text{Min}} \\ 1, & \text{otherwise.} \end{cases}$$
(2.3.14)

Here Q_t^{Min} and Q_t^{Max} are thresholds for the control error Q. The idea is to keep $|Q| < Q_t^{\text{Max}}$ so that the map parameters are close to their desired values, but to also keep $|Q| > Q_t^{\text{Min}}$ because an unnecessarily small |Q| means unnecessarily small timescales and therefore a large computational expense (because the apparent horizon must be found frequently). For binary black hole simulations where the holes have masses M_A and M_B , we find that the following choices work well:

$$Q_t^{\text{Max}} = \frac{2 \times 10^{-3}}{M_A/M_B + M_B/M_A}$$
(2.3.15)

$$Q_t^{\rm Min} = \frac{1}{4} Q_t^{\rm Max}.$$
 (2.3.16)

Once we have adjusted the timescale τ_d for every control system, we then use these timescales to choose the times t_{i+1} in Eq. (2.3.1) at which we update the polynomial coefficients λ_i^n in the map parameter expression $\lambda(t)$,

$$\Delta t_i := t_{i+1} - t_i = \alpha_d \min(\tau_d), \tag{2.3.17}$$

where typically $\alpha_d = 0.3$, and the minimum is taken over all map parameters (except for the $\ell = 0$ component of the horizon shape map). This ensures that the coefficients λ_i^n , are updated faster than the physical system is changing, and faster than the control system is damping. For α_d too large, we find that the control system becomes unstable.

We also use the timescale τ_d to choose the interval τ_m at which we measure the control error. For many map parameters, the associated control error is a function of apparent horizon quantities, which is why we desire τ_m to be as large as possible. But a certain number of measurements are needed for each Δt_i so that the control signals, defined in Eqs. (2.3.5) and (2.3.6), are sufficiently accurate and the control system is stable. We choose

$$\tau_m = \alpha_m \Delta t_i, \tag{2.3.18}$$

where α_m is typically between 0.25 and 0.3. In other words, we measure the control error three or four times before we update the control signal. This also ensures that the averaging timescale is greater than the measurement timescale, as we typically choose $\tau_{\text{avg}} \sim 0.25\tau_d$.

2.4 Control Systems for Maps

In a SpEC evolution, we transform the grid coordinates, x^i , into inertial coordinates, \bar{x}^i , through a series of elementary maps as in [1, 2, 3, 4, 5]. Several maps have been added and many improvements have been made since their original introduction. The full transformation is $\bar{x}^i = \mathcal{M}x^i$, where

$$\mathcal{M} = \mathcal{M}_{\text{Translation}} \circ \mathcal{M}_{\text{Rotation}} \circ \mathcal{M}_{\text{Scaling}}$$

$$\circ \mathcal{M}_{\text{Skew}} \circ \mathcal{M}_{\text{CutX}} \circ \mathcal{M}_{\text{Shape}}.$$
 (2.4.1)

Below we will describe each of these maps and how we measure the error in their parameters. Before we do this, however, we describe our domain decomposition and how we measure apparent horizons, because information from the grid and the horizons is used to determine the maps.

In the grid coordinates x^i , the domain decomposition looks like Fig. 4. There are two excision boundaries, A and B, which are spheres in grid coordinates. The grid-coordinate centers of these excision boundaries we will call C_H^i , where H is either A or B. The excision boundaries (and therefore C_H^i) remain fixed throughout the evolution. The purpose of many of the maps is to move the mapped centers $\bar{C}_H^i := \mathcal{M}(C_H^i)$ along with the centers of the apparent horizons. We measure the apparent horizons in an intermediate frame \hat{x}^i , which we call the *distorted* frame. This frame is connected to the grid frame by the map

$$\mathcal{M}_{\text{Distortion}} = \mathcal{M}_{\text{CutX}} \circ \mathcal{M}_{\text{Shape}}.$$
(2.4.2)

We will discuss exactly what $\mathcal{M}_{\text{Shape}}$ and $\mathcal{M}_{\text{CutX}}$ are below, but for now we need only demand that $\mathcal{M}_{\text{Distortion}}$ has two properties: The first is that it leaves the centers of the excision boundaries invariant, i.e.,

$$\hat{C}_H^i := \mathcal{M}_{\text{Distortion}}(C_H^i) = C_H^i.$$
(2.4.3)

The second property is that at each excision boundary H, $\mathcal{M}_{\text{Distortion}}$ leaves angles invariant. That is, if we define grid-frame polar coordinates (r_H, θ_H, ϕ_H) centered about excised region H in the usual way,

$$x^0 = C_H^0 + r_H \sin \theta_H \cos \phi_H, \qquad (2.4.4)$$

$$x^1 = C_H^1 + r_H \sin \theta_H \sin \phi_H, \qquad (2.4.5)$$

$$x^2 = C_H^2 + r_H \cos \theta_H, (2.4.6)$$

and similarly for polar coordinates $(\hat{r}_H, \hat{\theta}_H, \hat{\phi}_H)$ centered about \hat{C}_H^i in the \hat{x}^i frame, then the second property means that

$$\hat{\theta}_H = \theta_H,$$

$$\hat{\phi}_H = \phi_H,$$
(2.4.7)

on the excision boundary. Note that the index i on spatial coordinates ranges over (0, 1, 2) in this paper.

We find the apparent horizons in the frame \hat{x}^i by a fast flow method [20]. Each horizon is represented as a smooth surface in this frame, and is parameterized in terms of polar coordinates centered around \hat{C}_H^i : the radius of the horizon at each $(\hat{\theta}_H, \hat{\phi}_H)$ is given by

$$\hat{r}_{H}^{\text{AH}}(\hat{\theta}_{H}, \hat{\phi}_{H}) = \sum_{\ell m} \hat{S}_{\ell m}^{H} Y_{\ell m}(\hat{\theta}_{H}, \hat{\phi}_{H}), \qquad (2.4.8)$$

where again the index H labels which excision boundary is enclosed by the apparent horizon. Here $Y_{\ell m}(\hat{\theta}_H, \hat{\phi}_H)$ are spherical harmonics and $\hat{S}^H_{\ell m}$ are expansion coefficients that describe the shape of the apparent horizon.

We search for apparent horizons in the distorted frame \hat{x}^i rather than in the grid frame x^i or the inertial frame \bar{x}^i because this simplifies the formulae for the control systems. In particular, this choice decouples the control errors of the shape map $Q_{\ell m}$ (defined in Sec. 2.4.V, Eq. (2.4.49)) for different values of (ℓ, m) , and it decouples the errors $Q_{\ell m}$ from the control errors of the maps connecting the distorted and inertial frames.

For each surface H, we define the center of the apparent horizon

$$\hat{\xi}_{H}^{i} = \frac{\int_{H} \hat{x}^{i} (\hat{r}_{H}^{\text{AH}})^{2} d\hat{\Omega}}{\int_{H} (\hat{r}_{H}^{\text{AH}})^{2} d\hat{\Omega}},$$
(2.4.9)

where the integrals are over the surface. In the code, it suffices to use the following approximation of Eq. (2.4.9), which becomes exact as the surface becomes spherical:

$$\hat{\xi}_{H}^{0} = C_{H}^{0} - \sqrt{3/2\pi} \,\Re(\hat{S}_{11}^{H}), \qquad (2.4.10)$$

$$\hat{\xi}_{H}^{1} = C_{H}^{1} + \sqrt{3/2\pi} \,\Im(\hat{S}_{11}^{H}), \qquad (2.4.11)$$

$$\hat{\xi}_H^2 = C_H^2 + \sqrt{3/4\pi} \,\hat{S}_{10}^H. \tag{2.4.12}$$

Here we have used the property $\hat{C}_{H}^{i} = C_{H}^{i}$, Eq. (2.4.3). We distinguish the center C_{H}^{i} of the excision boundary from the center $\hat{\xi}_{H}^{i}$ of the corresponding apparent horizon. The former is fixed in time in the grid frame, but the latter will change as the metric quantities evolve. The purpose of several of the maps (namely $\mathcal{M}_{\text{Scaling}}$, $\mathcal{M}_{\text{Rotation}}$, and $\mathcal{M}_{\text{Translation}}$ described below) is to ensure that $\hat{\xi}_{H}^{i} - C_{H}^{i}$ is driven toward zero; i.e., to ensure that the centers of the excision boundaries track the centers of the apparent horizons.

We now describe each of the maps comprising \mathcal{M} and how we measure the error in their parameters. The order in which we discuss the maps is not the same as the order in which the maps are composed so that we can discuss the simplest maps first. To produce less cluttered equations in the following descriptions of the maps, we omit accents on variables that represent specific frames (i.e. we just write x instead of \hat{x} or \bar{x}) whenever the input or output frame of the map is unambiguous or explicitly stated.

I Scaling

The scaling map, $\mathcal{M}_{\text{Scaling}}$, causes the grid to shrink or expand (and the excision boundaries to move respectively closer together or farther apart) in the inertial frame, thereby allowing the grid to follow the two black holes as their separation changes.

This map transforms the radial coordinate with respect to the origin (i.e., the center of the outermost sphere in Fig. 4), such that the region near the black holes is scaled uniformly by a factor a and the outer boundary is scaled by a factor b,

$$R \mapsto aR + (b-a)R^3/R_{OB}^2.$$
 (2.4.13)

Here R is the radial coordinate, R_{OB} is the radius at the outer boundary, a is a parameter that will be determined by a control system, and b is another parameter that will be determined empirically.

In [1], the scaling map was simply $x^i \mapsto a(t)x^i$, which is recovered by Eq. (2.4.13) if b = a. In that case, as the black holes inspiral together and a decreases, the outer boundary of the grid decreases as well. For long evolutions, the outer boundary decreases so much that we can no longer extract gravitational radiation far from the hole. The addition of b to the map alleviates this difficulty, allowing the motion of the outer boundary to be decoupled from the motion of the holes. We choose b by an explicit functional form

$$b(t) = 1 - 10^{-6} t^3 / (2500 + t^2), \qquad (2.4.14)$$

which is designed to keep the outer boundary from shrinking rapidly, but to allow the boundary to move inward at a small speed, so that zero-speed modes are advected off the grid (and thus need no boundary condition imposed on them). We choose a cubic function of time in Eq. (2.4.14) because we have found that at least the first two time derivatives of b(t) must vanish initially, or else a significant ingoing pulse of constraint violations is produced at the outer boundary.

We control the scale factor a of this map using the error function

$$Q_a = a(dx^0 - 1), (2.4.15)$$

where

$$dx^{i} = \frac{\hat{\xi}_{A}^{i} - \hat{\xi}_{B}^{i}}{C_{A}^{0} - C_{B}^{0}}.$$
(2.4.16)

We assume that the separation vector $C_A^i - C_B^i$ in the grid frame is parallel to the x-axis. The idea is that the distorted-frame separation of the horizon centers along the x-axis is driven to be the same as the separation of the excision boundary centers.

To show that the control system for *a* obeys Eq. (2.3.3), we consider the change in Q_a under variations of *a* with all other maps held fixed, and with the inertial-coordinate centers $\bar{\xi}_{H}^{i}$ of the horizons held fixed.¹ This means that the distorted-frame centers of the horizons $\hat{\xi}_{H}^{i}$ appearing in Eq. (2.4.16) change with variations of *a*. The second term in Eq. (2.4.13) is small when evaluated near the horizon where $(R/R_{OB})^2 \ll 1$, so we can write the action of the scaling map as

$$\bar{\xi}^i_H = a\hat{\xi}^i_H, \qquad (2.4.17)$$

and therefore under variations δa ,

$$0 = \delta \bar{\xi}_H^i = \hat{\xi}_H^i \,\delta a + a \,\delta \hat{\xi}_H^i. \tag{2.4.18}$$

Taking variations of Eq. (2.4.15), using Eq. (2.4.18) to substitute for $\delta \hat{\xi}_{H}^{i}$, and noting that the excision-boundary centers C_{H}^{i} are constants and do not vary with a, we obtain

$$\delta Q_a = -\delta a, \tag{2.4.19}$$

¹ We consider inertial-frame quantities to be fundamental and determined by the solution of Einstein's equations plus gauge conditions, and therefore independent (modulo numerical truncation error) of the numerical grid and of the maps that we use to construct, move, and distort that grid.

so that Eq. (2.3.3) is satisfied.

To verify that the more general decoupling equation, Eq. (2.3.4), is satisfied for the scaling map, we must show, in addition to Eq. (2.4.19), that the quantity dx^0 defined in Eq. (2.4.16) is invariant under the action of the other maps, at least in the limit that all the control errors Q are small. We consider each map in turn. The translation map moves both apparent horizons together, so it leaves dx^0 invariant. Changes in the rotation map parameters will change dx^0 only by an amount proportional to the control errors of the rotation map (see Eqs. (2.4.25) and (2.4.26) below). The skew map (below) leaves the centers of the excision boundaries invariant. Because the intent of the rotation, translation, and scaling maps is to drive the centers of the apparent horizons toward the centers of the excision boundaries, this means that the skew map changes dx^0 by an amount proportional to the control errors of the rotation, translation, and scaling maps. Finally, the shape and CutX maps connect the distorted and grid frames, so they cannot affect dx^0 .

II Rotation

The rotation map, $\mathcal{M}_{\text{Rotation}}$, is a rigid 3D rotation about the origin that tracks the orbital phase and precession of the system,

$$x^{i} \mapsto \mathcal{R}^{i}_{j} x^{j}, \text{ where } \mathcal{R} = \begin{pmatrix} \cos \vartheta \cos \varphi & -\sin \vartheta & \cos \vartheta \sin \varphi \\ \sin \vartheta \cos \varphi & \cos \vartheta & \sin \vartheta \sin \varphi \\ -\sin \varphi & 0 & \cos \varphi \end{pmatrix}.$$
 (2.4.20)

The pitch and yaw map parameters (φ, ϑ) are controlled so as to align the line segment connecting the apparent horizon centers with the distorted-frame x-axis. Note that the map parameters (φ, ϑ) are functions of time, and are not to be confused with the polar coordinates (ϕ_H, θ_H) centered about each excision boundary. Then the control error is given by

$$Q_{\varphi} = -\frac{\hat{\xi}_{A}^{2} - \hat{\xi}_{B}^{2}}{\hat{\xi}_{A}^{0} - \hat{\xi}_{B}^{0}}$$
(2.4.21)

$$Q_{\vartheta} = \frac{\hat{\xi}_{A}^{1} - \hat{\xi}_{B}^{1}}{(\hat{\xi}_{A}^{0} - \hat{\xi}_{B}^{0})\cos\varphi}.$$
(2.4.22)

In the case of extreme precession where $\varphi \to \pi/2$, these equations are insufficient because Q_{ϑ} diverges. Our solution is to use quaternions, which avoid this singularity (for a complete discussion, see [21]).

The control errors Q_{ϑ} and Q_{φ} obey Eq. (2.3.4). To show this, consider variations of ϑ and φ with other maps held fixed, and with the inertial-coordinate centers $\bar{\xi}_{H}^{i}$ of the horizons held fixed. The rotation map, Eq. (2.4.20), implies that under these variations,

$$\delta \bar{\xi}_H^i = 0 = \mathcal{R}_j^i \delta \hat{\xi}_H^j + (\delta \mathcal{R})_j^i \hat{\xi}_H^j.$$
(2.4.23)

Multiplying this equation by \mathcal{R}^{-1} we obtain

$$\delta \hat{\xi}_H^i = -(\mathcal{R}^{-1} \delta \mathcal{R})_j^i \hat{\xi}_H^j, \qquad (2.4.24)$$

which yields

$$\frac{\partial \hat{\xi}_{H}^{i}}{\partial \varphi} = \mathcal{A}_{j}^{i} \hat{\xi}_{H}^{j}, \quad \text{where} \quad \mathcal{A} := -\mathcal{R}^{-1} \frac{\partial \mathcal{R}}{\partial \varphi} = \begin{pmatrix} 0 & 0 & -1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

$$\frac{\partial \hat{\xi}_{H}^{i}}{\partial \vartheta} = \mathcal{B}_{j}^{i} \hat{\xi}_{H}^{j}, \quad \text{where} \quad \mathcal{B} := -\mathcal{R}^{-1} \frac{\partial \mathcal{R}}{\partial \vartheta} = \begin{pmatrix} 0 & \cos \varphi & 0 \\ -\cos \varphi & 0 & -\sin \varphi \\ 0 & \sin \varphi & 0 \end{pmatrix}.$$

$$(2.4.25)$$

We can now verify Eq. (2.3.4) for the special case where the indices a and b in Eq. (2.3.4) are either ϑ and φ . This result is obtained in a straightforward way by differentiating Eqs. (2.4.21) or (2.4.22) with respect to ϑ or φ , and substituting Eqs. (2.4.25) or (2.4.26).

In addition, the control errors Q_{φ} and Q_{ϑ} are independent (to leading order in the control errors) of changes in the parameters of the other maps that make up \mathcal{M} : Variations of

Eqs. (2.4.21) and (2.4.22) with respect to the scaling map parameter a are zero, because both the numerator and denominator of Q_{φ} and Q_{ϑ} scale in the same way with a. Similarly, Q_{φ} and Q_{ϑ} are independent of the translation map, since both apparent horizons are translated by the same amount. The skew map can change Q_{φ} and Q_{ϑ} , but only by an amount proportional to control errors, because the skew map leaves C_{H}^{i} invariant and other maps ensure that $\hat{\xi}_{H}^{i}$ are close (within a control error) to C_{H}^{i} . The shape and CutX maps cannot affect Q_{φ} and Q_{ϑ} because those maps connect the grid and distorted frames (and therefore they cannot change the distorted-frame horizon centers $\hat{\xi}_{H}^{i}$).

III Translation

The translation map, $\mathcal{M}_{\text{Translation}}$, transforms the Cartesian coordinates, x^i , according to

$$x^i \mapsto x^i + f(R)T^i, \tag{2.4.27}$$

where f(R) is a Gaussian centered on the origin with a width set such that f(R) falls off to machine precision at the outer boundary radius, and $T^{i}(t)$ are translation parameters that are adjusted by a control system.

The translation map moves the grid to account for any drift of the "center of mass" of the system (as computed assuming point masses at the apparent horizon centers) in the inertial frame. This drift can be caused by momentum exchange between the black holes and the surrounding gravitational field [22, 23], by anisotropic radiation of linear momentum to infinity, or by linear momentum in the initial data. This is the third map (the other two are rotation and scaling) that drives the centers of the apparent horizons toward the centers of the excision boundaries. The apparent horizon centers $\hat{\xi}_A^i$ and $\hat{\xi}_B^i$ represent six degrees of freedom: one is fixed by the scaling map, two by rotation, and three by translation.

