


Accurate gravitational waveforms from binary black-hole systems

Thesis by

Michael Boyle

in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California
2009

(Defended October 20, 2008)

© 2009

Michael Boyle

No rights reserved

*Come, you lost Atoms, to your Centre draw,
And be the Eternal Mirror that you saw:
Rays that have wander'd into Darkness wide,
Return, and back into your Sun subside.*

From Farid ud-Din Attar's twelfth-century masterpiece

The Conference of the Birds

ACKNOWLEDGMENTS

My deep thanks go to Lee Lindblom, who has guided the first steps recounted herein with great patience and generosity. Kip Thorne has, when his busy schedule permitted, provided sage guidance and support, which have been crucial. Harald Pfeiffer and Mark Scheel have also been with me every step of the way, for which I am also very grateful. The four of them have lent their invaluable expertise to all my meanderings, and made sure I didn't get too far off course—no small task.

Much of my education these past five years has also come from my fellow TAPIRs. All the grad students who have given talks at our weekly lunch meetings have broadened my knowledge of astrophysics. Everyone who has come to the relativity group meetings has added to my insight. My thanks to all of them. I especially thank Yanbei Chen, who has not only made group meetings more interesting, but has also sat on my candidacy and thesis committees. Alan Weinstein also sat on both committees. I appreciate their efforts very much, because I'm sure it's not the most entertaining or productive use of their time, though their presence certainly was beneficial to me.

I thank JoAnn Boyd whose capable guidance got me through all sorts of bureaucratic obstacles without a hitch, whose green thumb saved at least one very endangered shamrock plant, and whose friendship brightened my time in Pasadena.

Many other people have also made my time here much more enjoyable, doing

their best to distract me from the oppressive southern-California weather. Sherry Suyu brought me together with Ilya Mandel, Cynthia Chiang, and Luc Bouten for some excellent hot pot. Nate Bode, Geoffrey Lovelace, Rob Owen, Jon Pritchard, and Chris Wegg were friends first and fellow nerds second. Chang-Kook Oh provided much light-hearted relief from studying, along with the tastiest Korean food I've ever had.

My undergraduate experience at the University of Chicago was undoubtedly the most transformative period of my life (so far). Much of that was due to Arunas Liulevicius, whose profound understanding and enthusiastic teaching opened me up to the depths and beauties of rigorous thought. That eccentricity was successfully counterbalanced by Dietrich Müller and Simon Swordy, who brought me on surprisingly diverse explorations of a more sublunary realm by way of cosmic-ray detectors. Meanwhile, my close friends Sarah and Phil Barbeau, Baird Allis, John Allread, Dave Brick, Matt Keeshin, Jason Laine, and Ben Linford kept me sane and happy.

I thank all my family members—Courtney, James, Mary Catherine, Claire, Catherine, Auntie Maureen, Uncle Peter, grandma, and grandpa—who not only put up with me, but encourage me when I need it, and always let me know they love me. Most of all, I thank my parents, who have sacrificed and worked hard all my life to give me (and my brother and sisters) the best education that they could, and have themselves taught me the greatest lessons just by the examples they set. Anything that's good about me came from them.

ABSTRACT

We examine various topics involved in the creation of accurate theoretical gravitational waveforms from binary black-hole systems.

In Chapter 2 a pseudospectral numerical code is applied to a set of analytic or near-analytic solutions to Einstein's equations which comprise a testbed for numerical-relativity codes. We then discuss methods for extracting gravitational-wave data from numerical simulations of black-hole binary systems, and introduce a practical technique for obtaining the asymptotic form of that data from finite simulation domains in Chapter 3. A formula is also developed to estimate the size of near-field effects from a compact binary. In Chapter 4 the extrapolated data is then compared to post-Newtonian (PN) approximations. We compare the phase and amplitude of the numerical waveform to a collection of Taylor approximants, cross-validating the numerical and PN waveforms, and investigating the regime of validity of the PN waveforms. Chapter 5 extends that comparison to include Padé and effective-one-body models, and investigates components of the PN models. In each case, a careful accounting is made of errors. Finally, we construct a long post-Newtonian–numerical hybrid waveform and evaluate the performance of LIGO's current data-analysis methods with it. We suggest certain optimizations of those methods, including extending the range of template mass ratios to unphysical ranges for certain values of the total mass, and a simple analytic cutoff frequency for the templates which results in nearly optimal matches for both Initial and Advanced LIGO.

CONTENTS

Acknowledgments	v
Abstract	vii
Contents	viii
List of figures	xi
List of tables	xviii
1 Introduction	1
1.1 Generating gravitational waves	2
1.2 Detecting gravitational waves	5
1.3 Modeling gravitational waveforms	7
1.4 This thesis	9
2 Testing the numerical code	13
2.1 Introduction	14
2.2 Solution method	16
2.3 Random initial data on flat space	22
2.4 Linear plane wave	27
2.5 Gauge wave	37
2.6 Gowdy spacetime	48
2.7 Discussion	52

3	Extrapolation	59
3.1	Introduction	60
3.2	Extrapolation in brief	66
3.3	Waveform extraction	72
3.4	Extrapolation technique applied to binary inspiral waveforms . . .	80
3.5	Near-field effects in the data	94
3.6	Conclusions	108
4	Comparing NR to PN	111
4.1	Introduction	112
4.2	Generation of numerical waveforms	119
4.3	Generation of post-Newtonian waveforms	160
4.4	PN–NR Comparison Procedure	172
4.5	Estimation of uncertainties	175
4.6	Results	191
4.7	Conclusions	212
5	Energy flux	217
5.1	Introduction	218
5.2	Computation of the numerical gravitational-wave energy flux . . .	225
5.3	Post-Newtonian approximants	235
5.4	Comparison with post-Newtonian approximants: Energy flux . . .	250
5.5	Estimation of (the derivative of) the center-of-mass energy	266
5.6	Comparing waveforms	271
5.7	Conclusions	282

5.8	Appendix: Padé approximants to the energy flux in the test particle limit	286
6	Data analysis	293
6.1	Introduction	294
6.2	Searches for gravitational waves from black-hole binaries	296
6.3	PN–NR hybrid waveform	303
6.4	Detection efficiency of gravitational-wave templates	313
6.5	Recommendations for improvements	327
A	Notation and conventions	335
A.1	Conventions in this thesis	335
A.2	Comparison with other references	344
B	Spin-weighted spherical harmonics	347
B.1	Spin-weighted functions	347
B.2	Behavior under rotation	350
B.3	Multipole decompositions	353
	Bibliography	355

LIST OF FIGURES

2.1	Constraints for Minkowski space with random noise	24
2.2	Error energy for Minkowski space with random noise	26
2.3	Phase error for 1-D sinusoidal linear wave	29
2.4	Constraints for 1-D Gaussian linear wave	33
2.5	Error energy for 1-D Gaussian linear wave	34
2.6	Constraints for 2-D linear waves	38
2.7	Error energy for 2-D linear waves	39
2.8	Constraints for high-amplitude 2-D gauge wave	41
2.9	Error energy for high-amplitude 2-D gauge wave	42
2.10	Constraints for shifted gauge wave	46
2.11	Error energy for shifted gauge wave	47
2.12	Constraints for expanding Gowdy spacetime	50
2.13	Error energy for expanding Gowdy spacetime	51
2.14	Constraints for collapsing Gowdy spacetime	53
2.15	Error energy for collapsing Gowdy spacetime	54
2.16	Comparing constraints for 1-D Gaussian linear waves	56
3.1	Data-analysis mismatch between finite-radius waveforms and the extrapolated waveform for Initial LIGO	63
3.2	Data-analysis mismatch between finite-radius waveforms and the extrapolated waveform for Advanced LIGO	64

3.3	Convergence of the amplitude of the extrapolated Ψ_4 , with increasing order of the extrapolating polynomial, N	84
3.4	Convergence of the phase of the extrapolated Ψ_4 , with increasing order of the extrapolating polynomial, N	85
3.5	Convergence of the phase of Ψ_4 , extrapolated with no correction for the dynamic lapse	86
3.6	Convergence of the amplitude of the extrapolated h , with increasing order of the extrapolating polynomial, N	88
3.7	Convergence of the phase of the extrapolated h , with increasing order of the extrapolating polynomial, N	89
3.8	Comparison of extrapolation of Ψ_4 using different sets of extraction radii	93
3.9	Relative amplitude difference between h data extracted at finite radius and data extrapolated to infinity, and calculated near-field effects . . .	102
3.10	Phase difference between h data extracted at finite radius and data extrapolated to infinity, and calculated near-field effects	103
3.11	Relative amplitude difference between Ψ_4 data extracted at finite radius and data extrapolated to infinity, and calculated near-field effects	106
3.12	Phase difference between Ψ_4 data extracted at finite radius and data extrapolated to infinity, and calculated near-field effects	107
4.1	Eccentricity removal	122
4.2	Spacetime diagram showing the region simulated by the numerical evolutions	129
4.3	Coordinate trajectories of the centers of the black holes	130

4.4	Deviation of total irreducible mass from its initial value	132
4.5	Gravitational waveform extracted at $r = 240M$	135
4.6	Normalized constraint violations of run 30c-1	136
4.7	Unnormalized constraint violations of run 30c-1	137
4.8	Convergence of the gravitational-wave phase without time shifting . .	143
4.9	Convergence of the gravitational-wave phase with time shifting	144
4.10	Convergence of the gravitational-wave amplitude	148
4.11	Convergence of the gravitational-wave amplitude	149
4.12	Difference between areal radius r_{areal} and coordinate radius r of selected extraction surfaces	151
4.13	Convergence of phase extrapolation with extrapolating-polynomial order	153
4.14	Convergence of extrapolated amplitude with extrapolating-polynomial order	155
4.15	Effect of choice of wave-extraction radii on extrapolated phase	156
4.16	Effect of choice of wave-extraction radii on extrapolated amplitude . .	157
4.17	Asymptotic behavior of the average lapse at large radii	181
4.18	Asymptotic behavior of higher angular moments of the lapse at large radii	182
4.19	Comparison of numerical simulation with TaylorT1 3.5/2.5 waveforms—phase difference	189
4.20	Comparison of numerical simulation with TaylorT1 3.5/2.5 waveforms—relative amplitude difference	190
4.21	Numerical and TaylorT1 3.5/2.5 waveforms	192

4.22 Comparison of numerical simulation with TaylorT2 3.5/2.5 waveforms— phase difference	193
4.23 Comparison of numerical simulation with TaylorT2 3.5/2.5 waveforms— relative amplitude difference	194
4.24 Comparison of numerical simulation with TaylorT3 3.5/2.5 waveforms— phase difference	197
4.25 Comparison of numerical simulation with TaylorT3 3.5/2.5 waveforms— relative amplitude difference	198
4.26 Comparison of numerical simulation with TaylorT4 3.5/2.5 waveforms— phase difference	201
4.27 Comparison of numerical simulation with TaylorT4 3.5/2.5 waveforms— relative amplitude difference	202
4.28 Numerical and TaylorT4 3.5/3.0 waveforms	204
4.29 TaylorT4 amplitude comparison for different PN orders	205
4.30 Phase comparison for various PN approximants	207
4.31 Late-time phase comparison for various PN approximants	208
4.32 Phase differences between numerical and post-Newtonian waveforms at $t = t_{(M\omega=-0.063)}$	210
4.33 Phase differences between numerical and post-Newtonian waveforms at $t = t_{(M\omega=-0.1)}$	211
5.1 Some aspects of the numerical simulation	224
5.2 Accuracy of the numerical flux	230
5.3 Contributions of various (l, m) modes to the total numerical gravita- tional wave flux	232

5.4	Accuracy of numerical $\dot{\omega}$	233
5.5	Ratio of GW frequencies ω and ϖ to orbital frequency Ω	252
5.6	Effect of choice of frequency type	253
5.7	Comparison of NR and PN energy flux	255
5.8	Early-time comparison of NR and PN energy flux	256
5.9	Comparison of normalized energy flux F/F_{Newt} for equal-mass systems	257
5.10	Cauchy convergence test of F/F_{Newt} for T- and P-approximants	260
5.11	Fitting several PN approximants to the numerical flux	262
5.12	GW frequency derivative $\dot{\omega}$ for the numerical relativity simulation and various PN approximants at 3.5PN order	265
5.13	Comparison of $\dot{\omega}$ for the numerical results and various PN approximants	267
5.14	Comparison of PN $\dot{\omega}$ with a heavily smoothed version of the numerical $\dot{\omega}$	269
5.15	Comparison of NR and PN $dE/d\omega$ versus GW frequency ω	272
5.16	Phase differences between the numerical waveform, and untuned, orig- inal EOB, untuned Padé, and Taylor waveforms, at two selected times close to merger	276
5.17	Phase differences between untuned and tuned P-approximants and NR waveforms	278
5.18	Phase accuracy of various PN approximants	281
5.19	Low-order normalized energy flux F/F_{Newt} versus GW frequency 2Ω in the test-mass limit	287
5.20	High-order normalized energy flux F/F_{Newt} versus GW frequency 2Ω in the test-mass limit	288

5.21	Convergence of the PN approximants in the test-mass limit	289
5.22	Cauchy convergence test of F/F_{Newt} in the test-mass limit for the T- and P-approximants	290
6.1	Convergence testing for numerical waveforms from a data-analysis per- spective	306
6.2	Amplitude and phase differences between the numerical and post- Newtonian waveforms blended to create the hybrid waveform	310
6.3	The last $t = 5000M_{\odot}$ of the hybrid waveform used in this analysis . . .	312
6.4	Hybrid Caltech–Cornell waveform scaled to various total masses shown against the Initial- and Advanced-LIGO noise curves	314
6.5	Histogram of overlaps found by 300 instances of the Amoeba algorithm	316
6.6	Overlaps between Caltech–Cornell hybrid waveforms and restricted stationary-phase pN waveforms for the Initial-LIGO PSD	320
6.7	Overlaps between Caltech–Cornell hybrid waveforms and restricted stationary-phase pN waveforms for the Advanced-LIGO PSD	321
6.8	Integrand of the inner product for a <i>TaylorF2</i> 3.5 pN waveform	323
6.9	Overlap between Caltech–Cornell waveform and restricted <i>TaylorF2</i> , 3.5 pN waveform as a function of cutoff frequency f_c	325
6.10	Maximum overlaps obtained by allowing η to range over unphysical values, compared to those obtained by restricting the range of η	326
6.11	Candidate f_c values for 3.5 pN templates with Initial LIGO	329
6.12	Candidate f_c values for 3.5 pN templates with Advanced LIGO	330

6.13 Loss in overlap when using our recommendations, compared to results
searching over all template families, masses, mass ratios, and cutoff
frequencies 334

LIST OF TABLES

4.1	Summary of the initial data sets	121
4.2	Overview of low-eccentricity simulations	128
4.3	Summary of uncertainties in the comparison between numerical relativity and post-Newtonian expansions	177
5.1	Summary of PN approximants	236
5.2	Normalized energy flux F/F_{Newt} for the T- and P-approximants at various PN orders and velocities v_{Ω}	258
5.3	Optimal a_5 and v_{pole} that minimize phase differences between tuned EOB models and the numerical simulation	280
5.4	Normalized energy flux F/F_{Newt} in the test-mass limit for the T- and P-approximants at different PN orders and at three different frequencies	291
6.1	Maximum overlaps between Caltech–Cornell hybrid waveforms and restricted SPA pN templates using the Initial-LIGO noise curve	318
6.2	Maximum overlaps between Caltech–Cornell hybrid waveforms and restricted SPA pN templates using the Advanced-LIGO noise curve	319
A.1	Comparison of sign conventions for geometric quantities.	346

CHAPTER

1

INTRODUCTION

Down a quiet road in the woods of Louisiana, the darkness is warm and comforting. The sky is spread wide, bejeweled above a rippling wind sculpture, restless in the solitude. Nearby, on a few delicate fibers, hangs an almost flawless set of mirrors—virtually motionless. An intense beam of light shines across them. With sublime sensitivity, the light echoes out sounds of the distant cosmos.

Almost a century ago, Einstein introduced an elegant theory to explain the nature of gravity. While the fundamentals are fairly simple and well understood, physicists have spent most of the past hundred years teasing out its implications. Gravitational waves, the Big Bang, and black holes—the broad strokes are all there; now we need to fill in the details. In recent years, we have begun to develop the tools that will allow us to investigate the most extreme environments in the universe. They will bring us deeper understanding of the theory’s consequences, and will be vital to the future of astronomy and physics.

Astronomy has always relied almost entirely on observations of light, pushing down to long-wavelength radio waves and up to high-energy gamma rays. A few

very interesting observations have also been made outside the electromagnetic spectrum, with neutrino telescopes and cosmic-ray detectors. As each new window has opened, more unexpected phenomena have been discovered. Now, with gravitational waves, Einstein's theory of general relativity is giving us a new way to extend our senses to the far reaches of the universe.

The signals, however, are extraordinarily subtle, squeezing and stretching the distance between objects near Earth by no more than a few parts in a billion billion, and typically far less.¹ Detecting them requires experiments of unsurpassed precision, and an accurate knowledge of just what the waves should look like. This thesis is an attempt to accurately model some of those waveforms. First, it will be helpful to review some of the most promising sources of gravitational radiation and the methods used to detect them.

1.1 Generating gravitational waves

Gravitational waves are ripples in the fabric of spacetime generated by accelerating masses. The greater the mass and the greater the acceleration, the larger the strain of the gravitational waves. Of course, like all radiation, the waves' amplitudes fall off inversely with the distance r from the source. But simple acceleration is not sufficient. The source needs to have a changing quadrupole moment,² the archetype of which is a simple binary—two bodies in orbit about each other. Indeed, the strongest astrophysical gravitational waves are expected

¹The quantity used to measure the size of a gravitational wave is the *strain*, usually denoted h . It is precisely the fractional change in lengths induced by the wave in free masses.

to be produced by simple binaries.

If we denote the total mass of the system by M , the reduced mass by μ , and the orbital angular velocity by Ω , the typical strain produced by the wave is, in order of magnitude,

$$h \approx \frac{G\mu}{c^2 r} \left(\frac{GM\Omega}{c^3} \right)^{2/3}, \quad (1.1)$$

In terms of familiar scales, this is

$$h \approx 10^{-21} \frac{\mu}{1 M_\odot} \frac{1 \text{ Mpc}}{r} \left(\frac{M}{1 M_\odot} \frac{\Omega}{1 \text{ Hz}} \right)^{2/3}. \quad (1.2)$$

The tiny coefficient gives us concern for the feasibility of detection. We need large masses, orbiting at high frequencies, as near to Earth as possible.

The closest contact binaries—pairs of more or less ordinary stars that are nearly touching—are found at roughly 100 pc from Earth, and orbit with periods as low as several hours [145]. If we assume typical masses of a few times the mass of the sun, this corresponds to a strain of about 10^{-20} . In addition to concern for the minuscule magnitude of the strain, we need to worry about the frequency band in which it is found. On Earth, many things happen on timescales of several hours. Observing such tiny fluctuations without coupling to tides, or thermal cycling, or any number of vibrations is well beyond the capabilities of current earthbound technology, and even at the limit of expected capabilities of planned space-based detectors.

Ordinary stars undergoing fusion are simply too large and loosely bound to accelerate very quickly without breaking apart, which would tend to “smear out” the waves, reducing their strength. More compact objects are needed to produce

²It is possible to have a purely octupole mode, however, the strongest sources expected in nature will be predominantly quadrupolar.

intense gravitational waves suitable for observations in the near future. White dwarfs, neutron stars, and black holes are the densest known concentrations of mass, and are known to be dynamic on timescales as brief as milliseconds [188]. This means that their frequencies can extend up to the kHz range,³ increasing the strain of the gravitational waves they emit, and placing those waves at frequencies that can be detected with high sensitivity.

Slightly elliptic spinning neutron stars would give off gravitational waves due to their changing nonsphericity. The waveforms emitted by such a star would be well modeled—basically a wave of constant frequency, modulated by the changing orientation and velocity of the detector. Similar signals are likely given off by the neutron star of a low-mass X-ray binary [30]. The waves could be detected by demodulating the signal, taking its Fourier transform, and essentially looking for excess power [64].⁴ The longer the observation time, the more sensitivity there is to be gained. On the other hand, the fraction of the star’s mass involved in the type of nonsymmetric motion that gives off gravitational waves would be quite small, meaning that the waves themselves would be correspondingly small.

Alternatively, a close encounter between a pair of compact objects involves essentially all of the mass in nonsymmetric motion. We will see in Chapter 6 that, for merging nonspinning, equal-mass black-hole binaries, $\left(\frac{M}{1M_{\odot}} \frac{\Omega}{1\text{Hz}}\right)$ reaches up to about 10^5 . Thus, at its peak amplitude, the signal from such a binary will be

$$h \approx 10^{-18} \frac{\mu}{1 M_{\odot}} \frac{1 \text{Mpc}}{r} . \quad (1.3)$$

While this signal is indeed tiny for realistic masses and distances, it is nonetheless

³White-dwarf binaries would merge or break up at somewhat lower frequencies; this scale is only valid for neutron stars and black holes.

⁴In fact, “Einstein@Home” is a distributed-computing project which allows users with home computers to donate idle time to searches doing just this type of analysis.

of a size that may be detected in the near future. The number of stellar-mass black holes in our galaxy is estimated to be in the hundreds of millions [193]. Though we have only a rough idea of how often these will meet and merge, it has been estimated that next-generation gravitational-wave detectors could detect merging black-hole binaries anywhere from dozens of times per year to a dozen times per day [102]. Even the lower number provides ample incentive to pursue detection of gravitational waves.

1.2 Detecting gravitational waves

Gravitational waves have been indirectly detected. Thirty years of observation of the binary pulsar B1913+16 have given us an accurate measurement of the system's tightening [247]. The rate at which the binary inspirals agrees with the prediction for energy loss in the form of gravitational waves to within 0.2%. In 1993, the Nobel Prize was awarded to Hulse and Taylor for this observation. To date, however, no *direct* observation has been made.

The main tool we have for direct measurement is a set of three instruments—two (LIGO [245]) located in Louisiana and Washington state in the U.S., and another (VIRGO [3]) in Italy.⁵ Their miles-long arms have mirrors at each end, reflecting high-power laser light back and forth. Using Michelson interferometry, the position of the mirrors is measured to within a small fraction of the width of a proton. As a gravitational wave passes the detectors, it stretches and shrinks the space between the mirrors. But the amplitude of that change is expected to be of roughly the same size as random fluctuations in the instrument, complicating

the detection.

In interferometric detectors, there are three fundamental limiting sources of noise [155]:

- Seismic noise. This is filtered heavily through a system of coupled oscillators and active controls. It is dominant at low frequencies, dropping off very quickly at higher frequencies.
- Thermal noise. Produced by the suspension and within the mirror itself, the power in this noise source typically falls off roughly as $f^{-4.5}$ and is dominant just after the seismic noise drops off.
- Shot noise. This is noise due to photon counting statistics. Increasing with frequency, it takes over from thermal noise, and climbs as f^2 .

Currently the most sensitive interferometers are the LIGO instruments.

With a low signal-to-noise ratio (SNR) it is simply not possible to just look at the data and see the signal. However, if we know what a possible signal would look like, we can test for its presence in the data. Though we can never be absolutely sure that a signal is or is not present, given a data stream $s(t)$ and a template waveform $h(t)$, we can derive some likelihood statistic describing whether or not h is contained in s . This depends on the level of noise in the detector, $S_n(f)$.⁶ We form the inner product

$$(s|h) := 4 \Re \int_0^\infty \frac{\tilde{s}(f) \tilde{h}^*(f)}{S_n(f)} df . \quad (1.4)$$

⁶Other detectors exist and are taking data. GEO [250] and TAMA [9] are other laser interferometers. Explorer [13], Allegro [192], AURIGA [91], NAUTILUS [14], and NIOBE [35] are cryogenic bar detectors sensitive to vibrations induced in the bars by gravitational waves. Each of these instruments has a lower optimal sensitivity than the LIGO–VIRGO instruments.

The larger this inner product is, the greater the likelihood that h is contained in s [65, 173], assuming the templates h are normalized to a constant. But this test assumes that we have a reasonably accurate template waveform h . If the template is inaccurate, it could match with noise just as well as it matches with the real data. Producing accurate waveforms and comparing them to models will be the main topic of this thesis.

1.3 Modeling gravitational waveforms

Science progresses by comparing its predictions to observations. General relativity's most intriguing results have only been encountered indirectly. Gravitational waves have been inferred to carry off energy from an inspiraling pair of stars. The existence of black holes has similarly been inferred from observations of the matter around them—evidence that so much mass is packed into such a small volume that we have no idea what might be in there other than a black hole. It would be comforting to have more direct observations, testing Einstein's theory in the most severe environments in the universe today. Of course, to do that, we need to understand exactly what those predictions are.

To describe the motion of the Moon around the Earth to high accuracy, we only need Newtonian gravity. To describe the motion of Mercury around the Sun, we need a first-order approximation from Einstein's theory [132]. The closer and more massive two orbiting objects are, the more we need general relativity. For a pair of black holes moving slowly, we can expand the terms in Einstein's equations in series, and solve approximately. Techniques for doing these “post-Newtonian

⁶We discuss these techniques in greater detail in Sec. 6.2.

approximations” have evolved over the past two decades, and have reached an impressive stage [42]. For some physical situations, they can accurately predict the waveform to within an orbit or two of the point of merger [59].⁷

However, as the black holes near each other, moving faster and faster, the approximations must eventually break down. We have basically no hope of solving the full, nonlinear Einstein’s equations analytically. Instead, we need to use computers to simulate the physical scenario. While post-Newtonian approximations break down at some point before merger, numerical simulations are costly, and cannot extend for long before merger. We need to extend the simulations to a domain in which we trust the post-Newtonian approximations, and check that the predicted waveforms agree on the overlap. Then, the final waveform will be a marriage of post-Newtonian and numerical methods. In the following chapters we have done just this, demonstrating how to produce an accurate gravitational waveform for binary black holes.

Of course, it is entirely possible that exotic sources other than black-hole binaries exist, and will only be discovered through gravitational-wave astronomy—much as pulsars, active galactic nuclei, and the cosmic microwave background were only discovered with the advent of radio and microwave astronomy. One enticement to gravitational-wave astronomy is exploration of the unknown, as well as observation of the known. But the only way we will be able to discover new phenomena in the data will be to understand the more mundane signals—like binary black holes.

⁷Post-Newtonian approximations are described in some detail in Secs. 4.3 and 5.3.

1.4 This thesis

This thesis examines various topics involved in the creation of accurate theoretical gravitational waveforms from binary black-hole systems. The data presented is from a simulation of an equal-mass, nonspinning system, evolved through 15 orbits, merger, and ringdown. It is used to compare with post-Newtonian approximations, and to evaluate data-analysis techniques for gravitational-wave detectors.

In Chapter 2 we present tests of the numerical code used to evolve black-hole binary systems. The code is applied to a set of exact and approximate solutions to Einstein’s equations, which comprise a testbed for numerical-relativity codes—the so-called “Mexico City Tests”. While the formulation of Einstein’s equations is different from the one used to evolve binaries, the underlying code infrastructure and basic numerical methods are the same. This chapter is extracted with minor revisions from Ref. [62], and was written in collaboration with Lee Lindblom, Harald P. Pfeiffer, Mark A. Scheel, and Lawrence E. Kidder, and published in 2007.

We then introduce a practical technique for obtaining the asymptotic form of gravitational-wave data from numerical simulations with finite-sized domains in Chapter 3. We discuss methods of extracting the data, and some of the errors that would be made in treating this data as the asymptotic waveform. A formula is developed to estimate the size of near-field effects encountered when extracting at finite radius. We show that these effects account for nearly all of the error in Regge–Wheeler–Zerilli data. This chapter (referred to as Ref. [63]) will be incorporated into a paper to be published with Abdul H. Mroué. The data comes

courtesy of Mark A. Scheel, Harald P. Pfeiffer, and Luisa Buchman.

In Chapter 4 data extrapolated from a highly accurate 15-orbit simulation is then compared to post-Newtonian (PN) approximations. We compare the phase and amplitude of the numerical waveform to a collection of Taylor approximants, cross-validating the numerical and PN waveforms, and investigating the regime of validity of the PN waveforms. We find one particular approximant which agrees with our numerical waveform to within the uncertainty throughout most of the inspiral. This chapter is extracted with minor revisions from Ref. [59], which was written in collaboration with Duncan A. Brown, Lawrence E. Kidder, Abdul H. Mroué, Harald P. Pfeiffer, Mark A. Scheel, Gregory B. Cook, and Saul A. Teukolsky, and published in 2007.

Chapter 5 extends that comparison to include Padé and effective-one-body models, and investigates components of the PN models. In each case, a careful accounting is made of errors. The waveforms themselves are also compared, showing that Padé and EOB waveforms do have high accuracy, though only slightly better than the best Taylor approximant. This chapter is extracted with minor revisions from Ref. [61], which was written in collaboration with Alessandra Buonanno, Lawrence E. Kidder, Abdul H. Mroué, Yi Pan, Harald P. Pfeiffer, and Mark A. Scheel. It has been submitted to *Physical Review D*, and is under review.

Finally, in Chapter 6, we use our numerical waveform—attached to a very long PN waveform—to study the detection efficiency of stationary-phase approximated post-Newtonian template waveforms currently used by ground-based gravitational-wave detectors to search for the coalescence of binary black holes.

We perform this study for the Initial and Advanced LIGO detectors. We make various recommendations to improve the application of the templates. We introduce a simple analytic formula which can be used to determine the frequency at which to end the template waveform, and show that using this cutoff achieves nearly optimal matches for all mass ranges in either detector. We suggest that this formula could also find use in searches for other physical situations (e.g., spinning and unequal-mass systems). This chapter (referred to as Ref. [60]) was written in collaboration with Duncan A. Brown and Larne Pekowsky. It has not yet been submitted for publication.

TESTING THE NUMERICAL-EVOLUTION CODE¹

The accuracy and stability of the Caltech-Cornell pseudospectral code is evaluated using the Kidder, Scheel, and Teukolsky (KST) representation of the Einstein evolution equations. The basic “Mexico City Tests” widely adopted by the numerical-relativity community are adapted here for codes based on spectral methods. Exponential convergence of the spectral code is established, apparently limited only by numerical roundoff error or by truncation error in the time integration. A general expression for the growth of errors due to finite machine precision is derived, and it is shown that this limit is achieved here for the linear plane-wave test.

¹This chapter is extracted with minor revisions from Ref. [62], which was written in collaboration with Lee Lindblom, Harald Pfeiffer, Mark Scheel, and Larry Kidder. Note that the evolution system used in this chapter (KST) is not the same system used for the successful binary evolutions presented in other chapters (generalized harmonic). The code infrastructure, however, is the same. The code was written by Harald Pfeiffer, Mark Scheel, and Larry Kidder. I wrote the analytic solutions in the code (except for the linear wave), ran the simulations, and analyzed the data. Crucial suggestions to improve the parameters of the tests and refine the analysis came from all of my co-authors. I led the writing of most of the text, with some exceptions, though all authors contributed significantly. In particular, the introduction is due mostly to Lee Lindblom, Sec. 2.2 is due largely to Mark Scheel, and Sec. 2.7 was in large part written by Harald Pfeiffer.

2.1 Introduction

A number of groups have now developed numerical-relativity codes sophisticated enough to evolve binary black-hole spacetimes [219, 21, 81, 126, 158, 230]. The gravitational waveforms predicted by these evolutions will play an important role in detecting and interpreting the physical properties of the sources of these waves, soon to be detected (we anticipate) by the community of gravitational-wave observers (e.g., LIGO, etc.). Therefore, such codes must be capable of performing stable and accurate simulations of very nonlinear and dynamical spacetimes.

Several years ago a large subset of the numerical-relativity community—the “Apples with Apples” collaboration [7]—proposed a series of basic code tests designed to verify the accuracy, stability, robustness, and efficiency of any code designed to find fully three-dimensional solutions to the Einstein evolution equations. These tests—often referred to as the “Mexico City Tests” because they were first formulated during a conference in Mexico City in May 2002—were designed to be analogous to the standard suite of tests used by the numerical-hydrodynamics community (e.g., tests to reproduce Sedov explosions, Sod shock tubes, blast waves, etc.) to commission new hydrodynamics codes. The Mexico City tests were designed to be applicable to any formulation of Einstein’s equations solved with any numerical method. All tests proposed so far concern bulk properties of the formulation and numerical method, and so all of the evolutions are carried out on a numerical grid with three-torus topology; no boundary conditions are needed (or tested). There are four basic tests, some of them in a number of variations: (a) the evolution of initial data with small, random departures from Minkowski spacetime; (b) the evolution of small-amplitude linear

plane-wave initial data; (c) the evolution of a nonlinear gauge-wave representation of flat spacetime; and (d) the evolution of initial data for a very dynamic and nonlinear Gowdy cosmological model.

The Mexico City tests have now been applied to a number of different numerical-relativity codes that use different formulations of the Einstein equations [7, 17]. But all of the codes tested so far use finite-difference numerical methods. In this paper we report the results of applying these tests to the code developed in collaboration between the Caltech and Cornell numerical-relativity groups. We use a first-order symmetric-hyperbolic formulation of the equations developed by Kidder, Scheel, and Teukolsky [181] (sometimes referred to as the KST formulation) and we solve the equations using pseudospectral numerical methods. The results reported here differ therefore from all previously tested cases both in the formulation of the Einstein equations and the numerical methods used to solve them.

In Sec. 2.2 we review the KST formulation of the equations, and the pseudospectral numerical methods we use to solve them. The remaining sections present the results of the various Mexico City tests, adapted somewhat to provide more challenging tests of a code based on spectral methods. In Sec. 2.3 we show that our code is stable when evolving small random perturbations of flat spacetime. In Sec. 2.4 we report the results of the small-amplitude plane-wave test. We demonstrate the convergence rates for different spatial resolutions and different time-step algorithms. We also derive an equation for the error introduced by finite machine precision, and show that it limits the convergence of our evolutions for small spatial and temporal resolutions. In Sec. 2.5 we investigate

the stability of our evolution code for nonlinear gauge waves. In this case, nonlinear terms give rise to an instability that is drastically reduced by suitably filtering the components of the spectral expansion. Section 2.6 shows the performance of our code for evolutions of the highly dynamical Gowdy spacetime, in which the exact analytical expressions for the components of the fields grow exponentially in time. Finally, we discuss and summarize our various results in Sec. 2.7.

2.2 Solution method

In this section we describe the formulation of the Einstein equations and the pseudospectral numerical solution method that we test. The Mexico City tests were designed with finite-difference methods in mind and were originally applied to formulations of the Einstein equations that are second-order in space and first-order in time. Both our numerical methods and our representation of the Einstein equations differ significantly from those in Ref. [7], so appropriate modifications to the Mexico City test suite (for example, the number of grid points used or the constraint quantities observed) are needed. These modifications are also described in this section.

2.2.1 KST formulation

The KST system [181] is a first-order symmetric hyperbolic generalization of York's representation of the ADM equations [252]. The dynamical variables of this system are the three-metric g_{ij} , the extrinsic curvature K_{ij} , and a new variable D_{kij} that is initially set equal to $\partial_k g_{ij}/2$. This last variable allows the system to

be put into first-order form. Its introduction results in two additional constraints:

$$C_{kij} := D_{kij} - \frac{1}{2} \partial_k g_{ij} , \quad (2.1)$$

$$C_{lkij} := \partial_{[l} D_{k]ij} . \quad (2.2)$$

The KST evolution equations are obtained from the ADM equations [252] by adding constant multiples of the various constraints to the evolution equations and by replacing the lapse with a lapse-density function. These changes do not affect the physical solutions of the system, but they do modify the unphysical constraint-violating solutions. The added constraint terms are proportional to constant parameters $\{\gamma_1, \gamma_2, \gamma_3, \gamma_4\}$, which are chosen to make the system symmetric hyperbolic [180]. The principal parts of the KST evolution equations, then, are given by:

$$\partial_t g_{ij} \simeq N^n \partial_n g_{ij} ; \quad (2.3)$$

$$\begin{aligned} \partial_t K_{ij} \simeq & N^n \partial_n K_{ij} - N \left[(1 + 2\gamma_0) g^{cd} \delta^n_{(i} \delta^b_{j)} \right. \\ & - (1 + \gamma_2) g^{nd} \delta^b_{(i} \delta^c_{j)} - (1 - \gamma_2) g^{bc} \delta^n_{(i} \delta^d_{j)} \\ & \left. + g^{nb} \delta^c_i \delta^d_j + 2\gamma_1 g^{n[b} g^{d]c} g_{ij} \right] \partial_n D_{bcd} ; \end{aligned} \quad (2.4)$$

$$\begin{aligned} \partial_t D_{kij} \simeq & N^n \partial_n D_{kij} - N \left[\delta^n_k \delta^b_i \delta^c_j - \frac{1}{2} \gamma_3 g^{nb} g_{k(i} \delta^c_{j)} \right. \\ & - \frac{1}{2} \gamma_4 g^{nb} g_{ij} \delta^c_k + \frac{1}{2} \gamma_3 g^{bc} g_{k(i} \delta^n_{j)} \\ & \left. + \frac{1}{2} \gamma_4 g^{bc} g_{ij} \delta^n_k \right] \partial_n K_{bc} . \end{aligned} \quad (2.5)$$

Here, the symbol \simeq indicates that terms algebraic in the fields (that is, nonprincipal terms) are not shown explicitly. The lapse function N is taken to be

$$N := g^{\gamma_0} e^Q , \quad (2.6)$$

and both the lapse density function Q and the shift N^i are assumed to be specified functions of the coordinates, rather than independent dynamical fields. Since each of the Mexico City tests involves reproducing either a known analytic solution of the Einstein equations or a small perturbation about a known solution, for all tests reported here we set the lapse density Q and the shift N^i from the appropriate analytic solution. We choose one set of the KST parameters for all the tests here: $\gamma_0 = 0.5$; $\gamma_1 = -0.21232$; $\gamma_2 = -0.00787402$; $\gamma_3 = -1.61994$; $\gamma_4 = -0.69885$. These values were chosen because they make the KST system symmetric hyperbolic and coincide with a set preferred by Owen [204] in his extension of the KST system.

To evaluate errors it is useful to look at constraint quantities. As mentioned above, the KST system has additional constraints, Eqs. (2.1) and (2.2), besides the usual Hamiltonian constraint C and momentum constraint C_i . To ensure that we are satisfying all the constraints, we monitor a single quantity \mathcal{C} that is zero if and only if all of the constraints vanish:

$$\mathcal{C} := \sqrt{C^2 + (C_i)^2 + (C_{kij})^2 + (C_{lkij})^2} , \quad (2.7)$$

where an object is squared using the evolved spatial metric. For example, $(C_i)^2 = g^{ij}C_iC_j$.

Likewise, when evaluating differences from analytically known solutions, we define an overall error quantity that includes the errors in all evolved variables g_{ij} , K_{ij} , and D_{kij} . Taking $\delta g_{ij} := g_{ij}^{\text{analytic}} - g_{ij}^{\text{evolved}}$, and similarly for other fundamental fields, this overall error quantity is given by

$$\delta \mathcal{U} := \sqrt{(\delta g_{ij})^2 + (\delta K_{ij})^2 + (\delta D_{kij})^2} . \quad (2.8)$$

Note that $\delta \mathcal{U}$ vanishes if and only if all evolved variables match the known solution.

For all error quantities \mathcal{Q} we display L_2 norms:

$$\|\mathcal{Q}\|_2 := \sqrt{\frac{1}{\text{Vol}} \int \mathcal{Q}^2 \sqrt{|g|} d^3x}, \quad (2.9)$$

where $\text{Vol} = \int \sqrt{|g|} d^3x$ is the volume of the domain. These norms are computed after each time step over the current $t = \text{const.}$ hypersurface. We refer to $\|\mathcal{C}\|_2$ as the constraint energy, and $\|\delta \mathcal{U}\|_2$ as the error energy.

The error quantities $\|\delta \mathcal{U}\|_2$ and $\|\mathcal{C}\|_2$ scale with the absolute magnitude of the fundamental fields and their derivatives, so it can be difficult to judge the significance of these error measures without knowing the overall scale of the variables in the problem. For this reason, we sometimes plot the *normalized* error energy $\|\delta \mathcal{U}\|_2 / \|\mathcal{U}\|_2$ and the *normalized* constraint energy $\|\mathcal{C}\|_2 / \|\partial \mathcal{U}\|_2$, where the normalization factors are defined by

$$\mathcal{U} := \sqrt{(g_{ij})^2 + (K_{ij})^2 + (D_{kij})^2}, \quad (2.10)$$

$$\partial \mathcal{U} := \sqrt{(\partial_i g_{jk})^2 + (\partial_i K_{jk})^2 + (\partial_i D_{jkl})^2}. \quad (2.11)$$

Note that $\|\delta \mathcal{U}\|_2 / \|\mathcal{U}\|_2$ and $\|\mathcal{C}\|_2 / \|\partial \mathcal{U}\|_2$ become of order unity when errors and constraint violations completely dominate the numerical solution. We display normalized error quantities only for tests involving the Gowdy spacetimes (Sec. 2.6), in which the fundamental variables vary exponentially in time. All other tests presented here involve perturbations of Minkowski spacetime, in which case the quantity $\|\partial \mathcal{U}\|_2$ is of order the size of the perturbation and is therefore inappropriate to use as a normalization factor. However, for perturba-

tions of Minkowski spacetime, the overall scale is of order unity so it suffices to display the unnormalized quantities $\|\delta \mathcal{U}\|_2$ and $\|\mathcal{C}\|_2$.

2.2.2 Pseudospectral methods

All of our numerical computations are carried out using pseudospectral methods; this is the first time the Mexico City tests have been applied to a pseudospectral code. A brief outline of our method is as follows. Given a system of partial differential equations

$$\partial_t u(\mathbf{x}, t) = \mathcal{F}[u(\mathbf{x}, t), \partial_i u(\mathbf{x}, t)] , \quad (2.12)$$

where u is a collection of dynamical fields, the solution $u(\mathbf{x}, t)$ is expressed as a time-dependent linear combination of N spatial basis functions $\phi_k(\mathbf{x})$:

$$u(\mathbf{x}, t) = \sum_{k=0}^{N-1} \tilde{u}_k(t) \phi_k(\mathbf{x}) . \quad (2.13)$$

Associated with the basis functions is a set of N_c collocation points \mathbf{x}_i . Given spectral coefficients $\tilde{u}_k(t)$, the function values at the collocation points $u(\mathbf{x}_i, t)$ are computed using Eq. (2.13). Conversely, the spectral coefficients are obtained by the inverse transform

$$\tilde{u}_k(t) = \sum_{i=0}^{N_c-1} w_i u(\mathbf{x}_i, t) \phi_k(\mathbf{x}_i) , \quad (2.14)$$

where w_i are weights specific to the choice of basis functions and collocation points. Thus it is straightforward to transform between the spectral coefficients $\tilde{u}_k(t)$ and the function values at the collocation points $u(\mathbf{x}_i, t)$.

To solve the differential equations, we evaluate spatial derivatives analytically using the known derivatives of the basis functions,

$$\partial_i u(\mathbf{x}, t) = \sum_{k=0}^{N-1} \tilde{u}_k(t) \partial_i \phi_k(\mathbf{x}) , \quad (2.15)$$

and we evaluate nonlinear terms using the values of $u(\mathbf{x}_i, t)$ at the collocation points. Thus we can write the partial differential equation, Eq. (2.12), as a set of *ordinary* differential equations for the function values at the collocation points,

$$\partial_t u(\mathbf{x}_i, t) = \mathcal{G}_i[u(\mathbf{x}_j, t)] , \quad (2.16)$$

where \mathcal{G}_i depends on $u(\mathbf{x}_j, t)$ for all j . We then integrate this system of ordinary differential equations in time, using (for example) a fourth-order Runge-Kutta algorithm.

Because the tests discussed here are periodic in all spatial dimensions, we use Fourier basis functions. If we choose a computational domain extending from $-1/2$ to $1/2$ in each of the x , y , and z directions, then each variable u is decomposed as

$$u(x, y, z) = \sum_{k=0}^{N_x-1} \sum_{l=0}^{N_y-1} \sum_{m=0}^{N_z-1} a_{klm} \phi_k(x) \phi_l(y) \phi_m(z), \quad (2.17)$$

where

$$\phi_k(x) = \begin{cases} 1 & k = 0 ; \\ \sin[\pi x(k+1)] & k > 0 \text{ (} k \text{ odd)} ; \\ \cos(\pi x k) & k > 0 \text{ (} k \text{ even)} . \end{cases} \quad (2.18)$$

For smooth solutions, the spectral approximation Eq. (2.13) converges exponentially (error $\sim e^{-\lambda N}$ for some $\lambda > 0$ which depends on the solution). This is much faster than the polynomial convergence (error $\sim 1/N^p$) obtained using

p th-order finite differencing. As a result, we run our tests at coarser resolutions than those recommended in Ref. [7] for finite-difference codes—typically we use $N_i = 9, 15, 21, 27,$ and 33 collocation points in the relevant directions. From Eqs. (2.17) and (2.18) we see that if we choose $N_x, N_y,$ or N_z to be an even integer, the highest-frequency component in our expansion will have a sine term but no matching cosine term. Consequently, the spatial derivative of this highest-frequency component will not be represented by our basis functions, causing a numerical instability. Therefore we choose $N_x, N_y,$ and N_z to be odd.

Because spectral methods so greatly reduce spatial-discretization errors, time-stepping error is often dominant. In order to make the time stepping and the spatial-discretization errors comparable in these tests, we use fourth-order Runge-Kutta ODE integration. The time-step sizes are chosen in an effort to use step sizes comparable to those used to test finite-difference methods in Ref. [7], while also ensuring that time-step errors do not dominate over our spatial-truncation errors. We use $\Delta t = \Delta x/20$ in the first test, and $\Delta t = \Delta x/40$ in all others, except where explicitly noted. Here, Δx is the minimum distance between collocation points.

2.3 Random initial data on flat space

Perhaps the simplest test of a numerical-relativity code is evolving standard Minkowski spacetime on a three-torus, T^3 . However, this test is *too* simple because all fundamental fields are spatially constant and most are identically zero, and hence most numerical methods will reproduce the correct solution exactly. This test can be made more discriminating by adding a small amount of random

noise to the initial data; the noise is intended to simulate the effect of finite numerical precision. A different random number is added to each component of each evolved variable, at each point in the domain. These random numbers are chosen to lie between -10^{-10} and 10^{-10} so that the system remains in the linear regime. If these small perturbations to a simple spacetime grow unstably, it is likely that the inevitable errors (e.g., discretization error or even numerical-roundoff error) that arise in any more complicated simulation will also grow unstably. For this test we vary the resolution in the x dimension, and we fix the resolution to three collocation points in each of the y and z dimensions.

If the perturbations in the fields are chosen to be of size ϵ , independent of resolution, then the perturbation in the n th spatial derivatives of these fields will be $\sim \epsilon(\Delta x)^{-n}$, where Δx is some measure of the distance between neighboring points. This means that error quantities involving derivatives (such as constraints) will be *larger* for finer resolutions.² This behavior is seen in the plot of the constraint energy in Fig. 2.1.

The purpose of this test is to establish that small constraint violations around flat space do not grow, and the KST system clearly passes this test. Whether or not constraint violations decay will depend on the evolution system and the numerical method. For example, artificial dissipation in the numerical method might cause all variations to decay, including constraint violations. Furthermore,

²The Mexico City collaboration [7] intended their Hamiltonian-constraint errors to be independent of resolution, so they chose the size of the perturbation ϵ to be *dependent* on resolution, $\epsilon \sim (\Delta x)^2$, which is the appropriate scaling for the second-order-in-space formulations of Einstein's equations they use. However, for the first-order-in-space formulation we use, the Hamiltonian constraint is computed using first derivatives of D_{kij} rather than second derivatives of g_{ij} , so the constraint will vary as $(\Delta x)^{-1}$. Note also that the $\epsilon \sim (\Delta x)^2$ scaling does not make the momentum constraint independent of resolution, as it depends on first derivatives of the fields. We simply choose ϵ to be independent of resolution.




Figure 2.1: Constraints for Minkowski space with random noise

Higher resolutions are expected to have larger constraints because more closely spaced points result in larger derivatives. The constraints do not grow in time.

if the evolution system contains constraint damping in some form, then the constraints should decay. Indeed, Owen has extended the KST system to include constraint damping [204]; running the same test, he observes exponential decay in the constraint quantities. The flat constraint violations observed in Fig. 2.1 indicate that the KST system with our parameter choice does not damp constraints and that the spectral method has insignificant artificial dissipation.

In Fig. 2.2 we see a linear growth of the error energy $\|\delta\mathcal{U}\|$ for this test. We find that the growth is caused solely by contributions from the metric g_{ij} ; the average values of K_{ij} and D_{kij} remain constant in time. We can understand this as follows. The average value of K_{ij} is determined by the random initial data and will in general be nonzero. The time derivative of K_{ij} , to first order in the amplitude of perturbations around flat space, involves only spatial derivatives of D_{kij} . (See Eq. (2.4).) These derivatives have zero average (up to roundoff errors $\sim 10^{-16}$), because the constant term in the Fourier expansion Eq. (2.18) is removed by differentiation, and therefore the average of K_{ij} will be constant in time. The time derivative of g_{ij} involves a term proportional to K_{ij} . Because the average of K_{ij} is constant in time and nonzero, the value of g_{ij} will therefore drift linearly in time. The average of K_{ij} is smaller for higher resolutions—because the average is taken over more random numbers—which means that the growth rate of g_{ij} should decrease with increasing resolution. Indeed, this is what we observe in Fig. 2.2.

We can verify that the nonzero average of K_{ij} is the only source of growth in g_{ij} by manually removing the average value of K_{ij} . We expect this will leave the norms of the components of g_{ij} approximately constant in time. This is

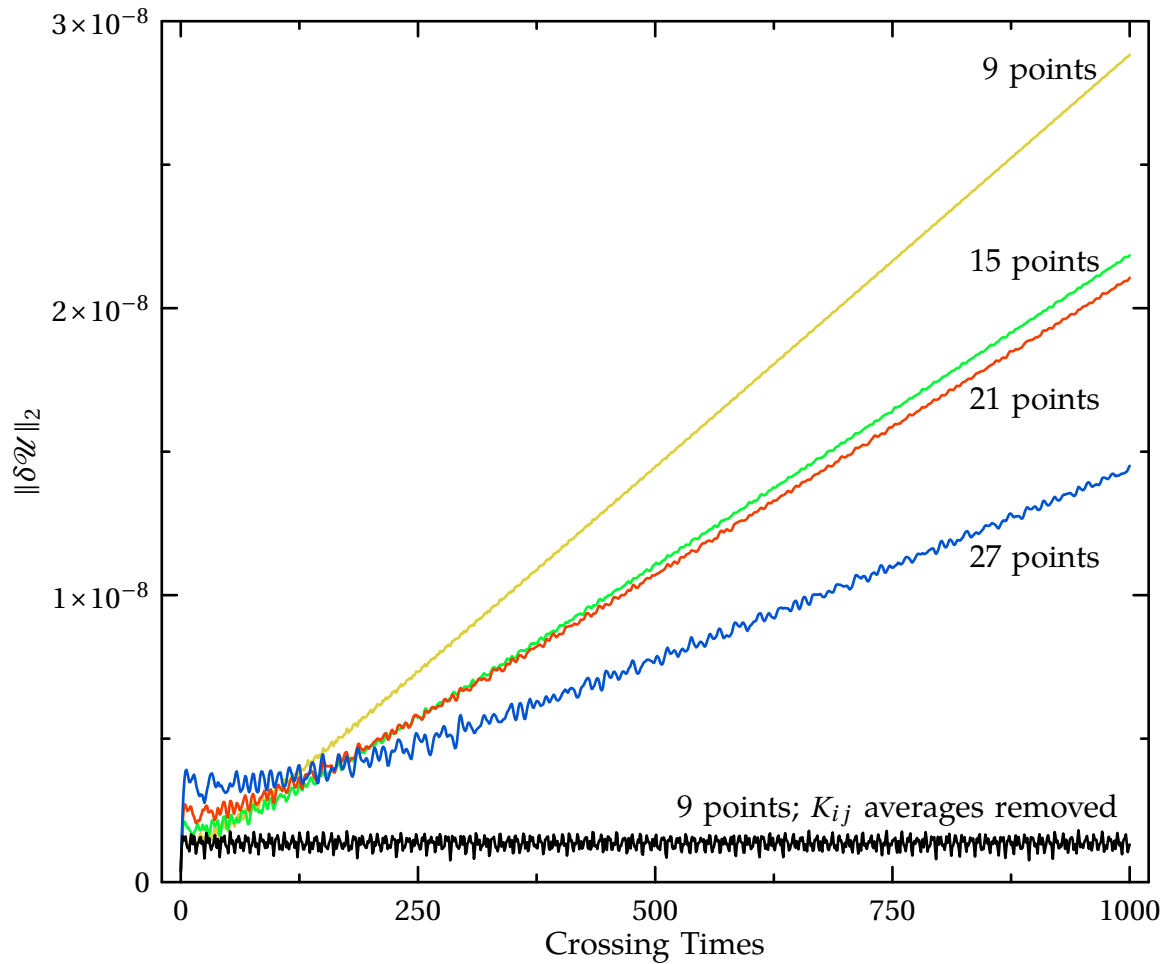


Figure 2.2: Error energy for Minkowski space with random noise

The linear increase in time is due to a nonzero average in the random noise added to K_{ij} . This average approaches zero as resolution is increased, since there are more points over which to average. The flat line shows the evolution when the average value of K_{ij} is set to zero in the initial data.

accomplished by setting the $k = 0$ spectral coefficients of all components of K_{ij} to zero in the initial data, after all the random numbers have been added. The flat line in Fig. 2.2 shows the result, indicating that the average offset in K_{ij} is the only source of growth in the evolved variables of the KST system for this test.

2.4 Linear plane wave

If the ultimate goal of simulating binary black hole mergers is to predict the gravitational-radiation waveforms for observations, an evolution system must at least be capable of propagating a simple linear plane wave through flat spacetime. The form suggested for the Mexico City tests in Ref. [7] is

$$ds^2 = -dt^2 + dx^2 + (1 + b)dy^2 + (1 - b)dz^2 , \quad (2.19)$$

where

$$b = b(x, t) = A \sin [2\pi(x - t)] . \quad (2.20)$$

This metric satisfies Einstein's equations only to linear order in the wave's amplitude A , so if the fully nonlinear numerical solution is compared to this approximate solution, there will be deviations of order A^2 that arise from our choice of "analytic solution" rather than from numerical errors. The amplitude A for the Mexico City tests is chosen to be 10^{-8} so that such deviations in the metric components g_{ij} are below machine precision. However, we still observe $\mathcal{O}(A^2)$ deviations in the variables K_{ij} and D_{kij} (which have values of order A), even with an amplitude of $A = 10^{-8}$, because the *relative* error is well above machine precision.

2.4.1 One-dimensional sinusoid

The sinusoidal waveform chosen in Eq. (2.20) is only a weak test for pseudospectral methods, because the Fourier basis functions defined in Eqs. (2.17) and (2.18) exactly resolve Eq. (2.20) at all times using only three basis functions; the only truncation errors are those associated with time discretization. Therefore, as a more challenging test, in Sec. 2.4.2 we repeat the plane wave evolution using a Gaussian-shaped wave. It is nevertheless instructive to evolve the sinusoid and study the resulting time-discretization errors. Since the dynamics involve no change in amplitude, but a change in phase, we expect the errors to be primarily phase errors, for reasonably small time steps.

This loss of temporal accuracy is particularly relevant in efforts to simulate sources for gravitational-wave observations, as the search for signals involves matching expected waveforms against observations. If there is significant error in the phase of the expected waveform, the overlap will be poor and detection will be more difficult. Although a constant overall scaling error in frequency—like the one found in this linear problem—could still result in detection, more complex situations would likely give rise to more complicated errors. The straightforward way to handle this problem is to minimize all time-stepping error.

In Fig. 2.3 we show the convergence of the phase error in the evolution of the sinusoidal linear wave. The solution is fully resolved on a $3 \times 1 \times 1$ grid. We keep this grid fixed, and decrease the size of the time step. Assuming that the only error is some phase error $\delta\phi$, the evolved g_{zz} will be given by

$$g_{zz} = 1 - 10^{-8} \sin [2\pi(x - t) + \delta\phi] . \quad (2.21)$$




Figure 2.3: Phase error for 1-D sinusoidal linear wave

Phase error at $t = 25$ crossing times for various time-step sizes, and several time-stepping algorithms. These tests were all run with three points in the x direction. The dashed line indicates the expected accuracy limit due to roundoff error. See Eq. (2.26).

At integer multiples of the light-crossing time for our computational domain, this can be written as

$$g_{zz} = 1 - A [\cos \delta \phi \sin (2\pi x) + \sin \delta \phi \cos (2\pi x)] . \quad (2.22)$$

That is, we can find the phase error easily from the $k = 1$ sine and cosine components of g_{zz} (which happen to be easily accessible quantities in our code).

For intermediate time-step sizes, we can see convergence toward zero phase error with decreasing time step. As expected, we observe second-order convergence for Iterated Crank-Nicholson stepping, and fourth- and sixth-order for the appropriate higher-order Runge-Kutta algorithms. At very small time-step sizes, a new effect is seen, causing the phase error to increase with decreasing time step. This effect can be understood as machine-roundoff error accumulating via a random walk process.

Suppose we have a variable $\mathcal{V}(t)$ that is evolved by adding the small changes needed to update its value at each time step. Each such operation will introduce a fractional error $\chi(t)$ caused by the finite machine precision. We assume that the standard time-step size is Δt , and that there are n intermediate operations in each time step. After an evolution through time T , the total error added in this way will be

$$\delta \mathcal{V} = \sum_{j=0}^{nT/\Delta t} \mathcal{V}(t_j) \chi(t_j) . \quad (2.23)$$

To avoid tracking each individual error contribution, we treat χ as a random variable taking values in some range, with some probability distribution.