The control errors will be more complicated than for the other control systems because translation and rotation do not commute. We define the control errors Q_T^i for each of the translation directions i = 0, 1, 2 as

$$Q_T^i = a\mathcal{R}_j^i \left[\hat{\xi}_B^j + \mathcal{P}_k^j (\hat{\xi}_A^k - \hat{\xi}_B^k) \right], \qquad (2.4.28)$$

where \mathcal{R} is the rotation matrix in Eq. (2.4.20), and \mathcal{P} is some matrix yet to be determined, which may depend on the rotation parameters (φ, ϑ) and on the constants C_H^i , but which may not depend on the translation parameters T^i . This control error must have the property that $Q_T^i = 0$ when $\hat{\xi}_H^i = C_H^i$, so our first restriction on \mathcal{P} is that it satisfies

$$0 = C_B^i + \mathcal{P}_j^i (C_A^j - C_B^j).$$
(2.4.29)

To check Eq. (2.3.4), we note that near the black holes we can neglect the last term in Eq. (2.4.13) and we can use $f(R) \sim 1$ in Eq. (2.4.27), so that the apparent horizon centers in the inertial and distorted frames are related by

$$a\mathcal{R}^{i}_{j}\hat{\xi}^{j}_{H} = \bar{\xi}^{i}_{H} - T^{i}.$$
 (2.4.30)

Inserting Eq. (2.4.30) into Eq. (2.4.28), we obtain

$$Q_T^i = \bar{\xi}_B^i - T^i - (\mathcal{RPR}^{-1})_j^i (\bar{\xi}_A^j - \bar{\xi}_B^j).$$
(2.4.31)

The second term in Eq. (2.4.31) is the only term that depends on the translation parameter T^{i} , so we have

$$\frac{\partial Q_T^i}{\partial T^j} = -\delta_{ij} \tag{2.4.32}$$

When varying map parameters, the inertial-frame horizon centers remain fixed, so the only other term in Eq. (2.4.31) that depends on map parameters is the last term, which depends on (ϑ, φ) because of the rotation matrices and because of the (ϑ, φ) dependence in \mathcal{P} . We can therefore write the variation of Q_T^i with respect to (ϑ, φ) as

$$\frac{\partial Q_T^i}{\partial W} = -\frac{\partial}{\partial W} (\mathcal{RPR}^{-1})^i_j (\bar{\xi}^j_A - \bar{\xi}^j_B), \qquad (2.4.33)$$

$$= -\frac{\partial}{\partial W} (\mathcal{RPR}^{-1})^i_j a \mathcal{R}^j_k (\hat{\xi}^k_A - \hat{\xi}^k_B), \qquad (2.4.34)$$

where W stands for either ϑ or φ , and where in the last line we have used Eq. (2.4.30). To obey Eq. (2.3.4), $\partial Q_T^i / \partial W$ must either be zero, or on the order of a control error Q. For i = 1, 2, the quantity $(\hat{\xi}_A^i - \hat{\xi}_B^i)$ is proportional to the control error of the rotation map, Eqs. (2.4.21) and (2.4.22), so these terms can be neglected in Eq. (2.4.34). However, for i = 0, $(\hat{\xi}_A^i - \hat{\xi}_B^i)$ is not proportional to a control error; instead $(\hat{\xi}_A^0 - \hat{\xi}_B^0)$ is driven to a constant finite value of $C_A^0 - C_B^0$ by the scaling control system. Therefore, in order to satisfy Eq. (2.3.4), we require

$$-\frac{\partial}{\partial W}(\mathcal{RPR}^{-1})\mathcal{R}\begin{pmatrix}1\\0\\0\end{pmatrix}=0.$$
 (2.4.35)

We find that we can satisfy both Eqs. (2.4.29) and (2.4.35) by choosing

$$\mathcal{P} = \frac{1}{C_B^0 - C_A^0} \begin{pmatrix} C_B^0 & -C_B^1 & -C_B^2 \\ C_B^1 & C_B^0 + C_B^2 \tan \varphi & 0 \\ C_B^2 & -C_B^1 \tan \varphi & C_B^0 \end{pmatrix}.$$
 (2.4.36)

Here we have again assumed that the separation between the centers of the excision boundaries is parallel to the x-axis, i.e., $C_A^1 = C_B^1$ and $C_A^2 = C_B^2$.

IV Skew

In Fig. 4, a prominent feature of the domain decomposition is a plane (a vertical line in the two-dimensional figure) that is perpendicular to the grid-frame x-axis and lies between the two excision boundaries A and B. We call this plane the "cutting plane".

The skew map, $\mathcal{M}_{\text{Skew}}$, acts on the distorted-frame coordinates, in which the cutting plane is still perpendicular to the *x*-axis. The skew map leaves the coordinates y, z unchanged, but changes the *x*-coordinate in order to give a skewed shape to the cutting plane, as shown in Fig. 5. Let x_C^i be the intersection point of the line segment connecting the excision boundary



Figure 4: Two-dimensional projection of the domain decomposition near the two black holes A and B. Shown are boundaries between subdomains. Each subdomain takes the shape of a spherical shell, a distorted cylindrical shell, or a distorted cylinder. Additional spherical-shell subdomains (not shown) surround the outer boundary of the figure and extend to large radius. This domain decomposition is explained in detail in the Appendix of [5]. The red, magenta, and blue surfaces are those for which the function $f_A(r_A, \theta_A, \phi_A)$ from Eq. (2.4.44) has a discontinuous gradient, as described in Sec. 2.4.V. centers and the cutting plane. The action of the skew map is defined as

$$x^{0} \mapsto x^{0} + W \sum_{j=1,2} \tan \left[\Theta^{j}(t)\right] \cdot (x^{j} - x_{C}^{j}),$$
 (2.4.37)

$$x^j \mapsto x^j, \text{ for } j = 1,2$$
 (2.4.38)

where W is a radial Gaussian function centered around x_C^i , and the angles $\Theta^j(t)$ are the time-dependent map parameters. For j = 1, 2 the parameter $\Theta^j(t)$ is the angle between the mapped and unmapped x^j -axis when projected into the $x^{3-j} = x_C^{3-j}$ plane. Note that the line intersecting x_C^i and parallel to the x-axis is left invariant by the skew map². The width of the Gaussian W is set such that W is below machine precision at the innermost wave extraction sphere. This implies that the spherical-shell subdomains used to evolve the metric in the wave zone will not be affected by the skew map.

The purpose of the skew map is to (as much as possible) align the cutting plane with the surfaces of the apparent horizons in the region where the surfaces are closest to the cutting plane. We derive the control system in charge of setting the parameters Θ^{j} by the following condition: We demand that the angle between the mapped cutting plane and the x-axis at x_{C}^{i} be driven to the (weighted) average of the angles at which the mapped apparent horizons intersect the same x-axis.

The input coordinates to the skew map are the coordinates in the distorted frame. For each horizon, we therefore measure an angle in the distorted frame as follows: Let x_H^{Int} be the distorted-frame x-coordinate at which apparent horizon H intersects the line segment connecting the centers of the excision boundaries. For j = 1, 2, we calculate the normal to the surface at this intersection point. This normal is projected into the $x^{3-j} = x_C^{3-j}$ plane. We define Θ_H^j as the angle between the projected normal and the x-axis. Thus, projecting the normal into the y = const. plane gives Θ_H^z and vice versa (see Fig. 6). We then compute

² The horizon centers $\hat{\xi}_{H}^{i}$ are not invariant under the skew map because they do not necessarily lie on this line. This deviation, which is quantified in Eqs. (2.4.10)–(2.4.12), leads to a coupling of $\mathcal{O}(Q)$ with the translation, rotation, and scaling maps.



Figure 5: Snapshot of the grid, viewed in the mapped frame, from a binary black hole evolution with $M_A = 8M_B$ shortly before the common horizon forms. Left: without $\mathcal{M}_{\text{Skew}}$. Right: with $\mathcal{M}_{\text{Skew}}$. The gray areas are excised regions. In the left panel, the grid is compressed as the excision boundaries track increasingly skewed apparent horizons, and the evolution is terminated because of excessively large constraint violation in the compressed region. This problem is resolved by the inclusion of $\mathcal{M}_{\text{Skew}}$.



Figure 6: A diagram of the skew map in the mapped frame. The horizon surfaces are drawn in black and the excision surfaces are drawn in red. The dotted red line represents the line segment connecting the excision surface centers, which is parallel to the *x*-axis. The normal to either horizon at the intersection point with this segment is represented by the straight black lines. The cutting plane is represented by the (skewed) blue curve, and the normal to this plane at x_C^i is represented by the central green line. The green lines near the two excision surfaces are constructed parallel to the central green line, where parallelism is defined assuming a Euclidean background. There are three angles involved in the skew control system – the angle between the normal to the skewed cutting plane and the *x*-axis, Θ^j , and the angle between the normal to either horizon and the normal to the skewed cutting plane, Θ_H^j . Skew control acts to minimize Θ_H^j . We measure Θ_H^j in the unmapped (distorted) frame, but Θ_H^j is invariant under the skew map for $W \approx 1$.

a weighted average of the Θ_H^j such that the horizon closer to the cutting plane has a larger effect on the skew angles,

$$\Theta_{\text{Avg}}^{j} = \frac{w_A \Theta_A^{j} + w_B \Theta_B^{j}}{w_A W(C_A) + w_B W(C_B)},$$
(2.4.39)

where w_A, w_B are averaging weights,

$$w_H = \exp\left[-\frac{x_C^0 - x_H^{\text{Int}}}{x_C^0 - C_H^0}\right].$$
 (2.4.40)

Here C_H^i is the center of the excision boundary H.

We measure Θ_{Avg}^{j} in the distorted frame, and in this frame the cutting plane is always normal to the *x*-axis. Therefore, thinking about the desired result in the distorted frame, we see that the control system for the skew map should drive Θ_{Avg}^{j} to zero. This leads us to consider the following control error for the skew angles: $Q_{\Theta}^{j} = \Theta_{Avg}^{j}$. Assuming that the function W is unity near the apparent horizons, we find that

$$\frac{\partial \Theta_{\text{Avg}}^{j}}{\partial \Theta^{j}} = -1 + \mathcal{O}(Q), \qquad (2.4.41)$$

in agreement with Eq. (2.3.3). In deriving Eq. (2.4.41), it is helpful to observe that the partial derivative in Eq. (2.4.41) is taken with the inertial-frame apparent horizon held fixed. For a fixed inertial-frame horizon, the only map that can change the shape (as opposed to merely the center) of the distorted-frame horizon (and thus Θ^{j}_{Avg}) is the skew map.

We do not use $Q_{\Theta}^{j} = \Theta_{Avg}^{j}$ for the entire evolution, however, because at early times, when the coordinate distance between the apparent horizons is larger than their combined radii, the skew map is not needed. Furthermore, the skew map can cause difficulties early during the run, especially during the "junk radiation" phase when the horizons are oscillating in shape. For this reason, we gradually turn on the skew map as the black holes approach each other. This is done by defining a roll-on function g that is zero when the horizons are far apart, and one when they are close together. This roll-on function is defined as

$$g = \frac{1}{2} \left[1 - \tanh\left(10\frac{x_A^{\text{Int}} - x_B^{\text{Int}}}{C_A^0 - C_B^0} - 5\right) \right].$$
 (2.4.42)

For values $g < 10^{-3}$ the skew map is turned off completely; this is not strictly necessary, but it saves some computation. Given the function g, we define the control error as

$$Q_{\Theta}^{j} = g\Theta_{\text{Avg}}^{j} - (1-g)\Theta^{j}.$$
(2.4.43)

This control error drives the skew angles to zero when the black holes are far apart and drives Θ_{Avg}^{j} to zero as they approach each other.

V Shape control

We define the shape map $\mathcal{M}_{\text{Shape}}$ as:

$$x^{i} \mapsto x^{i} \left(1 - \sum_{H} \frac{f_{H}(r_{H}, \theta_{H}, \phi_{H})}{r_{H}} \sum_{\ell m} Y_{\ell m}(\theta_{H}, \phi_{H}) \lambda^{H}_{\ell m}(t) \right).$$
(2.4.44)

The index H goes over each of the two excised regions A and B, and the map is applied to the grid-frame coordinates. The polar coordinates (r_H, θ_H, ϕ_H) centered about excised region H are defined by Eqs. (2.4.4)–(2.4.6), the quantities $Y_{\ell m}(\theta_H, \phi_H)$ are spherical harmonics, and $\lambda_{\ell m}^H(t)$ are expansion coefficients that parameterize the map near excision region H. The function $f_H(r_H, \theta_H, \phi_H)$ is chosen to be unity near excision region H and zero near the other excision region, so that the distortion maps for the two black holes are decoupled. Specifically, $f_H(r_H, \theta_H, \phi_H)$ is determined as illustrated in Fig. 4. For excision region A in the figure, $f_A(r_A, \theta_A, \phi_A) = 1$ between the excision boundary and the magenta surface, it falls linearly to zero between the magenta and red surfaces, and it is zero everywhere outside the red surface. This means that the gradient of $f_A(r_A, \theta_A, \phi_A)$ is discontinuous on the red surface, the magenta surface, and the blue surfaces in the figure. Because we ensure that these discontinuities occur on subdomain boundaries, they cause no difficulty with using spectral methods. Around excision region B, $f_B(r_B, \theta_B, \phi_B)$ is chosen similarly.

In previous implementations of the shape map [3], the functions $f_H(r_H, \theta_H, \phi_H)$ were chosen to be smooth Gaussians centered around each excision boundary rather than to be piecewise linear functions. We find smooth Gaussians to be inferior for two reasons. The first is that piecewise linear functions are easier and faster to invert (the inverse map is required for interpolation to trial solutions during apparent horizon finding). The second is that for smooth Gaussians, it is necessary to choose the widths of the Gaussians sufficiently narrow so that the Gaussian for excision region A does not overlap the Gaussian for excision region B and vice versa, so that the maps and control systems for A and B remain decoupled. However, decreasing the width of the Gaussians increases the Jacobians of the map, producing coordinates that are stretched and squeezed nonuniformly. We found that this form of "gridstretching" significantly increased the computational resources required to resolve the solution to a given level of accuracy. In other words, with smooth Gaussians we were forced to add computational resources just to resolve the large Jacobians.

A map very similar to Eq. (2.4.44) is also described in [4]. The difference compared to

this work is the choice of $f_H(r_H, \theta_H, \phi_H)$, which corresponds to a different choice of domain decomposition. The control system used to choose the map parameters $\lambda_{\ell m}^H(t)$ in [4] is also different than what is described here.

We control the expansion coefficients $\lambda_{\ell m}^{H}(t)$ of the shape map, Eq. (2.4.44), so that each excision boundary, as measured in the intermediate frame (t, \hat{x}^{i}) , has the same shape as the corresponding apparent horizon. In other words, we desire

$$\frac{\hat{r}_{\rm AH}}{\langle \hat{r}_{\rm AH} \rangle} = \frac{\hat{r}_{\rm EB}}{\langle \hat{r}_{\rm EB} \rangle},\tag{2.4.45}$$

where for brevity we have dropped the index H that labels the excision boundary. Here the angle brackets mean averaging over angles, $\hat{r}_{AH}(\hat{\theta}, \hat{\phi})$ is the radial coordinate of the apparent horizon defined in Eq. (2.4.8), and $\hat{r}_{EB}(\theta, \phi)$ is the radial coordinate of the excision boundary, which from Eq. (2.4.44) can be written

$$\hat{r}_{\rm EB} = r_{\rm EB} - \sum_{\ell m} Y_{\ell m}(\hat{\theta}, \hat{\phi}) \lambda_{\ell m}(t).$$
(2.4.46)

Here $r_{\rm EB}$ is the radius of the (spherical) excision boundary in the grid frame. In deriving Eq. (2.4.46) we have used the relations $\mathcal{M}_{\rm CutX} = \mathbf{1}$, $f(r, \theta, \phi) = 1$, $\hat{\theta} = \theta$, and $\hat{\phi} = \phi$, which hold on the excision boundary.

Combining Eqs. (2.4.8), (2.4.45), and (2.4.46) yields

$$\frac{r_{\rm EB} - \sum_{\ell m} Y_{\ell m}(\hat{\theta}, \hat{\phi}) \lambda_{\ell m}(t)}{r_{\rm EB} - Y_{00} \lambda_{00}(t)} = \frac{\sum_{\ell m} \hat{S}_{\ell m} Y_{\ell m}(\hat{\theta}, \hat{\phi})}{\hat{S}_{00} Y_{00}},$$
(2.4.47)

which we can satisfy by demanding that

$$\lambda_{\ell m}(t) + \hat{S}_{\ell m} \frac{r_{\rm EB} - Y_{00} \lambda_{00}(t)}{\hat{S}_{00} Y_{00}} = 0, \quad \ell > 0.$$
(2.4.48)

Therefore, given an apparent horizon and given a value of λ_{00} , a control system can be set up for each (ℓ, m) pair with $\ell > 0$, and the corresponding control errors are

$$Q_{\ell m} = -\lambda_{\ell m}(t) - \hat{S}_{\ell m} \frac{r_{\rm EB} - Y_{00}\lambda_{00}(t)}{\hat{S}_{00}Y_{00}}, \quad \ell > 0.$$
(2.4.49)

Driving these $Q_{\ell m}$ to zero produces an excision boundary that matches the shape of the corresponding apparent horizon.

Equation (2.4.49) determines $\lambda_{\ell m}(t)$ only for $\ell > 0$, and Eqs. (2.4.45)–(2.4.49) can be satisfied for arbitrary values of the remaining undetermined map coefficient $\lambda_{00}(t)$. Determination of $\lambda_{00}(t)$ is complicated enough that it is described in its own section, Sec. 2.5.

Note that Eq. (2.4.49) does not satisfy Eq. (2.3.4) because $\partial Q_{\ell m}/\partial \lambda_{00} = \hat{S}_{\ell m}/\hat{S}_{00}$, which does not vanish even if all the control errors are zero. For small distortions, this coupling between λ_{00} and $Q_{\ell m}$ seems to cause little difficulty. However, at times when the shapes and sizes of the horizons change rapidly (e.g., during the initial "junk radiation" phase, after the transition to a single excised region when a common apparent horizon forms, and after rapid gauge changes), simulations using Eq. (2.4.49) exhibit relatively large and high-frequency oscillations in $Q_{\ell m}$ that usually damp away but occasionally destabilize the evolution. Construction of a control system in which all $Q_{\ell m}$ are fully decoupled will be addressed in a future work.

VI CutX

The map $\mathcal{M}_{\text{CutX}}$ applies a translation along the grid-frame *x*-axis in the vicinity of the black holes, but without moving the excision boundaries (or the surrounding spherical shells) themselves. The action of the map is shown in Fig. 7.

The goal of this map is to allow for a slight motion of the cutting plane toward the smaller excision boundary. This is important for binary black hole systems with mass ratio $q \gtrsim 8$. As such a binary gets closer to merger, the inertial-frame coordinate distance between each excision boundary and the cutting plane decreases. Eventually, this distance falls to zero for the larger excision boundary, producing a coordinate singularity in which the Jacobian of one of the other maps (often the shape map) becomes infinite. By pushing the cutting plane toward the smaller excision boundary, the map \mathcal{M}_{CutX} avoids this singularity. Even for



Figure 7: Two-dimensional projection of the domain decomposition near the two black holes A and B. The function ρ in \mathcal{M}_{CutX} is zero on the magenta boundaries, one on the red boundaries, and goes linearly between zero and one in the "radial" direction between these boundaries. The gradient of ρ is discontinuous across the solid magenta, red, and blue boundaries. Dotted lines show the subdomain boundaries under the action of \mathcal{M}_{CutX} .

evolutions in which the inertial-coordinate distance between the larger excision boundary and the cutting plane remains finite but becomes small, the map $\mathcal{M}_{\text{CutX}}$ prevents large Jacobians from developing and thus increases numerical accuracy (because it is no longer necessary to add computational resources to resolve the large Jacobians).

Figure 8 shows the domain decomposition in the inertial frame for a binary black hole simulation with q = 8. The top panel shows the case without \mathcal{M}_{CutX} , and the compressed grid near the larger excision boundary is evident. The lower panel shows the case with \mathcal{M}_{CutX} , which removes the extreme grid compression. Looking at the Jacobian of the mapping from the inertial frame to the grid frame shortly before merger, we find that the infinity norm of the determinant of the Jacobian is twice as large in the case without \mathcal{M}_{CutX} .



Figure 8: Snapshot of the grid, viewed in the inertial frame, from a binary black hole evolution with $M_A = 8M_B$ shortly before the common horizon forms. Top: without $\mathcal{M}_{\text{CutX}}$. Bottom: with $\mathcal{M}_{\text{CutX}}$. The gray areas are excised regions. In the top panel, the grid is compressed as the larger excision boundary approaches the cutting plane; this is especially evident in the long narrow subdomains immediately adjacent to the larger excision boundary. Using $\mathcal{M}_{\text{CutX}}$ relieves the grid compression. The map \mathcal{M}_{CutX} is written as

$$x^{0} \mapsto x^{0} + \rho(x^{i})F(t),$$
 (2.4.50)

$$x^j \mapsto x^j, \qquad j = 1, 2, \tag{2.4.51}$$

where F(t) is adjusted by a control system.

We choose $\rho(x^i)$ to be zero within either of the spherical shell regions, i.e., inside the magenta spheres around A and B in Fig. 7, and it is also zero outside the outer magenta sphere in the same figure. We set $\rho(x^i) = 1$ on the solid red boundaries in Fig. 7, and everywhere between the two solid red boundaries that enclose excision boundary B. Every other region on the grid is bounded by a smooth red boundary on one side and a smooth magenta boundary on the other; in these regions ρ varies linearly between zero and one. The solid blue boundaries are locations (in addition to the solid red and solid magenta boundaries) in which the gradient of ρ is discontinuous. Full details for the calculation of ρ can be found in Appendix A2.

As with the skew map, the map $\mathcal{M}_{\text{CutX}}$ is inactive for most of the inspiral. Let x_H^{Exc} be the distorted-frame *x*-coordinate of the intersection of the excision boundary *H* with the line segment connecting the centers of the excision boundaries. This is similar to x_H^{Int} defined earlier; the difference is that x_H^{Int} refers to a point on the apparent horizon and x_H^{Exc} refers to a point on the excision boundary. The quantity x_H^{Exc} is time-dependent because it depends on the shape map.

The map \mathcal{M}_{CutX} is turned off completely as long as

$$\frac{x_C^0 - x_H^{\text{Exc}}}{x_H^{\text{Exc}} - C_H^0} \ge \frac{1}{2},$$
(2.4.52)

where C_H^i are the excision boundary centers, and x_C^i are the coordinates of the intersection of the cutting plane and the line segment connecting the centers of the excision boundaries, as introduced in Sec. 2.4.IV. Let t_0 be the coordinate time at which the map $\mathcal{M}_{\text{CutX}}$ is activated.³ We designate the target position $x = x_T$ of the cutting plane as

$$x_T = \left. \frac{\Delta x_A \cdot x_A^{\text{Exc}} + \Delta x_B \cdot x_B^{\text{Exc}}}{\Delta x_A + \Delta x_B} \right|_{t=t_0}, \qquad (2.4.53)$$

where

$$\Delta x_H = \left| x_H^{\text{Exc}} - C_H^0 \right|, \quad H = A, B.$$
(2.4.54)

Recall that x_H^{Exc} is time-dependent (as it is measured in the distorted frame), so we save the value of x_T as calculated at the time when the map $\mathcal{M}_{\text{CutX}}$ is activated, rather than recalculating x_T at every measurement time.

Now that we have a target x_T , we could designate $x_T - x_C^0$ as the target value to the function F(t) by setting $Q_F = -F + x_T - x_C^0$ and have the control system drive Q_F to zero, thus driving the x-coordinate of the cutting plane to x_T . However, because we turn on the CutX control system suddenly at time t_0 , we must be more careful. Turning on any control system suddenly will produce some transient oscillations, unless the control error and its relevant derivatives and integrals are all initially zero. In the case of the CutX map, which is turned on during a very dynamic part of the simulation when excision boundaries need to be controlled very tightly, these oscillations can prematurely terminate the run by, for example, pushing an excision boundary outside its accompanying apparent horizon.

To turn on $\mathcal{M}_{\text{CutX}}$ gradually, we replace x_T by a new time-dependent target function T(t) that gradually approaches x_T at late times but produces a control error Q_F with $Q_F = \partial Q_F / \partial t = 0$ at the activation time $t = t_0$. We start by estimating the time t_H^{XC} at which x_H^{Exc} will reach the cutting plane, where H is either A or B. This is done by the method

³ Both $\mathcal{M}_{\text{Skew}}$ and $\mathcal{M}_{\text{CutX}}$ are turned on late in the run, but the condition triggering their activation is different, given the different nature of the problems they address. $\mathcal{M}_{\text{Skew}}$ is needed for all runs where the horizons eventually intersect the line segment connecting their excision centers at an angle sufficiently different from $\pi/2$. This will happen essentially for all runs except the simplest head-on collisions. $\mathcal{M}_{\text{CutX}}$, on the other hand, is primarily for unequal-mass runs where the larger excision surface encroaches upon the cutting plane near merger.

described in Appendix A3. We then let t^{XC} be the minimum of t_A^{XC} and t_B^{XC} . The smooth target function T(t) is then defined by

$$T(t; t_0, x_T, t^{\text{XC}}) = x_T + \exp\left[-\left(\frac{t - t_0}{0.3 \cdot t^{\text{XC}}}\right)^2\right] \times \left[x_C^0 - x_T + F(t_0) + (t - t_0) \cdot \frac{\partial F(t)}{\partial t}\Big|_{t=t_0}\right].$$
(2.4.55)

Designating $T - x_C^0$ as the target for F(t)

$$Q_F = T - x_C^0 - F (2.4.56)$$

leads to

$$Q_F|_{t=t_0} = 0, \quad \left. \frac{\partial Q_F}{\partial t} \right|_{t=t_0} = 0,$$
 (2.4.57)

so at the activation time the control system does not produce transients. Furthermore, in the limit of small Q_F ,

$$x_C^0 + F(t)|_{t=t^{\rm XC}} \approx x_T,$$
 (2.4.58)

i.e., by the time the excision boundary would have touched the cutting plane (and formed a grid singularity), the smooth target function T(t) has approached x_T , and the cutting plane will have reached its designated target location, x_T . As the run proceeds, the behavior of the cutting plane is determined by the map $\mathcal{M}_{\text{CutX}}$ while the motion of the excision boundaries is determined by gauge dynamics and the behavior of the other control systems. We continue to monitor the distance between the cutting plane and the excision boundaries, and if it is predicted to touch within a time less than $\tau_d/0.15$, the $\mathcal{M}_{\text{CutX}}$ control system is reset, constructing a new target function $T(t; t_0, x_T, t^{\text{XC}})$, where t_0 is the time of the reset, and x_T, t^{XC} are also recalculated based on the state of the grid at this reset time.

Each time a new target function T is constructed, the damping time τ_d of the $\mathcal{M}_{\text{CutX}}$ control system is set to be $t^{\text{XC}}/2$.

The control systems responsible for \mathcal{M}_{CutX} and \mathcal{M}_{Shape} are decoupled, as \mathcal{M}_{CutX} controls the location of the cutting plane, leaving the excision boundary unchanged, while \mathcal{M}_{Shape} controls the shape of the excision boundary, leaving the location of the cutting plane unchanged. Recall that \mathcal{M}_{CutX} and \mathcal{M}_{Shape} define the mapping from the grid frame to the distorted frame. As the apparent horizons are found in the distorted frame, and the other maps only depend upon measurements of the horizons, \mathcal{M}_{CutX} and \mathcal{M}_{Shape} are decoupled from the other maps.

2.5 Size control

In this section we discuss how we control the spherical part of the map given by Eq. (2.4.44), namely, the coefficients λ_{00}^{H} for each excision boundary H. We apply the same method to each excision boundary, so for clarity, in this section we again drop the index H from the coefficients λ_{00}^{H} and $S_{\ell m}^{H}$, and from the coordinates $(r_{H}, \theta_{H}, \phi_{H})$.

I Characteristic speed control

Controlling the size of the excision boundary is more complicated than simply keeping the excision boundary inside the apparent horizon. This is because black hole excision requires conditions on the characteristic speeds of the system, and if these conditions are not enforced they are likely to be violated.

The minimum characteristic speed at each excision boundary is given by

$$v = -\alpha - \bar{\beta}^i \bar{n}_i - \bar{n}_i \frac{\partial \bar{x}^i}{\partial t}, \qquad (2.5.1)$$

where α is the lapse, $\bar{\beta}^i$ is the shift, and \bar{n}_i is the normal to the excision boundary pointing out of the computational domain, i.e., toward the center of the hole. Here (\bar{t}, \bar{x}^i) are the inertial frame coordinates. The first two terms in Eq. (2.5.1) describe the coordinate speed of the ingoing (i.e., directed opposite to \bar{n}_i) light cone in the inertial frame, and the last term accounts for the motion of the excision boundary (which is fixed in the grid frame) with respect to the inertial frame.

In our simulations, we impose no boundary condition whatsoever at each excision boundary. Therefore, well-posedness requires that all of the characteristic speeds, and in particular the minimum speed v, must be non-negative; in other words, characteristics must flow into the hole. In practice, if v becomes negative, the simulation is terminated, because a boundary condition is needed, but we do not have one to impose. This can occur even when the excision boundary is inside the horizon. In this case, one might argue that if the simulation is able to continue without crashing (e.g. by becoming unstable inside the horizon), that the solution outside the horizon would not be contaminated. We have not explored this possibility. Instead, we choose to avoid this situation by terminating the code if negative speeds are detected.

Therefore, we would like to control λ_{00} in such a way that v remains positive. We start by writing v in a way that separates terms that explicitly depend on $\dot{\lambda}_{00}$ from terms that do not. To do this we expand the derivative in the last term of Eq. (2.5.1) as

$$\frac{\partial \bar{x}^{i}}{\partial t} = \frac{\partial \bar{x}^{i}}{\partial \hat{x}^{a}} \frac{\partial \hat{x}^{a}}{\partial t}
= \frac{\partial \bar{x}^{i}}{\partial \hat{t}} \frac{\partial \hat{t}}{\partial t} + \frac{\partial \bar{x}^{i}}{\partial \hat{x}^{j}} \frac{\partial \hat{x}^{j}}{\partial t}
= \frac{\partial \bar{x}^{i}}{\partial \hat{t}} + \frac{\partial \bar{x}^{i}}{\partial \hat{x}^{j}} \frac{\partial \hat{x}^{j}}{\partial t},$$
(2.5.2)

where a in the first line of Eq. (2.5.2) is a four-dimensional spacetime index, and the last line of Eq. (2.5.2) follows from $\partial \hat{t}/\partial t = 1$. Inserting this into Eq. (2.5.1) yields

$$v = -\alpha - \bar{\beta}^i \bar{n}_i - \bar{n}_i \frac{\partial \bar{x}^i}{\partial \hat{t}} - \hat{n}_i \frac{\partial \hat{x}^i}{\partial t}, \qquad (2.5.3)$$

where \hat{x}^i is the frame that is obtained from the grid frame by applying the distortion map; see Eq. (2.4.2). We then use Eq. (2.4.44) to rewrite this as

$$v = -\alpha - \bar{\beta}^{i} \bar{n}_{i} - \bar{n}_{i} \frac{\partial \bar{x}^{i}}{\partial \hat{t}} + \hat{n}_{i} \frac{x^{i}}{r} \sum_{\ell m} Y_{\ell m}(\theta, \phi) \dot{\lambda}_{\ell m}(t), \qquad (2.5.4)$$

where we have used the relations $f(r, \theta, \phi) = 1$ and $\mathcal{M}_{CutX} = \mathbf{1}$ when evaluating the distortion map $\mathcal{M}_{Distortion}$ on the excision boundary.

By combining all the terms that do not explicitly depend on $\dot{\lambda}_{00}$ into a quantity v_0 , we obtain

$$v = v_0 + \hat{n}_i \frac{x^i}{r} Y_{00} \dot{\lambda}_{00}.$$
 (2.5.5)

Thus, the characteristic speed v can be thought of as consisting of two parts: one part, v_0 , that depends on the position and shape of the excision boundary and the values of the metric quantities there, and another part that depends on the average speed of the excision boundary in the direction of the boundary normal.

We now construct a control system that drives the characteristic speed v to some target speed v_T . This is a control system that controls the derivative quantity $\dot{\lambda}_{00}$, as opposed to directly controlling the map quantity λ_{00} . We choose

$$Q = (\min(v) - v_T)/\langle -\Xi \rangle, \qquad (2.5.6)$$

where

$$\Xi = \hat{n}_i \frac{x^i}{r} Y_{00}, \qquad (2.5.7)$$

the minimum is over the excision boundary, and the angle brackets in Eq. (2.5.6) refer to an average over the excision boundary. Note that $\Xi < 0$ because \hat{n}_i points radially inward and x^i/r points radially outward; this means that $\dot{Q} = -\ddot{\lambda}_{00}$, in accordance with our normalization choice for a control system on $\dot{\lambda}_{00}$.

As in our other controlled map parameters, we demand that λ_{00} is a function with a piecewise-constant second derivative. It is then easy to construct λ_{00} as a function with a piecewise-constant third derivative.

The control system given by Eq. (2.5.6) with a hand-chosen value of v_T has been used successfully [24, 25] in simulations of high-spin binaries. Figure 9 illustrates why characteristic speed control is crucial for the success of these simulations.



Minimum characteristic speeds

Figure 9: Minimum characteristic speed on one of the excision surfaces of an equal mass binary with equal aligned spins of $\chi = 0.97$, shown just before merger [25]. Two cases are shown: one without characteristic speed control (solid) and one that is restarted from the uncontrolled run at t = 6255M with characteristic speed control turned on and $v_T = 5 \times 10^{-5}$ (dashed). The uncontrolled run is terminated because of negative characteristic speeds shortly after t = 6290M, but the controlled run continues through merger and ringdown.

II Apparent horizon tracking

The characteristic speed control described above has the disadvantage that it requires a user-specified target value v_T . If v_T is chosen to be too small, then small fluctuations (due to shape control, the horizon finder, or simply numerical truncation error) can cause the characteristic speed to become negative and spoil the simulation.

If v_T is chosen to be too large, the simulation can also fail. To understand why, recall that characteristic speed control achieves the target characteristic speed by moving the excision boundary and thus changing the velocity term in Eq. (2.5.1). So if $v > v_T$ the control system moves the excision boundary radially inward, and if $v < v_T$ the control system moves the excision boundary radially outward. If v_T is too large, the control system can push the excision boundary outward until it crosses the apparent horizon. This halts the evolution because the apparent horizon can no longer be found.

One way to prevent the excision boundary from crossing the horizon is to drive the excision boundary to some constant fraction of the horizon radius, or in other words, drive the quantity $d/dt(\Delta r)$ to zero, where

$$\Delta r = 1 - \frac{\langle \hat{r}_{\rm EB} \rangle}{\langle \hat{r}_{\rm AH} \rangle} \tag{2.5.8}$$

is the relative difference between the average radius of the apparent horizon (in the intermediate frame) and the average radius of the excision boundary. Using Eqs. (2.4.46) and (2.4.8), we can write

$$\frac{d}{dt}\Delta r = \frac{\dot{\lambda}_{00}}{\hat{S}_{00}} + \frac{\dot{S}_{00}}{\hat{S}_{00}} \left(1 - \Delta r\right), \qquad (2.5.9)$$

and therefore a control system that adjusts $\dot{\lambda}_{00}$ to achieve $d/dt(\Delta r) = 0$ can be obtained by defining

$$Q = \dot{\hat{S}}_{00}(\Delta r - 1) - \dot{\lambda}_{00}.$$
 (2.5.10)

A slight generalization of this control system can be obtained by demanding that $d/dt(\Delta r) = \dot{r}_{\text{drift}}$, where \dot{r}_{drift} is some chosen constant,

$$Q = \dot{\hat{S}}_{00}(\Delta r - 1) - \dot{\lambda}_{00} + \hat{S}_{00}\dot{r}_{\text{drift}}.$$
 (2.5.11)

We have not found the control systems defined by Eqs. (2.5.10) and (2.5.11) to be especially useful on their own. One drawback of these systems is that they do not prevent the minimum characteristic speed v from becoming negative. Instead, we use the horizon tracking control systems described here as part of a more sophisticated control system discussed in the next section.

III Adaptive switching of size control

Here we introduce a means of controlling λ_{00} that combines the best features of characteristic speed control and horizon tracking. The idea is to continuously monitor the state of the system and switch between different control systems as the evolution proceeds.

At fixed intervals during the simulation that we call "measurement times," we monitor the minimum characteristic speed v and the relative distance between the horizon and the excision boundary Δr . The goal of the control system is to ensure that both of these values remain positive.

Consider first the characteristic speed v. Using the method described in Appendix A3, we determine whether v is in danger of becoming negative in the near future, and if so, we estimate the timescale τ_v on which this will occur. Similarly, we also determine whether Δr will soon become negative, and if so we estimate a corresponding timescale $\tau_{\Delta r}$. Having estimated both τ_v and $\tau_{\Delta r}$, we use these quantities to determine how to control $\dot{\lambda}_{00}$. In particular, we switch between horizon tracking, Eq. (2.5.10), and characteristic speed control, Eq. (2.5.6), based on τ_v and $\tau_{\Delta r}$.

We do this by the following algorithm, which favors horizon tracking over characteristic speed control unless the latter is essential. Assume that the control system for $\dot{\lambda}_{00}$ is currently tracking the horizon, Eq. (2.5.10), and that the current damping timescale is some value τ_d . If v is in danger of crossing zero according to the above estimate, and if $\tau_v < \tau_d$ and $\tau_v < \tau_{\Delta r}$, we then switch to characteristic speed control, Eq. (2.5.6), we set $v_T = 1.01v$ where v is the current value of the characteristic speed (the factor 1.01 prevents the algorithm from switching back from characteristic speed control to horizon tracking on the very next time step), and we reset the damping time τ_d equal to τ_v . Otherwise we continue to use Eq. (2.5.10), resetting $\tau_d = \tau_{\Delta r}$ if $\tau_{\Delta r} < \tau_d$.

Now assume that the control system for λ_{00} is controlling the characteristic speed, Eq. (2.5.6). If v is not in danger of crossing zero, so that we no longer need active control of the characteristic speed, then we switch to horizon tracking, Eq. (2.5.10), without changing the damping time τ_d . If v is in danger of crossing zero, and if Δr is either in no danger of crossing zero or if it will cross zero sufficiently far in the future such that $\tau_{\Delta r} \geq \sigma_1 \tau_d$, then we continue to use Eq. (2.5.6) with τ_d reset to min (τ_d, τ_v) so as to maintain control of the characteristic speed. Here σ_1 is a constant that we typically choose to be about 20.

The more complicated case occurs when both v and Δr are in danger of crossing zero and $\tau_{\Delta r} < \sigma_1 \tau_d$. In this case, we have two possible methods by which we attempt to prevent Δr from becoming negative. We choose between these two methods based on the values of $\tau_{\Delta r}$ and τ_d , and also by the value of a quantity that we call the *comoving characteristic speed*:

$$v_{c} = -\alpha - \bar{\beta}^{i} \bar{n}_{i} - \bar{n}_{i} \frac{\partial \bar{x}^{i}}{\partial \hat{t}} + \hat{n}_{i} \frac{x^{i}}{r} \left[Y_{00} \dot{\hat{S}}_{00} (\Delta r - 1) + \sum_{\ell > 0} Y_{\ell m}(\theta, \phi) \dot{\lambda}_{\ell m}(t) \right].$$

$$(2.5.12)$$

This quantity is the value that the characteristic speed v would have if horizon tracking (Eq. (2.5.10)) were in effect and working perfectly. Equation (2.5.12) is derived by assuming Q = 0 in Eq. (2.5.10), solving for $\dot{\lambda}_{00}$, and substituting this value into Eq. (2.5.4). The reason to consider v_c is that for min $(v_c) < 0$ (which can happen temporarily during a simulation), horizon tracking is to be avoided, because horizon tracking will drive min(v) toward a negative value, namely min (v_c) .