We estimate the accumulated error due to finite machine precision by taking suitable averages over an ensemble of random $\chi(t)$ and over a time interval T . Assuming there is no asymmetry between positive and negative values of $\chi(t)$,

this accumulated error would be zero. Of course, we expect almost never to see this case: the most likely outcome is an accumulated error comparable to the dispersion:

$$|\delta\mathcal{V}| \sim \sqrt{(\delta\mathcal{V})^2} \sim \sqrt{\sum_{j,k} \overline{\mathcal{V}(t_j)\mathcal{V}(t_k)\chi(t_j)\chi(t_k)}}, \quad (2.24)$$

where the overbar indicates the average over the ensemble of random errors $\chi(t)$. We can simplify this expression by assuming that $\chi(t)$ has no correlations between time steps, and further assuming that the probability distribution is constant in time and uniform, taking values in the range $[-\epsilon, \epsilon]$, where ϵ is the machine precision. This means that $\overline{\chi(t_j)\chi(t_k)} = \delta_{jk}\epsilon^2/3$. Finally, we approximate the discrete time sum as an integral, and obtain

$$|\delta\mathcal{V}| \sim \epsilon \sqrt{\frac{n}{3\Delta t} \int_{t_1}^{t_2} \mathcal{V}(t)^2 dt}. \quad (2.25)$$

We can test this formula by observing its effects in the case of phase error for the linear wave. Here, the only nontrivial evolved variable is $\mathcal{V} = g_{zz}$, which is very nearly 1; so the integral in Eq. (2.25) becomes simply the evolution time T , which has the value 25 for the results plotted in Fig. 2.3. If phase errors dominate, $\delta g_{xx} \sim A \sin \delta\phi$, so we have

$$|\delta\phi| \sim \frac{\epsilon}{A} \sqrt{\frac{25n}{3\Delta t}} \sim \frac{10^{-7}}{\sqrt{\Delta t}} \left(\frac{\epsilon}{10^{-16}} \right) \left(\frac{10^{-8}}{A} \right), \quad (2.26)$$

where n is assumed to be of order 10. This expression is plotted as the dashed line in Fig. 2.3, demonstrating that the addition of machine-precision errors causes the departure from the standard second-, fourth-, and sixth-order convergence we observed. From Eq. (2.26) we see that $|\delta\phi|$ is proportional to the ratio ϵ/A ; thus $|\delta\phi|$ is so large in this case because the wave amplitude is so small, $A = 10^{-8}$.

The phase error is only so clearly visible in these evolutions because the full solution is described precisely at each moment by the first three basis functions. This means that discretization error due to spatial differentiation is essentially at the level of machine precision. Indeed, using more than three points actually degrades the quality of these one-dimensional sinusoid evolutions. Power in higher-order basis functions can only be error, and hence will necessarily do worse than the low-resolution case. We omit plots of the error energy and constraints in the higher-resolution cases, as they are very nearly the same as those of the more complicated two-dimensional evolutions discussed in Sec. 2.4.3.

2.4.2 One-dimensional Gaussian

As a more challenging test of pseudospectral methods, we repeat the one-dimensional linear wave test using a periodic Gaussian-shaped wave:

$$b = A \sum_{j=-\infty}^{\infty} \exp \left[-\frac{(x-t+j)^2}{2w^2} \right], \quad (2.27)$$

with $A = 10^{-8}$. The summation over j ensures that the function is truly periodic at all times. In practice, j need only range over a few, depending on the width of the Gaussian. The width chosen here is $w = 0.05$ to ensure that features are not too sharp, while still presenting a nontrivial challenge to spectral differentiation.

We find behavior comparable to the sinusoidal case, although as expected, more collocation points in the x direction are needed to resolve the solution spatially (but we still use only a single point in each of the y and z directions). Note the exponential convergence of the constraints with increasing resolution

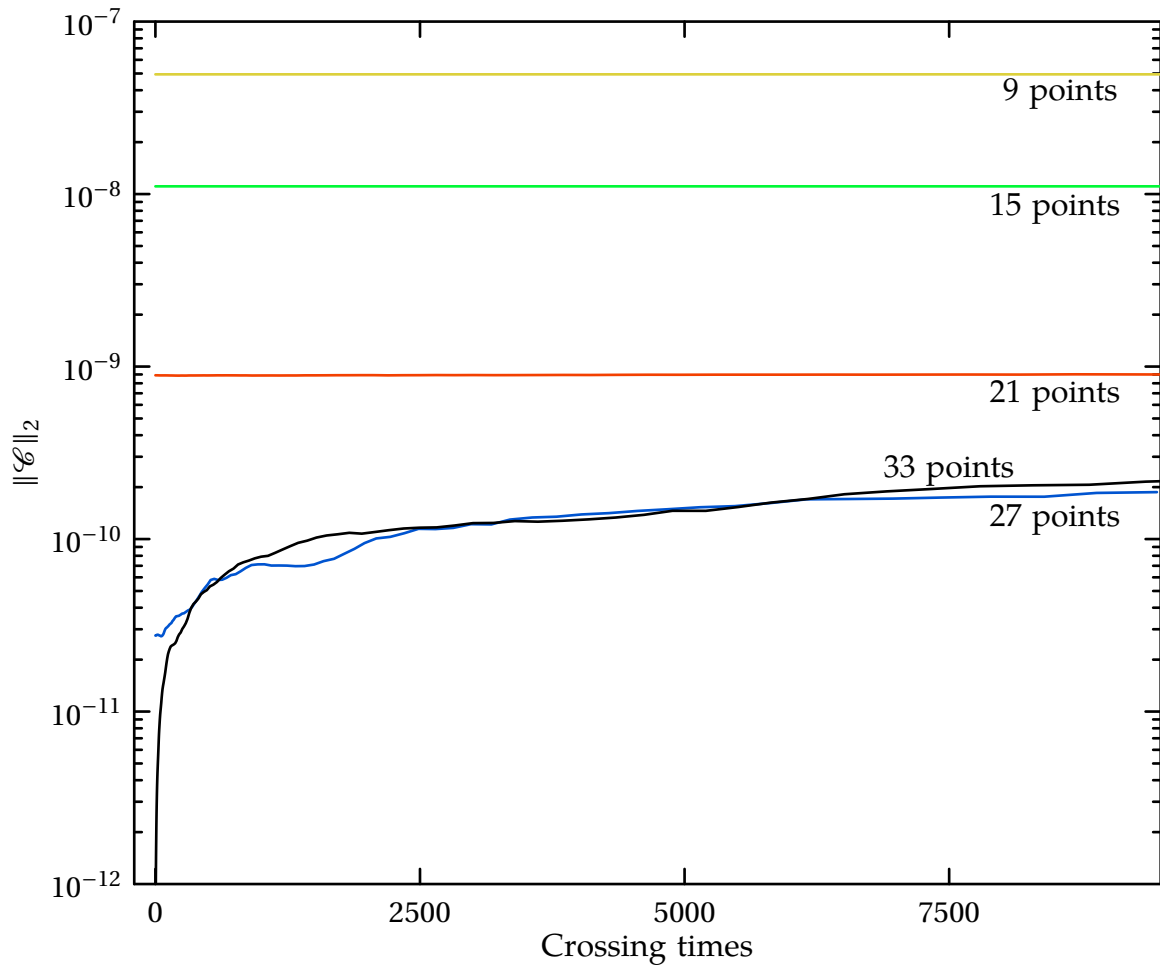


Figure 2.4: **Constraints for 1-D Gaussian linear wave**

Here we see the exponential convergence of the constraints with higher spatial resolution. At late times, the higher resolutions grow sublinearly in time, probably because of accumulated machine-roundoff error.




Figure 2.5: Error energy for 1-D Gaussian linear wave

The solid lines show the error energy at 1/2 crossing times, with clearly visible exponential convergence at early times. The dashed lines show the error energy at integer crossing times for the same resolutions. The smallness of the error energy at early times demonstrates the low dispersion of the numerical method; at later times, the error is dominated by the quadratic growth. Both effects are discussed in the text.

in Fig. 2.4. The constraint growth in the highest-resolution runs is slower than linear in time, and is probably caused by the accumulation of errors due to finite machine precision as discussed in Sec. 2.4.1.

Figure 2.5 presents the error energy for this run as the solid lines. At early times $\|\delta\mathcal{U}\|$ decreases exponentially to zero with increasing resolution, as one would expect. At late times, however, $\|\delta\mathcal{U}\|$ converges toward a parabola. The amplitude of this parabola scales in proportion to A^2 . In the rest of this section, we will first explain a subtlety arising when computing $\|\delta\mathcal{U}\|$, followed by a detailed explanation of why the terms $\mathcal{O}(A^2)$ manifest themselves in parabolic behavior of $\|\delta\mathcal{U}\|$.

The comparison of the computed solution with the analytic solution is performed at the collocation points. By virtue of the transformation Eqs. (2.13) and (2.14), the errors are initially exactly zero at the collocation points. The spatial-truncation error is nonzero of course, even at the initial time; it manifests itself as a deviation of the truncated series expansion from the analytic solution *between* collocation points. During the evolution, a linear wave will simply travel through the computational domain, returning to the original position after each light-crossing time. Since the spectral method has very small dispersion, the evolved shape remains the same. After each light-crossing time, therefore, the evolved solution again agrees to very high accuracy with the initial analytic solution *at* the collocation points. So, comparing the evolution with the analytic solution at integer multiples of the light-crossing time and at the collocation points will yield differences much smaller than spatial truncation error.³ Therefore, a fair compar-

³This is also true when comparing at intervals of $1/N$ of a crossing time if the number of collocation points in the direction of the wave's travel is divisible by N .

ison that includes the effects of spatial-truncation error must not be performed at integer light-crossing times. These considerations are evident from Fig. 2.5, where the solid lines show the “true” $\|\delta\mathcal{U}\|$ observed with 1/2 light-crossing interval offset, which suffices because the number of collocation points is always odd. The artificially small error energy observed at every complete light-crossing interval is shown as dashed lines, confirming the excellent low-dispersion property of our method.

At late times the differences between observation at full and 1/2 crossing times are swamped by the parabolic growth in $\|\delta\mathcal{U}\|$. Similar parabolic deviations of the evolution from the solution of the linearized equations are observed for the other two linear-wave evolutions, the 1-D and 2-D sinusoids (see Fig. 2.7). The growth in $\|\delta\mathcal{U}\|$ appears almost entirely due to growth in the $k = 0$ mode of diagonal terms in δg_{ij} . Using evolutions of waves with different amplitudes and wavelengths, we have verified that this growth is proportional to $A^2 t^2 / \lambda^2$, where A is the amplitude and λ the wavelength of the disturbance. The constant of proportionality depends directly on the KST parameter γ_1 appearing in Eq. (2.4). This parameter controls the addition of a term $\gamma_1 N g_{ij} C$ to the ADM evolution equation for K_{ij} . The Hamiltonian constraint C is roughly constant in time, and varies as A^2 / λ^2 . Since the $k = 0$ mode of $\gamma_1 N g_{ij} C$ is roughly $\gamma_1 \delta_{ij} C$, the $k = 0$ mode of K_{ii} will grow linearly with time in proportion to $\gamma_1 A^2 / \lambda^2$ for each i . That, in turn, will cause small quadratic growth in the $k = 0$ mode of g_{ii} . For the more well-behaved cases (highest resolutions for the Gaussians; all cases for the sinusoids) this model is an excellent fit for the observed error energy.

2.4.3 Two-dimensional linear waves

The linear wave tests above may be modified by rotating the coordinates by $\pi/4$ about the z axis, which gives a plane wave propagating along the x - y diagonal. By increasing the size of the domain by a factor of $\sqrt{2}$ in each direction, the rotated solution can be made periodic while maintaining the same wavelength. This converts the spacetime from essentially one-dimensional to essentially two-dimensional. The purpose of this test is to ensure that the symmetry of the one-dimensional version does not hide sources of error (although propagation along a diagonal obviously retains some symmetries). For these tests we use a single collocation point in the z dimension, and we vary the (equal) number of collocation points in the x and y directions. We run these tests to $t = 1000$ —ten times longer than is recommended by the Apples with Apples collaboration—to better observe the stability properties. As shown in Fig. 2.6, the constraints for the sinusoidal wave increase with increasing x and y resolution (still using only a single point in the z direction). The constraints for the Gaussian are very nearly the same as in the one-dimensional case. Again, the $A^2 t^2 / \lambda^2$ growth of $\|\delta \mathcal{U}\|$ is visible, as shown in Fig. 2.7.

2.5 Gauge wave

The next series of tests involves a simple but time-dependent gauge transformation of Minkowski space, in the form of a plane wave. The metric used for this




Figure 2.6: Constraints for 2-D linear waves

The solid lines represent the Gaussian wave, while the dashed lines represent the sinusoidal wave. The sinusoid is fully resolved spatially with 3 points. Going to higher resolutions merely introduces spatial errors in the unnecessary basis functions, which leads to an increase in the constraints with resolution.

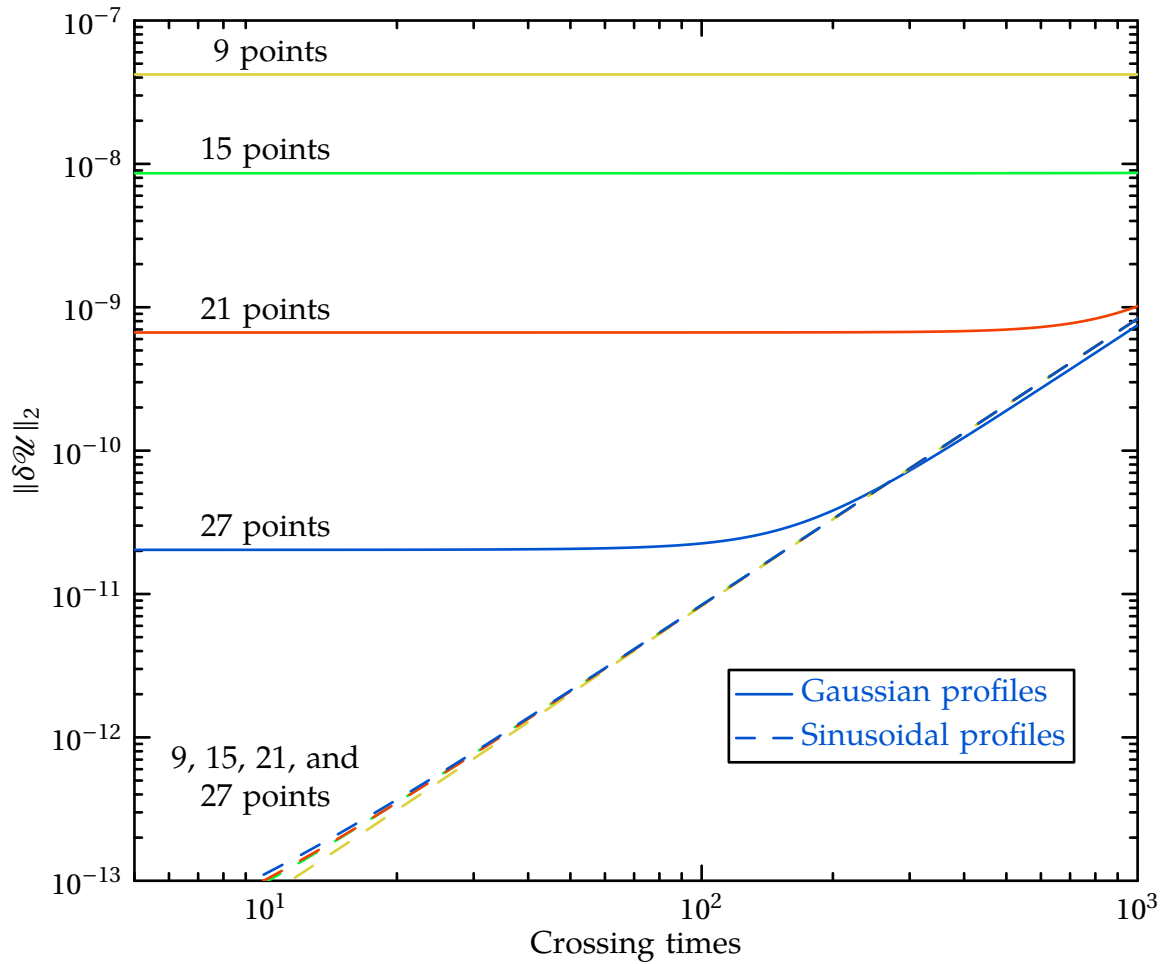


Figure 2.7: Error energy for 2-D linear waves

The solid lines represent the Gaussian wave, while dashed lines represent the sinusoidal wave, both observed at $1/2$ crossing times. As in the 1-D case (Fig. 2.5), both sets of evolutions converge to quadratic growth of the error caused by the Hamiltonian constraint, explained in the text.

Mexico City test has the form

$$ds^2 = -(1 + a)dt^2 + (1 + a)dx^2 + dy^2 + dz^2 , \quad (2.28)$$

$$a = A \sin [2\pi(x - t)] . \quad (2.29)$$

Two cases are considered: a low-amplitude case with $A = 0.01$, and a high amplitude case with $A = 0.1$. This is the first test for which the nonlinear terms in the equations play an important role.

For the linear plane wave test in Sec. 2.4, we found that because we use a Fourier basis, we were able to fully resolve the sinusoidal waveform using only three collocation points. This is not true for the gauge-wave test, because in this case the extrinsic curvature (one of our evolved variables) is not a simple sinusoid. Instead, its only nonzero component is

$$K_{xx} = -\pi \frac{A \cos [2\pi(x - t)]}{\sqrt{1 + A \sin [2\pi(x - t)]}} . \quad (2.30)$$

2.5.1 One-dimensional gauge wave

We ran the one-dimensional test described above using a single collocation point in each of the y and z directions, and varying the resolution in the x direction. We find that for both $A = 0.01$ and $A = 0.1$ our evolution is stable and convergent. Our error energy and constraint violations show no signs of instability and are strictly better than the filtered two-dimensional evolutions discussed below. We omit plots for this test because the two-dimensional test is more challenging and more discriminating.

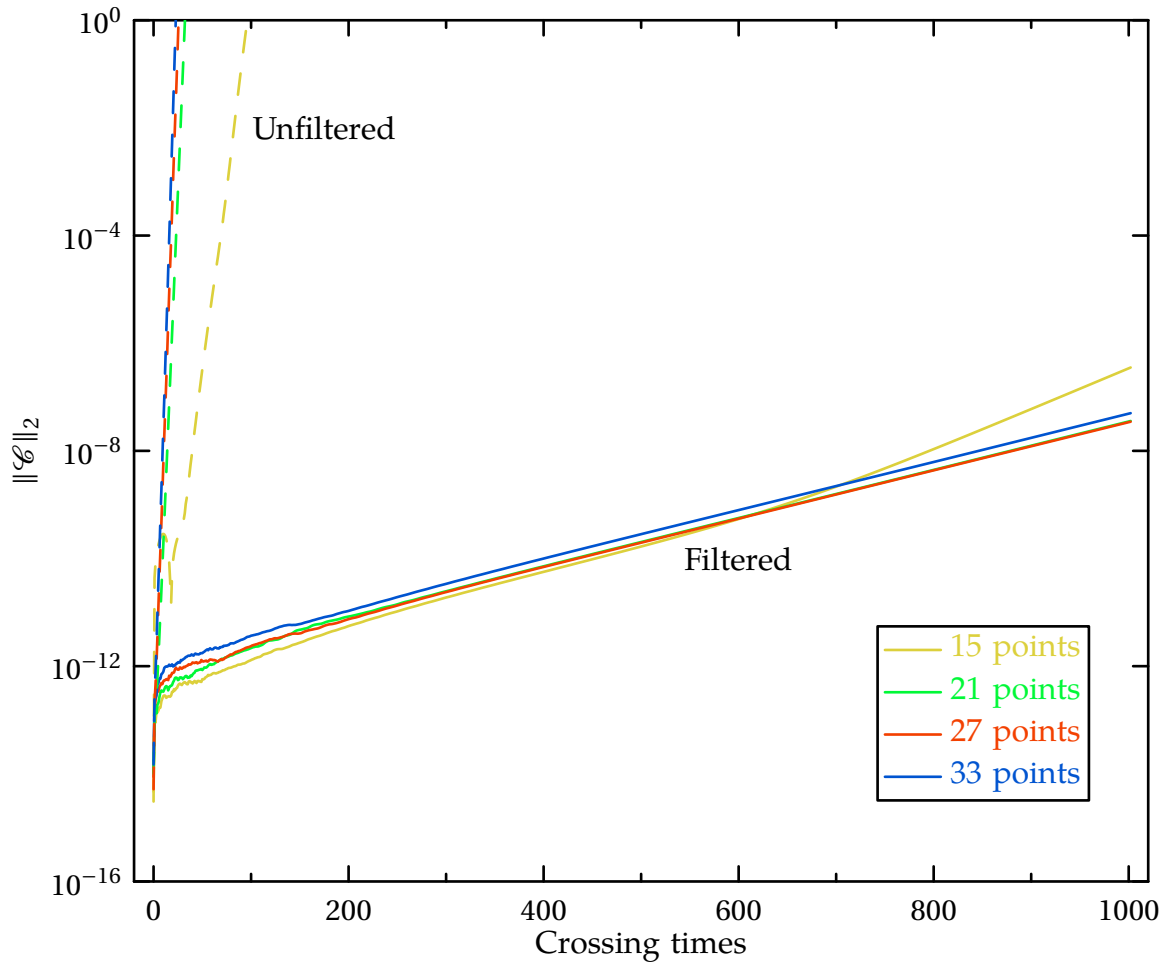


Figure 2.8: Constraints for high-amplitude 2-D gauge wave

Dashed lines indicate the unfiltered behavior; solid lines indicate the filtered behavior. Note that, despite an effective loss of resolution, filtering greatly improves the stability of the evolution.




Figure 2.9: Error energy for high-amplitude 2-D gauge wave

As in Fig. 2.8, dashed lines are unfiltered, and solid lines are filtered. The growth of the filtered error energy is exponential in time. For the highest resolution the time step was cut in half ($\Delta t = \Delta x/80$) to reduce time-discretization error to the same level as spatial-discretization error. The dotted line shows the same evolution with time step $\Delta t = \Delta x/40$, which is dominated by time-discretization error.

2.5.2 Two-dimensional gauge wave

A simple rotation of coordinates about the z axis makes the wave described by Eqs. (2.28) and (2.29) propagate along the x - y diagonal, as in the case of the linear wave. We use an equal number of collocation points in the x and y directions, and a single collocation point in z .

As for the one-dimensional gauge-wave test, we run at two different amplitudes: $A = 0.01$ and $A = 0.1$. For low amplitude, $A = 0.01$, our evolution of the 2-D gauge wave is stable and convergent. Again, we omit plots, as our results are strictly better than for the more interesting high-amplitude case.

For high amplitude, $A = 0.1$, we find an exponentially growing nonconvergent numerical instability, as seen in the curves labeled “unfiltered” in Figs. 2.8 and 2.9. This instability does not appear for the low-amplitude case, nor does it appear for either amplitude in the one-dimensional gauge-wave test.

The instability appears to be associated with aliasing caused by quadratic nonlinearities in the evolution equations; this is a well-known phenomenon that often occurs when applying spectral methods to nonlinear equations [58]. The basic mechanism for the instability can be understood by considering a truncated spectral expansion for some variable $u(x)$ in terms of N basis functions $\phi_k(x)$:

$$u(x) = \sum_{k=0}^{N-1} u_k \phi_k(x) . \quad (2.31)$$

The correct spectral expansion of the expression $u(x)^2$ can be obtained by squaring Eq. (2.31); for most basis functions—including the Fourier series of Eq. (2.18)—this yields a sum over a total of $2N$ basis functions, rather than just N . Of course, we keep only N basis functions (not $2N$) in our expansion, so the $k \geq N$ coefficients of the product must be eliminated. Ideally, these $k \geq N$ coefficients should

be simply discarded and the $k < N$ coefficients should remain untouched. But it turns out that for the pseudospectral method of evaluating nonlinear terms (i.e., Fourier transform to obtain values at spatial collocation points, square these values, then Fourier transform back to spectral space), the power in the extra $k \geq N$ coefficients of the product does not disappear, but instead appears as excess power in some of the $k < N$ coefficients (“aliasing”). Repeating this process on each time step builds up this excess power and produces the instability.

Fortunately, there is a well-known remedy for instabilities caused by aliasing in nonlinear terms: suppose that the upper half (i.e., those with $k \geq N/2$) of the coefficients u_k in Eq. (2.31) were all zero. Then the spectral expansion of $u(x)^2$ will have zeroes in all its $k \geq N$ coefficients, so there is no aliasing, and hence no instability. Therefore, we ensure that all coefficients with $k \geq k_{\text{cut}}$ are zero by removing those coefficients from the initial data and from the right-hand side of the evolution equations. It turns out (see, for example, Chapter 11.5 of Ref. [58]) that for a quadratic nonlinearity, it is sufficient to filter with $k_{\text{cut}} = 2N/3$ (and not $k_{\text{cut}} = N/2$) to eliminate aliasing. As mentioned in Sec. 2.2, the remaining number of nonzero coefficients must be odd, which is ensured by reducing k_{cut} by one if necessary.

The price we pay for stability via this filtering is that we must use a factor of 1.5 more spectral coefficients (and collocation points) than without filtering in order to achieve the same level of spatial-discretization error. Hence, we use more points for this test than for the previous ones: $N_i = 15, 21, 27,$ and 33 points. This leaves the effective resolutions at $\tilde{N}_i = 9, 13, 17,$ and 21 points, which are comparable to the resolutions we use in the unfiltered case. We see

from Figs. 2.8 and 2.9 that filtering dramatically reduces the instability. The initial constraint violations in these runs, $\|\mathcal{C}\| \approx 10^{-12}$, are at the level of the finite machine precision, so increasing the resolution causes *increased*—not decreased—constraint violations.

2.5.3 Shifted gauge wave

We also show the results of a new “shifted gauge wave” test suggested for addition to the “Apples with Apples” suite [17]. For this test we evolve flat space with the usual Minkowski coordinates $(\hat{t}, \hat{x}, \hat{y}, \hat{z})$ transformed to coordinates (t, x, y, z) via

$$\hat{t} = t - \frac{A}{4\pi} \cos[2\pi(x - t)] , \quad (2.32)$$

$$\hat{x} = x - \frac{A}{4\pi} \cos[2\pi(x - t)] , \quad (2.33)$$

$$\hat{y} = y , \quad (2.34)$$

$$\hat{z} = z . \quad (2.35)$$

This test includes the effects of a nonvanishing shift vector. We use the same computational domain and KST parameters as in the standard gauge wave tests above. The amplitude suggested in Ref. [17] is $A = 0.5$. We also run simulations with $A = 0.1$.

At high amplitude, $A = 0.5$, we see exponentially growing nonconvergent instabilities. Without filtering, the code crashes after just a few crossing times. By filtering out the top 1/3 spectral coefficients as described above, the evolution can be extended as far as $t = 60$. No other choice of filtering seems to improve this

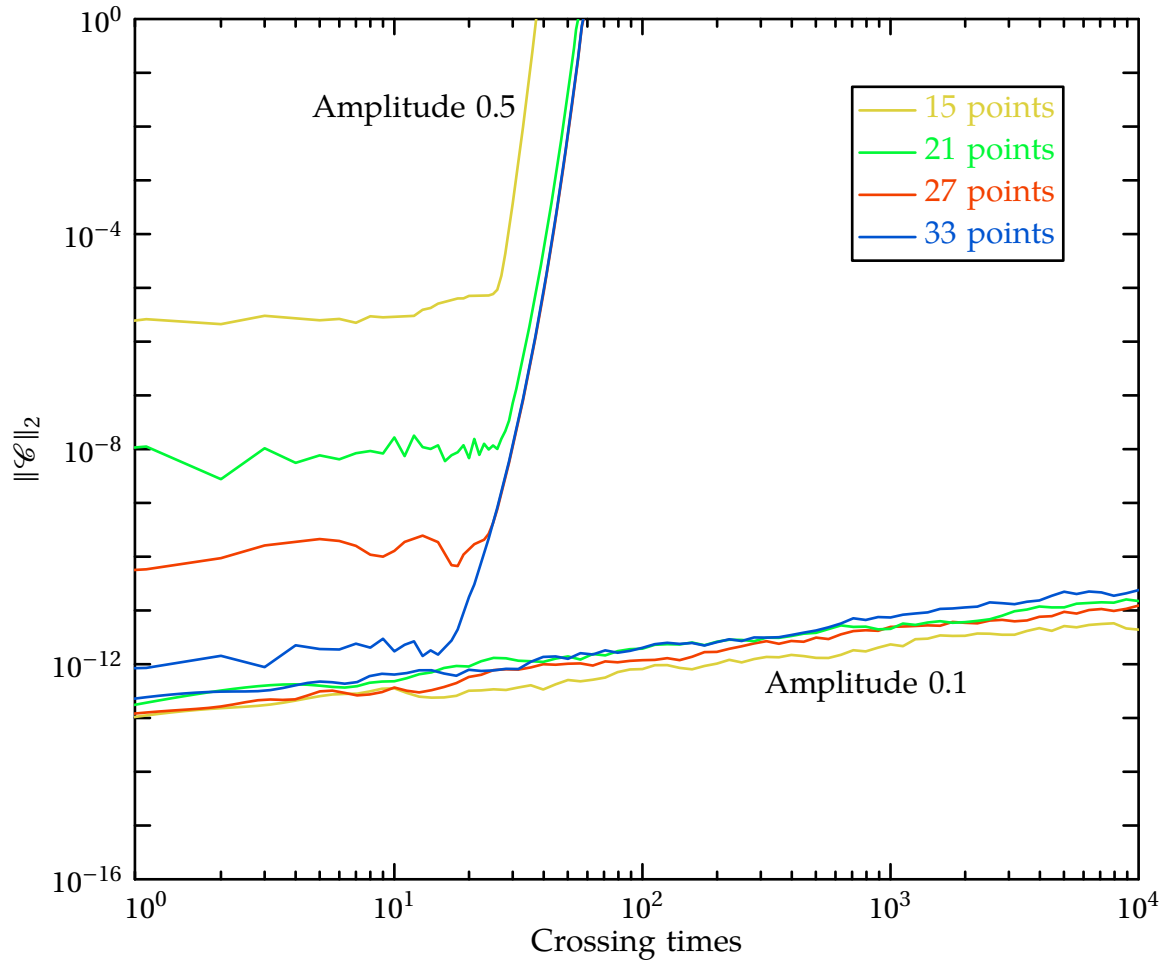


Figure 2.10: Constraints for shifted gauge wave

Solid lines indicate $A = 0.5$, and dashed lines indicate $A = 0.1$. For both amplitudes we filter out the top 1/3 spectral coefficients.