The first method of preventing Δr from becoming negative is to switch to horizon tracking, Eq. (2.5.10); we do this if min $(v_c) > 0$ and if $\tau_{\Delta r} < \sigma_2 \tau_d$, where σ_2 is a constant that we typically choose to be 5. The second method, which we employ if $\sigma_2 \tau_d < \tau_{\Delta r} < \sigma_1 \tau_d$ or if min $(v_c) < 0$, is to continue to use characteristic speed control, Eq. (2.5.6), but to reduce the target characteristic speed v_T by multiplying its current value by some fraction η (typically 0.25). By reducing the target speed v_T , we reduce the outward speed of the excision boundary, and this increases $\tau_{\Delta r}$. The reason we use the latter method for min $(v_c) < 0$ is that the former method, horizon tracking, will drive v toward v_c and we wish to keep v positive. The reason we use the latter method for $\sigma_2 \tau_d < \tau_{\Delta r} < \sigma_1 \tau_d$ is that in this range of timescales the control system is already working hard to keep v positive, therefore a switch to horizon tracking is usually followed by a switch back to characteristic speed control in only a few time steps, and rapid switches between different types of control make it more difficult for the control system to remain locked. The constants η , σ_1 , and σ_2 are adjustable, and although the details of the algorithm (i.e., which particular quantity is controlled at which particular time) are sensitive to the choices of these constants, the overall behavior of the algorithm is not, provided $1 < \sigma_2 < \sigma_1$.

The overall behavior of this algorithm is to cause $\dot{\lambda}_{00}$ to settle to a state in which both v > 0 and $\Delta r > 0$ for a long stretch of the simulation. Occasionally, when either v or Δr threatens to cross zero as a result of evolution of the metric or gauge, the algorithm switches between characteristic speed control and horizon tracking several times and then settles into a new near-equilibrium state in which both v > 0 and $\Delta r > 0$.

The actual value of Δr and of v in the near-equilibrium state is unknown a priori, and in particular this algorithm can in principle settle to values of Δr that are either small enough that control system timescales must be very short in order to prevent Δr from becoming negative, or large enough that excessive computational resources are needed to resolve the metric quantities deep inside the horizon. To prevent either of these cases from occurring, we introduce a constant correction velocity $\dot{r}_{\rm corr}$, a nominal target value $\Delta r_{\rm target}$, and a nominal target characteristic speed $v_{\rm target}$. We currently choose these to be $\dot{r}_{\rm corr} = 0.005$, $\Delta r_{\rm target} = 0.08$, and $v_{\rm target} = 0.08$. These quantities are used only to decide whether to replace Eq. (2.5.10) by Eq. (2.5.11) when the algorithm chooses horizon tracking, as described below; these quantities are not used when the algorithm chooses characteristic speed control.

First we consider the case where Δr is too large during horizon tracking. If $\Delta r > 1.1\Delta r_{\text{target}}$, then we replace Eq. (2.5.10) with Eq. (2.5.11), and we use $\dot{r}_{\text{drift}} = -\dot{r}_{\text{corr}}$. We continue to use $\dot{r}_{\text{drift}} = -\dot{r}_{\text{corr}}$ until $\Delta r < \Delta r_{\text{target}}$, at which point we set $\dot{r}_{\text{drift}} = 0$.

The case where Δr is too small during horizon tracking is more complicated because we want to be careful so as to not make the characteristic speeds decrease, since it is more difficult to recover from decreasing characteristic speeds when Δr is small. In this case, if $\Delta r < 0.9\Delta r_{\text{target}}$, $v < 0.9v_{\text{target}}$, and $\langle n^k \partial_k v_c \rangle > 0$, then we replace Eq. (2.5.10) with Eq. (2.5.11), and we use $\dot{r}_{\text{drift}} = v_{\text{corr}}$. Here v_c is the comoving characteristic speed defined in Eq. (2.5.12) and the angle brackets refer to an average over the excision boundary. The condition on the gradient of v_c ensures that v_c will increase if Δr is increased; otherwise a positive $\dot{r}_{\rm drift}$ will threaten the positivity of the characteristic speeds. We continue to use $\dot{r}_{\rm drift} = v_{\rm corr}$ until either $v > 0.9v_{\rm target}$, $\Delta r > 0.9\Delta r_{\rm target}$, $\langle n^k \partial_k v_c \rangle < 0$, or if v will cross $v_{\rm target}$ or Δr will cross $\Delta r_{\rm target}$ within a time less than the current damping time τ_d (where crossing times are determined by the procedure of Appendix A3). When any of these conditions occur, we switch to $\dot{r}_{\rm drift} = 0$, and we prohibit switching back to a positive $\dot{r}_{\rm drift}$ until either $v > 0.99v_{\rm target}$ or $\Delta r > 0.99\Delta r_{\rm target}$, or until both v and Δr stop increasing in time. These last conditions prevent the algorithm from oscillating rapidly between positive and zero values of $\dot{r}_{\rm drift}$.

2.6 Merger and ringdown

The maps and control systems in Sec. 2.4 were discussed in the context of a domain decomposition with two excision boundaries, one inside each apparent horizon. The skew and CutX maps were introduced specifically to handle difficulties that occur when two distorted excision boundaries approach each other. At some point in the evolution as the two black holes become sufficiently close, a common apparent horizon forms around them. After this occurs, the simulation can be simplified considerably by constructing a new domain decomposition with a single excision boundary placed just inside the common horizon. The evolved variables are interpolated from the old to the new domain decomposition, and the simulation then proceeds on the new one.

The algorithm for transitioning to a new grid with a single excision boundary is fundamentally the same as that described in [3, 4], but there have been improvements in the maps and the control systems, so for completeness we describe the procedure here.

At some time $t = t_m$ shortly after a common horizon forms, we define a new, post-merger
set of grid coordinates \check{x}^i and a corresponding new domain decomposition composed only of spherical shells. In \check{x}^i coordinates, the new excision boundary is a sphere centered at the origin. We also define a post-merger map $\bar{x}^i = \mathcal{M}_{\text{Ringdown}}\check{x}^i$, where

$$\mathcal{M}_{\text{Ringdown}} = \mathcal{M}_{\text{TranslationRD}} \circ \mathcal{M}_{\text{Rotation}} \circ \mathcal{M}_{\text{ScalingRD}} \circ \mathcal{M}_{\text{ShapeRD}}.$$
 (2.6.1)

Note that the post-merger grid coordinates differ from the pre-merger grid coordinates, but there is only one set of inertial coordinates \bar{x}^i . Recall that in the dual-frame picture [1], the inertial-frame components of tensors are stored and evolved, so the evolved variables are continuous at $t = t_m$ even though the grid coordinates are not.

The post-merger translation map $\mathcal{M}_{\text{TranslationRD}}$ and its control system are the same as discussed in Sec. 2.4.III, except that $\mathcal{M}_{\text{TranslationRD}}$ translates with respect to the origin of the \check{x}^i coordinate system, which is different from the origin of the x^i coordinate system. The post-merger translation control system keeps the common apparent horizon centered at the origin in the post-merger grid frame. Note that the center of the common apparent horizon does not necessarily lie on the same line as the centers of the individual apparent horizons (see [23] for an example). This means that $\mathcal{M}_{\text{TranslationRD}}$ is not continuous with the pre-merger translation map $\mathcal{M}_{\text{Translation}}$, except near the outer boundary where both maps are the identity.

To evolve a distorted single black hole we do not need a rotation map. However, if rotation is turned off suddenly at $t = t_m$, then the outer boundary condition (which is imposed in inertial coordinates) is not smooth in time, and we see a pulse of gauge and constraint violating modes propagate inward from the outer boundary because of this sharp change in the boundary condition. Therefore, we use the same rotation map $\mathcal{M}_{\text{Rotation}}$ before and after merger, and we slow down the rotation gradually. To accomplish this, instead of adjusting the map parameters (ϑ, φ) by a control system, for $t \geq t_m$ we set these parameters equal to prescribed functions that approach a constant at late times. We set

$$\vartheta(t) = A + [B + C(t - t_m)] e^{-(t - t_m)/\tau_{\text{roll}}},$$
(2.6.2)

where the constants A, B, and C are determined by demanding that $\vartheta(t)$ and its first two time derivatives are continuous at $t = t_m$. The parameter φ obeys an equation of the same form. We choose $\tau_{\text{roll}} = 20M$.

The post-merger scaling map $\mathcal{M}_{\text{ScalingRD}}$ is written

$$\breve{R} \mapsto \left[1 + \frac{b(t) - b(t_m)}{b(t_m)} \sin^2\left(\frac{\pi \breve{R}}{2\breve{R}_{\rm OB}}\right)\right] \breve{R},$$
(2.6.3)

where \check{R} is the post-merger grid-coordinate radius, and \check{R}_{OB} is the outer boundary radius in the post-merger grid coordinates. The function b(t) is given by Eq. (2.4.14) before and after merger. We set $\check{R}_{OB} = b(t_m)R_{OB}$ so that at $t = t_m$, $\mathcal{M}_{ScalingRD}$ matches $\mathcal{M}_{Scaling}$ at the outer boundary. Note that the post-merger scaling map is the identity near the merged black hole; the only purpose of this map is to keep the inertial-coordinate location of the outer boundary smooth in time at $t = t_m$.

The post-merger shape map $\mathcal{M}_{\text{ShapeRD}}$ and its control system are the same as discussed in Secs. 2.4.V and 2.5, except the map is centered about the origin of the \check{x}^i coordinate system, the sum over excised regions in Eq. (2.4.44) runs over only one region (which we call excision region C), and the function $f_C(r_C, \theta_C, \phi_C)$ appearing in Eq. (2.4.44) is

$$f_C(r_C, \theta_C, \phi_C) = \begin{cases} (r_C - r_{\max}) / (r_{\rm EB} - r_{\max}), & r_{\rm EB} \le r_C < r_{\max} \\ 0, & r_C \ge r_{\max} \end{cases}$$
(2.6.4)

where r_{max} is the radius of a (spherical) subdomain boundary that is well inside the wave zone. We typically choose $r_{\text{max}} = 32r_{\text{EB}}$.

Once the post-merger map parameters are initialized, as discussed below, we interpolate all variables from the pre-merger grid to the post-merger grid. The inertial coordinates \bar{x}^i are the same before and after merger, so this interpolation is easily accomplished using both the pre-merger and post-merger maps, e.g. $F(x^i) = F(\mathcal{M}^{-1}\mathcal{M}_{\text{Ringdown}}\check{x}^i)$ for some function F. After interpolating all variables onto the post-merger grid, we continue the evolution. This entire process — from detecting a common apparent horizon to continuing the evolution on the new domain — is done automatically.

I Initialization

All that remains for a full specification of $\mathcal{M}_{\text{Ringdown}}$ is to describe how parameters of the two discontinuous maps $\mathcal{M}_{\text{ShapeRD}}$ and $\mathcal{M}_{\text{TranslationRD}}$ are initialized at $t = t_m$. To do this, we first construct a temporary distorted frame with coordinates \dot{x}^i , defined by

$$\bar{x}^i = \mathcal{M}_{\text{Translation}} \circ \mathcal{M}_{\text{Rotation}} \circ \mathcal{M}_{\text{ScalingRD}} \dot{x}^i.$$
 (2.6.5)

The coordinates \dot{x}^i are the same as the post-merger distorted coordinates except that the pre-merger translation map $\mathcal{M}_{\text{Translation}}$ is used in Eq. (2.6.5). We then represent the common apparent horizon (which is already known in \bar{x}^i coordinates) in \dot{x}^i coordinates:

$$\dot{x}^{i}_{\mathrm{AH}}(\dot{\theta}_{C},\dot{\phi}_{C},t) = \dot{x}^{i}_{\mathrm{AH0}}(t) + \dot{n}^{i} \sum_{\ell m} \dot{S}^{C}_{\ell m}(t) Y_{\ell m}(\dot{\theta}_{C},\dot{\phi}_{C}).$$
(2.6.6)

The expansion coefficients $\hat{S}_{\ell m}^{C}(t)$ and the horizon expansion center $\dot{x}_{AH0}^{i}(t)$ are computed using the (inertial frame) common horizon surface plus the (time-dependent) map defined in Eq. (2.6.5). Here $(\dot{\theta}_{C}, \dot{\phi}_{C})$ are polar coordinates centered about \dot{x}_{AH0}^{i} , and \dot{n}^{i} is a unit vector corresponding to the direction $(\dot{\theta}_{C}, \dot{\phi}_{C})$. Equation (2.6.6) is the same expansion as Eq. (2.4.8), except here we write the expansion to explicitly include the center $\dot{x}_{AH0}^{i}(t)$. In Eq. (2.6.6) we write $\dot{S}_{\ell m}^{C}(t)$ and $\dot{x}_{AH0}^{i}(t)$ as functions of time because we compute them at several discrete times surrounding the matching time t_{m} . By finite-differencing in time, we then compute first and second time derivatives of $\dot{S}_{\ell m}^{C}$ and \dot{x}_{AH0}^{i} at $t = t_{m}$.

Once we have $\dot{S}_{\ell m}^C$ and its time derivatives, we initialize $\lambda_{\ell m}^C(t_m) = -\dot{S}_{\ell m}^C(t_m)$, where $\lambda_{\ell m}^C(t)$ are the map parameters appearing in Eq. (2.4.44), and the minus sign accounts for the sign difference between the definitions of Eqs. (2.4.44) and (2.6.6). We do this for all (ℓ, m) except for $\ell = m = 0$: we set $\lambda_{00}^C(t_m) = 0$, thus defining the radius of the common apparent horizon in the post-merger grid frame. For all (ℓ, m) including $\ell = m = 0$ we set $d\lambda_{\ell m}^C/dt \mid_{t=t_m} = -d\dot{S}_{\ell m}^C/dt \mid_{t=t_m}$, and similarly for the second derivatives.

Note that the temporary distorted coordinates \dot{x}^i are incompatible with the assumptions of our control system, because the center of the excision boundary $\dot{x}^i_{AH0}(t)$ in these coordinates is time-dependent. We therefore define new post-merger distorted coordinates \hat{x}^i by

$$\bar{x}^i = \mathcal{M}_{\text{TranslationRD}} \circ \mathcal{M}_{\text{Rotation}} \circ \mathcal{M}_{\text{ScalingRD}} \hat{x}^i,$$
 (2.6.7)

where

$$\hat{x}^{i} = \dot{x}^{i} - \dot{x}^{i}_{\rm AH0}(t).$$
(2.6.8)

If we denote the map parameters in $\mathcal{M}_{\text{Translation}}$ by T_0^i , the map parameters in $\mathcal{M}_{\text{TranslationRD}}$ by T^i , and if we denote $\mathcal{M}_{\text{Rotation}} \circ \mathcal{M}_{\text{ScalingRD}}$ by the matrix M_j^i , then Eqs. (2.6.7) and (2.6.8) require

$$T^{i}(t) = T_{0}^{i}(t) + M_{j}^{i}(t)\dot{x}_{AH0}^{j}(t).$$
(2.6.9)

Here we have assumed that f(R) appearing in Eq. (2.4.27) is unity in the vicinity of the horizon, so we can treat the translation map as a rigid translation near the horizon. We initialize T^i and its first two time derivatives at $t = t_m$ according to Eq. (2.6.9). Note that the distorted-frame horizon expansion coefficients $\hat{S}_{\ell m}^C$ and hence the initialization of $\lambda_{\ell m}^C(t)$ are unchanged by the change in translation map because \hat{x}^i and \hat{x}^i differ only by an overall translation, Eq. (2.6.8), which leaves angles invariant.

2.7 Control systems for efficiency

Although the most important use of control systems in SpEC is to adjust parameters of the mapping between the inertial and grid coordinates, another situation for which control systems are helpful is the approximation of functions that vary slowly in time, are needed frequently during the simulation, but are expensive to compute.

For example, the average radius of the apparent horizon, r_{AH} , is used in the control system for $\dot{\lambda}_{00}$, which is evaluated at every time step. Evaluation at each time step is necessary to allow the control system for $\dot{\lambda}_{00}$ to respond rapidly to sudden changes in the characteristic speed or the size of the excision boundary, as may occur after regridding, after mesh refinement changes, or when other control systems (such as size control) are temporarily out of lock. However, computing the apparent horizon (and thus its average radius $r_{\rm AH}$) at every time step is prohibitively expensive. To significantly reduce the expense, we define a function with piecewise-constant second derivative, $r_{\rm AH}^{\rm appx}(t)$, we define a control error

$$Q = r_{\rm AH} - r_{\rm AH}^{\rm appx}, \qquad (2.7.1)$$

and we define a control system that drives Q to zero. We then pass $r_{AH}^{appx}(t)$ instead of r_{AH} to those functions that require the average apparent horizon radius at each time step. The control error Q, and thus the expensive computation of the apparent horizon, needs to be evaluated only infrequently, i.e., on the timescale on which the average horizon radius is changing, which may be tens or hundreds of time steps.

2.8 Summary

In simulations of binary black holes, we use a set of time-dependent coordinate mappings to connect the asymptotically inertial frame (in which the black holes inspiral about one another, merge, and finally ringdown) to the grid frame in which the excision surfaces are stationary and spherical. The maps are described by parameters that are adjusted by a control system to follow the motion of the black holes, to keep the excision surface just inside the apparent horizons of the holes, and also to prevent grid compression. We take care to decouple the control systems and choose stable control timescales.

The scaling, rotation, and translation maps are used to track the overall motion of the black holes in the inertial frame. These are the most important maps during the inspiral phase of the evolution, as the shapes of the horizons remain fairly constant after the relaxation of the initial data. As the binary approaches merger, the horizon shapes begin to distort, and shape and size control start to become important. Shape and size control are especially crucial for unequal-mass binaries and for black holes with near-extremal spins. In the latter case, the excision surface (which must remain an outflow boundary with respect to the characteristics of the evolution system) needs to be quite close to the horizon.

In order to decouple the shape maps of the individual black holes, our grid is split by a cutting plane at which the shape maps reduce to the identity map. As the black holes merge, a skew map is introduced in order to align the cutting plane with the excision surfaces (see Fig. 5) and to minimize grid compression between the two black holes.

Finally, the implementation of CutX control was necessary to complete the merger of high mass-ratio configurations. In these systems, the excision boundaries approach the cutting plane asymmetrically as the black holes merge, thus compressing the grid between the cutting plane and the nearest excision boundary. CutX control translates the cutting plane to keep it centered between the two excision boundaries, alleviating this grid compression.

We find that we need all of the maps described in this paper with a sufficiently tight control system in order to robustly simulate a wide range of the parameter space of binary black hole systems [4, 24, 5]. Many of these maps and control systems are also used in our simulations of black hole – neutron star binaries [26, 27, 28].

A1 Implementation of exponentially-weighted averaging

Our exponentially-weighted averaging scheme uses an averaging timescale τ_{avg} to smooth noisy quantities F that are measured at intervals $\tau_m := t_k - t_{k-1}$, where F is the measured value of the control error Q, its integral, or its derivatives. Recall from Sec. 2.3.III that we typically choose $\tau_{avg} \sim 0.25\tau_d$ and $\tau_m \sim 0.075 \min(\tau_d)$.

The averaging is implemented by solving the following system of ordinary differential

equations (ODEs):

$$\frac{d}{dt}W = 1 - \frac{W}{\tau_{\text{avg}}},\tag{A1.1}$$

$$\frac{d}{dt}(W\tau) = t - \frac{W\tau}{\tau_{\text{avg}}},\tag{A1.2}$$

$$\frac{d}{dt}(WF_{\rm avg}) = F - \frac{WF_{\rm avg}}{\tau_{\rm avg}},\tag{A1.3}$$

where the evolved variables are a weight factor W, the average value F_{avg} , and the effective time τ at which F_{avg} is calculated. The ODEs are solved approximately using backward Euler differencing, which results in the recursive equations

$$W_{k} = \frac{1}{D} \left(\tau_{m} + W_{k-1} \right), \tag{A1.4}$$

$$\tau_k = \frac{1}{DW_k} \left(\tau_m t_k + W_{k-1} \tau_{k-1} \right), \tag{A1.5}$$

$$F_{\text{avg}}^{k} = \frac{1}{DW_{k}} \left[\tau_{m} F(t_{k}) + W_{k-1} F_{\text{avg}}^{k-1} \right],$$
(A1.6)

where $D = 1 + \tau_m / \tau_{avg}$. The recursion is initialized with $W_0 = 0$, $\tau_0 = t_0$, and $F_{avg}^0 = F(t_0)$.

In the control law equations, Eqs. (2.3.5) and (2.3.6), we want to use the averaged value at t_k instead of τ_k . Therefore, to adjust for the offset induced by averaging, $\delta t_k = t_k - \tau_k$, we evaluate at t_k an approximate Taylor series for F_{avg}^k expanded about τ_k

$$F_{\text{avg}}(t_k) = \sum_{m=0}^{N-n} \frac{\delta t^m}{m!} \left(\frac{d^m F}{dt^m}\right)_{\text{avg}}^k,$$
(A1.7)

where n is defined by $F := d^n Q/dt^n$ (for convenience, F represents the integral of Q when n = -1), and N is the same as in Eq. (2.3.1), where it is defined as the number of derivatives used to represent the map parameter. Note that this approximation matches the true Taylor series inasmuch as W is constant.

For the highest order derivative of Q, where n = N, we can no longer directly adjust for the offset because Eq. (A1.7) reduces to

$$F_{\text{avg}}(t_k) = F_{\text{avg}}^k. \tag{A1.8}$$

We would need to take additional derivatives of Q to provide a meaningful correction, but this is exactly what we want to avoid because of the noisiness of derivatives. Therefore, to circumvent taking higher order derivatives, we substitute $F_{\text{avg}}(t_k)$ into the control law, e.g. Eq. (2.3.11) for PID, and then solve for $d^N Q/dt^N$.

A2 Details on computing the CutX map weight function

As discussed in Sec. 2.4.VI, we use the map $\mathcal{M}_{\text{CutX}}$ to control the location of the cutting plane in the last phase of the merger, as the distance between the excision boundaries and the cutting plane becomes small. The value of the weight function $\rho(x^i)$ in Eq. (2.4.50) is zero in the spherical shells describing the wave zone and in the shells around the excision boundaries (see Fig. 10). The weight function is unity on the cutting plane, and on the "cut-sphere" surfaces around either excision boundary. For the smaller excision surface, we have two such cut-sphere surfaces, and the value of the weight function is one within the region enclosed by these two cut-spheres and the cutting plane. The weight function transitions linearly from zero to one (from the magenta curves to the red curves in Fig. 10). In these transitional regions the value of $\rho(x^i)$ is computed as described below.

Consider the region spanning the volume between the spherical shells around the excision boundary H (with H = A, B) and the cutting plane $x = x_C^0$, as shown in Fig. 11. (This is referred to in the code as the M region.) In order to calculate the value of $\rho(x^i)$ in this region, we shoot a ray from C_H^i in the direction of x^i . This ray intersects both the outer spherical boundary of the shells and the cutting plane. The x-coordinate of the intersection with the shells will be

$$x_M = C_H^0 + \frac{(x^0 - C_H^0)R_H}{\sum_i (x^i - C_H^i)^2},$$
(A2.1)

where R_H is the outer radius of the spherical region around C_H^i . Then, for a point x^i in the



Figure 10: Schematic diagram indicating the way the $\mathcal{M}_{\text{CutX}}$ weight function is defined. The value of the weight function $\rho(x^i)$ is one on the red curves and in the black region enclosed by red curves, it vanishes in the white regions (inside the two inner magenta curves and outside the outer magenta curve) and it changes linearly from zero to one in the gray regions between the red and magenta curves.



- Figure 11: Schematic diagram illustrating the algorithm used to compute the $\mathcal{M}_{\text{CutX}}$ weight function inside the cut-sphere. The point x^i represents an arbitrary point in the M region. R_H is the radius of the magenta sphere, and x_M is the x^0 -component of the point on the magenta sphere that intersects the ray pointing from C_H^i toward x^i . The weight function $\rho(x^i)$ is zero at the intersection of the ray with the magenta curve, it is unity at the intersection of the ray with the cutting plane, and changes linearly in between. Similarly, for a point in the E region, between the magenta sphere and the spherical part of the red cut-sphere, $\rho(x^i)$ changes linearly from zero (at the intersection of the ray with the inner, magenta sphere) to unity (at the intersection with the outer, red cut-sphere.)
- M region ρ is given by

$$\rho(x^i) = \frac{x^0 - x_M}{x_C^0 - x_M}.$$
(A2.2)

All other transitional regions are delimited on both ends by spheres. Our computational domain has two zones of this type; one is depicted in Fig. 11 (labeled as the E region), while the other is seen in Fig. 12. Let x_P^i denote the "center of projection," i.e., the point toward which the unmapped radial grid lines of the given zone converge, and let x_S^i denote the center of a sphere of radius R (noting that the inner and outer spheres may have different centers). In Fig. 11 the center of projection coincides with C_H^i . In Fig. 12 we choose the center of projection to be on the cutting plane, at its intersection with the line segment connecting the centers of the two excision surfaces, x_C^i . To compute $\rho(x^i)$ in this type of region, we shoot a



Figure 12: Schematic diagram illustrating the algorithm used to compute the \mathcal{M}_{CutX} weight function in the region between the red cut-sphere and the outer, magenta sphere.

ray from x_P^i in the direction of x^i . This ray intersects the two spheres that delimit the region. We define r_0 to be the distance between x_P^i and the point of intersection with the (magenta) sphere, for which $\rho = 0$; similarly, we define r_1 to be the distance between x_P^i and the point of intersection with the (red) sphere, for which $\rho = 1$. Then, $\rho(x^i)$ is given by

$$\rho(x^i) = \frac{r - r_0}{r_1 - r_0},\tag{A2.3}$$

where $r = |x^i - x_P^i|$. To compute r_0 and r_1 , we must solve

$$\sum_{i} \left[x_{P}^{i} + \eta (X^{i} - x_{P}^{i}) - x_{S}^{i} \right]^{2} = R^{2}$$
(A2.4)

for the associated intersection point X^i . The parameter $\eta(x_P^i, x_S^i, X^i, R)$ is defined as the

positive solution of the quadratic system,

$$\eta = B + \sqrt{B^2 - C},\tag{A2.5}$$

where

$$B = \frac{\sum_{i} (X^{i} - x_{P}^{i})(x_{S}^{i} - x_{P}^{i})}{\sum_{i} (X^{i} - x_{P}^{i})^{2}},$$
(A2.6)

$$C = \frac{\sum_{i} (x_{S}^{i} - x_{P}^{i})^{2} - R^{2}}{\sum_{i} (X^{i} - x_{P}^{i})^{2}}.$$
(A2.7)

The input coordinates to $\mathcal{M}_{\text{CutX}}$ are the coordinates \tilde{x}^i that are obtained from the grid coordinates x^i by the shape map, $\tilde{x}^i = \mathcal{M}_{\text{Shape}}(x^i)$. As seen above, computation of the weight function at any point requires knowledge of the region containing that point, which in turn requires knowledge of its grid-frame coordinates x^i . Therefore, to compute the weight function one must first compute $x^i = \mathcal{M}_{\text{Shape}}^{-1}(\tilde{x}^i)$. When computing the forward CutX map, this inverse is done only once. However, the situation is more complicated when computing the inverse CutX map, because the grid-frame coordinates depend upon the inverse CutX map itself, which depends on the weight function. This leads to an iterative algorithm where we must call the inverse map function of $\mathcal{M}_{\text{Shape}}$ from each iteration of the inverse map function of $\mathcal{M}_{\text{CutX}}$. This could easily become an efficiency bottleneck. We mitigate this by keeping $\mathcal{M}_{\text{CutX}}$ inactive for most of the run, only activating it when it becomes crucial in order to avoid a singular map.

A3 Estimation of zero-crossing times

Several of the control systems described here are designed to ensure that some measured quantity remains positive. For example, in Sec. 2.5.III, we demand that both the characteristic speed v at the excision boundary and the difference Δr between the horizon radius and the excision boundary radius remain positive. Similarly, in Sec. 2.4.VI we demand that each excision boundary does not cross the cutting plane.

Part of the algorithm for ensuring that some quantity q remains positive is estimating whether q is in danger of becoming negative in the near future, and if so, estimating the timescale τ on which this will occur. In this Appendix we describe a method of obtaining this estimate.

We assume the quantity q is measured at a set of (not necessarily equally-spaced) measurement times, and we remember the values of q at several (typically 4) previous measurement times. At each measurement time t_0 we fit these remembered values to a line $q(t) = a + b(t-t_0)$. The fit gives us not only the slope b and the intercept a, but also error bars for these two quantities δa and δb . Assuming that the true q(t) lies within the error bars, the earliest time at which q will cross zero is $t = t_0 + (-a + \delta a)/(b - \delta b)$, and the latest time is $t = t_0 + (-a - \delta a)/(b + \delta b)$. If both of these times are finite and in the future, then we regard q as being in danger of crossing zero, and we estimate the timescale on which this will happen as $\tau = -a/b$.

Acknowledgments

We would like to thank Harald Pfeiffer for useful discussions. We would like to thank Abdul Mroué for performing simulations that led to the development and improvement of some of the mappings presented in this paper. We gratefully acknowledge support from the Sherman Fairchild Foundation; from NSF grants PHY-0969111 and PHY-1005426 at Cornell, and from NSF grants PHY-1068881 and PHY-1005655 at Caltech. Simulations used in this work were computed with the SpEC code [17]. Computations were performed on the Zwicky cluster at Caltech, which is supported by the Sherman Fairchild Foundation and by NSF award PHY-0960291; on the NSF XSEDE network under grant TG-PHY990007N; and on the GPC supercomputer at the SciNet HPC Consortium [29]. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund–Research Excellence; and the University of Toronto.

References

- 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.
- [2] 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(12):124038, 2007.
- [3] M. A. Scheel, M. Boyle, T. Chu, L. E. Kidder, K. D. Matthews and H. P. Pfeiffer. High-accuracy waveforms for binary black hole inspiral, merger, and ringdown. *Phys. Rev. D*, 79:024003, 2009.
- [4] Béla Szilágyi, Lee Lindblom, and Mark A. Scheel. Simulations of Binary Black Hole Mergers Using Spectral Methods. *Phys. Rev. D*, 80:124010, 2009.
- [5] Luisa T. Buchman, Harald P. Pfeiffer, Mark A. Scheel, and Béla Szilágyi. Simulations of unequal mass binary black holes with spectral methods. *Phys. Rev. D*, 86:084033, 2012.
- [6] 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.
- [7] 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.
- [8] Béla Szilágyi, H.-O. Kreiss, and J. Winicour. Modeling the black hole excision problem. Phys. Rev. D, 71:104035, 2005.

- [9] Frans Pretorius. Evolution of binary black hole spacetimes. Phys. Rev. Lett., 95:121101, 2005.
- [10] Mohammad Motamed, M. Babiuc, Béla Szilágyi, and H-O. Kreiss. Finite difference schemes for second order systems describing black holes. *Phys. Rev. D*, 73:124008, 2006.
- [11] Béla Szilágyi, Denis Pollney, Luciano Rezzolla, Jonathan Thornburg, and Jeffrey Winicour. An explicit harmonic code for black-hole evolution using excision. *Class. Quantum Grav.*, 24:S275–S293, 2007.
- [12] Gregory B. Cook, Mijan F. Huq, Scott A. Klasky, Mark A. Scheel, et al. Boosted three-dimensional black-hole evolutions with singularity excision. *Phys. Rev. Lett.*, 80(12):2512–2516, 1998.
- [13] Steve Brandt, Randall Correll, Roberto Gomez, Mijan Huq, Pablo Laguna, Luis Lehner, Pedro Marronetti, Richard A. Matzner, David Neilsen, Jorge Pullinand Erik Schnetter, Deirdre Shoemaker, and Jeffrey Winicour. Grazing collisions of black holes via the excision of singularities. *Phys. Rev. Lett.*, 85:5496–5499, 2000.
- [14] Miguel Alcubierre and Bernd Brügmann. Simple excision of a black hole in 3+1 numerical relativity. *Phys. Rev. D*, 63(10):104006, 2001.
- [15] D. Shoemaker, K. Smith, U. Sperhake, P. Laguna, E. Schnetter, and D. Fiske. Moving black holes via singularity excision. *Class. Quantum Grav.*, 20:3729–3744, 2003.
- [16] 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*, 80:124051, 2009.
- [17] http://www.black-holes.org/SpEC.html.
- [18] Stuart Bennett. A History of Control Engineering, 1930-1955. Peter Peregrinus Ltd., London, United Kingdom, 1993.
- [19] Eduardo D. Sontag. Mathematical Control Theory: Deterministic Finite Dimensional

Systems. Number 6 in Textbooks in Applied Mathematics. Springer, New York, second edition, 1998.

- [20] Carsten Gundlach. Pseudospectral apparent horizon finders: An efficient new algorithm. *Phys. Rev. D*, 57(2):863–875, Jan 1998.
- [21] Serguei Ossokine, Lawrence E. Kidder, and Harald P. Pfeiffer. Precession-tracking coordinates for simulations of compact-object-binaries. arXiv:1304.3067, 2013.
- [22] 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.
- [23] 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. *Phys. Rev. D*, 82:064031, 2010.
- [24] Geoffrey Lovelace, Mark. A. Scheel, and Béla Szilágyi. Simulating merging binary black holes with nearly extremal spins. *Phys. Rev. D*, 83:024010, 2011.
- [25] Geoffrey Lovelace, Michael Boyle, Mark A. Scheel, and Béla Szilágyi. High-accuracy gravitational waveforms for binary-black-hole mergers with nearly extremal spins. *Class. Quant. Grav.*, 29:045003, 2012.
- [26] Matthew D. Duez, Francois Foucart, Lawrence E. Kidder, Harald P. Pfeiffer, Mark A. Scheel, and Saul A. Teukolsky. Evolving black hole-neutron star binaries in general relativity using pseudospectral and finite difference methods. *Phys. Rev. D*, 78:104015, 2008.
- [27] Francois Foucart, Matthew D. Duez, Lawrence E. Kidder, and Saul A. Teukolsky. Black hole-neutron star mergers: effects of the orientation of the black hole spin. *Phys. Rev.*, D83:024005, 2011.
- [28] Francois Foucart, Matthew D. Duez, Lawrence E. Kidder, Mark A. Scheel, Béla Szilágyi,

and Saul A. Teukolsky. Black hole-neutron star mergers for $10M_{\odot}$ black holes. *Phys. Rev. D*, 85(4):044015, February 2012.

[29] Chris Loken, Daniel Gruner, Leslie Groer, Richard Peltier, Neil Bunn, Michael Craig, Teresa Henriques, Jillian Dempsey, Ching-Hsing Yu, Joseph Chen, L Jonathan Dursi, Jason Chong, Scott Northrup, Jaime Pinto, Neil Knecht, and Ramses Van Zon. SciNet: Lessons Learned from Building a Power-efficient Top-20 System and Data Centre. J. Phys.: Conf. Ser., 256:012026, 2010.

CHAPTER 3

FINAL SPIN AND RADIATED ENERGY IN NUMERICAL SIMULATIONS OF BINARY BLACK HOLES WITH EQUAL MASSES AND EQUAL, ALIGNED OR ANTI-ALIGNED SPINS

Contents		
3.1	Introduction	
3.2	Simulation methods	
3.3	Simulations	
3.4	Results	
3.5	Conclusions	
$\mathbf{A1}$	Method for constructing our fitting formulae	

The behavior of merging black holes (including the emitted gravitational waves and the properties of the remnant) can currently be computed only by numerical simulations. This paper introduces ten numerical relativity simulations of binary black holes with equal masses and equal spins aligned or anti-aligned with the orbital angular momentum. The initial spin magnitudes have $|\chi_i| \leq 0.95$ and are more concentrated in the aligned direction because of the greater astrophysical interest of this case. We combine these data with five previously reported simulations of the same configuration, but with different spin magnitudes, including the highest spin simulated to date, $\chi_i \approx 0.97$. This data set is sufficiently accurate to enable us to offer improved analytic fitting formulae for the final spin and for the energy radiated by

gravitational waves as a function of initial spin. The improved fitting formulae can help to improve our understanding of the properties of binary black hole merger remnants and can be used to enhance future approximate waveforms for gravitational wave searches, such as Effective-One-Body waveforms.

3.1 Introduction

Binary black holes are an important source for gravitational-wave detectors such as the Laser Interferometer Gravitational-Wave Observatory (LIGO), GEO, and Virgo [1, 2, 3]. Searches for gravitational-wave signals have been able to constrain the event rate for binary black hole mergers, but a direct detection of gravitational waves has not yet been made [4, 5]. These searches require predictions ("templates") of the expected gravitational waves; so far, only non-spinning templates have been included [5]. However, there is evidence that spin is relevant in astrophysical black holes, from both theoretical predictions [6, 7, 8] and observational data [9, 10, 11].

Therefore, LIGO and other gravitational-wave detectors need to include spin as a parameter in their template waveforms; otherwise, the search space (and thus the detection rate) is reduced because of an insensitivity to spinning sources [12, 13]. Accurate simulations of spinning binary black hole mergers are also needed to infer the properties (e.g. masses and spins) of binaries from the detected waveforms ("parameter estimation") [14].

For both detection and parameter estimation, numerical simulations are too computationally expensive to generate waveforms for the entire parameter space of binary black hole mergers. Instead, numerical simulations are used to calibrate and validate the approximate, analytic models that are actually used to generate template waveforms. For instance, the Effective-One-Body (EOB) model, calibrated using numerical simulations that include merging black holes with spins aligned or anti-aligned with the orbital angular momentum [15], is used by the LIGO collaboration to estimate how sensitive their search is to waveforms from spinning systems [5]. However, Ref. [15] has shown that the EOB model poorly predicts configurations with large aligned spins, and that more numerical relativity simulations are needed in this region of spin parameter space to improve the calibration of the model.

Binary black holes whose spins are aligned (or anti-aligned) with the orbital angular momentum involve far fewer parameters than generic binaries with arbitrary spin directions, but nevertheless they can be used to construct templates capable of detecting a sizeable fraction of precessing binaries [12]. Furthermore, aligned-spin systems are astrophysically motivated by studies including observations of the micro-quasar XTE J1550-564 [16], models of gas-rich galaxy mergers [17], and population synthesis models [18].

In this paper, we introduce ten new simulations of binary black holes with equal masses and equal spins aligned or anti-aligned with the orbital angular momentum. We use the notation $S_{|\chi|}^{\pm\pm}$ to refer to specific cases, where the subscript is approximately the dimensionless spin magnitude at t = 0, and the superscripts indicate whether each black hole has the aligned (+) or anti-aligned (-) spin orientation. The new simulations are $S_{0.95}^{++}$, $S_{0.9}^{++}$, $S_{0.85}^{++}$, $S_{0.8}^{++}$, $S_{0.6}^{++}$, $S_{0.2}^{+-}$, $S_{0.2}^{--}$, $S_{0.8}^{--}$, and $S_{0.9}^{--}$.

To more fully cover the aligned-spin space, we include data previously reported for $S_{0.97}^{++}$ [19], $S_{0.95}^{--}$ [20], $S_{0.0}^{--}$ [21], $S_{0.44}^{--}$ [22], and $S_{0.44}^{++}$ [23] in our analysis. The $S_{0.95}^{++}$ case joins the two simulations from Refs. [20, 19] as the only simulations to date of merging black holes with spin magnitudes above $\chi \approx 0.93$ (the "Bowen-York limit") [24, 25, 26]. We use this combined dataset to improve on prior phenomenological fitting formulae for the final spin of the merger remnant [27, 28, 29] and the radiated energy from inspiral through ringdown [29, 30, 31]. These improved formulae can be used to reduce EOB waveform phase errors in the ringdown (see Eq. (19) and surrounding text in Ref. [15]) and therefore provide more accurate templates for gravitational-wave searches [31].