Figure 2.11: Error energy for shifted gauge wave

Solid lines indicate $A = 0.5$, while dashed lines indicate $A = 0.1$. The growth in the $A = 0.1$ runs is roughly linear in time, accelerating to quadratic at later times. The dotted line indicates the standard time step ($\Delta t = \Delta x/40$) with 33 points, which is dominated by temporal discretization error, while the blue dashed curve uses $\Delta t = \Delta x/80$.

further. We also run the test with an amplitude of $A = 0.1$. For this amplitude, the evolutions are stable with filtering but unstable without. Figs. 2.10 and 2.11 show the constraints and error energy for these evolutions. The initial constraint violations in these runs, $\|\mathcal{C}\| \approx 10^{-13}$, are at the level of the finite machine precision, so increasing the resolution causes increased, rather than decreased, constraint violations. The growth in $\|\delta\mathcal{U}\|$ seen in Fig. 2.11 is linear in time for $t < 100$, becoming quadratic at late times. The quadratic in time growth is dominated by time-stepping error, which tests show is convergent. (Reducing this error to the level of spatial-truncation error would require a prohibitive amount of computing time at the higher resolutions.)

2.6 Gowdy spacetime

The Gowdy spacetimes are dynamic cosmological solutions that present a serious challenge to any numerical relativity code. The Gowdy spacetimes are vacuum cosmological models having two spatial Killing fields (planar symmetry) that expand from (or, when time-reversed, contract toward) a curvature singularity. Two particular examples of these spacetimes with relatively simple analytical forms were chosen for the Mexico City tests: one in which the spacetime expands away from the singularity; another in which it collapses toward the singularity.

2.6.1 Expanding form

The metric chosen for the expanding case is

$$ds^2 = t^{-1/2} e^{\frac{\lambda-\lambda_0}{2}} (-dt^2 + dz^2) + t(e^P dx^2 + e^{-P} dy^2), \quad (2.36)$$

where

$$P(t, z) = J_0(2\pi t) \cos(2\pi z) , \quad (2.37)$$

$$\begin{aligned} \lambda(t, z) = & -2\pi t J_0(2\pi t) J_1(2\pi t) \cos^2(2\pi z) \\ & + 2\pi^2 t^2 [J_0^2(2\pi t) + J_1^2(2\pi t)] , \end{aligned} \quad (2.38)$$

$\lambda_0 = \lambda(1, 1/8)$, and J_n is a Bessel function. Asymptotically, P approaches zero as time increases, and λ increases linearly with time. Because the metric components are singular at $t = 0$, the Mexico City test begins the evolution at $t = 1$ and proceeds forward in time.

The time step Δt required for numerical stability is roughly given by the Courant condition $\Delta t \lesssim \Delta x / v$, where Δx is the spacing between collocation points and v is the coordinate speed of wave propagation, which in this case is the coordinate speed of light. For the Gowdy metric the coordinate speed of light in the z -direction is constant in time, but in the x - and y -directions it varies roughly like $t^{3/4} e^{t/2}$. Therefore, the maximum allowed time step decreases in time like $t^{-3/4} e^{-t/2}$, so for any fixed time step, the evolution will soon become numerically unstable if there is any perturbation in the x - or y -directions. This problem can be circumvented by running the simulation with just one point in the transverse directions, effectively eliminating any perturbation that could seed the instability.

Another difficulty with evolving the expanding Gowdy metric is that the metric components and derivatives become enormous very quickly. By $t \sim 725$ the numbers become larger than 10^{310} , so the evolution cannot be easily handled using standard 64-bit floating-point arithmetic. Our evolutions do not actually crash until $t = 700$; unfortunately errors dominate our evolutions long before this

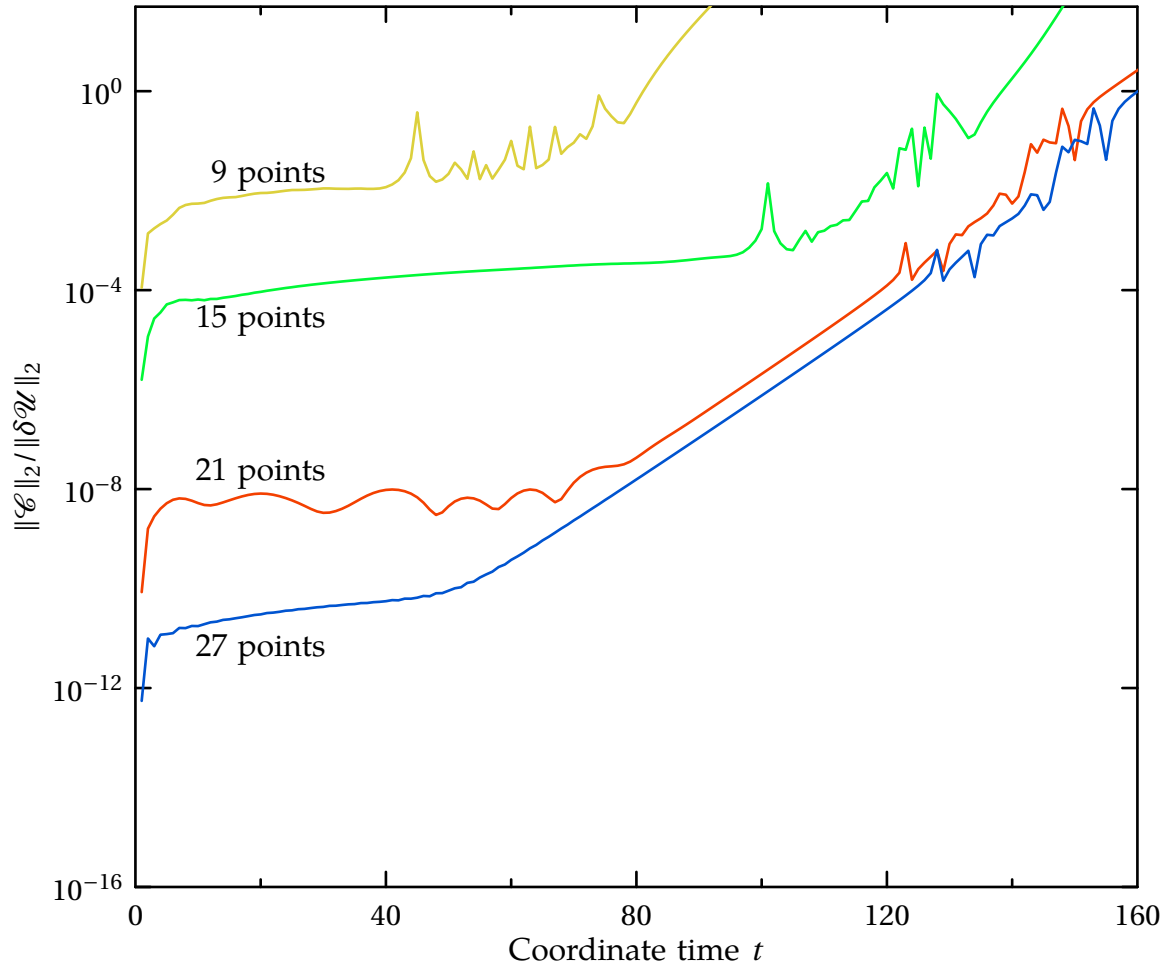


Figure 2.12: Constraints for expanding Gowdy spacetime

At early times, the exponential convergence of spectral methods is clearly visible. Soon, however, the evolutions are dominated by constraints growing roughly as $e^{t/5}$.




Figure 2.13: Error energy for expanding Gowdy spacetime

The error energy converges with increased spatial resolution, but $\|\delta\mathcal{U}\|_2/\|\mathcal{U}\|_2$ grows like $e^{t/5}$.

time, as seen in Fig. 2.13. The normalized error energy—along with the constraints shown in Fig. 12—grows roughly as $e^{t/5}$, and accuracy is completely lost in these evolutions by $t \sim 150$.

2.6.2 Collapsing form

The time coordinate in the Gowdy metric given above can be transformed so that the initial singularity is approached only asymptotically in the past. The new time coordinate, τ , is defined by $\tau := \frac{1}{c} \ln(t/k)$, where $c = 0.0021195119214607454$, and $k = 9.6707698127640558$. The spacetime can be evolved backwards indefinitely without reaching the singularity; that is, the time step is chosen to be negative. For purposes of convenience, the evolution is begun at an initial time of $\tau = \tau(t_0)$, where $t_0 = 9.8753205829098263$, which is a zero of $J_0(2\pi t)$.

This evolution is far less challenging than the expanding case. This is because the lapse function is essentially an exponential in τ , so that the spacetime is becoming *less* dynamical as the simulation progresses and τ becomes more negative. The main challenge in this test is resolving the spatial features of the solution. For spectral methods, the convergence should be exponential with increasing resolution, which is indeed the behavior shown in Figs. 2.14 and 2.15.

2.7 Discussion

We have applied the full suite of Mexico City tests [7]—modified suitably—to a pseudospectral implementation of the KST formulation of Einstein’s equations.




Figure 2.14: Constraints for collapsing Gowdy spacetime

Note that the simulation starts at $\tau \sim 9.875$, and proceeds backwards.




Figure 2.15: Error energy for collapsing Gowdy spacetime

The simulation starts at $\tau \sim 9.875$, and proceeds backwards.

We have also implemented the shifted gauge-wave test suggested by Babiuc, et al. [17], and suggested a number of minor changes to the tests that make them better challenges for pseudospectral methods. These tests reveal that the KST equations with pseudospectral methods demonstrate excellent convergence and accuracy, along with very good stability in all but a few cases. We have derived a fundamental limit Eq. (2.25) for the time-step accuracy possible in a method-of-lines numerical simulation, and have shown that our implementation is capable of quickly achieving that limit in the simple case of a sinusoidal linear wave. We have also shown that the use of filtering is very effective in reducing nonlinear aliasing instabilities.

The Mexico City tests provide a basic set of benchmarks for evaluating any numerical relativity code: allowing direct comparisons between different codes that use different numerical techniques and different formulations of the Einstein equations. However, the tests in their present form make too many implicit assumptions about the evolution system and the numerical methods. Since the creation of the tests, numerical relativity codes have become more diverse: using a variety of improved numerical techniques (fixed and adaptive mesh refinement, higher-order finite-differencing, multi-block methods, spectral methods) and at least two evolution systems (generalized harmonic and BSSN) capable of successfully evolving binary black hole spacetimes.

To accommodate the wide range of numerical methods and evolution systems now being used, future tests need to be formulated in more abstract terms. We recommend the following specific changes to the statement of the tests:

1. A code should demonstrate convergence, both spatial and temporal, appro-




Figure 2.16: **Comparing constraints for 1-D Gaussian linear waves**

Hamiltonian constraint norms $\|\mathcal{H}\|_2$ are much smaller than the constraint energies $\|\mathcal{C}\|_2$ for this test, so by themselves are not good diagnostics of constraint violations. In Fig. 2.4, we showed the constraint energies.

priate for the numerical method used, for each of the tests (gauge wave, linear wave, Gowdy spacetime, etc.).

The number of grid points or the time step needed to achieve a given accuracy is highly dependent on the numerical implementation. Therefore, the test specifications should not dictate a certain number of grid points or a certain time-step size as the original formulation of the tests did.

2. The combined error of all evolution variables, and the combined constraint violation (including all constraints of an evolution system, see Eqs. (2.7) and (2.8)), should be reported for each of the tests.

Prescriptions for examining errors of particular variables or constraints, such as those given in the original Mexico City tests, are not applicable to evolution systems that do not evolve those particular variables or constraints (e.g., tetrad or generalized harmonic evolution systems). In addition such prescriptions may not encompass all variables or constraints (as in the KST system), and may therefore fail to detect errors that accumulate only in a subset of the evolved variables. To illustrate this point, Fig. 2.16 shows both the total constraint energy $\|\mathcal{C}\|$, and the Hamiltonian constraint for the Gaussian linear wave (see Fig. 2.4). The Hamiltonian constraint turns out to be anomalously small for the KST system in this case, and so is not a good overall error indicator.

3. Use periodic Gaussian wave spatial profiles in the linear and gauge wave tests.

The sinusoidal spatial profiles specified in the original Mexico City tests with periodic boundary conditions provide an artificial advantage for spectral techniques.

Periodic Gaussian profiles are no more difficult for finite difference codes, and provide a significantly greater challenge for spectral methods. Finally,

4. Output data at generic times, not at integer multiples of the light-crossing time.

Outputting data at exact integer multiples of the light-crossing time significantly underestimates the errors in codes with very small dissipation (such as spectral codes).

We believe these recommendations will make it easier to apply the Mexico City tests fairly to a far wider class of numerical relativity codes, and so facilitate apples-with-apples comparisons between these codes. We have learned a great deal about the subtle properties of our code by carefully running and analyzing these simple tests. We encourage other groups to make their results from these tests public so that meaningful and objective comparisons between codes can be made.

EXTRAPOLATING GRAVITATIONAL-WAVE DATA FROM NUMERICAL SIMULATIONS¹

Techniques are developed for obtaining the asymptotic form of gravitational-wave data from numerical simulations. The extrapolation techniques are discussed in the context of Newman–Penrose and Regge–Wheeler–Zerilli data, and applied to extrapolation of waveforms from an equal-mass, non-spinning black-hole binary simulation. The sources of discrepancies between waveforms extracted at finite radius and those extrapolated to infinite radius are discussed. A formula is developed to estimate the size of near-field effects. It is shown that these near-field effects are indeed the largest source of the discrepancies during inspiral for Regge–Wheeler–Zerilli data, and that effects of similar magnitude and scaling dominate for Newman–Penrose data. Thus it is demonstrated that during the inspiral, techniques to improve the extraction method, while probably quite useful, are not entirely sufficient and

¹This chapter will be incorporated into a paper to be published with Abdul H. Mroué. The data from the binary simulations come courtesy of Mark Scheel, Harald Pfeiffer, and Luisa Buchman. The work contained in this chapter is my own, with the exception of the general argument for the retarded time in Sec. 3.2.2, which is due to Lee Lindblom. Lee and Mark also gave the text thorough readings and suggested revisions, which have improved it substantially.

must be supplemented with extrapolation.

3.1 Introduction

As numerical simulations of black-hole binaries improve, the criterion for success moves past the ability of a code to merely persist through many orbits of inspiral, merger, and ringdown. Accuracy becomes the goal, as related work in astrophysics and analysis of data from gravitational-wave detectors begin to rely more heavily on results from numerical relativity. The greatest challenge in the field today is to find and eliminate systematic errors that could pollute results built on numerics. Though there are many possible sources of such error, one stands out as being particularly easy to manage and—as we show—a particularly large effect: the error made by extracting gravitational waveforms from a simulation at finite radius, and treating these waveforms as though they were the asymptotic form of the radiation.

The desired waveform is the one to which post-Newtonian approximations aspire, and the one sought by gravitational-wave observatories: the asymptotic waveform. In typical numerical simulations, data extraction takes place at a distance of order $100M$ from the black holes. At this radius, the waves are still rapidly changing due to real physical effects. Near-field effects are plainly evident, scaling with powers of the ratio of the reduced wavelength to the radius, $(\lambda/r)^k$.² Extraction methods aiming to eliminate the influence of gauge alone (e.g., improved Regge–Wheeler–Zerilli or quasi-Kinnersley techniques) will not be able to account for these physical changes.

²We use the standard notation $\lambda := \lambda/2\pi$.

Even using a rather naive, gauge-dependent extraction method, we find that the near-field effects dominate the error in extracted waves throughout the inspiral. For extraction at $r = 50 M$, in the early stages of a 16-orbit equal-mass binary inspiral, these effects can account for an error of more than 50% in amplitude, or a phase difference of more than one radian. Crucially, the amplitude and phase differences change most rapidly during the merger, meaning that coherence is lost between the inspiral and merger/ringdown segments of the waveform. For the matched-filtering technique used to analyze data from gravitational-wave detectors, this can lead to serious error.

Matched filtering compares two signals, $s_1(t)$ and $s_2(t)$. It does this by Fourier transforming each into the frequency domain, taking the product of the signals, weighting each inversely by the noise—which is a function of frequency—and integrating over all frequencies. This match is optimized over the time and phase offsets of the input waveforms. For appropriately normalized waveforms, the result is a number between 0 and 1, denoted $\langle s_1 | s_2 \rangle$.³ If we imagine that s_1 is the extrapolated waveform and that s_2 is the waveform extracted at finite radius, we can evaluate the match between them. If the extrapolated waveform accurately represents the “true” physical waveform, the mismatch (defined as $1 - \langle s_1 | s_2 \rangle$) shows us the loss of signal in data analysis if we were to use the finite-radius waveforms to search for physical waveforms in detector data.

The waveforms have a simple scaling with the total mass of the system, which sets the frequency scale relative to the noise present in the detector. In Figs. 3.1 and 3.2, we show mismatches for a range of masses of interest to LIGO data analysis, using the Initial- and Advanced-LIGO noise curves, respectively, to weight

³For a more precise discussion of matched filtering, see Refs. [136, 137], or Chapter 6.

the matches. The value of r denotes the radius of extraction for the finite-radius waveform.

These figures demonstrate that the mismatch can be several percent when extracting at a radius of $r = 50 M$. For extraction at $r = 225 M$, the mismatch is never more than about 0.1%. Lindblom et al. [186] cite a target mismatch of less than 0.5% between the physical waveform and a class of model templates to be used for detection of events in detector data. Thus, for example, if these numerical waveforms were to be used in construction of template banks, the waveform extracted at $r = 50 M$ would be entirely insufficient, though the $r = 225 M$ waveform may be acceptable. Estimating the parameters of the waveform—masses and spins of the black holes, for instance—requires still greater accuracy. For the loudest signals expected to be seen by Advanced LIGO, the required mismatch may be roughly 10^{-4} [186]. In this case, even extraction at $r = 225 M$ would be insufficient; some method must be used to obtain the asymptotic waveform.

To remove the finite-radius effects—as we show below—we cannot simply rely on improved extraction methods, like quasi-Kinnersley techniques [29]. While those methods will no doubt improve the data quality, they cannot obtain the asymptotic waveform by extracting data at just one radius. An interesting method has been suggested [1, 2] for using the waveform and—essentially—its radial derivatives at a single radius, and matching to an analytic near-field expansion of the waveform. This method merits attention and may be very useful, though it imposes many assumptions about the behavior of the solution. In particular, it essentially assumes the correctness of the near-field formulas discussed below.

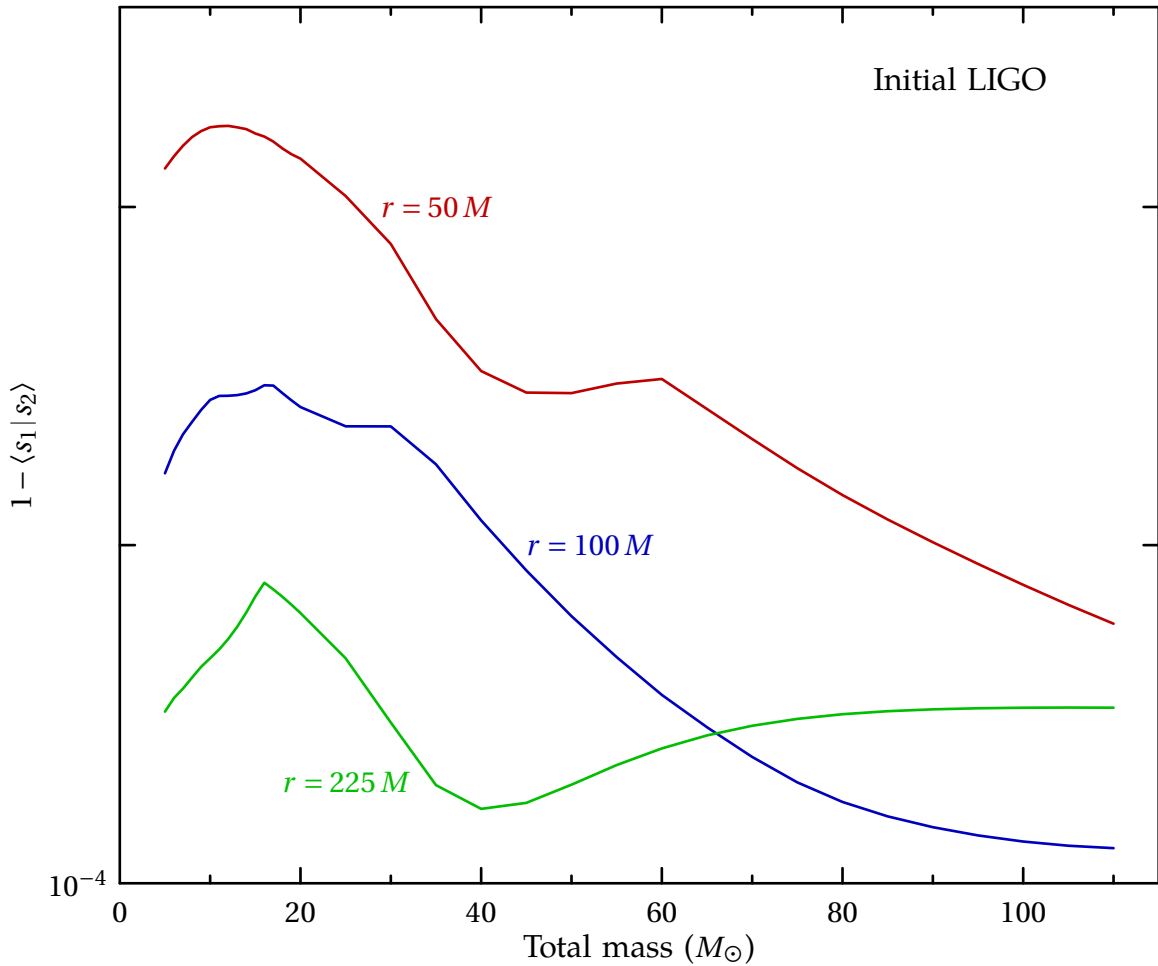


Figure 3.1: **Data-analysis mismatch between finite-radius waveforms and the extrapolated waveform for Initial LIGO**

This plot shows the mismatch between extrapolated waveforms and waveforms extracted at several finite radii, scaled to various values of the total mass of the binary system, using the Initial-LIGO noise curve. The waveforms are shifted in time and phase to find the maximum match.




Figure 3.2: **Data-analysis mismatch between finite-radius waveforms and the extrapolated waveform for Advanced LIGO**

This plot shows the mismatch between extrapolated waveforms and waveforms extracted at several finite radii, scaled to various values of the total mass of the binary system, using the Advanced-LIGO noise curve. The waveforms are shifted in time and phase to find the maximum match.

While we show this assumption to be basically correct for Regge–Wheeler–Zerilli waveforms, we prefer not to impose the assumption from the beginning.

Instead, we propose extraction of waveform data using more naive methods, at a series of radii—whether on a series of concentric spheres, or at various radii along an outgoing null ray. These data can then be expressed as functions of extraction radius and retarded time using a simple method we describe. For each value of retarded time, the waveforms can then be fit to a polynomial in inverse powers of the extraction radius. The asymptotic waveform is simply the first nonzero term in the polynomial. Though this method also imposes certain assumptions, they amount to assuming that the data behave as radially propagating waves, and that the metric itself is asymptotically Minkowski.

Extrapolation is, by its very nature, a dangerous procedure. The final result may be numerically unstable, in the sense that it will fail to converge as the order of the extrapolating polynomial is increased. This is to be expected, as the size of the effects to be removed falls below the size of noise in the waveform data. There are likely better methods of determining the asymptotic form of gravitational waves produced by numerical simulations. Characteristic evolution is an example of a promising technique that may become common in the near future [33, 165, 15, 16]. Nonetheless, extrapolation does provide a rough and ready technique which can easily be implemented by numerical-relativity groups using existing frameworks.

This chapter presents a simple method for implementing the extrapolation of gravitational-wave data from numerical simulations, and the motivation for doing so. We begin, in Sec. 3.2, by briefly outlining the basic method for ex-

trapolating generic data, and suggesting a simple method for treating the crucial problem of associating data from different space and time locations. In Sec. 3.3 we review two standard methods for extracting waveforms from general simulations: Regge–Wheeler–Zerilli and Newman–Penrose methods. We outline the extrapolation process in greater detail in Sec. 3.4, applying the general method to a particular binary black-hole simulation, along with details of the algorithm. We then present the results of extrapolation using data from a long simulation of the inspiral and merger of equal-mass black holes for both Regge–Wheeler–Zerilli data and Newman–Penrose data. In Sec. 3.5, we calculate expected near-field effects, and compare the extrapolated waveform to data extracted at finite radius. We then show that—for Regge–Wheeler–Zerilli data, at least—most of the effect of extrapolation involves removing the influence of near-field effects. Finally we conclude, in Sec. 3.6, with a discussion of remaining issues in extrapolation.

3.2 Extrapolation in brief

There are many types of data that can be extracted from a numerical simulation of an isolated source of gravitational waves. Below, we discuss two common methods of extracting gravitational waveforms—using the Newman–Penrose Ψ_4 quantity, or the metric perturbation h extracted using Regge–Wheeler–Zerilli techniques. Even if we focus on a particular type of waveform, the data can be extracted at a series of points along the z axis, for example, or decomposed into multipole components and extracted on a series of spheres around the source. To simplify this introductory discussion of extrapolation, we ignore the variety of particular types of waveform data. Instead, we generalize to some abstract quan-

tity f , which encapsulates the quantity to be extrapolated and behaves roughly as a radially outgoing wave. We assume that f travels along outgoing null cones, which we parametrize by a retarded time t_{ret} . Along each of these null cones, we further assume that f can be expressed as a convergent series in $1/r$ —where r is some radial coordinate—for all radii of interest. That is, we assume

$$f(t_{\text{ret}}, r) = \sum_{k=0}^{\infty} \frac{f_{(k)}(t_{\text{ret}})}{r^k}, \quad (3.1)$$

for some functions $f_{(k)}$. The asymptotic behavior of f is given by the lowest nonzero $f_{(k)}$.⁴

Given data for such an f at a set of retarded times, and a set of radii $\{r_i\}$, it is a simple matter to fit the data for each value of t_{ret} to a polynomial in $1/r$. That is, for each value of t_{ret} , we take the set of data $\{f(t_{\text{ret}}, r_i)\}$ and fit it to a finite polynomial so that

$$f(t_{\text{ret}}, r_i) \simeq \sum_{k=0}^N \frac{f_{(k)}(t_{\text{ret}})}{r_i^k}. \quad (3.2)$$

Standard algorithms [218] can be used to accomplish this fitting. Of course, because we are truncating the series of Eq. (3.1) at $k = N$, some of the effects from $k > N$ terms will appear at lower orders. We will need to choose N appropriately, checking that the extrapolated quantity has converged sufficiently with respect to this order.

3.2.1 Choice of r parameter

One subtlety to be considered is the choice of r parameter to be used in the extraction and fitting. For numerical simulation of an isolated system, one simple

⁴For example, if $f = r\Psi_4$, then $f_{(0)}$ gives the asymptotic behavior; if $f = \Psi_4$, then $f_{(1)}$ gives the asymptotic behavior.

and obvious choice is the coordinate radius r_{coord} used in the simulation. Alternatively, if the data is measured on some spheroidal surface, it is possible to define an areal radius r_{areal} by measuring the area of the sphere A along with f , and setting $r_{\text{areal}} := \sqrt{A/4\pi}$. Still other choices are certainly possible.