The remainder of this paper is organized as follows. In Sec. 3.2, we discuss the numerical

methods that we employ in our simulations. In Sec. 3.3, we report on the values and convergence of the constraint violations, masses, and spins. In Sec. 3.4, we use the horizon data to improve the phenomenological fitting formulae for final spin and radiated energy as a function of initial spin. Section 3.5 contains our conclusions, and Appendix A1 details our method for constructing the fitting formulae.

3.2 Simulation methods

All simulations used in this paper were generated with the Spectral Einstein Code (SpEC) [32]. In this section, we describe the methods for the ten new simulations. For detailed methods of the previously reported SpEC simulations, see Refs. [19, 20, 21, 22, 23] and references therein. Throughout this paper, we use units where G = c = 1, and we report lengths and times in units of M, the total Christodoulou mass in the initial data.

To produce initial data, we solve the extended conformal thin-sandwich equations with quasi-equilibrium boundary conditions [33, 34, 35, 36, 37, 38]. We adopt free data based on a weighted superposition of two Kerr-Schild black holes, which enables us to construct initial data containing black holes with nearly extremal spins [39, 40]. The constraint equations are solved using a spectral elliptic solver [41], and the free parameters are iterated until the target masses and spins are achieved to within some tolerance.

We evolve the initial data on a "cut-spheres" domain [42] using spectral adaptive mesh refinement, which will be detailed in a forthcoming paper. On a timescale of 50M, we change smoothly from the initial data gauge to damped harmonic gauge [43, 44, 45], which helps prevent coordinate singularities. We use a fifth-order Dormand-Prince dense adaptive timestepper.

To reduce eccentricity, we first evolve each system for 2.5 orbits beyond the time when the spurious "junk" radiation is sufficiently far from the black holes so as to have a negligible effect on the black hole trajectories. Then we fit the time derivative of the orbital frequency $\dot{\Omega}$ to find improved initial angular and radial velocities (Ω_0 and \dot{d}_0/d_0) [46]. We iterate this procedure until an eccentricity below 10^{-3} is achieved.

We use the dual-frames technique to do spectral excision of the singularities [47]. As described in other papers reporting high spin simulations using SpEC [19, 20], the most important aspect of this excision technique is careful control of the excision boundary. This must accomplish three tasks. First, it must distort the shape of the boundary so that it matches the shape of the apparent horizon. Second, it must regulate the fractional separation between the excision surface and the apparent horizon — if the separation is too small, then the horizon falls out of the computational domain, but if the separation is too large, then the excision surface falls far inside the horizon, where large gradients are computationally expensive to resolve. Third, it must keep all characteristic speeds on the excision surface positive; i.e., the excision surface must be a pure outflow boundary. This is because we do not impose boundary conditions on the excision surface. Instead, we monitor the characteristic speeds; if they ever become negative, then our evolution system becomes ill-posed, and we terminate the simulation. These three tasks are challenging for high spin systems in part because of the additional distortion of the horizons (see Fig. 13), and they are especially challenging for large aligned spins because such systems spend more time in the dynamic regime before merger.

Using the fast-flow method described in Ref. [48], we find the apparent horizons as an expansion in spherical harmonics, truncated at a given maximum order ℓ . As the system evolves, we adaptively change ℓ to satisfy accuracy criteria for the resolution of the horizon. After a common horizon is found during merger, the evolution continues on a new domain with a single excision surface that subsumes the two individual excision regions [45].

We measure the quasi-local spin, S, on each horizon using the approximate Killing vector



Figure 13: Effect of spin on horizon geometry. This image shows the intrinsic Ricci scalar on the apparent horizons in simulation $S_{0.85}^{++}$. The proper separation of the horizons along the line connecting their centers is about 1.7*M*. Both spin effects (gradients as a function of polar angle) and tidal bulges (dark red regions near the intersection with the line connecting the horizon centers) can be seen.

method described in Ref. [39]. The dimensionless spin is then

$$\chi = \frac{S}{M_{\rm Ch}^2},\tag{3.2.1}$$

where $M_{\rm Ch}$ is the Christodoulou mass,

$$M_{\rm Ch}^2 = M_{\rm irr}^2 + \frac{S^2}{4M_{\rm irr}^2},\tag{3.2.2}$$

and $M_{\rm irr}$ is the irreducible mass, which is a function of the horizon area,

$$M_{\rm irr} = \sqrt{\frac{A_{\rm AH}}{16\pi}}.$$
(3.2.3)

With these definitions, $\chi = 1$ represents an extremal black hole [49].

We choose an integer k to characterize the resolution of each simulation. We call k the resolution level (or "Lev"). It sets the resolution by defining the target maximum truncation error for the adaptive mesh refinement and adaptive horizon finding as

$$\epsilon_{\max} = 10^{-4} e^{-k}. \tag{3.2.4}$$

Around the excision boundary, where the most resolution is required, we reduce the maximum truncation error for adaptive mesh refinement by a factor of 10^{-2} .

χ_0	$\dot{Md_0}/d_0 \times 10^4$	$M\Omega_0$	$N_{\rm orbits}$
0.95	7.26420673	0.01395360	25.4
0.9	5.48222492	0.01419573	24.9
0.85	4.33347923	0.01437107	24.7
0.8	3.54332917	0.01450430	24.2
0.6	1.65215665	0.01487274	22.8
0.2	0.09507527	0.01525060	19.9
-0.2	-0.69081937	0.01538827	17.2
-0.6	-1.95883097	0.01527384	14.6
-0.8	-3.60252091	0.01501397	13.3
-0.9	-5.42657163	0.01474328	12.8

Table 1: Initial data parameters (radial velocity \dot{d}_0 and angular velocity Ω_0) at separation $d_0 = 15.366M$ for the ten new simulations with target spin, χ_0 . Also included is the approximate number of orbits until merger. Here M is the sum of the Christodoulou masses at time t = 0.

3.3 Simulations

There are ten new simulations presented in this paper: $S_{0.95}^{++}$, $S_{0.9}^{++}$, $S_{0.85}^{++}$, $S_{0.8}^{++}$, $S_{0.6}^{++}$, $S_{0.2}^{++}$, $S_{0.2}^{--}$, $S_{0.2}^{--}$, $S_{0.6}^{--}$, $S_{0.8}^{--}$, $S_{0.9}^{--}$. Initial data were generated with a target Christodoulou mass for each hole $M_0 = 0.5$, target spin for each hole χ_0 , and target ADM linear momentum $p_0^i = 0$. We fix the initial separation at $d_0 = 15.366M$, and then we iterate as summarized in Sec. 3.2 to obtain the initial radial velocity \dot{d}_0/d_0 and angular velocity Ω_0 . The targets are met to within an absolute error of $\mathcal{O}(10^{-8})$, and the resulting initial data parameters are reported in Table 1. We construct our initial data with a target total Christodoulou mass of M = 1 so that our evolution code units are essentially interchangeable with units of M.

At least three different resolutions were evolved for each case to check convergence. Figure 14 shows convergence of the (normalized) volume-averaged L^2 -norm of the generalized harmonic constraint energy [50] for a representative case. Each time the domain structure is changed to alleviate grid compression, the constraints jump because of interpolation errors, but then slowly decay back to their baseline levels.

Additional simulations are used in our analysis of masses and spins in Sec. 3.4: $S_{0.97}^{++}$ [19],



Figure 14: Normalized constraint violations for $S_{0.9}^{--}$. For each resolution level, k, we plot $||\mathcal{C}||$, the volume-averaged L^2 -norm of the generalized harmonic constraint energy divided by the volume-averaged L^2 -norm of the dynamical field gradients. This measure is defined in Eq. (71) of Ref. [50]. As the resolution level increases, the constraints decrease. Jumps in the constraints are attributed to changes in the domain structure, and the spike around $t \sim 3500M$ corresponds to the merger.

 $S_{0.44}^{++}$ [23], $S_{0.0}^{--}$ [21], $S_{0.44}^{--}$ [22], and $S_{0.95}^{--}$ [20]. Although these have been previously reported, we include them below for completeness. It should be noted that $S_{0.44}^{++}$ and $S_{0.44}^{--}$ are older simulations and therefore used different initial data and evolution machinery than described in Sec. 3.2. Simulations $S_{0.95}^{--}$, $S_{0.0}^{--}$, and $S_{0.97}^{++}$ used the initial data methods of Sec. 3.2, but earlier implementations of the evolution methods.

I Mass and spin

We define the initial spin, χ_i , to be the spin after the system has relaxed from the initial data and the junk radiation at the apparent horizon has become negligible. The spin before this time is not physically relevant to the rest of the evolution. There are subtle issues to consider when choosing the time to measure χ_i . If we choose too early a time, then junk radiation effects will still be present. If we are overly cautious and choose too late a time, then the system will have emitted enough gravitational radiation to significantly change the spin.



Figure 15: Plots of the apparent horizon quantities as a function of time for a representative case, $S_{0.9}^{++}$. The top panels display the dimensionless spin and the bottom panels display the Christodoulou mass. From left to right, the panels display the inspiral, merger, and ringdown. We normalize the *y*-scales separately so that the differences between each resolution can be clearly seen. The discontinuity in the middle panel indicates where we begin to measure the mass and spin on the common horizon. The dots in the early inspiral identify our choice of t_i for each resolution level.

Case	$ \chi_i $	$ \chi_f $
$S_{0.95}^{}$	0.949053(-30)	0.37567(-18)
$S_{0.9}^{}$	0.899569(-11)	0.392748(-12)
$S_{0.8}^{}$	0.7997602(59)	0.4268932(30)
$S_{0.6}^{}$	0.59993163(71)	0.4942327(-31)
$S_{0.44}^{}$	0.437568970(-10)	0.547851(20)
$S_{0.2}^{}$	0.1999802(-40)	0.6242202(-61)
$S_{0.0}^{}$	$64(-29) \times 10^{-8}$	0.686445(-52)
$S_{0.2}^{++}$	0.200035(-19)	0.7464314(-96)
$S_{0.44}^{++}$	0.4365505(95)	0.8140(10)
$S_{0.6}^{++}$	0.5999635(14)	0.857808(15)
$S_{0.8}^{++}$	0.7998737(-44)	0.907526(14)
$S_{0.85}^{++}$	0.849826(15)	0.919088(30)
$S_{0.9}^{++}$	0.8997371(-15)	0.930212(23)
$S_{0.95}^{++}$	0.9495863(-25)	0.940852(29)
$S_{0.97}^{++}$	0.969504(13)	0.944964(11)

Table 2: Dimensionless spins measurements. For each case, we provide the initial spinmagnitude of each hole and the final spin magnitude of the remnant at the highestresolution. Adding the number in parentheses to the last two significant digitsgives the value at the next highest resolution.

We use a histogram method to determine χ_i . Let $\{\chi_I\}$ be the set of spin measurements during the inspiral. The range of $\{\chi_I\}$ is split uniformly into N bins, where N is the size of $\{\chi_I\}$, and then each element of $\{\chi_I\}$ is put into the appropriate bin.¹ We choose $\chi_i = \chi(t_i)$, where t_i is the latest time when the spin is in the bin containing the most measurements. In the initial relaxation, the spin is oscillating, and during the inspiral, the spin changes more rapidly as the holes approach each other. Under these conditions, this method selects the spin just after the junk radiation, when the spin is nearly constant for a long interval. Figure 15 shows the evolution of the mass and spin as a function of time for a representative case, and identifies our choice of t_i by the dots in the early inspiral.

We compute t_i from χ because the behavior of the mass is not as simple during the inspiral. The histogram method applied to mass will pick out the local maximum late in the inspiral that is present in most of our cases (see Fig. 15). We define the initial mass to be $M_i = M(t_i)$,

¹ If the time interval between spin measurements is not equally spaced, we weight each measurement by the average of the two adjacent time intervals.

Case	M_i	M_f	$E_{\rm rad}(\%)$
$S_{0.95}^{}$	0.999856(68)	0.968134(33)	3.1727(33)
$S_{0.9}^{}$	1.00016197(73)	0.967909(-27)	3.2248(28)
$S_{0.8}^{}$	1.0000859(-11)	0.9665941(-16)	3.348894(50)
$S_{0.6}^{}$	1.00002292(-78)	0.963769(-14)	3.6253(13)
$S_{0.44}^{}$	2.2470608(-22)	2.159561(-49)	3.8940(21)
$S_{0.2}^{}$	0.999956(26)	0.9564388(84)	4.3519(17)
$S_{0.0}^{}$	0.9999971(-43)	0.9516182(-74)	4.83791(33)
$S_{0.2}^{++}$	0.999961(22)	0.945471(16)	5.44923(46)
$S_{0.44}^{++}$	2.2451548(28)	2.10099(-44)	6.421(20)
$S_{0.6}^{++}$	1.00001907(-96)	0.926868(-19)	7.3149(18)
$S_{0.8}^{++}$	1.0000765(-22)	0.911275(-28)	8.8794(26)
$S_{0.85}^{++}$	1.000108(-12)	0.906168(-73)	9.3931(62)
$S_{0.9}^{++}$	1.0001513(-29)	0.900366(-48)	9.9770(46)
$S_{0.95}^{++}$	1.00021743(77)	0.893703(-65)	10.6492(66)
$S_{0.97}^{++}$	1.0002384(-94)	0.890691(-22)	10.9521(14)

Table 3: Christodoulou mass measurements. For each case, we provide the total initial mass of the black holes, the final mass of the remnant, and the radiated energy computed from Eq. (3.3.1) at the highest resolution. Adding the number in parentheses to the last two significant digits gives the value at the next highest resolution.

the sum of the Christodoulou masses at time t_i .

The final spin and Christodoulou mass, χ_f and M_f , are measured at the last observation time, when the merger remnant is in quasi-equilibrium and approximates a Kerr black hole. We report the initial and final spins in Table 2, and the initial and final Christodoulou masses in Table 3.

From the initial and final Christodoulou masses, we can infer the fraction of the black hole energy that is radiated in gravitational waves during the evolution:

$$E_{\rm rad} = 1 - \frac{M_f}{M_i}.$$
 (3.3.1)

We expect mass and spin measurements at higher resolutions to be more accurate. For this reason, we weight the uncertainty of a particular measurement by a function of resolution level k in our analysis in Sec. 3.4. However, as illustrated by the comparisons in Fig. 16, these quantities are not strictly convergent in a number of cases. As described in Sec. 3.2, the



Figure 16: Differences in the final masses and spins between resolution levels. For each case, we compare χ_f and M_f of the highest resolution to the two lower resolutions. Note that, except for the older $S_{0.44}^{\pm\pm}$ simulations, all differences are $\lesssim 10^{-4}$. Differences in the initial masses and spins behave similarly.

most stringent resolution requirements occur in the vicinity of the apparent horizons, but the accuracy may be dominated by short, under-resolved segments of the evolution. The initial masses and spins appear to be randomly perturbed by the junk radiation as the initial data relaxes, and the final masses and spins appear to be affected by the details of the coalescence, where we see a spike in constraint violations (Fig. 14). Apart from these under-resolved segments, we do see convergence in the time derivatives of the masses and spins, but the absolute values remain offset from one another.

We have investigated other potential sources of uncertainty, but found them to lie below the resolution level uncertainty. For example, one source of uncertainty in the masses and spins is the resolution of the surface of the horizon. In Fig. 17, we show a representative plot of error in final spin as a function of ℓ of the horizon finder. Let $\Delta \chi_k$ be the resolution level error between the two highest resolutions, and let $\Delta \chi_\ell$ be the resolution error between the ℓ



Figure 17: Convergence of the final spin for $S_{0.9}^{--}$ as a function of the ℓ of the horizon finder. We plot the difference between $\chi_f(\ell)$ and $\chi_f(\ell = 20)$ for each resolution. The adaptive horizon finder for this case chose $\ell = 8$ at the final time of the simulation.

chosen by the adaptive horizon finder and the ℓ for which the horizon is fully resolved. At the final time of the simulation, we find that, in all cases, $\Delta \chi_k > \Delta \chi_\ell$ by several orders of magnitude.

A source of uncertainty in the radiated energy is the energy that would have been radiated by the binary as it proceeds from infinite separation to the separation d_0 at which we start the simulation. As discussed in Ref. [19], Alvi's formula [51] estimates that the energy radiated from $d = \infty$ to $d = d_0$ is one part in 10⁶. Since this is smaller than our resolution level uncertainty, it is safe to ignore this difference and we can think of $E_{\rm rad}$ as the total radiated energy from infinite separation through ringdown.

3.4 Results

Much effort has been put into constructing phenomenological formulae for the final spin [52, 27, 28, 29] and radiated energy [52, 29, 30, 31] as a function of initial spin. Because the SpEC code has the capability to generate and evolve initial data of black holes with spins above the "Bowen-York limit" of $\chi \approx 0.93$ [26], we are able to provide new data points to test and improve these formulae.

We use a Bayesian nonlinear measurement error model (described in Appendix A1) to fit and compare new parametric formulae. This approach (1) accounts for uncertainties in both the initial spin data and the output data (i.e., final spin or radiated energy); (2) accounts for the expected improvement in accuracy of results as the resolution level increases; and (3) includes a simple systematic error component quantifying misfit between a chosen formula and the curve the data are converging toward.² The framework lets us predict an output as a function of initial spin, with prediction uncertainties that account for the uncertainties in the parameters of the chosen formula (including correlated uncertainties) and the typical scale of the systematic error.

The new fitting formulae that we provide here are only applicable to equal mass binary black hole configurations with equal spins aligned or anti-aligned with the orbital angular momentum. More general formulae exist (see e.g. Refs. [53, 54, 28, 29]), but they are less accurate at high spins because of the scarcity of simulations with both unequal masses and high spins in random orientations.

 $^{^{2}}$ To keep the calculations analytically tractable, the systematic error component accounts only for the typical magnitude of misfit (essentially, the root-mean-square of the residuals), and does not account for correlations or patterns in the residuals.

I Final spin

Using the data from Table 2, we construct a new fitting formula for the final spin as a function of initial spin. We fit to a fourth-order polynomial,

$$\tilde{\chi}_f = a_0 + a_1 \chi_i + a_2 \chi_i^2 + a_3 \chi_i^3 + a_4 \chi_i^4.$$
(3.4.1)

The best fit to our data has the parameters a_n and associated covariance Σ_a :

$$a = \begin{pmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \\ a_4 \end{pmatrix} = \begin{pmatrix} 0.686402(60) \\ 0.30660(14) \\ -0.02684(33) \\ -0.00980(19) \\ -0.00499(35) \end{pmatrix}$$
(3.4.2)
$$\Sigma_a = \begin{pmatrix} 3.6 & 0.31 & -14 & -0.45 & 11 \\ 0.31 & 21 & -4.8 & -26 & 6.0 \\ -14 & -4.8 & 110 & 7.1 & -110 \\ -0.45 & -26 & 7.1 & 36 & -9.5 \\ 11 & 6.0 & -110 & -9.5 & 120 \end{pmatrix} \times 10^{-9}$$
(3.4.3)

The uncertainty in a_n , given in parentheses in Eq. (3.4.2), is estimated by $\sqrt{\Sigma_a^{nn}}$. However, the parameter estimates are highly correlated; therefore, the full covariance matrix is used in the computation of the fit uncertainty σ_f in Eq. (A1.12). In Fig. 18, we show the fit and residuals using Eq. (3.4.1) with the parameters from Eq. (3.4.2).