One objective in choosing a particular r parameter is to ensure the physical relevance of the final extrapolated quantity. If we try to detect the wave, for example, we may want to think of the detector as being located at some constant value of r . Or, we may want r to asymptotically represent the luminosity distance. These conditions may be checked by inspecting the asymptotic behavior of the metric components in the given coordinates. For example, if the metric components in a coordinate system including r asymptotically approach those of the standard Minkowski metric, it is not hard to see that an inertial detector could follow a path of constant r parameter.

Suppose we have two different parameters r and \tilde{r} which can be related by a series expansion

$$r = \tilde{r} [1 + a/\tilde{r} + \dots] . \quad (3.3)$$

In the data presented in this paper, we can show that r_{coord} and r_{areal} are related in this way. Introducing the expansion coefficients $\tilde{f}_{(k)}$, we can write

$$f(t_{\text{ret}}, r) = \sum_{k=0}^{\infty} \frac{f_{(k)}(t_{\text{ret}})}{r^k} = \sum_{k=0}^{\infty} \frac{\tilde{f}_{(k)}(t_{\text{ret}})}{\tilde{r}^k} . \quad (3.4)$$

Inserting Eq. (3.3) into this formula, Taylor expanding, and equating terms of equal k , this shows that $f_{(0)} = \tilde{f}_{(0)}$ and $f_{(1)} = \tilde{f}_{(1)}$. Thus, if the asymptotic behavior of f is given by $f_{(0)}$ or $f_{(1)}$, the final extrapolated data should not depend on whether r or \tilde{r} is used. On the other hand, in practice we truncate the series in Eq. (3.1) at finite order. This means that higher-order terms could “pollute” $f_{(0)}$

or $f_{(1)}$. The second objective in choosing an r parameter, then, is to ensure fast convergence of the series in Eq. (3.1) or Eq. (3.2). If the extrapolated quantity does not converge quickly as the order of the extrapolating polynomial N is increased, it may be due to a poor choice of r parameter.

The coordinate radius used in a simulation may be subject to large gauge variations which are physically irrelevant, and hence are not reflected in the wave's behavior. That is, the wave may not fall off nicely in inverse powers of that coordinate radius. For the data discussed later in this paper, we find that using the coordinate radius of extraction spheres is indeed a poor choice, while using the areal radius of those extraction spheres improves the convergence of the extrapolation.

3.2.2 Choice of retarded time parameter

Similar considerations must be made for the choice of retarded time t_{ret} to be used in extrapolation. It may be possible to evolve null geodesics in numerical simulations, and use these to define the null curves on which data is to be extracted. While this is an interesting possibility that deserves investigation, we use a simpler method here based on an approximate retarded time constructed using the coordinates of the numerical simulation.

Again, we have two motivations for choosing a retarded time parameter. First is the physical suitability in the asymptotic limit. For example, we might want the asymptotic t_{ret} to be (up to an additive term constant in time) the proper time along the path of a detector located at constant r . Again, checking the asymptotic behavior of the metric components with respect to t_{ret} and r should be a sufficient

test of the physical relevance of the parameters. Second, we wish to have rapid convergence of the extrapolation series using the chosen parameter, which also needs to be checked.

As before, we can also show the equivalence of different choices for the t_{ret} parameter. Suppose we have two different approximations t_{ret} and \tilde{t}_{ret} which can be related by a series expansion

$$t_{\text{ret}} = \tilde{t}_{\text{ret}} [1 + b/r + \dots] . \quad (3.5)$$

Using the new expansion coefficients $\tilde{f}_{(k)}$, we can write

$$f(t_{\text{ret}}, r) = \sum_{k=0}^{\infty} \frac{f_{(k)}(t_{\text{ret}})}{r^k} = \sum_{k=0}^{\infty} \frac{\tilde{f}_{(k)}(\tilde{t}_{\text{ret}})}{r^k} . \quad (3.6)$$

Now, however, we need to assume that the functions $f_{(k)}$ can be well-approximated by Taylor series. If this is true, we can again show that $f_{(0)} = \tilde{f}_{(0)}$ or, if we have $f_{(0)} = \tilde{f}_{(0)} = 0$, that $f_{(1)} = \tilde{f}_{(1)}$. The condition that f be well-approximated by a Taylor series is nontrivial, and can help to inform the choice of f . Similarly, the speed of convergence of the extrapolation can help to inform the choice of a particular t_{ret} parameter.

Since we will be considering radiation from an isolated compact source, our basic model for t_{ret} comes from the Schwarzschild spacetime; we assume that the system in question approaches this spacetime at increasing distance. In analogy with the time-retardation effect on outgoing null rays in a Schwarzschild spacetime [92], we define a ‘‘tortoise coordinate’’ r_* by:

$$r_* := r + 2M_{\text{ADM}} \ln \left(\frac{r}{2M_{\text{ADM}}} - 1 \right) , \quad (3.7)$$

where M_{ADM} is the ADM mass of the initial data.⁵ In standard Schwarzschild coordinates, the appropriate retarded time would be given by $t_{\text{ret}} = t - r_*$. It is not hard to see that the exterior derivative dt_{ret} is null with respect to the Schwarzschild metric.

Taking inspiration from this, we can attempt to account for certain differences from a Schwarzschild background. Let T and R denote the simulation's coordinates, and suppose that we extract the metric components g^{TT} , g^{TR} , and g^{RR} from the simulation. We seek a $t_{\text{ret}}(T, R)$ such that

$$dt_{\text{ret}} = \frac{\partial t_{\text{ret}}}{\partial T} dT + \frac{\partial t_{\text{ret}}}{\partial R} dR \quad (3.8)$$

is null with respect to these metric components. That is, we seek a t_{ret} such that

$$g^{TT} \left(\frac{\partial t_{\text{ret}}}{\partial T} \right)^2 + 2g^{TR} \left(\frac{\partial t_{\text{ret}}}{\partial T} \right) \left(\frac{\partial t_{\text{ret}}}{\partial R} \right) + g^{RR} \left(\frac{\partial t_{\text{ret}}}{\partial R} \right)^2 = 0. \quad (3.9)$$

We introduce the ansatz $t_{\text{ret}} = t - r_*$, where t is assumed to be a slowly varying function of R ,⁶ and r_* is given by Eq. (3.7) with R in place of r on the right side.

If we ignore $\partial t / \partial R$ and insert our ansatz into the last equation, we have

$$g^{TT} \left(\frac{\partial t}{\partial T} \right)^2 - 2g^{TR} \left(\frac{\partial t}{\partial T} \right) \left(\frac{1}{1 - 2M_{\text{ADM}}/R} \right) + g^{RR} \left(\frac{1}{1 - 2M_{\text{ADM}}/R} \right)^2 = 0. \quad (3.10)$$

We can solve this for $\partial t / \partial T$:

$$\frac{\partial t}{\partial T} = \frac{1}{1 - 2M_{\text{ADM}}/R} \frac{g^{TR} \pm \sqrt{(g^{TR})^2 - g^{TT} g^{RR}}}{g^{TT}}. \quad (3.11)$$

⁵ Kocsis and Loeb [182] pointed out that the propagation of a roughly spherical gravitational wave should be affected primarily by the amount of mass *interior* to the wave. Because the waves from a merging binary can carry off a significant fraction (typically a few percent) of the binary's mass, this suggests that we should allow the mass in this formula to vary in time, falling by perhaps a few percent over the duration of the waveform. However, this is a small correction of a small correction; we have not found it necessary. Perhaps with more refined methods, this additional correction would be relevant.

Substituting the Schwarzschild metric components shows that we should choose the negative sign in the numerator of the second term. Finally, we can integrate (numerically) to find

$$t = \int_0^T \frac{1}{g^{TT}} \frac{g^{TR} - \sqrt{(g^{TR})^2 - g^{TT} g^{RR}}}{1 - 2M_{\text{ADM}}/R} dT' . \quad (3.12)$$

Now, in the case where g^{TR} is small compared to 1, we may wish to ignore it, in which case we have

$$t = \int_0^T \frac{\sqrt{-g^{RR}/g^{TT}}}{1 - 2M_{\text{ADM}}/R} dT' . \quad (3.13)$$

It is not hard to see that this correctly reduces to $t = T$ in the Schwarzschild case.

For the data discussed later in this paper, we make further assumptions that $g^{RR} = 1 - 2M_{\text{ADM}}/R$, and that $R = r_{\text{areal}}$. That is, we define the corrected time

$$t_{\text{corr}} := \int_0^T \sqrt{\frac{-1/g^{TT}}{1 - 2M_{\text{ADM}}/r_{\text{areal}}}} dT' \quad (3.14a)$$

and the retarded time

$$t_{\text{ret}} := t_{\text{corr}} - r_* . \quad (3.14b)$$

We will show that this corrected time leads to a significant improvement over the naive choice of $t(T) = T$.

3.3 Waveform extraction

The first task in analyzing data from a numerical simulation is simply extracting the relevant waveform. There are two common methods used to achieve this

⁶More specifically, we need $|\partial t/\partial R| \ll |\partial r_*/\partial R|$. This condition needs to be checked for all radii used, at all times in the simulation. For the data presented below, we have checked this, and shown it to be a valid assumption, at the radii used for extrapolation.

aim. One uses Regge–Wheeler–Zerilli functions to extract the metric perturbation directly. This method explicitly decomposes the full spacetime into a spherically symmetric background and a perturbation. Particular combinations of the perturbations are then taken, making the final result invariant *to first-order* under gauge perturbations. The second method uses Newman–Penrose scalars to decompose the Weyl tensor, directly reading off the curvature. We now examine these two techniques in more detail.

3.3.1 Regge–Wheeler–Zerilli functions

Both methods of extracting gravitational-wave information have their origins in treatments of perturbations of Schwarzschild and Kerr black holes. While these analyses are valid right down to—or even within—the black hole’s horizon, we only need them at some distance from the center of our domain. The physical situation of interest to simulations of binary black holes is an asymptotically flat spacetime. In this case, the spacetime in the outer parts of a simulation can be thought of as a perturbation of Schwarzschild.

Regge and Wheeler [222] first introduced a formalism for studying odd-parity perturbations of a Schwarzschild black hole in Schwarzschild coordinates, while investigating the stability of black holes. Zerilli [255] extended their formalism to allow for even-parity perturbations. The Regge–Wheeler–Zerilli (RWZ) methods are gauge dependent, however, requiring the gauge to be fixed before any interpretation can be made of the physical content of their results. Moncrief [196] suggested an improvement to the combined formalism, making it *linearly* invariant to gauge perturbations. Sarbach and Tiglio [228] have extended this further,

with a geometric treatment that can be used to study perturbations on an arbitrary slicing of a spherically symmetric spacetime.⁷ These generalized RWZ functions obey wave equations, and can be used to extract the (linearly) gauge-invariant metric perturbation. Second-order methods have also been developed [146, 83, 203, 147].

We begin by assuming the existence of a spherically symmetric background metric $\underline{g}_{\mu\nu}$ that can be decomposed into a time–radius sector and a spherical sector:

$$\underline{g}_{\mu\nu} dx^\mu dx^\nu = \check{g}_{ab} dx^a dx^b + r^2 \mathring{\Omega}_{AB} dx^A dx^B . \quad (3.15)$$

Here, lowercase Latin indices refer to the t – r sector, as does the check (suggesting the null-cone) over the metric. Similarly, uppercase Latin indices refer to the ϑ – φ sector, as does the ring (suggesting the sphere) over the metric. The metric $\mathring{\Omega}_{AB}$ is assumed to be the standard unit-sphere metric. We will also need the derivative operator $\mathring{\nabla}$ and antisymmetric tensor $\mathring{\epsilon}$ appropriate to the spheroidal submanifold spanned by ϑ, φ .

Next, we define the metric perturbation $\delta g_{\mu\nu}$ as the full metric minus the background:

$$\delta g_{\mu\nu} := g_{\mu\nu} - \underline{g}_{\mu\nu} . \quad (3.16)$$

We then expand the perturbation in (tensor, vector, and scalar) spherical harmon-

⁷Useful reviews of RWZ methods can be found in Refs. [198, 71, 226].

ics:

$$\delta g_{ab} = \sum_{l,m} \tilde{P}_{ab}^{l,m} Y_{l,m} , \quad (3.17a)$$

$$\delta g_{Ab} = \sum_{l,m} \left[\tilde{Q}_b^{l,m} \overset{\circ}{\nabla}_A Y_{l,m} + \tilde{R}_b^{l,m} \overset{\circ}{e}^C{}_A \overset{\circ}{\nabla}_C Y_{l,m} \right] , \quad (3.17b)$$

$$\delta g_{AB} = \sum_{l,m} \left[\tilde{S}^{l,m} r^2 \overset{\circ}{\Omega}_{AB} Y_{l,m} + \tilde{T}^{l,m} \overset{\circ}{\nabla}_{(A} \overset{\circ}{\nabla}_{B)} Y_{l,m} + \tilde{U}^{l,m} \overset{\circ}{e}^C{}_{(A} \overset{\circ}{\nabla}_{B)} \overset{\circ}{\nabla}_C Y_{l,m} \right] . \quad (3.17c)$$

The components $\tilde{P}_{ab}^{l,m}, \dots, \tilde{U}^{l,m}$ will be gauge dependent. Even in flat spacetime, we could introduce a gauge perturbation which mimics a gravitational wave in terms of its effect on some of the components. For example, if we only look at $\tilde{T}^{l,m}$ and ignore other components, we may easily be deceived by a gauge perturbation. Moncrief, however, pointed out that it is possible to take certain combinations of the components to obtain components that are gauge invariant to first order.

For example, take the magnetic-parity part of the metric perturbation:

$$\delta g_{Ab} = \tilde{R}_b^{l,m} \overset{\circ}{e}^C{}_A \overset{\circ}{\nabla}_C Y_{l,m} , \quad \delta g_{AB} = \tilde{U}^{l,m} \overset{\circ}{e}^C{}_{(A} \overset{\circ}{\nabla}_{B)} \overset{\circ}{\nabla}_C Y_{l,m} . \quad (3.18)$$

A general magnetic-parity gauge transformation looks like

$$\xi_a = 0 \quad \xi_A = \zeta(t, r) \overset{\circ}{e}^C{}_A \overset{\circ}{\nabla}_C Y_{l,m} , \quad (3.19)$$

for some function $\zeta(t, r)$. If this is an *infinitesimal* gauge transformation, the components of the metric perturbation transform as

$$\tilde{R}_b^{l,m} \rightarrow \tilde{R}_b^{l,m} + \nabla_b \zeta , \quad \tilde{U}^{l,m} \rightarrow \tilde{U}^{l,m} + \zeta . \quad (3.20)$$

Looking at these expressions for a moment suggests a way to cancel the gauge dependence:

$$R_b^{l,m} := \tilde{R}_b^{l,m} - \nabla_b \tilde{U}^{l,m} . \quad (3.21)$$

The new component $R_b^{l,m}$ is invariant to linear order under gauge transformations. Analogous—though more complicated—expressions may be obtained for the other components of the metric perturbation [228, 224].

These invariant components are then used to reconstruct the metric, which can be decomposed in the standard way [42] by defining the polarization tensors

$$\varepsilon_+^{\alpha\beta} := \vartheta^\alpha \vartheta^\beta - \varphi^\alpha \varphi^\beta \quad \text{and} \quad \varepsilon_\times^{\alpha\beta} := \vartheta^\alpha \varphi^\beta + \varphi^\alpha \vartheta^\beta, \quad (3.22)$$

and using these to define

$$h_+ := \frac{1}{2} \varepsilon_+^{\alpha\beta} \delta g_{\alpha\beta} \quad \text{and} \quad h_\times := \frac{1}{2} \varepsilon_\times^{\alpha\beta} \delta g_{\alpha\beta}. \quad (3.23)$$

Here, ϑ^α and φ^α are the standard coordinate vectors, which do not necessarily have unit length. These two quantities are commonly combined into a single complex field:

$$h := h_+ - i h_\times. \quad (3.24)$$

Details of the procedure used for the data presented here are given in [224]. In the notation of that paper, the metric perturbation is given by [231]

$$h = \frac{1}{r} \sum_{l=2}^{\infty} \sum_{m=-l}^l \sqrt{(l+2)(l+1)l(l-1)} \left(\Phi_{l,m}^{(+)} + i \Phi_{l,m}^{(-)} \right) {}_{-2}Y_{l,m}, \quad (3.25)$$

where ${}_{-2}Y_{l,m}$ are the spin-weight -2 spherical harmonics discussed in Sec. 3.3.3.

3.3.2 Newman–Penrose scalars

Newman and Penrose introduced a useful decomposition of the Weyl tensor, among other geometric quantities [201]. At each point in space, they define a complex null tetrad $(l^\alpha, n^\alpha, m^\alpha, \bar{m}^\alpha)$, where \bar{m}^α is simply the complex conjugate

of m^α . The tetrad is assumed to obey the orthonormality conditions $l^\alpha n_\alpha = -1$, $m^\alpha \bar{m}_\alpha = 1$, with all other products being zero, and the vector m^α assumed to be a complex combination of spacelike vectors. This tetrad is then used to define five complex functions of space and time:⁸

$$\Psi_0 := C_{\alpha\beta\gamma\delta} l^\alpha m^\beta l^\gamma m^\delta ; \quad (3.26a)$$

$$\Psi_1 := C_{\alpha\beta\gamma\delta} l^\alpha m^\beta l^\gamma n^\delta ; \quad (3.26b)$$

$$\Psi_2 := C_{\alpha\beta\gamma\delta} l^\alpha m^\beta \bar{m}^\gamma n^\delta ; \quad (3.26c)$$

$$\Psi_3 := C_{\alpha\beta\gamma\delta} l^\alpha n^\beta \bar{m}^\gamma n^\delta ; \quad (3.26d)$$

$$\Psi_4 := C_{\alpha\beta\gamma\delta} n^\alpha \bar{m}^\beta n^\gamma \bar{m}^\delta . \quad (3.26e)$$

These ten real degrees of freedom correspond to the ten degrees of freedom in the Weyl tensor itself, effectively encoding the curvature of a vacuum spacetime. Indeed, it is not difficult to re-express the Weyl tensor in terms of these five functions and the tetrad—essentially inverting the equations above. (See, e.g., Chandrasekhar’s Eq. (1.298) [92].) Because we will always be working in vacuum, we note that $C_{\alpha\beta\gamma\delta} = R_{\alpha\beta\gamma\delta}$.

A general spacetime, naturally, has nonzero values for each of these quantities. Moreover, because the choice of tetrad is not unique, we can generally mix components by rotating the tetrad. However, it can be shown [227] that—in vacuum, asymptotically flat spacetimes—there exists a tetrad for which Ψ_4 is $\mathcal{O}(r^{-1})$ along an outgoing null ray, and all other Newman–Penrose scalars fall off more quickly. This is essentially the so-called *peeling theorem* [235]. It shows that, if

⁸Note that Newman and Penrose use the opposite metric-signature convention and convention for the Riemann tensor, as compared to the one used here (see Sec. A.2). The equations they use to define Ψ_n also have opposite signs, so overall, the signs of the Ψ_n used here should agree with those of Newman and Penrose.

we hope to find radiation from an isolated source using the Newman–Penrose formalism, Ψ_4 is—at least—a crucial element.

In fact, there exist many tetrads such that Ψ_4 is $\mathcal{O}(r^{-1})$, and all other Newman–Penrose scalars fall off more quickly. This variety—and the ambiguity it brings about—makes the issue of choosing a particular such tetrad in a general spacetime a delicate one. Here, we simply define the tetrad components by the coordinate basis used in an evolution. Explicitly, we define the tetrad components by

$$l^\alpha := \frac{1}{\sqrt{2}} (t^\alpha + r^\alpha) , \quad (3.27a)$$

$$n^\alpha := \frac{1}{\sqrt{2}} (t^\alpha - r^\alpha) , \quad (3.27b)$$

$$m^\alpha := \frac{1}{\sqrt{2}} (\vartheta^\alpha + i\varphi^\alpha) . \quad (3.27c)$$

(When evaluation on the z axis is necessary, we always take the limit with $\varphi = 0$.) The vector t^α is a unit-magnitude vector, normal to the spacelike hypersurface of constant time in the simulation. The vector r^α is a unit vector within that spacelike hypersurface which is normal to the extraction sphere at that time. Both of the latter are normalized by the full spacetime metric. Again, the vectors ϑ^α and φ^α are the standard coordinate vectors. This choice actually leaves the orthonormality conditions for m^α and \bar{m}^α unsatisfied in general; they do not have unit magnitude, and are not orthogonal. Nonetheless, when the metric in these coordinates asymptotically approaches the standard Minkowski metric, we can expect that the orthonormality conditions will be satisfied asymptotically. In turn, we can also expect that this will suffice because the physically relevant part of the resulting Ψ_4 will be selected by extrapolation.

Relation of Ψ_4 to h

In the linear theory, we can construct a Minkowskian background, and small perturbations on top of it. In this case, we define the tetrad above with respect to the background, orthonormalizing the tetrad with respect to the background metric. Then, for plane waves propagating along the n^α vector, it is not difficult to show that⁹

$$\Psi_4 = -\ddot{h} = -(\ddot{h}_+ - i\ddot{h}_\times) \quad \text{in the linear approximation,} \quad (3.28)$$

where double dots denote double time derivatives. Note that h , h_+ , and h_\times were defined in Eqs. (3.23) and (3.24). This gives us further confidence that Ψ_4 is a physically relevant quantity to extract. Unfortunately, it also shows us that we need to integrate Ψ_4 twice, if we wish compute h from Ψ_4 . Various integration techniques have been developed, the most successful of which will be discussed in later chapters.

It is worth emphasizing that this relation is true only in the linear approximation, and only with a particular choice of tetrad. In highly nonlinear spacetimes, there is no reason to expect this relation to hold. Indeed, even in the linear approximation, this need not hold for general tetrads satisfying the conditions set out by Newman and Penrose; this simple expression is *only* obtained when the wave is propagating along the n^α vector and when $(\varepsilon_+ - i\varepsilon_\times)^{\alpha\beta} = 2\bar{m}^\alpha \bar{m}^\beta$. (Recall the definitions of ε_+ and ε_\times from Eq. (3.22).) For example, if we were to rotate the m^α vector by an angle θ , we would have $\Psi_4 = -e^{-2i\theta} \ddot{h}$.

⁹There are many sign choices that need to be made before deriving this relation: the sign of Riemann; the sign of Ψ_4 ; the metric signature; the sign of metric perturbations; etc. Equation (3.28) is correct for the signs chosen here. For a review of those signs, see Appendix A.

3.3.3 Spin-weighted spherical harmonics

Newman and Penrose [202] introduced a useful set of functions similar to the standard spherical harmonics, designed to decompose functions like Ψ_4 —the spin-weighted spherical harmonics (SWSHs), ${}_s Y_{l,m}$.¹⁰ In particular, Ψ_4 and h can be decomposed in terms of SWSHs of spin weight $s = -2$:

$$\Psi_4^{l,m}(t,r) := \int_0^{2\pi} \int_0^\pi \Psi_4(t,r,\vartheta,\varphi) {}_{-2}\bar{Y}_{l,m}(\vartheta,\varphi) \sin\vartheta d\vartheta d\varphi, \quad (3.29)$$

$$h^{l,m}(t,r) := \int_0^{2\pi} \int_0^\pi h(t,r,\vartheta,\varphi) {}_{-2}\bar{Y}_{l,m}(\vartheta,\varphi) \sin\vartheta d\vartheta d\varphi. \quad (3.30)$$

The waveform data used throughout the rest of this chapter will consist solely of the $(l,m) = (2,2)$ mode. When we present that data, we will generally drop the l,m indices.

3.4 Extrapolation technique applied to binary inspiral waveforms

In this section, we briefly discuss the data to be extracted from the simulation and the form in which they will be used. We show the results of the extrapolation. Then, we discuss the order of the extrapolating polynomial, and the trade-offs involved in moving to higher or lower orders. Finally, we discuss how this impacts and is impacted by the choice of radii at which to extract data.

¹⁰For further discussion of the spin-weighted spherical harmonics, see Appendix B and references therein.

3.4.1 Data to be extracted from the simulation

To begin the extrapolation procedure, we extract Ψ_4 and h data on a set of spheres at constant coordinate radius in the simulation. In the black-hole binary simulations used here (the same as those discussed in Refs. [59, 61, 60, 229]), these spheres are located every $\Delta r_{\text{coord}} = 10M_{\text{irr}}$ from an inner radius of $r_{\text{coord}} = 50M_{\text{irr}}$ to an outer radius of $r_{\text{coord}} = 240M_{\text{irr}}$, where M_{irr} denotes the total apparent-horizon mass of the two holes at the beginning of the simulation. This extraction occurs at time steps of $\Delta t_{\text{coord}} \approx 0.5M_{\text{irr}}$ throughout the simulation. We also measure the areal radius, r_{areal} , of those spheres by integrating the induced area element over the sphere to find the area A , and defining $r_{\text{areal}} := \sqrt{A/4\pi}$. This differs from the coordinate radius r_{coord} by roughly $M_{\text{irr}}/r_{\text{coord}}$. Because of gauge effects, the areal radius of a coordinate sphere changes as a function of time, so we measure this as a function of time. Finally, we measure the average lapse $N := 1/\sqrt{-g^{TT}}$ as a function of coordinate time on the extraction spheres to correct for dynamic lapse. The areal radius and lapse are used to compute the retarded time t_{ret} defined in Eq. (3.14).

The gravitational-wave data Ψ_4 and h , the areal radius r_{areal} , and the lapse N are all measured as functions of the code coordinates t_{coord} and r_{coord} . We can use these to construct the retarded time defined in Eq. (3.14), using r_{areal} in place of r . This, then, will also be a function of the code coordinates. The mapping between $(t_{\text{ret}}, r_{\text{areal}})$ and $(t_{\text{coord}}, r_{\text{coord}})$ is invertible, so we can rewrite Ψ_4 and h as functions of t_{ret} and r_{areal} .

As noted in Sec. 3.2.2, we need to assume that the extrapolated functions are well approximated by Taylor series. Because the real and imaginary parts of Ψ_4

and h are rapidly oscillating in the data presented here, we prefer to use the same data in smoother form. We define the complex amplitude A and phase ϕ of the wave:

$$r_{\text{areal}} M_{\text{irr}} \Psi_4 := A e^{i\phi} \quad (3.31a)$$

or

$$r_{\text{areal}} h / M_{\text{irr}} := A e^{i\phi} , \quad (3.31b)$$

where A and ϕ are functions of t_{ret} and r_{areal} . Note that this definition factors out the dominant $1/r$ behavior of the amplitude. Because of the ambiguity in this definition of the phase, we remove discontinuities of 2π at successive times. The continuous phase is easier to work with for practical reasons, and is certainly much better approximated by a Taylor series.

A slight complication arises in the relative phase offset between successive radii. Noise in the early parts of the waveform make the overall phase offset go through multiples of 2π essentially randomly. We choose some fairly noise-free time and ensure that phases corresponding to successive extraction spheres are matched at that (retarded) time, by simply adding multiples of 2π to the phase of the entire waveform—that is, we add multiples of 2π to the phase at all times.

Extrapolation of the waveform, then, basically consists of finding the asymptotic forms of these functions, A and ϕ as functions of time. We apply the general technique discussed in Sec. 3.2 to A and ϕ . That is, we fit the data to polynomials

in $1/r$ for each value of retarded time:

$$A(t_{\text{ret}}, r_{\text{areal}}) \simeq \sum_{k=0}^N \frac{A_{(k)}(t_{\text{ret}})}{r_{\text{areal}}^k}, \quad (3.32a)$$

$$\phi(t_{\text{ret}}, r_{\text{areal}}) \simeq \sum_{k=0}^N \frac{\phi_{(k)}(t_{\text{ret}})}{r_{\text{areal}}^k}. \quad (3.32b)$$

The asymptotic waveform is fully described by $A_{(0)}$ and $\phi_{(0)}$.

3.4.2 Results

We show the results of extrapolating Newman–Penrose data in Figs. 3.3 and 3.4. The first plot shows the relative amplitude difference between waveforms extrapolated with different orders of extrapolating polynomial; the second plot shows the phase difference for the same data. We see fair convergence for the orders shown here, though there are clearly problems with noise, and other features near merger. Ignoring high-frequency features, we estimate that the $N = 4$ waveform is correct to within roughly 1% in amplitude, or 0.01 radians throughout all times shown here.

For comparison, we also show the result of extrapolation without the correction for dynamical lapse of Eq. (3.14), in Fig. 3.5. As noted earlier, a sharp gauge pulse (roughly a 1% change in lapse over a time span of $20M_{\text{irr}}$) occurs near merger, which is denoted by the dotted vertical line. Without the correction, we see non-convergence of the extrapolated waveform after the merger; the difference between successive orders becomes larger with higher order. We are led to the conclusion that the correction is crucial to convergent extrapolation for the gauge used in these simulations.




Figure 3.3: **Convergence of the amplitude of the extrapolated Ψ_4 , with increasing order of the extrapolating polynomial, N**

This figure shows the convergence of the relative amplitude of the extrapolated Newman–Penrose waveform, as the order N of the extrapolating polynomial is increased. (See Eq. (3.32).) That is, we subtract the amplitudes of the two waveforms, and normalize at each time by the amplitude of the second waveform. We see that increasing the order tends to amplify the apparent noise during the early and late parts of the waveform. Nonetheless, the broad (low-frequency) trend is towards convergence. The dotted vertical line denotes merger, defined as the time of peak amplitude of the wave. Also note that the differences decrease as the system nears merger; this is a first indication that the extrapolated effects are due to near-field influences. Compare the convergence of h extrapolation in Fig. 3.6.




Figure 3.4: **Convergence of the phase of the extrapolated Ψ_4 , with increasing order of the extrapolating polynomial, N**

This figure is much the same as Fig. 3.3, except that it shows the convergence of phase. Again, increasing the extrapolation order tends to amplify the apparent noise during the early and late parts of the waveform, though the broad (low-frequency) trend is towards convergence. The dotted vertical line denotes merger, defined as the time of peak amplitude of the wave. Compare the convergence of h extrapolation in Fig. 3.7.




Figure 3.5: **Convergence of the phase of Ψ_4 , extrapolated with no correction for the dynamic lapse**

This figure is just the same as Fig. 3.4, except that no correction is done to account for the dynamic lapse. (See Eq. (3.14) and surrounding discussion.) Observe that the convergence is very poor after merger (indicated by the dotted vertical line). This corresponds to the time after which a sharp pulse in the lapse is observed. We conclude from this graph and comparison with the previous graph that the correction is crucial to convergence of Ψ_4 extrapolation through merger and ringdown.

Finally, we also show the results of extrapolation of Regge–Wheeler–Zerilli data, in Figs. 3.6 and 3.7. Convergence here is worse than in the Newman–Penrose case, and is completely absent after merger. This suggests that the wave does not propagate nicely as a function of t_{ret} , or does not fall off in powers of $1/r$. It is entirely possible that gauge effects are responsible—we currently have no estimate for the size of gauge effects in our RWZ data.¹¹ Alternatively, it is possible that near-field effects are all removed by $N = 1$ or $N = 2$ extrapolation, and going to higher orders simply allows the polynomials to fit to numerical noise. More work needs to be done improving extraction of these data, or—perhaps—improving the extrapolation. Nevertheless, for the rest of this paper we will treat the ($N = 4$) RWZ data as being trustworthy before the merger, at the level of roughly 4% in amplitude, and 0.04 radians in phase, as judged from the convergence plots. These are smaller errors than the actual difference between the extrapolated waveform and the waveforms extracted at finite radius. This suggests that even in this case, there are real benefits to be gained by extrapolating.

Another way to check the accuracy of the extrapolated waveforms is to differentiate the h data twice in time, and compare to the Ψ_4 data. When we do this we find agreement to within the 4% and 0.04 radians suggested above. Because extraction occurs in two very different ways, this gives us confidence that details of the extraction method are not polluting our final results—before merger, at least.

¹¹It has been suggested that second-order corrections to RWZ functions could provide error estimates, while also improving extraction. See Refs. [146, 83, 203], and especially [147] for discussion.




Figure 3.6: **Convergence of the amplitude of the extrapolated h , with increasing order of the extrapolating polynomial, N**

This figure shows the convergence of the relative amplitude of the extrapolated Regge–Wheeler–Zerilli waveform, as the order N of the extrapolating polynomial is increased. Note the poor convergence at higher orders, and the lack of convergence after merger (denoted by the dotted vertical line).




Figure 3.7: **Convergence of the phase of the extrapolated h , with increasing order of the extrapolating polynomial, N**

This figure shows the convergence of the phase of the extrapolated Regge–Wheeler–Zerilli waveform, as the order N of the extrapolating polynomial is increased. As in Fig. 3.6, we see poor convergence—or even a lack of convergence beyond $N = 3$ —before merger and complete lack of convergence after merger for any order. Presumably, these indicate the presence of effects in the data which we have not suggested, and which do not fall off in powers of $1/r$ or do not propagate as functions of t_{ret} .

3.4.3 Choosing the order of extrapolation

Deciding on an appropriate order of extrapolation to be used for a given purpose requires balancing competing effects. As we see in Fig. 3.3, for example, there is evidently some benefit to be gained from using higher-order extrapolation; there is clearly some convergence for each of the orders shown. On the other hand, higher-order methods amplify the apparent noise in the waveform.¹²

The optimal order depends on the accuracy needed. For some applications, little accuracy is needed, so a low-order extrapolation (or even no extrapolation) is preferable. If high-frequency noise is not considered an issue, then simple high-order extrapolation should suffice. Of course, if both high accuracy and low noise are required, data may easily be filtered, mitigating the problem of noise amplification.¹³ There is some concern that this may introduce subtle inaccuracy; filtering is more art than science, and it is difficult to establish precise error bars for filtered data.

Blending

Another consideration is the part of the physical waveform from which the data are collected. Because near-field effects scale inversely with powers of the fre-

¹²So-called “junk radiation” is a ubiquitous feature of initial data for current numerical simulations of binary black-hole systems. It is clearly evident in simulations as large-amplitude, high-frequency waves which die out as the simulation progresses. While it is astrophysically extraneous, it is nevertheless a real and fully resolved artifact of the initial data. Better initial data would, presumably, decrease its magnitude. This is the source of what looks like noise in the waveforms at early times. It is less apparent in h data than in Ψ_4 data because Ψ_4 effectively amplifies high-frequency components, due to the relation $\Psi_4 \approx -\ddot{h}$.

¹³For example, the MATLAB function `filtfilt`, with a low-pass sixth-order Butterworth filter with cutoff frequency just above the noise frequency applied to the early parts of both the input and output data is satisfactory, reducing the apparent noise by more than an order of magnitude. This filtering may be applied to either the complex data, or to its amplitude and phase, which allows for a lower cutoff frequency.

quency, the accuracy of a given order of extrapolation will improve as the system nears merger, while noise during and after merger will be a greater problem for higher-order methods. These factors combine to dictate different optimal extrapolation orders for different portions of the waveform. For example, fourth-order extrapolation may be necessary during inspiral, whereas second-order extrapolation is sufficient and less noisy during merger and ringdown. It is desirable to combine the different portions—fourth-order data for inspiral, and second-order for merger and ringdown—into a single waveform.

This blending can be achieved by a simple linear transition function, for example. We may define the amplitude and phase of the blended waveform by

$$A_{\text{blend}}(t) = \tau(t) A_{N=2} + [1 - \tau(t)] A_{N=4} , \quad (3.33a)$$

$$\phi_{\text{blend}}(t) = \tau(t) \phi_{N=2} + [1 - \tau(t)] \phi_{N=4} , \quad (3.33b)$$

where the blending function τ is given by

$$\tau(t) = \begin{cases} 0 & \text{if } t < t_1, \\ \frac{t-t_1}{t_2-t_1} & \text{if } t_1 \leq t < t_2, \\ 1 & \text{if } t_2 \leq t. \end{cases} \quad (3.34)$$

Here t_1 and t_2 are chosen for the desired accuracy and noise level. We separate the two times so that the transition between the two waveforms is more gradual, and the result is smoother.

3.4.4 Choosing extraction radii

Another decision needs to be made regarding the number and location of extraction surfaces. Choosing the number is fairly easy, because there is typically little cost in increasing the number of extraction radii (especially relative to the cost of—say—running a simulation). The only restriction is that the number of data points needs to be larger than the order of the extrapolating polynomial; more can hardly hurt. More careful consideration needs to be given to the location of the extraction surfaces.

For the extrapolations shown in Figs. 3.3 and 3.4, data was extracted on spheres spaced by roughly $10M_{\text{irr}}$, from $r = 75M_{\text{irr}}$ to $r = 225M_{\text{irr}}$. The outer radius of $225M_{\text{irr}}$ was chosen simply because this is the largest radius at which data exists throughout the simulation. In choosing the inner radius, there are two competing motives.

On one hand, we want the largest spread possible between the inner and outer extraction radii to stabilize the extrapolation. A rough rule of thumb [218] says that the distance to be extrapolated should be no greater than the distance covered by the data. Because the extrapolating polynomial is a function of $1/r$, the distance to be extrapolated is $1/r_{\text{outer}} - 1/\infty = 1/r_{\text{outer}}$. The distance covered by the data is $1/r_{\text{inner}} - 1/r_{\text{outer}}$, so if the rule of thumb is to be satisfied, the inner extraction radius should be no more than half of the outer extraction radius.

On the other hand, we would like the inner extraction radius to be as far out as possible. Extracting data near the violent center of the simulation is a bad idea for many reasons. Coordinate ambiguity, tetrad errors or nonlinear gauge effects, near-field effects—all are more severe near the center of the simulation.

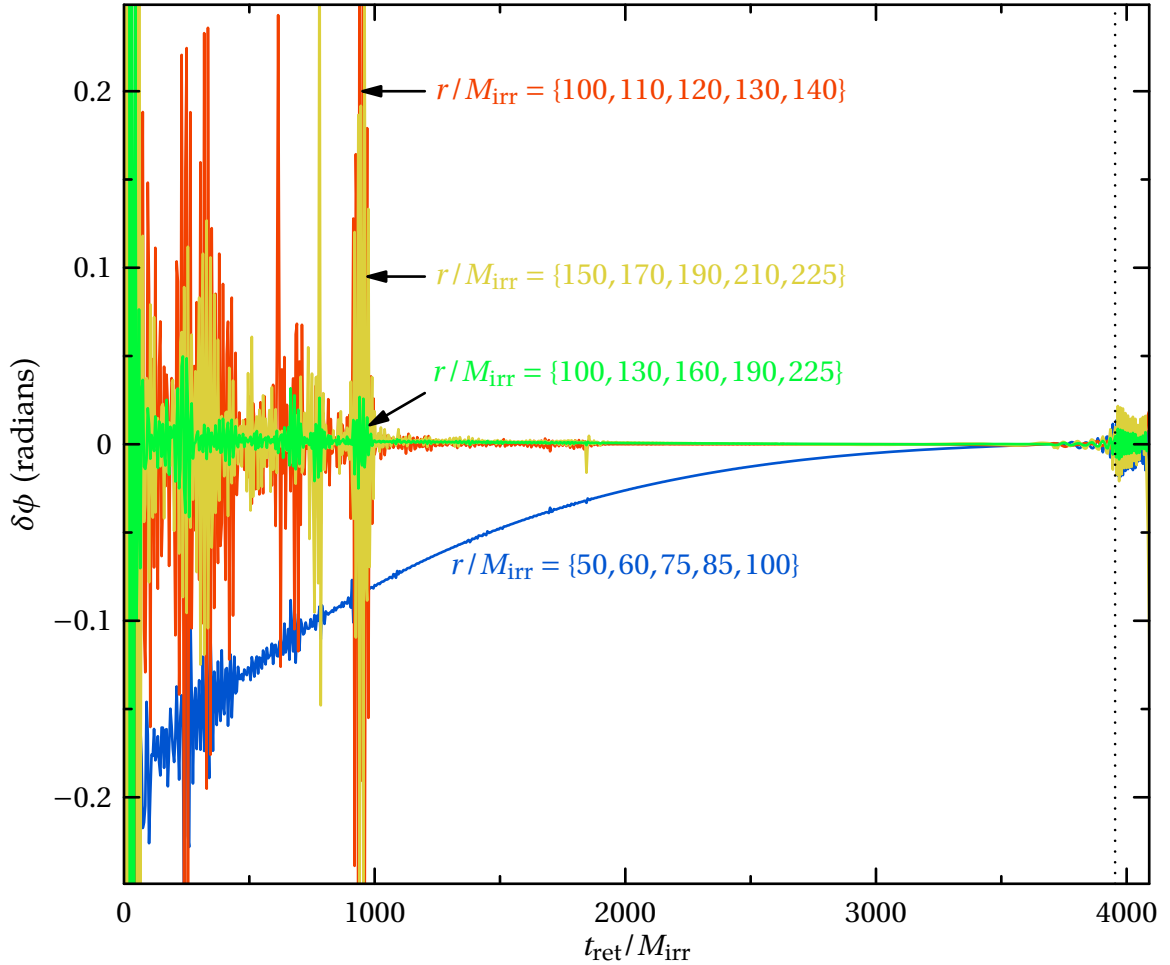


Figure 3.8: **Comparison of extrapolation of Ψ_4 using different sets of extraction radii**

This figure compares the phase of waveforms extrapolated with various sets of radii. All comparisons are with respect to the data set used elsewhere in this chapter, which uses extraction radii $r/M_{\text{irr}} = \{75, 85, 100, 110, 120, \dots, 200, 210, 225\}$. The order of the extrapolating polynomial is $N = 3$ in all cases. Roughly the same behavior is seen for h .

The larger these errors are, the more work the extrapolation needs to do. This effectively means that higher-order extrapolation is needed if data are extracted at low radii. The exact inner radius needed for extrapolation depends on the desired accuracy and, again, the part of the simulation from which the waveform is needed.

We can compare data extrapolated using different sets of radii. Fig. 3.8 shows a variety, compared to the data used elsewhere in this chapter. The extrapolation order is $N = 3$ in all cases. Note that the waveforms labeled $r/M_{\text{irr}} = \{50, \dots, 100\}$ and $r/M_{\text{irr}} = \{100, \dots, 225\}$ both satisfy the rule of thumb that the inner radius should be at most half of the outer radius, while the other two waveforms do not; it appears that violation of the rule of thumb leads to greater sensitivity to noise. One waveform is extrapolated using only data from small radii, $r/M_{\text{irr}} = \{50, \dots, 100\}$. It is clearly not converged, and would require higher-order extrapolation if greater accuracy is demanded. The source of the difference is presumably the near-field effect discussed in the next section, which is proportionally larger at small radii.

3.5 Near-field effects in the data

We can calculate the approximate near-field effects expected from a binary, in terms of the asymptotic field. This allows us to use the extrapolated waveforms to estimate the near-field effects. Comparing these estimates to the observed difference between extrapolated waveforms and waveforms extracted at finite radius will allow us to test our calculations. This will show that the field at a finite radius really is essentially different from the asymptotic field, even if we ignore

coordinate effects and various other “delicate issues” in wave extraction [185]. While our near-field calculation only gives an order-of-magnitude estimate of the error in Ψ_4 , it gives a much more accurate accounting for the error in h .

3.5.1 Calculating near-field effects

The familiar electrodynamics problem of finding the field of an idealized oscillating electric dipole provides a simple prototype for near-field effects. If the frequency of oscillation is ω , the exact solution is easily found, and expressed as a quadratic polynomial in $1/r\omega$, where r is the distance to the dipole [246, 172]. The solution is generally divided into two regions, where each of a pair of opposing approximations is valid: the near zone with $r\omega \ll 1$, and the far zone (or radiation zone) with $r\omega \gg 1$. In the near zone, the field includes contributions that will not be found at larger radii. The rough boundary between the two zones depends, of course, on the frequency of the source. At lower frequencies, the near zone extends farther out. Though the analogy is clearly imperfect, radiation from a black-hole binary shows similar signs of dependence on the source’s frequency [242]. In the remainder of this section, we will explicitly calculate those near-field effects, and estimate their size in simulations.

Calculating near-field effects in h

We define the symbol $h^{\alpha\beta} := \eta^{\alpha\beta} - \sqrt{-g} g^{\alpha\beta}$, where $g := \det(g_{\alpha\beta})$. In harmonic coordinates, in the linear approximation, this quantity obeys a wave equation. Explicitly, in vacuum, we have

$$\square h^{\alpha\beta} = 0, \tag{3.35}$$

where \square is just the flat-space wave operator [242]. Assuming that solutions to this equation can be expressed in the form of a multipolar expansion, we can write down the general outgoing solution as¹⁴

$$h^{\alpha\beta}(t, r, \vartheta, \varphi) = \sum_{l=0}^L \partial_{I_l} \left[\frac{1}{r} Q^{\alpha\beta I_l}(t-r) \right], \quad (3.36)$$

where each $Q^{\alpha\beta I_l}$ is just some multipole tensor, I_l is a spatial multi-index with l Cartesian components (that is, $I_l := i_1 \cdots i_l$), and summation over the multi-index is implied. Also note the obvious notation $\partial_{I_l} := \partial_{i_1} \cdots \partial_{i_l}$. Though it is not made explicit here, the right side implicitly depends on ϑ and φ , in that the differential operators depend on those coordinates, making the dependence on ϑ and φ from the left side nontrivial. For technical reasons, we must assume the existence of some L , representing the maximum moment of interest. Its value, of course, is arbitrary, so this should be of no practical importance.

Obviously, the uncontracted indices of the multipole components must have the same symmetries as the tensors they represent. For instance, the $Q^{\alpha\beta I_l}$ must be symmetric in α and β . As for the multi-index, we may as well suppose that the I_l are symmetric, because the partial derivatives commute, so only the symmetric part will be used. Also, we can impose the condition that the multipole be trace-free on each pair of indices in I_l . This makes the $Q^{\alpha\beta I_l}$ unique, and allows us to interpret each l as corresponding—as usual—to the weight of an irreducible representation of the rotation group. That is, each term in the sum in

¹⁴Sachs [227] and Pirani [213] introduced expressions of this form for various solutions to the flat-space wave equation. Thorne [242] gives the same expression in a different form; his is the decomposition of the free indices into time–time, time–space, and space–space components, each written as a pure-trace, a pure-divergence, etc., coupled to symmetric, trace-free tensors. Finally, Blanchet and Damour [44] proved that this solution encompasses all possible solutions, subject to a few basic conditions.

Eq. (3.36) corresponds to the usual 2^l -pole moment. (This interpretation of the rank- l symmetric, trace-free tensors is discussed, for example, by Pirani [213] and Thorne [242].)

The (nonrestrictive) assumption that the multipole tensor is trace-free is crucial to simplifying the differentiation in Eq. (3.36). In this case, we can show¹⁵ that the correct formula is

$$h^{\alpha\beta} = \sum_{l=0}^L h_{(l)}^{\alpha\beta} \quad \text{where} \quad h_{(l)}^{\alpha\beta} = \sum_{k=0}^l \frac{(-1)^l (l+k)!}{2^k k! (l-k)!} \frac{1}{r^{k+1}} {}^{(l-k)}Q^{\alpha\beta I_l}(t-r) N_{I_l}. \quad (3.37a)$$

Here, $N_{I_l} := n_{i_1} \cdots n_{i_l}$ is a collection of unit normal vectors, and ${}^{(m)}f(\tau) := \frac{d^m}{d\tau^m} f(\tau)$. (This expression is far more complicated if we do not assume that the multipole moment is symmetric and trace-free on the multi-index.) In particular, for the $l=2$ component, we have

$$h_{(l=2)}^{\alpha\beta} = \frac{1}{r} \left[\frac{d^2}{dt^2} + \frac{3}{r} \frac{d}{dt} + \frac{3}{r^2} \right] Q^{\alpha\beta ij}(t-r) n_i n_j. \quad (3.37b)$$

The first term in braces corresponds to the quantity obtained by extrapolation, and is the vastly dominant component of any gravitational waves from any equal-mass binary inspiral that will likely be detected in Earth's neighborhood.

We can use the last formula to estimate the size of the near-field terms in the $l=2$ component. Because we do not know (nor do we particularly care about) all the components of the multipole $Q^{\alpha\beta ij}$, let us focus on the asymptotic part of h :

$$h_{(l=2)}^{\text{asy}} = \frac{1}{r} (\varepsilon_+ - i\varepsilon_\times)_{\alpha\beta} n_i n_j \frac{d^2}{dt^2} Q^{\alpha\beta ij}(t-r), \quad (3.38)$$

¹⁵The formula is verified by double induction on l and k , dropping terms involving the Kronecker delta because of the trace-free property of the multipole with which they are contracted. The components are assumed to be Cartesian, meaning that there is no complication from covariant differentiation to be concerned about.

where the ε_+ and ε_\times tensors were defined in Eq. (3.23). In terms of this, we can rewrite Eq. (3.37b) as

$$h_{(l=2)} = \left[1 + \frac{3}{r} \int dt + \frac{3}{r^2} \iint dt dt' \right] h_{(l=2)}^{\text{asy}}, \quad (3.39)$$

interpreting the integrals as operators acting on $h_{(l=2)}^{\text{asy}}$. Performing these integrations requires setting one or two complex integration constants. Doing this reliably is difficult. Instead of attempting to actually do that integration, then, we will make a simple approximation. If $h_{(l=2)}^{\text{asy}}$ can be approximated as $h_{(l=2)}^{\text{asy}} \approx A e^{i\omega t}$, for real constants A and ω , we can re-express the above as

$$h_{(l=2)} \approx \left[1 - \frac{3i}{r\omega} - \frac{3}{r^2\omega^2} \right] h_{(l=2)}^{\text{asy}}. \quad (3.40)$$

Here we directly see the near-field terms behaving as powers of $1/r\omega = \lambda/r$.

As we will see in subsequent chapters, comparisons of gravitational waves generally decompose the wave into (time-varying) amplitude and phase. Because the second term in this formula is purely imaginary, it contributes—for small values of $1/r\omega$ —almost solely to the phase of the waveform, rather than its amplitude. Specifically, the $l = 2$ component of h measured at a finite radius will differ from the asymptotic waveform in phase by roughly $-3/r\omega$ radians. The fractional difference in amplitude will be roughly $3/2r^2\omega^2$ according to this formula. However, that is a second-order effect, and we will see the amplitude difference is poorly predicted by this formula.

Nonetheless, there is some insight to be gained from this approximate formula. Note the different behavior with radius, indicating that amplitude errors can be mitigated rapidly by moving to a larger extraction radius (or are particularly problematic when extracting at small radii), while phase errors diminish more

slowly. This can inform our choice of radii at which to extract data for a desired amount of precision.

Dominant gravitational-wave frequencies in numerical simulations are typically no more than about $1/2M$, and some times as small as $1/30M$ in the early inspiral segment of current numerical simulations. Extraction radii in current simulations are of order $100M$, meaning that phase errors induced by near-field effects may be a full cycle in the early inspiral, falling to a fraction of that near merger. Because precise phasing is crucial to successful detection via matched filtering, this effect can potentially spoil the accuracy of numerical waveforms, and needs to be accounted for.

Note that the assumptions that have gone into the derivation above include the assumption of linearity in the metric perturbation, and the assumption that the coordinates are harmonic. Nonlinearities and non-harmonic features of the coordinates will act as sources in Eq. (3.35). Assuming that our h data is reasonably gauge invariant, of course, the assumption of harmonic coordinates may be less important [2].

Calculating near-field effects in Ψ_4

A very similar analysis directly shows just the same near-zone behavior for Ψ_4 , without using the approximation $\Psi_4 \approx -\ddot{h}$ everywhere. The crucial assumption in this argument is that the background curvature and the length scales on which it changes are much larger than the reduced wavelength. In that case, it is not hard to show that the Riemann tensor obeys a homogeneous wave equation in

vacuum [57]:

$$\square R_{\alpha\beta\gamma\delta} = 0 . \quad (3.41)$$

The extra indices simply carry through into the expression for the general outgoing solution:

$$R_{\alpha\beta\gamma\delta}(t, r, \vartheta, \varphi) = \sum_{l=0}^L \partial_{I_l} \left[\frac{1}{r} Q_{\alpha\beta\gamma\delta}{}^{I_l}(t-r) \right] , \quad (3.42)$$

where the multi-index is, again, fully symmetric and trace-free. The differentiation follows in the same way as for Eq. (3.37a):

$$R_{\alpha\beta\gamma\delta} = \sum_{l=0}^L \sum_{k=0}^l \frac{(-1)^l (l+k)!}{2^k k! (l-k)!} \frac{1}{r^{k+1}} {}^{(l-k)} Q_{\alpha\beta\gamma\delta}{}^{I_l}(t-r) N_{I_l} . \quad (3.43a)$$

The expression for the $l=2$ component looks exactly the same as before:

$$R_{\alpha\beta\gamma\delta}^{(l=2)} = \frac{1}{r} \left[\frac{d^2}{dt^2} + \frac{3}{r} \frac{d}{dt} + \frac{3}{r^2} \right] Q_{\alpha\beta\gamma\delta}{}^{ij}(t-r) n_i n_j . \quad (3.43b)$$

Again, we contract this equation and extract the asymptotic field

$$\Psi_4^{\text{asy}} = \frac{1}{r} l^\alpha \bar{m}^\beta l^\gamma \bar{m}^\delta n_i n_j \frac{d^2}{dt^2} Q_{\alpha\beta\gamma\delta}{}^{ij}(t-r) , \quad (3.44)$$

rewrite the above as

$$\Psi_4 = \left[1 + \frac{3}{r} \int dt + \frac{3}{r^2} \iint dt dt' \right] \Psi_4^{\text{asy}} , \quad (3.45)$$

and approximate as before:

$$\Psi_4 \approx \left[1 - \frac{3i}{r\omega} - \frac{3}{r^2\omega^2} \right] \Psi_4^{\text{asy}} . \quad (3.46)$$

Now, we did not directly integrate in Eq. (3.39) (the equivalent of Eq. (3.45)) because of difficulties in setting the integration constants that would be needed.

In this case, however, if we have also extracted h , then we can use the linear relation $\Psi_4 \approx -\ddot{h}$. Because these are asymptotic quantities, the linear approximation should hold precisely. So, we can also rewrite the above as

$$\Psi_4 = \Psi_4^{\text{asy}} - \frac{3}{r} \dot{h}^{\text{asy}} - \frac{3}{r^2} h^{\text{asy}} . \quad (3.47)$$

A useful check of the validity of the approximation made in going from Eq. (3.45) to Eq. (3.46) is to compare that result with the result of Eq. (3.47). We have done this, and find that there is very little difference—roughly within the uncertainty estimated above—until just before merger.

Note that the only approximation we have made to arrive at Eq. (3.47) is that Riemann obeys a homogeneous wave equation, which relies on the assumption that the reduced wavelength λ is much less than the radius of curvature of the background spacetime \mathcal{R} and less than the length scale on which it changes \mathcal{L} . At a distance r from a Schwarzschild black hole, $\mathcal{R} \approx \sqrt{r^3/M}$ and $\mathcal{L} \lesssim r$, so we are assuming roughly $\lambda \ll r$.

3.5.2 Measuring near-field effects

We can now compare the extrapolated waveforms with data extracted at finite radius. The difference will show us how sensitive the extraction method is to the near-field effects just discussed, and to the effects of gauge discussed in Sec. 3.2. It will also give us an idea of the error we would make by using data extracted at finite radius rather than extrapolating. Figures 3.9 and 3.10 show as solid lines the relative amplitude and phase differences of h extracted from the simulation directly, compared to the data extrapolated with a fourth-order polynomial. The




Figure 3.9: **Relative amplitude difference between h data extracted at finite radius and data extrapolated to infinity, and calculated near-field effects**

For each extraction radius, we show the relative amplitude difference between the extracted and extrapolated data as solid lines. The dashed lines show the near-field effects, calculated using the extrapolated waveforms along with Eq. (3.40).




Figure 3.10: **Phase difference between h data extracted at finite radius and data extrapolated to infinity, and calculated near-field effects**

For each extraction radius, we show the phase difference between the extracted and extrapolated data as solid lines. The dashed lines show the near-field effects, calculated using the extrapolated waveforms along with Eq. (3.40).

dashed lines show the calculated near-field effects, which are given in terms of the asymptotic fields by Eq. (3.40).

The first feature we note is the size of the errors we would make if we were to treat waves extracted at finite radius as the asymptotic waveform. With extraction at $50 M_{\text{irr}}$, the measured amplitude difference begins at over 80%, falling to zero as the binary nears merger. These differences fall off roughly as $1/r^2$; the values at a given time decrease by a factor of roughly 4 between successive lines. For extraction at $r = 200 M_{\text{irr}}$, the differences from the extrapolated waveform are never more than a few percent. The phase error scales inversely with the radius, and can be as much as 1.6 radians for the simulation shown here, with extraction at $50 M_{\text{irr}}$. Most importantly, the phase errors are changing during the simulation. A constant offset would be largely irrelevant for many purposes; a *changing* error means that, for example, coherence with a matched-filtering template would be lost.

We also observe that the near-field effect gives a very good estimate of the size of the difference between data extracted at finite radius and data extrapolated to infinity. Recalling from Sec. 4.6 that the RWZ extrapolation can only be trusted to roughly 4% in amplitude, or about 0.04 radians, the calculated near-field effect is, in fact, within the error bounds. From this, we can conclude that the dominant effect on waveforms extracted with the RWZ methods used here is the near-field effect. Note that the near-field effects are fairly low order in λ/r , implying that extrapolation with $N > 2$ is unnecessary—at least for the data presented here.

The same comparison is made for Ψ_4 data, in Figs. 3.11 and 3.12. Again, solid lines show the relative amplitude and phase differences of Ψ_4 extracted from

the simulation directly, compared to the data extrapolated with a fourth-order polynomial. The dashed lines show the calculated near-field effects, which are given in terms of the asymptotic fields by Eq. (3.47).