We fit to a fourth-order polynomial because the high accuracy of our dataset enables us to identify significant third- and fourth-order trends in the residuals of a fit to a second-order polynomial, which is the fitting function used in Refs. [27, 28, 29]. The difference between the logarithm of the marginalized likelihood function (LML) for the best-fit fourth-order and second-order polynomials is ~40, indicating that the fourth-order polynomial provides a significantly better fit. If the two additional degrees of freedom were fitting noise, rather than



Figure 18: In the top panel, we plot our preferred fitting formula (solid line), the fourth-order polynomial in Eq. (3.4.1), and a comparison with a second-order polynomial (dashed line) for χ_f as a function of χ_i . Our data points are plotted as polygons, where more sides indicates higher resolution level. In the bottom panel, we plot the residuals of the fourth-order polynomial. We indicate our fit parameter (dotted line) and total prediction (dashed line) uncertainties (defined in Appendix A1), which in this case are nearly identical. Note that the residuals for the two lower resolution runs for $S_{0.44}^{++}$ are too large to fit in this panel.

some underlying structure in the data, then we would only expect a change in maximum LML of $\mathcal{O}(1)$.³

The estimated systematic error magnitude $\hat{\sigma}_{\Delta}$ (defined in Appendix A1) for the fourthorder polynomial formula is negligibly small, suggesting that the significant behavior is captured as well as could be expected. However, the residuals, especially at large aligned spins, display trends suggesting that there is additional structure not captured by the fourthorder polynomial (such trends are ignored by our simple systematic error model). This encouraged us to explore a fifth-order polynomial formula, but it did not reduce the residuals enough to justify the additional degree of freedom. This does not rule out the possibility that a different formula could capture the behavior even more accurately.

We compare our data to existing fitting formulae for final spin in Fig. 19. The χ_f data corroborate the existing fitting formulae, but indicate deviations at large spins (especially in the aligned direction). This is an expected consequence of the scarcity of high spin numerical relativity data heretofore. In the figure, we provide a quantity for each fit, r, that measures how much larger its systematic error is than that of our fourth-order polynomial. This is essentially a ratio of the root-mean-square residuals (more precisely, r is the ratio of the $\hat{\sigma}_{\Delta}$ values). The previously reported formulae have roughly 100 to 250 times as much systematic error as our fourth-order polynomial fit. While the formula in Tichy 2008 [29] performs best, we note that it has a large uncertainty.

II Radiated energy

Following the procedure in Sec 3.4.I, we use the data from Table 3 to construct a new fitting formula for the radiated energy fraction, $E_{\rm rad}$, as a function of initial spin. We fit to a

³ The leading-order term in the maximum LML is proportional to a chi-squared-like quantity, so for nested models, such as the second- and fourth-order polynomials, the change in the maximum LML should roughly mimic the asymptotic statistics of likelihood ratio tests, as given by Wilks's theorem [55]. Two models are said to be nested if the simpler one is a special case of the more complicated one.



Figure 19: Final spin as a function of initial spin. In the left panels we plot our data (circles) along with the fitting formula (red line) with error estimates (dashed) from several other studies. The top panel is from Ref. [29], the middle is from Ref. [27], and the bottom is from Ref. [28]. In the right panels we plot the difference between our data and the corresponding fitting formula on the left. The value r quantifies the size of the systematic error compared to the fourth-order polynomial.

hyperbolic function,

$$\tilde{E}_{\rm rad} = b_0 + \frac{b_1}{b_2 + \chi_i}.$$
(3.4.4)

The best fit to our data has the parameters b_n and associated covariance Σ_b :

$$b = \begin{pmatrix} b_0 \\ b_1 \\ b_2 \end{pmatrix} = \begin{pmatrix} 0.00258(29) \\ -0.07730(79) \\ -1.6939(59) \end{pmatrix}$$
(3.4.5)
$$\Sigma_b = \begin{pmatrix} 0.83 & 2.2 & 16 \\ 2.2 & 6.2 & 46 \\ 16 & 46 & 350 \end{pmatrix} \times 10^{-7}$$
(3.4.6)

The uncertainty in b_n , given in parentheses in Eq. (3.4.5), is estimated by $\sqrt{\Sigma_b^{nn}}$. In Fig. 20, we show the fit and residuals using Eq. (3.4.4) with the parameters from Eq. (3.4.5).

We use a hyperbolic fitting function instead of a second-order polynomial (as in Refs. [29, 30]) or a constrained second-order polynomial, e.g. $\tilde{E}_{rad} = c_0 + c_1 \chi_i + (c_1/4) \chi_i^2$ (as in Ref. [31]). Parabolic fits show visible offsets in various regions of the initial spin space, which can be seen in plots of the residuals in [31] and in the comparison plot in the top panel of Fig. 20. The difference between the maximum LML for the 3-parameter hyperbola and the second-order polynomial is ~36, indicating that the hyperbola is a dramatically better fit to the data.

In Fig. 21, we compare our data to existing fitting formulae for $E_{\rm rad}$. All previous formulae suffer from the same systematic deficiencies as the best second-order polynomial fit to our data shown in Fig. 20. The ratio of the systematic error magnitude in these formulae to its magnitude in our 3-parameter hyperbolic fit, r, is shown in the figure and ranges from roughly 40 to 130. Note that it is not meaningful to compare these r values to those shown in Fig. 19, because we have not added any additional degrees of freedom in our $E_{\rm rad}$ model compared to a second-order polynomial (unlike in our final spin model, which adds two degrees of freedom).



Figure 20: In the top panel, we plot our preferred fitting formula (solid line), the hyperbolic function in Eq. (3.4.4), and a comparison with a second-order polynomial (dashed line) for $E_{\rm rad}$ as a function of χ_i . Our data points are plotted as polygons, where more sides indicates higher resolution level. In the bottom panel, we plot the residuals of the hyperbolic function. We indicate our fit parameter (dotted line) and total prediction (dashed line) uncertainties (defined in Appendix A1).


Figure 21: $E_{\rm rad}$ as a function of initial spin. In the left panels we plot our data (circles) along with the fitting formula (red line) with error estimates (dashed) from several other studies. The top panel is from Ref. [29], the middle is from Ref. [30], and the bottom is from Ref. [31]. In the right panels we plot the difference between our data and the corresponding fitting formula on the left. The value r quantifies the size of the systematic error compared to the 3-parameter hyperbola.

III Extremality

An important aspect of these fitting formulae is their ability to predict remnant properties for nearly extremal initial spins. How much of the initial mass can be radiated as gravitational waves, and how fast can the remnant hole spin? Our prediction for the radiated energy and final spin for an extremal initial spin configuration, $\chi_i = 1$, is

$$\tilde{E}_{\rm rad}(1) = 0.11397(18),$$
 (3.4.7)

$$\tilde{\chi}_f(1) = 0.951383(85),$$
(3.4.8)

where the uncertainty (in parentheses) is σ_{tot} , the total prediction uncertainty defined in Eq. (A1.13), evaluated at $\chi_i = 1$. The highest radiated energy predicted by any of the formulae we compare against in this paper is $E_{rad}(1) = 0.0995(8)$.

Previous estimates of $E_{\rm rad}$ underestimated the mass loss for large, aligned initial spins. The most extreme data point in this paper, $S_{0.97}^{++}$, was identified as a potential outlier [31] and $E_{\rm rad}$ was expected to be $\leq 10\%$ for an extremal, aligned configuration inspiraling from infinity. Additional data presented here, most notably $S_{0.95}^{++}$ and $S_{0.9}^{++}$, suggest that $S_{0.97}^{++}$ is not an outlier. Furthermore, these cases indicate that even a $\chi_i = 0.9$ inspiral is capable of radiating $\geq 10\%$ of its initial mass.

Simulations with $\chi_i > 0.93$ are an important factor in our fitting formulae. This high-spin regime is not accessible with the most popular initial data methods for binary black hole evolutions, which assume conformal flatness (cf. Ref. [39] and references therein). To assess the impact of the high spin simulations, we compare our best fits to fits of a subset of the data, omitting cases $S_{0.95}^{--}$, $S_{0.95}^{++}$, and $S_{0.97}^{++}$. We identify several key results.

For the χ_f formula, we find that $\Delta \tilde{\chi}_f(1) \gtrsim 2.5 \sigma_{\text{tot}}^{\text{sub}}(\chi_f, 1)$. That is, the prediction of the final spin with the full dataset differs from the prediction with the subset by more than 2.5 times the total prediction uncertainty in the fit to the subset. The parameter uncertainty in the full dataset is smaller by a factor $\sigma_f/\sigma_f^{\text{sub}} \approx 0.6$, which is much smaller than would

be expected from adding 3 data points randomly distributed in initial spin (the expected improvement based on the root-N rate would be $\sqrt{12/15} \approx 0.9$). In a random sampling context, one would typically have to more than double the size of the dataset to get such a reduction in parameter uncertainties. Of course, we have not chosen the subset randomly. Note that because the systematic error magnitude is negligible, $\hat{\sigma}_{\Delta} \approx 0$, the total prediction uncertainty of Eq. (A1.13) has the same behavior as the parameter uncertainty.

For the $E_{\rm rad}$ formula, we find that $\Delta \tilde{E}_{\rm rad}(1) \gtrsim 3.5 \sigma_{\rm tot}^{\rm sub}(E_{\rm rad}, 1)$; the lower spin subset poorly predicts the extremal $E_{\rm rad}$. Parameter uncertainties decrease only slightly faster than the expected root-N rate for adding 3 randomly placed data points, $\sigma_f/\sigma_f^{\rm sub} \approx 0.85$. However, the total prediction uncertainty at $\chi_i = 1$ increases, $\sigma_{\rm tot}/\sigma_{\rm tot}^{\rm sub} \approx 1.15$, because the additional high spin data deviates most from the fitting formula that is based on lower spin data. That is, $E_{\rm rad}$ for $\chi_i > 0.93$ is unanticipated by the fit to the lower spin subset, causing the systematic error magnitude $\hat{\sigma}_{\Delta}$ to increase. While this highlights the importance of the high spin data in assessing the predictive power of the fitting formula for near-extremal initial spins, it also suggests that we are unlikely to capture the behavior of $E_{\rm rad}$ much better with our simple fitting formula. Furthermore, both manual and algorithmic searches [56] have not identified any better formulae, which leads us to believe that for the best predictive results at high spins, a non-parametric approach may be preferred. Such an approach could be implemented using, for example, a correlated Gaussian process [57], which would provide a way to predict final masses and spins without the use of a parametric fitting formula.

The analysis comparing the subset to the full dataset does not change in any appreciable way if we use a near-extremal spin, e.g. $\chi_i = 0.97$, instead of the most extreme case, $\chi_i = 1$.

3.5 Conclusions

In this paper, we present and analyze a family of numerical relativity simulations performed using SpEC in order to construct improved fitting formulae for the final spin and radiated energy as a function of initial spin. We consider a physically motivated, one-dimensional subset of the binary black hole parameter space, in which the black holes have equal masses and equal spins aligned or anti-aligned with the orbital angular momentum. The improvement in these fitting formulae is most dramatic in the regime where the initial spin is above the "Bowen-York limit," since for the first time data from simulations above this limit have been included in the fits.

For the final spin, we improve on the second-order polynomial fitting formula by using a fourth-order polynomial to capture the statistically significant cubic and quartic features. For the radiated energy, we find that a 3-parameter hyperbolic fitting formula is greatly preferred to a second-order polynomial. The qualitatively different behavior at large, aligned spins in the new fit to $E_{\rm rad}$ implies that there is somewhat more power in gravitational waves from nearly extremal sources than previously thought, perhaps because of higher-order effects that become relevant at very high spins.

We have shown that performing more nearly extremal simulations is the most effective way to reduce the uncertainty in the fitting formula parameters. However, we have also observed that the systematic uncertainty in $E_{\rm rad}$ may prohibit a simple fitting formula from providing any further significant improvement to the prediction uncertainty of $E_{\rm rad}$ for high, aligned spins.

Analytic models, such as the aligned-spin EOB model, are needed to generate templates for gravitational-wave detectors (e.g. LIGO), because of the prohibitive expense of generating sufficient numerical relativity data to adequately cover the parameter space. The fitting formulae we define in this paper can be used to better calibrate these models, and therefore improve future template waveforms.

A1 Method for constructing our fitting formulae

Our goal is to find a convenient but reasonably accurate function that predicts the final black hole spin, χ_f , or the fractional radiated energy, $E_{\rm rad}$, as a function of the initial spin, χ_i . We will specify one or more simple parametric candidate functions, find the best parameter values, quantify uncertainties in the parameters and predictions, and compare rival candidate functions. To treat both the χ_f and $E_{\rm rad}$ problems in generic notation, we let ξ denote the *predictor* (i.e., χ_i), and η denote the *response* we seek to predict (i.e., χ_f or $E_{\rm rad}$). We have one or more parametric models for the relationship, $\eta \approx f(\xi; \theta)$, with parameters θ (we sometimes suppress the parameter dependence below to simplify notation).

The data for the analysis are from post-processing outputs from deterministic numerical calculations of the binary black hole merger. A complex computation produces initial data (ID) targeting a specified value of ξ , but the actual value of ξ that the generated ID corresponds to necessarily differs from the target value. A processing algorithm estimates the actual value to be x. Evolution of the ID produces high-dimensional outputs that are processed to produce the computed response, y, that estimates the result, η , that would be obtained by solving the PDEs exactly. A set of (x, y) pairs constitutes the basic data we must use to find $f(\xi; \theta)$.

A variety of parameters govern the accuracy of the ID, evolution, and processing algorithms. These are summarized via a resolution level k (defined in Sec. 3.2) assigned to each (x, y) pair, with the x and y values likely to be closer to the ξ and η values when k is larger. For every choice of ID, we have results for multiple values of k, comprising repeated measurements of (ξ, η) of varying accuracy. We have developed a Bayesian nonlinear measurement error model for the analysis.⁴ Letting the index n label the choice of ID, the model specification is:

$$x_{nk} = \xi_n + \epsilon_{nk},\tag{A1.1}$$

$$y_{nk} = \eta_n + \delta_{nk} \tag{A1.2}$$

$$= f(\xi_n; \theta) + \Delta_n + \delta_{nk}, \tag{A1.3}$$

for N ID cases (n = 1 to N), and $k \in L_n$ for ID case n, where L_n denotes a set of levels run for case n (for most cases, $L_n = \{2, 3, 4\}$, but for runs targeting $\chi_i = \pm 0.44$, $L_n = \{1, 2, 3\}$). Here ϵ_{nk} and δ_{nk} denote *level error* terms reflecting the difference between numerical results at finite resolution and the actual solution to the differential equations we are studying. For Eq. (A1.3) we set $\eta_n = f(\xi_n; \theta) + \Delta_n$, where Δ_n is a *discrepancy* term representing the difference between the true response and the prediction based on the fitting function.

To complete the model we must assign (prior) probability density functions (PDFs), i.e. *priors*, to a number of random variables: the level errors, ϵ_{nk} and δ_{nk} ; the latent predictor variables, ξ_n ; and the latent discrepancy variables, Δ_n .

We assign independent, zero-mean normal distributions to the level error terms, ϵ_{nk} and δ_{nk} , with standard deviations σ_x/α_k and σ_y/α_k (respectively). We assign α_k scale factors to capture the notion that we expect the errors to be smaller (on average) for higher levels. For the calculations here, we took $\alpha_k = (1/2)^{4-k}$, so the standard deviations for the highest-resolution k = 4 results are σ_x and σ_y , and the error scales double for each decreasing level. We did not explore this assignment except to verify that this choice has a much higher likelihood than taking $\alpha_k = 1$, i.e., the data themselves show clear evidence for convergence as k grows. Although in principle we could let the error scale be different for each ID case, for simplicity we assign a common error scale across ID cases; the modest amount of data we have do not indicate a strong variation of error scale with ID. We adopt normal

⁴ For introductions aimed at physicists, see Refs. [58, 59, 60] for Bayesian inference and Ref. [61] for multilevel Bayesian modeling. Multilevel measurement error models inspiring our approach here are covered in Refs. [62, 63].

distributions, partly for convenience, but also because we are modeling relationships between scalar quantities calculated from high-dimensional computational outputs with complicated algorithms. Presuming the final errors result from numerous additive contributions whose uncertainties have finite variance, the central limit theorem motivates the normal choice.

We assign informative but relatively broad priors for the ξ_n values, reflecting the ability to produce ID corresponding to a ξ_n value close to a desired target value, μ_n . The priors are normal with means μ_n (equal to the target value for ID case n) and common standard deviation w = 0.002, reflecting the typical change in mass and spin as a result of the initial relaxation (as seen in Fig. 15). These values do not strongly impact the results.

We also assign independent, zero-mean normal distribution PDFs to the discrepancy terms, with common standard deviation σ_{Δ} . The quantity σ_{Δ} represents the typical scale of systematic error magnitude in the model. A more flexible and realistic choice would be to assign a correlated Gaussian process prior over the space of discrepancy functions, $\Delta(\xi_n)$, and to identify $\Delta_n = \Delta(\xi_n)$. This would resemble the practice in the literature on Bayesian emulation of input/output response surfaces, the prevailing approach in the literature on the statistical analysis of the results of deterministic numerical simulations (see, e.g., [64, 65]). But the goal of that literature is not to find simple and tractable fitting functions; it instead builds nonparametric emulators that, while simpler than the simulators being emulated, are still computationally nontrivial. Moreover, the vast majority of existing work on emulation addresses cases with precisely known inputs, which is not the case here; uncertainty in the predictor significantly complicates implementation of Gaussian process regression [66, 67]. The independent normal PDF for Δ_n will enable us to invoke a simple approximation leading to analytical results.

Finally, we adopt flat priors for the fitting function model parameters, θ .

The conditional dependency structure of such a multilevel model can be represented by a directed acyclic graph (DAG). A graphical model of this type can be readily coded in a



Figure 22: Directed acyclic graph displaying the conditional dependence structure of the Bayesian nonlinear measurement error model adopted for the fitting function analysis. See text for a detailed description.

DAG-oriented statistical modeling language (e.g., WinBUGS or JAGS) to enable Bayesian computation via Markov chain Monte Carlo (MCMC) posterior sampling. Here the focused goal (finding a simple fitting function) and the small uncertainties in the level error terms (well below 1% for nonzero spins) motivated an analytical approach based on linearization of $f(\xi)$. This lets us avoid the complexity of MCMC, producing a fast algorithm that is relatively simple to use.

Figure 22 shows the DAG for our model. Circles denote random variables (RVs, uncertain quantities with assigned or computed PDFs). Shaded circles are the data (x, y), and shaded squares are fixed constants that help define the model. We marginalize over the error RVs $(\epsilon_{nk}, \delta_{nk}, \text{ and } \Delta_n)$ and the uncertain input variables (ξ_n) , and then we solve for the remaining non-shaded variables simultaneously. The plates (enclosing boxes) denote parts of the graph that are replicated as indicated by the quantity in the lower right corner of each plate.