Curiously, we see that there is a much smaller relative amplitude error, as compared to RWZ data. With extraction at $50 M_{\text{irr}}$, the measured amplitude difference begins below -30% , goes to zero, and even switches sign, leveling off at over 10% near merger and during ringdown. These differences also seem to fall off roughly as $1/r^2$, and—for extraction at $r = 200 M_{\text{irr}}$ —are never more than a few percent. Again, the phase error scales roughly inversely with the radius, and can be as much as 1.4 radians for the data shown here, with extraction at $50 M_{\text{irr}}$.

Unfortunately, the successes of our near-field estimates for RWZ are not matched by our estimates for Ψ_4 data. We see that the estimated near-field effects reflect the general form of the errors, though they poorly describe the size. In fact, the amplitude estimate seems to be wrong by very nearly a factor of -2 . An improved choice of tetrad would not likely remove these effects, unless there is some subtle dependence of the tetrad on the waves' frequency; the time dependence of the errors suggests that we have the right functional form for the near-field effects with the wrong coefficients.

We do expect some error in our near-zone calculation. The crucial approximation was that the reduced wavelength is much less than the radius of curvature of the background spacetime, and the scale on which the background curvature changes. This approximation, of course, is not well satisfied early in the simula-




Figure 3.11: **Relative amplitude difference between Ψ_4 data extracted at finite radius and data extrapolated to infinity, and calculated near-field effects**

For each extraction radius, we show the relative amplitude difference between the extracted and extrapolated data as solid lines. The dashed lines show the near-field effects, calculated using the extrapolated waveforms along with Eq. (3.47), though Eq. (3.46) gives essentially the same result.

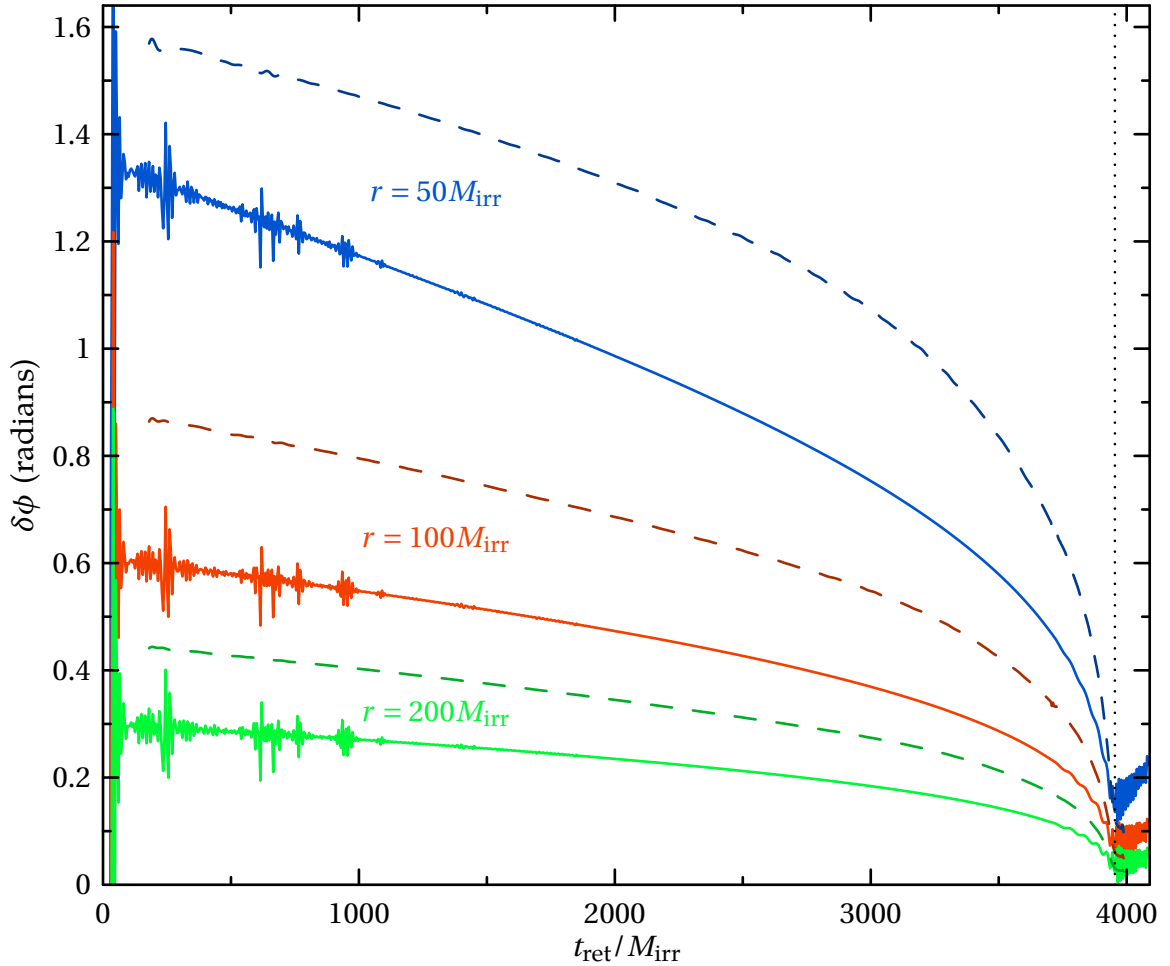


Figure 3.12: **Phase difference between Ψ_4 data extracted at finite radius and data extrapolated to infinity, and calculated near-field effects**

For each extraction radius, we show the phase difference between the extracted and extrapolated data as solid lines. The dashed lines show the near-field effects, calculated using the extrapolated waveforms along with Eq. (3.47), though Eq. (3.46) gives essentially the same result.

tion, close to the binary; the initial reduced wavelength of the gravitational waves is $\lambda \approx 30 M_{\text{irr}}$. On the other hand, it is not clear why this approximation would be so poor for the near-field effects in Riemann, while the linearity approximation for the near-field effects in h would be so good, by comparison.

3.6 Conclusions

We have developed a simple method for extrapolating gravitational-wave data from numerical simulations. We have discussed two methods commonly used to extract waveforms from simulations, and applied the extrapolation technique to those methods. Convergence tests indicate that the extrapolated waveforms are probably accurate to within 4% in amplitude and 0.04 radians in phase up to the merger for Regge–Wheeler–Zerilli data, and about 1/4 of that for Newman–Penrose data, including merger and ringdown. For reasons that remain unclear, extrapolation of the Regge–Wheeler–Zerilli waveform fails just after merger, though extrapolation of the Newman–Penrose waveform seems trustworthy well into the ringdown. We also compared our calculated near-field effects to the observed effects by comparing waveforms extracted at finite radius to the extrapolated waveforms. The calculated near-field effects seem to be very accurate for Regge–Wheeler–Zerilli data, but very poor for Newman–Penrose data. The lack of agreement in the latter case is an interesting problem for future work.

As with any type of extrapolation, a note of caution is naturally in order. It is entirely possible that the “true” function being extrapolated bears no resemblance to the approximating function we choose, outside of the domain on which we have data. We may, however, have reason to believe that the true function takes

a certain form. If the data in question are generated by a homogeneous wave equation, for instance, we know that well-behaved solutions fall off in powers of $1/r$. In any case, there is a certain element of faith that extrapolation is a reasonable thing to do. While that faith may indeed be misplaced, there are methods of checking whether or not it is: goodness-of-fit statistics, error estimates, and convergence tests. To be of greatest use, goodness-of-fit statistics and error estimates for the output waveform require error estimates for the input waveforms. We leave this for future work.

We still do not know the correct answers to the questions numerical relativity considers. Thus, large systematic errors could be hidden in plain view. To eliminate them, we need to use multiple, independent methods for our calculations. For example, we might extract Ψ_4 directly by calculating the Riemann tensor and contracting appropriately with our naive coordinate tetrad, and extract the metric perturbation using the formalism of Regge–Wheeler–Zerilli and Moncrief. By differentiating the latter result twice and comparing to Ψ_4 , we could verify that details of the extraction methods are not producing systematic errors. Nonetheless, it is possible that infrastructure used to find both could be leading to errors.

In the same way, simulations need to be performed using different gauge conditions, numerical techniques, code infrastructures, boundary conditions, and even different extrapolation methods. Only when multiple schemes arrive at the same result can we be truly confident in any of them. But to arrive at the same result, the waveforms from each scheme need to be processed as carefully as possible. We have shown that extrapolation—or more advanced techniques—is

crucial for highly accurate gravitational waveforms.

COMPARING NUMERICAL-RELATIVITY WAVEFORMS TO POST-NEWTONIAN APPROXIMATIONS¹

Numerical simulations of 15 orbits of an equal-mass binary black-hole system are presented. Gravitational waveforms from these simulations, covering more than 30 cycles and ending about 1.5 cycles before merger, are compared with those from quasi-circular zero-spin post-Newtonian (PN) approximations. The cumulative phase uncertainty of these comparisons is about 0.05 radians, dominated by effects arising from the small residual spins of the black holes and the small residual orbital eccentricity in the simulations. Matching numerical results to PN waveforms early in the run yields excellent agreement (within 0.05 radians) over the first ~ 15 cycles, thus validating the numerical simulation and establishing a regime where PN theory is accurate. In the last 15 cycles before merger, however, generic time-domain Taylor ap-

¹ This chapter is extracted with minor revisions from Ref. [59], which was written in collaboration with Duncan A. Brown, Lawrence E. Kidder, Abdul H. Mroué, Harald P. Pfeiffer, Mark A. Scheel, Gregory B. Cook, and Saul A. Teukolsky. I was responsible for extrapolating the data from the simulations, generating the post-Newtonian waveforms used in this paper, and establishing the matches between the waveforms, while Abdul checked my work with different methods. I also double-checked the error estimates, and shared in writing the text.

proximants build up phase differences of several radians. But—apparently by coincidence—one specific post-Newtonian approximant (TaylorT4 at 3.5PN order) agrees much better with the numerical simulations, having accumulated phase differences of less than 0.05 radians over the 30-cycle waveform. Gravitational-wave amplitude comparisons are also done between numerical simulations and post-Newtonian, and the agreement depends on the post-Newtonian order of the amplitude expansion: the amplitude difference is about 6–7% for 0.0PN order and becomes smaller for increasing order. A newly derived 3.0PN amplitude correction improves agreement significantly (<1% amplitude difference throughout most of the run, increasing to 4% near merger) over the previously known 2.5PN-amplitude terms.

4.1 Introduction

The last two years have brought tremendous progress in simulations of black-hole binaries, starting with the first stable simulation of orbiting and merging black holes [219, 221], development of the moving-puncture method [84, 21] and rapid progress by other groups [86, 158, 126, 230, 234, 69, 191, 134, 238]. Since then, an enormous amount of work has been done on the late inspiral and merger of black-hole binaries, among them studies of the universality of the merger waveforms [20, 19], investigations into black-hole kicks [22, 150, 183, 89, 149, 160, 233, 93, 82, 68, 18, 159, 158, 232] and spin dynamics [88, 87, 85], comparisons to post-Newtonian models [75, 6, 32], and applications to gravitational-wave data analysis [206, 78, 27].

Compared to the intense activity focusing on simulations close to merger,

there have been relatively few simulations covering the inspiral phase. To date, only three simulations [24, 23, 209, 157, 166] cover more than five orbits. Long inspiral simulations are challenging for a variety of reasons. First, the orbital period increases rapidly with separation, so that simulations must cover a significantly longer evolution time. In addition, the gravitational waveform must be extracted at larger radius (and the simulation must therefore cover a larger spatial volume) because the gravitational wavelength is longer. Furthermore, gravitational-wave data analysis requires small *absolute* accumulated phase uncertainties in the waveform, so the relative phase uncertainty of the simulation must be smaller.

Gravitational-wave detectors provide a major driving force for numerical relativity (NR). The first-generation interferometric gravitational-wave detectors, such as LIGO [26, 245], GEO600 [162], and VIRGO [4, 3], are now operating at or near their design sensitivities. Furthermore, the advanced generation of detectors are entering their construction phases. This new generation of interferometers will improve detector sensitivity by a factor of ~ 10 and hence increase expected event rates by a factor of ~ 1000 [143]. One of the most promising sources for these detectors is the inspiral and merger of binary black holes (BBHs) with masses $m_1 \sim m_2 \sim 10\text{--}20 M_\odot$ [139]. These systems are expected to have circularized long before their gravitational waves enter the sensitive frequency band of ground-based detectors [208].

A detailed and accurate understanding of the gravitational waves radiated as the black holes spiral towards each other will be crucial not only to the initial detection of such sources, but also to maximize the information that can be ob-

tained from signals once they are observed. When the black holes are far apart, the gravitational waveform can be accurately computed using a post-Newtonian (PN) expansion. As the holes approach each other and their velocities increase, the post-Newtonian expansion is expected to diverge from the true waveform. It is important to quantify any differences between theoretical waveforms and the true signals, as discrepancies will cause a reduction of search sensitivity. Several techniques have been proposed to address the problem of the breakdown of the post-Newtonian approximation [106, 76, 73], but ultimately, the accuracy of the post-Newtonian waveforms used in binary-black-hole gravitational-wave searches can only be established through comparisons with full numerical simulations.

Unfortunately, comparing post-Newtonian approximations to numerical simulations is not straightforward, the most obvious problem being the difficulty of producing long and sufficiently accurate numerical simulations, as explained above. In addition, post-Newtonian waveforms typically assume circular orbits, and most astrophysical binaries are expected to be on circular orbits late in their inspiral, so the orbital eccentricity within the numerical simulation must be sufficiently small.² Another factor that complicates comparisons is the variety of post-Newtonian approximants available, from several straightforward Taylor expansions to more sophisticated Padé resummation techniques and the effective one-body approach (see, e.g. [108, 110, 106, 76, 77, 114, 103, 105, 72], as well as Sec. 4.3.5 below). While all post-Newtonian approximants of the same order should agree sufficiently early in the inspiral (when neglected higher-order terms are small), they begin to disagree with each other during the late inspiral when

²Unfortunately, this circularization occurs on extremely long timescales [208]—thousands of orbits—making it impossible to run the numerical simulation long enough to radiate the eccentricity away.

the post-Newtonian approximation starts to break down—exactly the regime in which NR waveforms are becoming available.

Finally, agreement (or disagreement) between NR and PN waveforms will also depend very sensitively on the precise protocol used to compare the waveforms. Are PN and NR waveforms matched early or late in the inspiral? Is the matching done at a particular time, or is a least-squares fit performed over some region of the waveform? Does one compare frequencies $\omega(t)$ or phases $\phi(t)$? Are comparisons presented as functions of time, or of frequency? Up to which cutoff frequency does one compare PN with NR?

Despite these difficulties, several comparisons between NR and PN have been done for the last few orbits of an equal-mass, nonspinning black-hole binary. The first such study was done by Buonanno et al. [75] based on simulations performed by Pretorius [219] lasting somewhat more than 4 orbits (~ 8 gravitational-wave cycles). This comparison performs a least-squares fit over the full waveform, finds agreement between the numerical evolution and a particular post-Newtonian approximant (in our language TaylorT3 3.0/0.0³) and notes that another approximant (TaylorT4 3.5/0.0) will give similarly good agreement. However, as the authors note, this study is severely limited by numerical resolution, sizable initial eccentricity (~ 0.015), close initial separation of the black holes, and coordinate artifacts; for these reasons, the authors do not quantify the level of agreement.

More recently, Baker et al. [24, 23] performed simulations covering the last ~ 14 cycles before merger. These simulations have an orbital eccentricity ~ 0.008 [24],

³We identify post-Newtonian approximants with three pieces of information: the label introduced by [108] for how the orbital phase is evolved; the PN order to which the orbital phase is computed; and the PN order that the amplitude of the waveform is computed. See Sec. 4.3.5 for more details.

forcing the authors to use a fitted smooth (“de-eccentrized”) gravitational-wave phase to obtain a monotonically increasing gravitational-wave frequency. Comparing to TaylorT4 3.5/2.5, they find agreement between numerical and post-Newtonian gravitational-wave phase to within their numerical errors, which are about 2 radians. The authors also indicate that other post-Newtonian approximants do not match their simulation as well as TaylorT4, but unfortunately, they do not mention whether any disagreement is significant (i.e., exceeding their numerical errors). Pan et al. [206] performed a more comprehensive analysis of the numerical waveforms computed by Pretorius [75] and the Goddard group [24, 23], confirming that TaylorT4 3.5/0.0 matches the numerical results best.

The most accurate inspiral simulation to date was performed by the Jena group and presented in Husa et al. [166] and Hannam et al. [157]. This simulation covers 18 cycles before merger and has an orbital eccentricity of ~ 0.0018 [167]. Discarding the first two cycles, which are contaminated by numerical noise, and terminating the comparison at a gravitational-wave frequency $M\omega = -0.1$ (see Eq. (4.15) for the precise definition), their comparison extends over 13 cycles. We discuss the results of Ref. [157] in more detail in Sec. 4.6.1.

This paper presents a new inspiral simulation of a nonspinning equal-mass black-hole binary. This new simulation more than doubles the evolution time of the simulations in Refs. [24, 23, 157, 166], resulting in a waveform with 30 gravitational-wave cycles, ending ~ 1.5 cycles before merger, and improves numerical-truncation errors by one to two orders of magnitude over those in Refs. [24, 23, 157, 166]. The orbital eccentricity of our simulations is $\sim 6 \times 10^{-5}$; this low eccentricity is achieved using refinements of techniques described in Ref. [209]. We

present a detailed analysis of various effects which might influence our comparisons to post-Newtonian waveforms for nonspinning black-hole binaries on circular orbits. These effects result in an uncertainty of ~ 0.05 radians out of the accumulated ~ 200 radians. Perhaps surprisingly, the largest uncertainty arises from the residual orbital eccentricity, despite its tiny value. The second largest effect arises due to a potential residual spin on the black holes, which we bound by $|S|/M_{\text{irr}}^2 < 5 \times 10^{-4}$.

We compare the numerical waveforms with four different time-domain post-Newtonian Taylor-approximants [108, 110, 73] and we match PN and NR waveforms at a specific time during the inspiral. We explore the effects of varying this matching time. When matching ~ 9 cycles after the start of our evolution, all post-Newtonian approximants of 3.0PN and 3.5PN order in orbital phase agree with our simulation to within ~ 0.03 radians over the first 15 cycles. This agreement is better than the combined uncertainties of the comparison, thus validating our simulations in a regime where the 3.5PN truncation error of post-Newtonian theory is comparable to the accuracy of our simulations. Lower-order post-Newtonian approximants (2.0PN and 2.5PN order), however, accumulate a significant phase difference of ~ 0.2 radians over this region.

Extending the comparison toward merger (as well as when matching closer to merger), we find, not surprisingly, that the agreement between PN and NR at late times depends strongly on exactly what post-Newtonian approximant we use [108, 110]. Typical accumulated phase differences are on the order of radians at frequency $M\omega = -0.1$. One particular post-Newtonian approximant, TaylorT4 at 3.5PN order in phase, agrees with our NR waveforms far better than the other

approximants, the agreement being within the phase uncertainty of the comparison (0.05 radians) until after the gravitational-wave frequency passes $M\omega = -0.1$ (about 3.5 cycles before merger). It remains to be seen whether this agreement is fundamental or accidental, and whether it applies to more complicated situations (e.g., unequal masses, nontrivial spins).

We also compare the post-Newtonian gravitational-wave amplitude to the numerical amplitude, where we estimate the uncertainty of this comparison to be about 0.5%. Restricted waveforms (i.e., 0PN order in the amplitude expansion) are found to disagree with the numerical amplitudes by 6–7%. An amplitude expansion of order 2PN shows significantly better agreement than the expansion at order 2.5PN. A newly derived 3PN amplitude [178] is found to give much better agreement than the 2.0PN amplitude.

This paper is organized as follows: Section 4.2 discusses our numerical techniques. In particular, we describe how we construct binary-black-hole initial data, evolve these data for 15 orbits, extract gravitational-wave information from the evolution, and produce a gravitational waveform as seen by an observer at infinity. Section 4.3 details the generation of post-Newtonian waveforms, including details of how we produce the four approximants that we compare against NR. We describe our procedure for comparing NR and PN waveforms in Sec. 4.4, and present a detailed study of various sources of uncertainty in Sec. 4.5. The comparisons between NR and PN are presented in Sec. 4.6. This section is split into two parts. First, we compare each PN approximant separately with the numerical simulation. Second, we show some additional figures which facilitate comparisons between the different PN approximants. Finally, we present some

concluding remarks in Sec. 4.7. The impatient reader primarily interested in PN–NR comparisons may wish to proceed directly to Table 4.3, summarizing the uncertainties of our comparisons, and then continue to Sec. 4.6, starting with Fig. 4.19.

4.2 Generation of numerical waveforms

In order to perform a quantitative comparison between numerical and post-Newtonian waveforms, it is important to have a code capable of starting the black holes far enough apart to be in a regime where we believe the post-Newtonian approximation is valid, track the orbital phase extremely accurately, and do so efficiently, so that the simulation can be completed in a reasonable amount of time. Furthermore, the gravitational waves from such a simulation must be extracted in a manner that preserves the accuracy of the simulation and predicts the waveform as seen by a distant observer, so that a comparison with the post-Newtonian waveform can be made. In this section we describe the techniques we use to do this, as well as the results of a simulation starting more than 15 orbits prior to merger.

When discussing numerical solutions of Einstein’s equations, we write all dimensionful quantities in terms of some mass scale M , which we choose to be the sum of the irreducible masses of the two black holes in the initial data:

$$M = M_{\text{irr},1} + M_{\text{irr},2} . \quad (4.1)$$

The irreducible mass of a single hole is defined as

$$M_{\text{irr}} := \sqrt{A/16\pi} , \quad (4.2)$$

where A is the surface area of the event horizon; in practice we take A to be the surface area of the apparent horizon. More generally, it is more appropriate to use the Christodoulou mass of each black hole,

$$M_{\text{Chr}}^2 = M_{\text{irr}}^2 + \frac{S^2}{4M_{\text{irr}}^2}, \quad (4.3)$$

instead of the irreducible mass. Here S is the spin of the hole. However, for the case considered in this paper, the spins are sufficiently small that there is little difference between M_{Chr} and M_{irr} .

4.2.1 Initial data

Initial data are constructed within the conformal thin sandwich formalism [253, 212] using a pseudo-spectral elliptic solver [210]. We employ quasi-equilibrium boundary conditions [96, 97] on spherical excision boundaries, choose conformal flatness and maximal slicing, and use Eq. (33a) of Ref. [90] as the lapse boundary condition. The spins of the black holes are made very small via an appropriate choice of the tangential shift at the excision surfaces, as described in [90].

As the most accurate post-Newtonian waveforms available assume adiabatic inspiral of quasi-circular orbits, it is desirable to reduce the eccentricity of the numerical data as much as possible. Using techniques developed in [209], each black hole is allowed to have a nonzero initial velocity component towards the other hole. This small velocity component v_r and the initial orbital angular velocity Ω_0 are then fine-tuned in order to produce an orbit with very small orbital eccentricity.⁴ We have improved our eccentricity-reduction procedure since the

⁴An alternative method of producing low-eccentricity initial data, based on post-Newtonian

Table 4.1: Summary of the initial data sets

Name	d	Ω_0	f_r	$\nu_r \times 10^4$	$M\omega_0$	M_{ADM}/M	J_{ADM}/M^2	s_0/m	$e_{\text{ds}/dt}$
30a	30	0.0080108	0.939561	0.00	-0.01664793	0.992333	1.0857	17.37	1.0×10^{-2}
30b	30	0.0080389	0.939561	-4.90	-0.0167054	0.992400	1.0897	17.37	6.5×10^{-4}
30c	30	0.0080401	0.939561	-4.26	-0.0167081	0.992402	1.0898	17.37	5.0×10^{-5}
24a	24	0.0110496	0.92373	-8.29	-0.0231947	0.990759	1.0045	14.15	1.1×10^{-3}
24b	24	0.0110506	0.923739	-8.44	-0.0231967	0.990767	1.0049	14.15	1.5×10^{-4}

The first block of numbers (d , Ω_0 , f_r , and ν_r) represent raw parameters entering the construction of the initial data. The second block gives some properties of each initial data set: M denotes the sum of the irreducible masses, M_{ADM} and J_{ADM} the ADM energy and angular momentum, and s_0 the initial proper separation between the horizons. The last column lists the eccentricity computed from Eq. (4.7). The initial data set 30c is used for all evolutions (except for consistency checks) described in this paper.




Figure 4.1: **Eccentricity removal**

The quantity shown is the time derivative of the proper separation between the black holes in each of three simulations. This is shown for short evolutions of the $d = 30$ initial data sets 30a, 30b, and 30c (see Table 4.1). These three data sets represent zero through two iterations of our eccentricity-reduction procedure. The orbital eccentricity is reduced significantly by each iteration. The final curve (30c) corresponds to the simulation used in the rest of this paper.

version described in [209], so we summarize our new iterative procedure here:

We start with a quasi-circular (i.e., $v_r = 0$) initial data set at coordinate separation $d = 30$, where Ω_0 is determined by equating Komar mass with Arnowitt–Deser–Misner (ADM) mass [90]. We then evolve these data for about 1.5 orbits, corresponding to a time $t/M \approx 600$. From this short evolution, we measure the proper separation s between the horizons by integration along the coordinate axis connecting the centers of the black holes. We fit the time derivative ds/dt in the interval $100 \lesssim t/M \lesssim 600$ to the function

$$\frac{ds}{dt} = A_0 + A_1 t + B \cos(\omega t + \delta) , \quad (4.4)$$

where we vary all five parameters A_0, A_1, B, ω , and δ to achieve the best fit. The desired smooth inspiral is represented by the part $A_0 + A_1 t$; the term $B \cos(\omega t + \delta)$ corresponds to oscillations caused by orbital eccentricity.

For a *Newtonian* orbit with radial velocity $B \cos(\omega t + \delta)$ at initial separation s_0 , it is straightforward to determine the changes to the orbital frequency and the radial velocity which make the orbit perfectly circular, namely

$$\Omega_0 \rightarrow \Omega_0 + \frac{B \sin \delta}{2s_0} , \quad (4.5)$$

$$v_r \rightarrow v_r - \frac{B \cos \delta}{2} . \quad (4.6)$$

For Newtonian gravity, Eq. (4.6) will of course result in a circular orbit with $v_r = 0$. In general relativity, Ω_0 and v_r will be different from their Newtonian values, for instance $v_r < 0$ to account for the inspiral of the two black holes. Nevertheless, we assume that small perturbations around the zero-eccentricity

ideas, is developed in [167]. While that technique is computationally more efficient than ours, it merely reduces orbital eccentricity by a factor of ~ 5 relative to quasi-circular initial data, which is insufficient for the comparisons presented here. (see Sec. 4.5.5).

inspiral trajectory behave similarly to small perturbations around a Newtonian circular orbit. Therefore, we apply the same formulas, Eqs. (4.5) and (4.6), to obtain improved values for Ω_0 and ν_r for the black hole binary, where s_0 is the initial proper separation between the horizons. We then use the new values of Ω_0 and ν_r to construct a new initial-data set, again evolve for two orbits, fit to Eq. (4.4), and update Ω_0 and ν_r . We continue iterating this procedure until the eccentricity is sufficiently small.

We estimate the eccentricity for each iteration from the fit to Eq. (4.4) using the formula

$$e_{ds/dt} = \frac{B}{s_0 \omega} , \quad (4.7)$$

which is valid in Newtonian gravity for small eccentricities. Successive iterations of this procedure are illustrated in Fig. 4.1 and yield the initial-data sets 30a, 30b, and 30c summarized in Table 4.1. Eccentricity decreases by roughly a factor of 10 in each iteration, with 30c having $e_{ds/dt} \approx 5 \times 10^{-5}$. The evolutions used during eccentricity reduction need not be very accurate and need to run only for a short time, $t \sim 600M$. One iteration of this procedure at our second lowest resolution requires about 250 CPU-hours. For completeness, Table 4.1 also lists parameters for initial data at smaller separation; these data will be used for consistency checks below. Apart from these consistency checks, the remainder of this paper will focus exclusively on evolutions of the low-eccentricity initial data set 30c.

4.2.2 Evolution of the inspiral phase

The Einstein evolution equations are solved with the pseudo-spectral evolution code described in Ref. [230]. This code evolves a first-order representation [187]

of the generalized-harmonic system [141, 144, 220]. We handle the singularities by excising the black-hole interiors from our grid. Our outer boundary conditions [187, 223, 225] are designed to prevent the influx of unphysical constraint violations [237, 142, 25, 239, 80, 240, 180] and undesired incoming gravitational radiation [70], while allowing the outgoing gravitational radiation to pass freely through the boundary.