Dashed circles indicate RVs that play the role of *hyperparameters*, i.e., parameters defining prior PDFs for lower-level RVs. Formally, we could account for uncertainty in the hyperparameters by assigning them priors of their own and marginalizing over them (the *hierarchical Bayes* approach). As a simpler approximation, we optimized these hyperparameters (the *empirical Bayes* approach), using constant prior PDFs for them. Directed edges (arrows) in a Bayesian network DAG are used to indicate the dependency structure. The top-level RVs have no dependencies (incoming arrows); their PDFs would be specified a priori for a full hierarchical Bayesian analysis. The conditional PDFs of lower-level RVs depend only on the values of their dependencies. The full joint PDF for all RVs is the product of the prior and conditional PDFs. Therefore, Fig. 22 indicates that the joint PDF for the RVs comprising our model may be written

$$p(\theta, \sigma_x, \sigma_y, \sigma_\Delta, \xi, \Delta, \epsilon, \delta, x, y) = p(\theta) p(\sigma_x) p(\sigma_y) p(\sigma_\Delta)$$

$$\times \prod_{n=1}^N \left[p(\xi_n | \mu_n, w) \, p(\Delta_n | \sigma_\Delta) \right]$$

$$\times \prod_{k \in L_n} p(\epsilon_{nk} | \sigma_x) p(x_{nk} | \xi_n, \epsilon_{nk})$$

$$\times p(\delta_{nk} | \sigma_y) p(y_{nk} | \theta, \xi_n, \Delta_n, \delta_{nk}) \right], \qquad (A1.4)$$

where $(\xi, \Delta, \epsilon, \delta, x, y)$ is shorthand notation for the indexed collections of those variables. Since we are adopting a constant prior PDF for θ , and an empirical Bayes treatment of the hyperparameters $\psi = (\sigma_x, \sigma_y, \sigma_\Delta)$, the quantity of interest is the conditional PDF for the data, (x, y), and the latent parameters, $(\xi, \Delta, \epsilon, \delta)$, given the fitting function parameters and the hyperparameters,

$$p(\xi, \Delta, \epsilon, \delta, x, y | \theta, \psi) = \prod_{n=1}^{N} \left[p(\xi_n | \mu_n, w) \, p(\Delta_n | \sigma_\Delta) \right]$$
$$\times \prod_{k \in L_n} p(\epsilon_{nk} | \sigma_x) p(x_{nk} | \xi_n, \epsilon_{nk})$$
$$\times p(\delta_{nk} | \sigma_y) p(y_{nk} | \theta, \xi_n, \Delta_n, \delta_{nk}) \right],$$
(A1.5)

The model Eqs. (A1.1) and (A1.3) imply that the conditional PDFs for x_{nk} and y_{nk} in these equations are δ -functions. This lets us trivially marginalize over ϵ and δ , giving a

marginal PDF for the remaining variables,

$$p(\xi, \Delta, x, y | \theta, \psi) = \prod_{n=1}^{N} \left[p(\xi_n | \mu_n, w) p(\Delta_n | \sigma_\Delta) \right]$$
$$\times \prod_{k \in L_n} p(\epsilon_{nk} = x_{nk} - \xi_n | \sigma_x)$$
$$\times p(\delta_{nk} = y_{nk} - f(\xi_n; \theta) - \Delta_n | \sigma_y) \right].$$
(A1.6)

Marginalizing over ξ and Δ gives the marginal likelihood function (the probability for the data, conditioned on parameter values) for the fitting parameters and hyperparameters,

$$\mathcal{L}_{M}(\theta,\psi) = \prod_{n=1}^{N} \int d\xi_{n} \int d\Delta_{n} \left[\mathcal{N}(\xi_{n}|\mu_{n},w) \mathcal{N}(\Delta_{n}|\sigma_{\Delta}) \times \prod_{k\in L_{n}} \mathcal{N}(x_{nk}-\xi_{n}|0,\sigma_{x}) \mathcal{N}(y_{nk}-f(\xi_{n};\theta)-\Delta_{n}|0,\sigma_{y}) \right],$$
(A1.7)

where $\mathcal{N}(z|\mu,\sigma)$ denotes the normal distribution PDF for z with mean μ and standard deviation σ ,

$$\mathcal{N}(z|\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(z-\mu)^2/2\sigma^2}.$$
(A1.8)

When $f(\xi)$ is a nonlinear function of ξ , the ξ integral in Eq. (A1.7) is in general intractable. However, the x and y errors are small, so we expect a local linear approximation of $f(\xi)$ to be very accurate over regions of ξ_n that have significant probability density. So we use

$$f(\xi_n;\theta) \approx f(\tilde{\xi}_n;\theta) + (\xi_n - \tilde{\xi}_n)f'(\tilde{\xi}_n;\theta)$$
(A1.9)

in Eq. (A1.7), where $f'(\xi; \theta)$ denotes the derivative of the fitting function with respect to ξ , and $\tilde{\xi}_n$ is a fixed reference value of ξ_n based on the x_{nk} values for a particular n (we use a weighted mean of the x_{nk}). With this linearization, the integrals in the marginal likelihood function can be performed analytically.

We estimate the parameters for a candidate fitting function by maximizing the marginal likelihood function over both θ and ψ :

$$(\hat{\theta}, \hat{\psi}) = \arg \max \mathcal{L}_M(\theta, \psi).$$
 (A1.10)

For the fitting functions studied here, the θ dependence of the marginal likelihood function is approximately multivariate Gaussian. To quantify the θ uncertainties, we calculate the observed Fisher information matrix (with ψ fixed at $\hat{\psi}$),

$$I_{\alpha\beta} = \frac{\partial^2}{\partial \theta_{\alpha} \partial \theta_{\beta}} \mathcal{L}_M(\theta, \psi) \Big|_{\hat{\theta}, \hat{\psi}}, \tag{A1.11}$$

where θ_{α} denotes the α th parameter of the fitting function. The posterior PDF for θ (conditional on $\hat{\psi}$) is then approximately a multivariate normal PDF with mean $\hat{\theta}$ and covariance matrix $\Sigma = I^{-1}$.

To predict the value of the response at a specified value of ξ , we calculate an approximate predictive distribution (also conditioned on $\hat{\psi}$) using the multivariate normal PDF and the delta method (propagation of errors). The model assumes the response is given by the sum of the fitting function and a discrepancy term with zero mean. The most probable value of the response is simply $f(\xi; \hat{\theta})$. There are two components to the uncertainty in the prediction. One comes from propagating the θ uncertainty (accounting for correlations between the parameters, which can be large). The resulting standard deviation in the fitting function evaluated at ξ is $\sigma_f(\xi)$, satisfying

$$\sigma_f^2(\xi) = \sum_{\alpha\beta} \frac{\partial f(\xi;\theta)}{\partial \theta_\alpha} \Sigma_{\alpha\beta} \frac{\partial f(\xi;\theta)}{\partial \theta_\alpha}.$$
 (A1.12)

The full uncertainty in the predicted response must also account for the uncertainty in the discrepancy term, which is given by the hyperparameter σ_{Δ} that we estimate from the data. The full uncertainty in the prediction is

$$\sigma_{\rm tot}(\xi) = \sqrt{\sigma_f^2(\xi) + \hat{\sigma}_{\Delta}^2}.$$
 (A1.13)

This calculation ignores the uncertainty in the value of σ_{Δ} , but that uncertainty is relatively small in our calculations.

To compare rival parametric fitting functions, a formal model comparison could be implemented, e.g., using Bayes factors (which would require assigning normalized priors to θ for each candidate fitting function, and integrating the product of the prior and the marginal likelihood function over θ), or an information criterion such as the Bayesian information criterion (BIC) or the Akaike information criterion (AIC). The BIC and AIC rank models according to their maximum likelihoods, penalized by a term depending on the number of parameters in each model (and the sample size in the case of the BIC). These criteria were developed for comparing simple parametric models, not multilevel models with many latent parameters. We adopt a less formal approach here. We simply calculate the logarithm of ratios of the maximum marginal likelihood function. For the models we consider, the log-ratio for the best model vs. the next-best competitor is large (well over 10), far larger than the typical penalty terms in information criteria, so the choice of best model is unambiguous.

Acknowledgments

We thank Matt Giesler, Tony Chu, and Bryant Garcia for providing data from their simulations. We gratefully acknowledge support from the Sherman Fairchild Foundation; from NSF grants PHY-0969111 and PHY-1005426 at Cornell; and from NSF grants PHY-1068881, PHY-1005655, and DMS-1065438 at Caltech. Loredo's work for this project was supported by NSF grant AST-0908439 and NASA grant NNX09AK60G. Simulations used in this work were computed with SpEC [32]. Computations were performed on SHC at Caltech, which is supported by the Sherman Fairchild Foundation; on the Zwicky cluster at Caltech, which is supported by the Sherman Fairchild Foundation and by NSF award PHY-0960291; on the NSF XSEDE network under grant TG-PHY990007N; and on the GPC supercomputer at the SciNet HPC Consortium [68]. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund–Research Excellence; and the University of Toronto.

REFERENCES

- B. Abbott et al. LIGO: The Laser interferometer gravitational-wave observatory. *Rept. Prog. Phys.*, 72:076901, 2009.
- [2] H. Grote. The status of GEO 600. Class. Quant. Grav., 25:114043, 2008.
- [3] F. Acernese, M. Alshourbagy, P. Amico, F. Antonucci, S. Aoudia, et al. Virgo status. *Class. Quant. Grav.*, 25:184001, 2008.
- [4] J. Abadie et al. Search for gravitational waves from binary black hole inspiral, merger and ringdown. *Phys. Rev. D*, 83:122005, 2011.
- [5] The LIGO Scientific Collaboration, the Virgo Collaboration, J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, and et al. Search for Gravitational Waves from Binary Black Hole Inspiral, Merger and Ringdown in LIGO-Virgo Data from 2009-2010. ArXiv e-prints, September 2012.
- [6] Kip S. Thorne. Disk-accretion onto a black hole. ii. evolution of the hole. Astrophys. J., 191:507, 1974.
- [7] Charles F Gammie, Stuart L Shapiro, and Jonathan C McKinney. Black hole spin evolution. *Astrophys. J.*, 602:312, 2004.
- [8] Michael Kesden, Guglielmo Lockhart, and E. Sterl Phinney. Maximum black-hole spin from quasi-circular binary mergers. *Phys. Rev. D*, 82:124045, 2010.
- [9] Jeffrey E McClintock, Rebecca Shafee, Ramesh Narayan, Ronald A Remillard, Shane W

Davis, and Li-Xin Li. The spin of the near-extreme Kerr black hole GRS 1915+105. *Astrophys. J.*, 652:518, 2006.

- [10] Jian-Min Wang, Yan-Mei Chen, Luis C Ho, and Ross J McLure. Evidence for rapidly spinning black holes in quasars. Astrophys. J., 642:L111, 2006.
- [11] L.W. Brenneman, C.S. Reynolds, M.A. Nowak, R.C. Reis, M. Trippe, et al. The Spin of the Supermassive Black Hole in NGC 3783. Astrophys. J., 736:103, 2011.
- [12] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaría, and J. Seiler. Inspiral-merger-ringdown waveforms for black-hole binaries with nonprecessing spins. *Phys. Rev. Lett.*, 106:241101, June 2011.
- [13] Chris Van Den Broeck, Duncan A. Brown, Thomas Cokelaer, Ian Harry, Gareth Jones, et al. Template banks to search for compact binaries with spinning components in gravitational wave data. *Phys.Rev.*, D80:024009, 2009.
- [14] J. Aasi et al. Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network. 2013.
- [15] Andrea Taracchini, Yi Pan, Alessandra Buonanno, Enrico Barausse, Michael Boyle, et al. Prototype effective-one-body model for nonprecessing spinning inspiral-merger-ringdown waveforms. *Phys. Rev. D*, 86:024011, 2012.
- [16] James F. Steiner and Jeffrey E. McClintock. Modeling the Jet Kinematics of the Black Hole Microquasar XTE J1550-564: A Constraint on Spin-Orbit Alignment. Astrophys. J., 745:136, 2012.
- [17] Tamara Bogdanovic, Christopher S. Reynolds, and M. Coleman Miller. Alignment of the spins of supermassive black holes prior to merger. *Astrophys.J.Lett.*, 2007.
- [18] T. Fragos, M. Tremmel, E. Rantsiou, and K. Belczynski. Black Hole Spin-Orbit Misalignment in Galactic X-ray Binaries. 2010.
- [19] Geoffrey Lovelace, Michael Boyle, Mark A. Scheel, and Béla Szilágyi. High-accuracy

gravitational waveforms for binary-black-hole mergers with nearly extremal spins. *Class. Quant. Grav.*, 29:045003, 2012.

- [20] Geoffrey Lovelace, Mark. A. Scheel, and Béla Szilágyi. Simulating merging binary black holes with nearly extremal spins. *Phys. Rev. D*, 83:024010, 2011.
- [21] Bryant Garcia, Geoffrey Lovelace, Lawrence E. Kidder, Michael Boyle, Saul A. Teukolsky, Mark A. Scheel, and Béla Szilágyi. Are Different Approaches to Constructing Initial Data for Binary Black Hole Simulations of the Same Astrophysical Situation Equivalent? *Phys. Rev. D*, 86:084054, 2012. Accepted for publication in Phys. Rev. D.
- [22] 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*, 80:124051, 2009.
- [23] Tony Chu. Numerical simulations of black-hole spacetimes. PhD thesis, California Institute of Technology, 2012.
- [24] Gregory B. Cook and James W. York, Jr. Apparent horizons for boosted or spinning black holes. *Phys. Rev. D*, 41(4):1077–1085, 1990.
- [25] Sergio Dain, Carlos O. Lousto, and Ryoji Takahashi. New conformally flat initial data for spinning black holes. *Phys. Rev. D*, 65:104038, 2002.
- [26] Mark Hannam, Sacha Husa, and Niall. Ó Murchadha. Bowen-york trumpet data and black-hole simulations. *Phys. Rev. D*, 80:124007, 2009.
- [27] Luciano Rezzolla et al. Spin Diagrams for Equal-Mass Black-Hole Binaries with Aligned Spins. Astrophys. J., 679:1422–1426, 2008.
- [28] Enrico Barausse and Luciano Rezzolla. Predicting the direction of the final spin from the coalescence of two black holes. *Astrophys.J.*, 704:L40–L44, 2009.
- [29] Wolfgang Tichy and Pedro Marronetti. The Final mass and spin of black hole mergers. *Phys. Rev. D*, 78:081501, 2008.

- [30] Christian Reisswig, Sascha Husa, Luciano Rezzolla, Ernst Nils Dorband, Denis Pollney, et al. Gravitational-wave detectability of equal-mass black-hole binaries with aligned spins. *Phys. Rev. D*, 80:124026, 2009.
- [31] Enrico Barausse, Viktoriya Morozova, and Luciano Rezzolla. On the mass radiated by coalescing black-hole binaries. *Astrophys.J.*, 758:63, 2012.
- [32] http://www.black-holes.org/SpEC.html.
- [33] James W. York. Conformal "thin-sandwich" data for the initial-value problem of general relativity. *Phys. Rev. Lett.*, 82(7):1350–1353, Feb 1999.
- [34] Gregory B. Cook. Corotating and irrotational binary black holes in quasicircular orbits. *Phys. Rev. D*, 65(8):084003, Mar 2002.
- [35] Gregory B. Cook and Harald P. Pfeiffer. Excision boundary conditions for black-hole initial data. *Phys. Rev. D*, 70(10):104016, Nov 2004.
- [36] 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.
- [37] Eric Gourgoulhon, Philippe Grandclément, and Silvano Bonazzola. Binary black holes in circular orbits. I. A global spacetime approach. *Phys. Rev. D*, 65:044020, 2002.
- [38] 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.
- [39] 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.
- [40] Geoffrey Lovelace. Reducing spurious gravitational radiation in binary-black-hole simulations by using conformally curved initial data. *Class. Quantum Grav.*, 26:114002, 2009.

- [41] 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.
- [42] Luisa T. Buchman, Harald P. Pfeiffer, Mark A. Scheel, and Béla Szilágyi. Simulations of unequal mass binary black holes with spectral methods. *Phys. Rev. D*, 86:084033, 2012.
- [43] Lee Lindblom and Béla Szilágyi. An improved gauge driver for the GH Einstein system. Phys. Rev. D, 80:084019, 2009.
- [44] Matthew W. Choptuik and Frans Pretorius. Ultra Relativistic Particle Collisions. Phys. Rev. Lett., 104:111101, 2010.
- [45] Béla Szilágyi, Lee Lindblom, and Mark A. Scheel. Simulations of Binary Black Hole Mergers Using Spectral Methods. *Phys. Rev. D*, 80:124010, 2009.
- [46] Abdul H. Mroué, Harald P. Pfeiffer, Lawrence E. Kidder, and Saul A. Teukolsky. Measuring orbital eccentricity and periastron advance in quasi-circular black hole simulations. *Phys. Rev. D*, 82:124016, 2010.
- [47] Daniel A. Hemberger, Mark A. Scheel, Lawrence E. Kidder, Béla Szilágyi, Geoffrey Lovelace, Nicholas W. Taylor, and Saul A. Teukolsky. Dynamical Excision Boundaries in Spectral Evolutions of Binary Black Hole Spacetimes. *Class. Quantum Grav.*, 30(11):115001, 2013.
- [48] Carsten Gundlach. Pseudospectral apparent horizon finders: An efficient new algorithm. Phys. Rev. D, 57(2):863–875, Jan 1998.
- [49] Ivan Booth and Stephen Fairhurst. Extremality conditions for isolated and dynamical horizons. Phys. Rev. D, 77(8):084005, 2008.
- [50] 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.
- [51] Kashif Alvi. Energy and angular momentum flow into a black hole in a binary. Phys. Rev. D, 64(10):104020, Oct 2001.

- [52] Manuela Campanelli, Carlos O. Lousto, and Yosef Zlochower. Spin-orbit interactions in black-hole binaries. *Phys. Rev. D*, 74:084023, 2006.
- [53] Carlos O. Lousto, Manuela Campanelli, Yosef Zlochower, and Hiroyuki Nakano. Remnant masses, spins, and recoils from the merger of generic black-hole binaries. *Class. Quant. Grav.*, 27:114006, 2010.
- [54] Laszlo A. Gergely and Peter L. Biermann. The typical mass ratio and typical final spin in supermassive black hole mergers. 2012.
- [55] S. S. Wilks. The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. The Annals of Mathematical Statistics, 9(1):60–62, 1938.
- [56] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. Science, 324(5923):81–85, 2009.
- [57] Carl Edward Rasmussen and Christopher K. I. Williams. Gaussian processes for machine learning. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, 2006.
- [58] E. T. Jaynes and G. L. Bretthorst. *Probability Theory*. April 2003.
- [59] D. S. Sivia. *Data analysis*. Oxford Science Publications. Oxford University Press, Oxford, second edition, 2006. A Bayesian tutorial, With J. Skilling.
- [60] P. Gregory. Bayesian Logical Data Analysis for the Physical Sciences. May 2010.
- [61] Thomas J. Loredo. Bayesian astrostatistics: A backward look to the future. In Joseph M. Hilbe, editor, Astrostatistical Challenges for the New Astronomy, volume 1 of Springer Series in Astrostatistics, pages 15–40. Springer New York, 2013.
- [62] Raymond J. Carroll, David Ruppert, Leonard A. Stefanski, and Ciprian M. Crainiceanu. Measurement error in nonlinear models, volume 105 of Monographs on Statistics and Applied Probability. Chapman & Hall/CRC, Boca Raton, FL, second edition, 2006. A modern perspective.

- [63] D.M. Giltinan. *Nonlinear models for repeated measurement data*. Monographs on Statistics and Applied Probability. Chapman and Hall, 1995.
- [64] Marc C. Kennedy and Anthony O'Hagan. Bayesian calibration of computer models. J. R. Stat. Soc. Ser. B Stat. Methodol., 63(3):425–464, 2001.
- [65] Dave Higdon, Marc Kennedy, James C. Cavendish, John A. Cafeo, and Robert D. Ryne. Combining field data and computer simulations for calibration and prediction. SIAM J. Sci. Comput., 26(2):448–466, 2004.
- [66] Andrew McHutchon and Carl Edward Rasmussen. Gaussian process training with input noise. In J. Shawe-Taylor, R.S. Zemel, P. Bartlett, F.C.N. Pereira, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems 24, pages 1341–1349. 2011.
- [67] Patrick Dallaire, Camille Besse, and Brahim Chaib-draa. An approximate inference with gaussian process to latent functions from uncertain data. *Neurocomputing*, 74(11):1945– 1955, May 2011.
- [68] Chris Loken, Daniel Gruner, Leslie Groer, Richard Peltier, Neil Bunn, Michael Craig, Teresa Henriques, Jillian Dempsey, Ching-Hsing Yu, Joseph Chen, L Jonathan Dursi, Jason Chong, Scott Northrup, Jaime Pinto, Neil Knecht, and Ramses Van Zon. SciNet: Lessons Learned from Building a Power-efficient Top-20 System and Data Centre. J. Phys.: Conf. Ser., 256:012026, 2010.