The code uses a fairly complicated domain decomposition to achieve maximum efficiency. Each black hole is surrounded by several (typically six) concentric spherical shells, with the inner boundary of the innermost shell (the excision boundary) just inside the horizon. A structure of touching cylinders (typically 34 of them) surrounds these shells, with axes along the line between the two black holes. The outermost shell around each black hole overlaps the cylinders. The outermost cylinders overlap a set of outer spherical shells, centered at the origin, which extend to a large outer radius. External boundary conditions are imposed only on the outer surface of the largest outer spherical shell. We vary the location of the outer boundary by adding more shells at the outer edge. Since all outer shells have the same angular resolution, the cost of placing the outer boundary farther away (at full resolution) increases only linearly with the radius of the boundary. External boundary conditions are enforced using the method of Bjorhus [34], while inter-domain boundary conditions are enforced with a penalty method [153, 161].

We employ the dual-frame method described in Ref. [230]: we solve the equations in an *inertial frame* that is asymptotically Minkowski, but our domain decomposition is fixed in a *comoving frame* that rotates with respect to the inertial frame

and also shrinks with respect to the inertial frame as the holes approach each other. The positions of the holes are fixed in the comoving frame; we account for the motion of the holes by dynamically adjusting the coordinate mapping between the two frames. Note that the comoving frame is referenced only internally in the code as a means of treating moving holes with a fixed domain. Therefore all coordinate quantities (e.g., black-hole trajectories, wave-extraction radii) mentioned in this paper are inertial-frame values unless explicitly stated otherwise.

One side effect of our dual-frame system is that the outer boundary of our domain (which is fixed in the comoving frame) moves inward with time as observed in the inertial frame. This is because the comoving frame shrinks with respect to the inertial frame to follow the motion of the holes. In Refs. [230, 209] the inertial-frame coordinate radius r (with respect to the center of mass) and the comoving coordinate radius r' are related by a simple scaling

$$r = a(t)r' . \quad (4.8)$$

The expansion parameter $a(t)$ is initially set to unity and decreases dynamically as the holes approach each other, so that the comoving-frame coordinate distance between the holes remains constant. The outer boundary of the computational grid is at a fixed comoving radius R'_{bdry} , which is mapped to the inertial-coordinate radius $R_{\text{bdry}}(t) = a(t)R'_{\text{bdry}}$. Because we wish to accurately compute the gravitational radiation as measured far from the holes, it is desirable to have a moderately large outer boundary ($R_{\text{bdry}}(t) \gtrsim 200M$) throughout the run. For the linear mapping, Eq. (4.8), this requires a very distant outer boundary early in the run, $R_{\text{bdry}}(0) \simeq 1000M$. Computationally this is not very

expensive. However, the initial junk radiation contaminates the evolutions for a time interval proportional to the light-crossing time to the outer boundary, and for $R_{\text{bdry}}(0) \simeq 1000M$ it would be necessary to discard a significant portion of the evolution.

We therefore use the mapping

$$r = \left[a(t) + (1 - a(t)) \frac{r'^2}{R_0'^2} \right] r' , \quad (4.9)$$

for some constant R_0' which is chosen to be roughly the radius of the outer boundary in comoving coordinates. This mapping has the following properties: (1) At the initial time $t = 0$, the map reduces to the identity map because $a(0) = 1$. Thus we do not need to re-map our initial data before evolving. (2) For small radii (e.g., at the locations of the black holes), the map reduces to the linear map, $r = a(t)r' + \mathcal{O}(r'^3)$. This allows use of the control system without modifications. (3) The moving radius $r' = R_0'$ is mapped to a *constant* inertial radius: $r(R_0') = R_0'$. This allows us to keep the inertial radius of the outer boundary constant (or nearly constant⁵) in time rather than shrinking rapidly.

In total, we have run three evolutions of the 30c initial data set; these use different combinations of outer-boundary radius and radial mapping between inertial and moving coordinates. Some properties of these evolutions are summarized in Table 4.2. We also performed extensive convergence testing, running the same evolution on up to six distinct resolutions, N1 to N6. The coarsest resolution 30c-1/N1 uses approximately 41^3 grid points (summing all grid points

⁵In practice, we choose R_0' somewhat larger than the outer boundary, so that the outer boundary of the computational domain slowly contracts in inertial coordinates. This makes the zero-speed characteristic fields *outgoing* there, avoiding the need to impose boundary conditions on those fields.

Table 4.2: Overview of low-eccentricity simulations

Name	ID	N_{orbits}	R_{bdry}	Radial map	Resolutions
30c-1	30c	15.6	462M	Eq. (4.9)	N1, N2, ..., N6
30c-2	30c	15.6	722M	Eq. (4.8)	N2, N4, N6
30c-3	30c	15.6	202M	Eq. (4.8)	N2, N3, ..., N6
24b-1	24b	8.3	160M	Eq. (4.8)	N2, N3, N4

R_{bdry} is the initial coordinate radius of the outer boundary; this radius changes during the evolution according to the choice of radial map between inertial and comoving coordinates. The last column lists the different resolutions run for each evolution, N6 being the highest resolution. Evolution 30c-1/N6 forms the basis of our post-Newtonian comparisons, and is used in all figures unless noted otherwise.

in all the subdomains), while the most accurate evolution, 30c-1/N6, uses about 67^3 grid points. The run 30c-1/N2 required about 2,500 CPU-hours and run 30c-1/N6 about 19,000, where our simulations do not take advantage of symmetries. The distance to the outer boundary is adjusted by adding or removing outer spherical shells to an otherwise unmodified domain-decomposition. Run 30c-1 has 20 such outer spherical shells, while 30c-2 utilizes 32 and 30c-3 only 8. Thus, the total number of grid points varies slightly between runs, e.g., about 71^3 for 30c-2/N6. Figure 4.2 indicates the different behavior of the outer-boundary location for these three evolutions.

For all of the evolutions 30c-1/2/3, the coordinate trajectories of the centers of the apparent horizons appear as in Fig. 4.3. The regular inspiral pattern without noticeable oscillations once again indicates that our evolutions indeed have very low eccentricity.

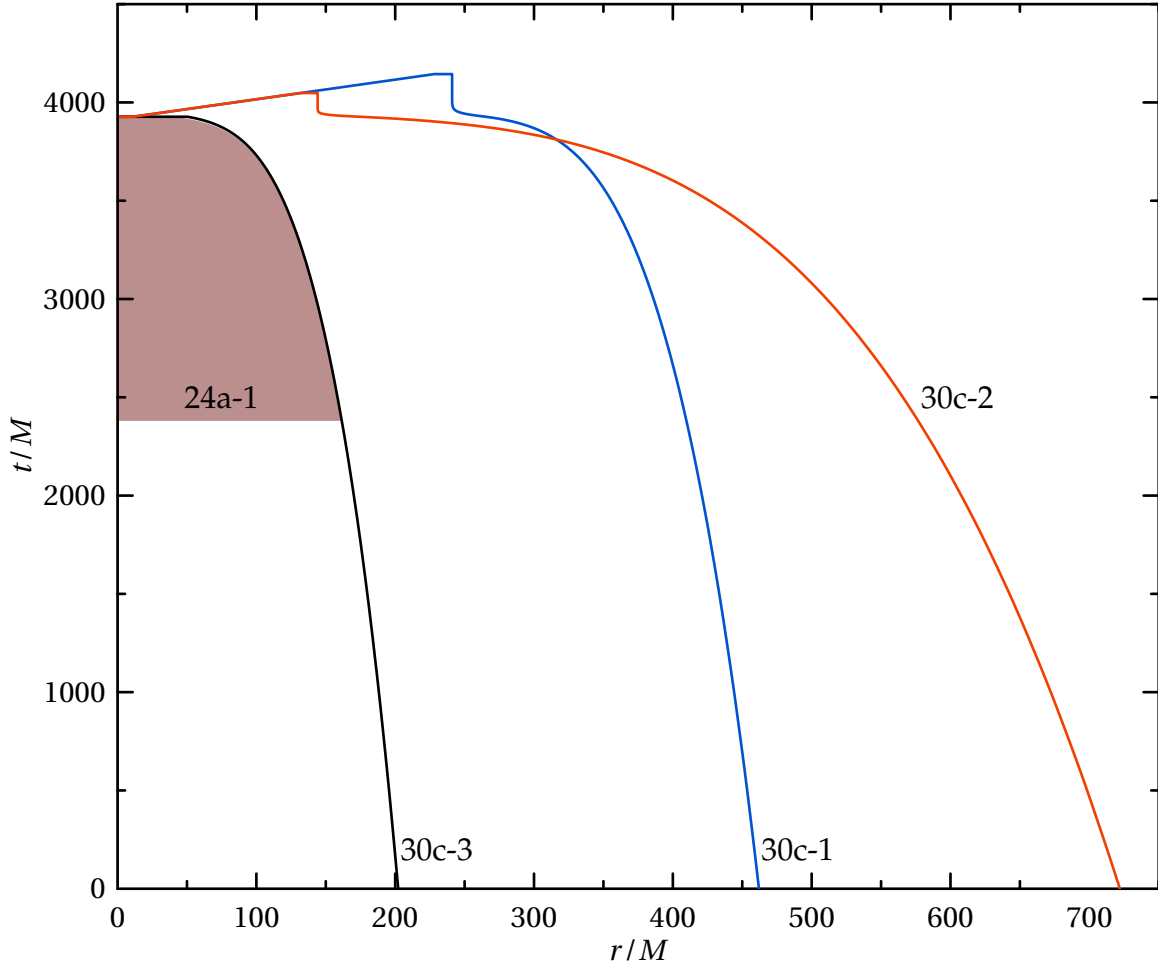


Figure 4.2: **Spacetime diagram showing the region simulated by the numerical evolutions**

The positions of the boundaries of various simulations are shown, as functions of time, with labels corresponding to names used in Table 4.2. The various 30c runs correspond to the identical initial data, but different outer boundaries. Note that the simulation crashes at $t \sim 3930M$. However, the source of the problem is near the center of the simulation. Useful gravitational-wave data is still available at larger radii. Thus, the inner boundary of the simulation is moved outwards at roughly the speed of light. This allows us to “escort” the useful data to useful extraction radii, as seen in simulations 30c-1 and 30c-2. (See Sec. 4.2.3.)




Figure 4.3: **Coordinate trajectories of the centers of the black holes**

The small circles and ellipses show the apparent horizons at the initial time and at the time when the simulation ends and wave escorting begins. The inset shows an enlargement of the dashed box.

Figure 4.4 demonstrates the convergence of the black hole mass $M(t)$ with spatial resolution for run 30c-1. The mass $M(t)$ is computed as the sum of the irreducible masses of both black holes, as defined in Eq. (4.2). At the highest resolution, $M(t)$ deviates by only a few parts in 10^6 from its initial value of M .

Our apparent-horizon finder works by expanding the radius of the apparent horizon as a series in spherical harmonics up to some order L . We utilize the fast flow methods developed by Gundlach [154] to determine the expansion coefficients; these are significantly faster than our earlier minimization algorithms [28, 211]. The apparent horizon is almost spherical during the inspiral, so that the expansion in L converges exceedingly fast: $L = 8$ results in a relative error of the irreducible mass of better than 10^{-8} . The distortion of the horizons becomes more pronounced toward the end of the evolution when the black holes approach each other rapidly. This results in an error of 10^{-6} in the $L = 8$ apparent-horizon expansion for the last $10M$ of the evolution.

We also measure the quasi-local spin using coordinate rotation vectors projected into the apparent-horizon surfaces [67, 11, 12]. Only the z component of the spin is nonzero (i.e., the spins are aligned with the orbital angular momentum). The spin starts at $S_z/M_{\text{irr}}^2 \approx -6 \times 10^{-5}$ and increases slowly to -5×10^{-4} during the evolution, where the negative sign indicates that the black-hole spin is anti-aligned with the orbital angular momentum. Thus it appears the black holes' spins move further from a corotational state. We believe this effect is caused by the use of coordinate rotation vectors when calculating the quasi-local spin, rather than more sophisticated approximate Killing vectors [128, 98, 205]. Preliminary results with approximate Killing vectors find the initial spin to be less than 10^{-6} ,




Figure 4.4: **Deviation of total irreducible mass from its initial value**

Plotted are the six different resolutions of run 30a-1.

and slowly increasing during the evolution to a final value of 2×10^{-5} at the end of the comparison interval to post-Newtonian theory. Given the preliminary character of these results, we will take here the conservative bound $|\vec{S}|/M_{\text{irr}}^2 \leq 5 \times 10^{-4}$ obtained from coordinate rotation vectors.

4.2.3 Escorting gravitational waves

The simulation presented in Fig. 4.3 ends when the horizons of the black holes become too distorted just before merger. At this time, most of the domain (all regions except for the immediate vicinity of the two holes) is still well resolved, and the spacetime contains gravitational radiation that has not yet propagated out to the large radii where we perform wave extraction. So instead of losing this information, which consists of several gravitational-wave cycles, we evolve only the outer portions of our grid beyond the time at which the code crashes in the center, effectively “escorting” the radiation out to the extraction radii.

To do this, we first stop the evolution shortly before it crashes, and we introduce a new spherical excision boundary that surrounds both black holes and has a radius of roughly three times the black hole separation. This new excision boundary moves radially outward at slightly faster than the speed of light so that it is causally disconnected from the interior region where the code is crashing, and so that no boundary conditions are required on this boundary. We then continue the evolution on the truncated spherical-shell domain that extends from the new excision boundary to the outer boundary. To move both boundaries appropriately, we employ a new radial coordinate mapping

$$r = A(t)r'(r') + B(t) , \quad (4.10)$$

where $r(r')$ is given by Eq. (4.9). The functions $A(t)$ and $B(t)$ are chosen to satisfy three criteria. First, the inner boundary of the spherical shell moves outward with coordinate speed of unity, which turns out to be slightly superluminal. Second, the outer boundary location $R_{\text{bdry}}(t)$ has continuous first and second time derivatives at the time we transition to the truncated domain. And finally,

the outer boundary location $R_{\text{bdry}}(t)$ approaches some fixed value at late times. Fig. 4.2 shows the motion of the inner and outer radii for evolutions 30c-1 and 30c-2 (we did not perform wave escorting for 30c-3). For 30c-1, wave escorting extends the evolution for an additional time $220M$ beyond the point at which the simulation stops in the center.

Figure 4.5 shows the gravitational waveform extracted at inertial-coordinate radius $R = 240M$ for the run 30c-1. The brown vertical line indicates the time when wave escorting starts. Wave escorting allows us to extract another 4 cycles of gravitational waves. When computing the gravitational-wave strain $h(t)$ from the Newman–Penrose scalar Ψ_4 (see Eq. (4.11) below), one must choose integration constants during the time integration. These integration constants were chosen such that $h(t)$ has zero average and first moment [209], which is sufficiently accurate for the illustrative Fig. 4.5. To avoid errors caused by the choice of integration constants, the comparison to post-Newtonian waveforms below is based entirely on Ψ_4 .

In the lower two panels of Fig. 4.5 there is a significant amount of noise near the beginning of the run, at $t < 250M$. This noise is (barely) evident in the top panel of Fig. 4.5 as well. The noise is a manifestation of “junk radiation”—a pulse of radiation often seen at the beginning of numerical-relativity simulations, caused by the initial data not being precisely a snapshot of an evolution that has been running for a long time. Among the effects that produce junk radiation are incorrect initial distortions of the individual holes, so that each hole radiates as it relaxes to its correct quasi-equilibrium shape.




Figure 4.5: **Gravitational waveform extracted at $r = 240M$**

From top panel to bottom: the real part of the (2,2) component of $rM\Psi_4$; the gravitational-wave strain, obtained by two time integrals of $\Re[rM\Psi_4]$; the frequency of the gravitational wave, Eq. (4.15); the wavelength, $\lambda = |2\pi/\omega|$. The vertical brown line at $t \approx 3930M$ indicates the time when wave escorting starts.




Figure 4.6: **Normalized constraint violations of run 30c-1**

This plot shows the L^2 norm of all constraints, normalized by the L^2 norm of the spatial gradients of all dynamical fields. Norms are taken only in the regions outside apparent horizons. The discontinuous jumps before $t/M = 4000$ represent the time at which the innermost regions (which contain the largest constraint violations) are excised. The resolutions N2 through N6 are convergent.




Figure 4.7: **Unnormalized constraint violations of run 30c-1**

This plot shows the L^2 norm of all constraints, with no normalization (compare Fig. 4.6). Norms are taken only in the regions outside apparent horizons. The discontinuous jumps before $t/M = 4000$ represent the time at which the innermost regions (which contain the largest constraint violations) are excised. The resolutions N2 through N6 are convergent.

Our evolution code does not explicitly enforce either the Einstein constraints or the secondary constraints that arise from writing the system in first-order form. Therefore, examining how well these constraints are satisfied provides a useful consistency check. Figures 4.6 and 4.7 show the constraint violations for run 30c-1. The first plot shows the L^2 norm of all the constraint fields of our first-order generalized-harmonic system, normalized by the L^2 norm of the spatial gradients of the dynamical fields (see Eq. (71) of Ref. [187]). The second plot shows the same quantity, but without the normalization factor (i.e., just the numerator of Eq. (71) of Ref. [187]). The L^2 norms are taken over the entire computational volume that lies outside of apparent horizons. At early times, $t < 500M$, the constraints converge rather slowly with resolution because the junk radiation contains high frequencies. Convergence is more rapid during the smooth inspiral phase, after the junk radiation has exited through the outer boundary. The constraints increase around $t \sim 3900M$ as the code begins to fail near the two merging holes, but then the constraints decrease again after the failing region is excised for wave escorting. The normalized constraint violations are less than 10^{-4} until just before the peak (which occurs at $t = 3930M$ for all but the lowest resolutions). The size of the peak causes some concern that the waveforms at late times may be contaminated by constraint violations to a non-negligible degree. However, near the peak, the constraint violations are large only in the inner regions of the domain near the black holes. (Note that the curves in Figs. 4.6 and 4.7 decrease by two orders of magnitude immediately after these inner regions are excised at $t = 3930M$.) Because all constraint quantities propagate at the speed of light or slower for the formulation of Einstein's equations that we use, any influence

that the constraint peak has on the extracted waveform occurs after the constraint violations have had time to propagate out to the wave-extraction zone. This is very late in the waveform, well after the gravitational-wave frequency reaches $M\omega = -0.1$, as can be seen from the right panel of the spacetime diagram in Fig. 4.2.

4.2.4 Waveform extraction

Gravitational waves are extracted using the Newman–Penrose scalar Ψ_4 , using the same procedure as in [209]. To summarize, given a spatial hypersurface with timelike unit normal n^μ , and given a spatial unit vector r^μ in the direction of wave propagation, the standard definition of Ψ_4 is the following component of the Weyl curvature tensor,

$$\Psi_4 = -C_{\alpha\mu\beta\nu}\ell^\mu\ell^\nu\bar{m}^\alpha\bar{m}^\beta, \quad (4.11)$$

where $\ell^\mu := \frac{1}{\sqrt{2}}(n^\mu - r^\mu)$, and m^μ is a complex null vector (satisfying $m^\mu\bar{m}_\mu = 1$) that is orthogonal to r^μ and n^μ . Here an overbar denotes complex conjugation.

For (perturbations of) flat spacetime, Ψ_4 is typically evaluated on coordinate spheres, and in this case the usual choices for n^μ , r^μ , and m^μ are

$$n^\mu = \left(\frac{\partial}{\partial t}\right)^\mu, \quad (4.12a)$$

$$r^\mu = \left(\frac{\partial}{\partial r}\right)^\mu, \quad (4.12b)$$

$$m^\mu = \frac{1}{\sqrt{2}r} \left(\frac{\partial}{\partial\vartheta} + i\frac{1}{\sin\vartheta}\frac{\partial}{\partial\varphi}\right)^\mu, \quad (4.12c)$$

where (r, ϑ, φ) denote the standard spherical coordinates. With this choice, Ψ_4

can be expanded in terms of spin-weighted spherical harmonics of weight -2 :

$$\Psi_4(t, r, \vartheta, \varphi) = \sum_{l, m} \Psi_4^{l, m}(t, r) {}_{-2}Y_{l, m}(\vartheta, \varphi) , \quad (4.13)$$

where the $\Psi_4^{l, m}$ are expansion coefficients defined by this equation.

For curved spacetime, there is considerable freedom in the choice of the vectors r^μ and m^μ , and different researchers have made different choices [75, 138, 29, 199, 79, 81, 69] that are all equivalent in the $r \rightarrow \infty$ limit. We choose these vectors by first picking an extraction two-surface \mathcal{E} that is a coordinate sphere ($r^2 = x^2 + y^2 + z^2$ using the global asymptotically Cartesian coordinates employed in our code) centered on the center of mass of the binary system, i.e., the point of symmetry. We choose r^μ to be the outward-pointing spatial unit normal to \mathcal{E} (that is, we choose r_i proportional to $\nabla_i r$ and raise the index with the spatial metric). Then we choose m^μ according to Eq. (4.12c), using the standard spherical coordinates ϑ and φ defined on these coordinate spheres. Finally we use Eqs. (4.11) and (4.13) to define the $\Psi_4^{l, m}$ coefficients.

Note that the m^μ vector used here is not exactly null nor exactly of unit magnitude at finite r . The resulting $\Psi_4^{l, m}$ at finite r will disagree with the asymptotic waveforms. Our definition does, however, agree with the standard definition given in Eqs. (4.11)–(4.13) as $r \rightarrow \infty$. Because we extrapolate the extracted waves to find the asymptotic radiation field (see Section 4.2.6), these effects should not play a role in our PN comparisons; relative errors in $\Psi_4^{l, m}$ introduced by using the simple coordinate tetrad fall off like $1/r$, and thus should vanish after extrapolating to obtain the asymptotic behavior. While more careful treatment of the extraction method—such as those discussed in [200, 207, 185]—may improve the quality of extrapolation and would be interesting to explore in the future, the

naive choice made here should be sufficient to ensure that the waveform after extrapolation is correct to the accuracy needed for these simulations.

In this paper, we focus on the $(l, m) = (2, 2)$ mode. Following common practice (see, e.g. [20, 69]), we split the extracted waveform into a real phase ϕ and a real amplitude A , defined by

$$\Psi_4^{2,2}(r, t) = A(r, t)e^{i\phi(r, t)} . \quad (4.14)$$

The gravitational-wave frequency is defined as

$$\omega := \frac{d\phi}{dt} . \quad (4.15)$$

Note that these definitions (along with the definition of Ψ_4) result in *decreasing* phase and negative frequencies when the binary rotates in the usual sense, with orbital angular momentum in the positive z direction. Equation (4.14) defines ϕ only up to additive multiples of 2π . These multiples of 2π are chosen to make ϕ continuous through each evolution, still leaving an overall multiple of 2π undetermined. We will consider only phase differences in this paper, so the choice of this overall phase offset is irrelevant.

4.2.5 Convergence of extracted waveforms

In this section we examine the convergence of the gravitational waveforms extracted at fixed radius, without extrapolation to infinity. This allows us to study the behavior of our code without the complications of extrapolation. The extrapolation process and the resulting extrapolated waveforms are discussed in Sec. 4.2.6.

Figure 4.8 shows the convergence of the gravitational-wave phase ϕ with numerical resolution for the run 30c-1. For this plot, the waveform is extracted at a fixed radius $R = 77M$. Each line shows the difference between ϕ computed at some particular resolution and ϕ computed from our highest-resolution run 30c-1/N6. When subtracting results at different resolutions, no time or phase adjustment has been performed. The difference in ϕ between the two highest-resolution runs is smaller than 0.03 radians throughout the run, and it is smaller than 0.02 radians between $t = 1000M$ and the point at which $M\omega = -0.1$.

At times before $1000M$, the phase convergence of our simulation is limited to about 0.05 radians because of effects of junk radiation (described at the end of Section 4.2.3). The sharp pulse of junk radiation has comparatively large numerical-truncation error, and excites all characteristic modes at the level of truncation error, including waves that propagate back toward the origin. Generation of these secondary waves stops when the pulse of junk radiation leaves through the outer boundary (i.e., after one light-crossing time). Because we use the improved outer boundary conditions of Rinne et al. [225], there are no significant reflections when the junk radiation passes through the outer boundary. However, the waves produced before the junk radiation leaves remain in the computational domain for two additional light-crossing times, until they eventually leave through the outer boundary.

Figure 4.9 shows phase comparisons between different waveforms after we perform a time shift and phase shift so that the waveforms agree at $M\omega = -0.1$. Our procedure for time shifting and phase shifting is the same as the shifting




Figure 4.8: **Convergence of the gravitational-wave phase without time shifting**

Data is extracted (not extrapolated) at the radius $R = 77M$. All lines show differences with respect to our highest resolution run, 30c-1/N6. No time or phase shifts have been performed.




Figure 4.9: **Convergence of the gravitational-wave phase with time shifting**

Data is extracted (not extrapolated) at the radius $R = 77M$. All lines show differences with respect to our highest-resolution run, 30c-1/N6. The various simulations have been aligned at $M\omega = -0.1$ by a time and phase shift. The thin vertical line indicates the time at which $M\omega = -0.1$ for 30c-1/N6. (Compare Fig. 4.8, for which no shifts have been applied.)

procedure we use to compare NR with PN waveforms (see Sec. 5.6.1), so that the error estimates we extract from Fig. 4.9 are relevant for our PN–NR comparison.

There are three different types of comparisons shown in Fig. 4.9: phase differences between runs with the same initial data but with different outer boundary locations, phase differences between runs with different initial data, and phase differences between different numerical resolutions of the same run (this last comparison is the same as that shown in the previous figure, except in the latter figure, the waveforms are shifted in time and phase). We will discuss all three of these in turn.

First, we compare the phase difference of 30c-1/N6 with runs that have different outer boundary locations. Run 30c-2 (with more distant outer boundary) agrees to within 0.002 radians with run 30c-1, but run 30c-3 (with closer outer boundary), has a much larger phase difference compared to 30c-1. We believe that this is because run 30c-3 has a very small ratio of outer-boundary location to gravitational wavelength; R/λ is about 1.1 for the first two-thirds of the run, and remains less than 2 for the entire run.

We can explain the order of magnitude of these phase differences using the analysis of Buchman and Sarbach [70]. Our outer-boundary conditions are not perfectly absorbing, but instead they reflect some fraction of the outgoing radiation.⁶ The ratio of the amplitude of curvature perturbations (i.e., Ψ_4) of the reflected wave to that of the outgoing wave is

$$q \approx \frac{3}{2(2\pi)^4} \left(\frac{\lambda}{R}\right)^4. \quad (4.16)$$

⁶Note that, in a comparison of various boundary conditions [225], the boundary conditions used for this data produced smaller reflections than other boundary conditions commonly used in numerical relativity.

The incoming reflected waves grow like $1/r$ as they travel inward just like the outgoing waves decrease by $1/r$ as they propagate outward. Therefore, the ratio of amplitudes of incoming and outgoing waves will have approximately the same value, q , at smaller radii, and we assume for the sake of this rough argument that this ratio remains equal to q even in the vicinity of the black holes (where it is no longer technically meaningful to talk about “radiation”). Now consider the second time derivative of the gravitational-wave phase, $\ddot{\phi}$. This is nonzero only because of gravitational-wave emission, so $\ddot{\phi}$ is proportional to some power of the outgoing wave amplitude. To get the correct power, we can use Eq. (4.48) to find $\dot{x} \sim x^5$, so Eq. (4.39) yields $-\ddot{\phi} \sim x^{11/2}$ (we assume gravitational-wave phase is twice the orbital phase). The amplitude of Ψ_4 scales like x^4 , so $-\ddot{\phi} \sim A^{11/8}$. Let us assume for the sake of this rough error estimate that the change in $\ddot{\phi}$ due to the *ingoing* reflected wave scales similarly with amplitude, $-\ddot{\phi} \sim \bar{A}^{11/8}$, where $\bar{A} = qA$ is the amplitude of the reflected ingoing wave. Therefore the unphysical gravitational-wave force acting back on the system due to boundary reflections will cause fractional errors in the second derivative of the phase of about $q^{11/8}$. That is, the magnitude of the error $\delta\phi$ caused by the improper boundary condition will be given by

$$\frac{d^2\delta\phi}{dt^2} = q^{11/8} \frac{d^2\phi}{dt^2}. \quad (4.17)$$

Integrating this yields $\delta\phi = q^{11/8}\phi$, where ϕ is the total gravitational-wave phase accumulated during the evolution. For 30c-3, $\lambda/R \sim 0.9$, so $q \sim 6 \times 10^{-4}$, which yields $\delta\phi \sim 0.08$ radians for an accumulated gravitational-wave phase of about 200 radians. This rough estimate agrees in order of magnitude with the phase difference between 30c-3 and 30c-1 as shown in Fig. 4.9. The run 30c-1 has an

outer boundary about 2.5 farther away, reducing the reflection coefficient by a factor $2.5^4 \approx 40$, so for 30c-1 this estimate of the phase error gives $\delta\phi = 5 \times 10^{-4}$ radians. Therefore, we expect reflection of the outgoing radiation at the outer boundary to be insignificant for 30c-1. This is confirmed by the excellent agreement between runs 30c-1 and 30c-2 (the latter having even larger outer boundary).

The second comparison shown in Fig. 4.9 is the phase difference between 30c-1/N6 and 24b-1/N4, a shorter 8-orbit evolution started from a separate initial-data set (set 24b in Table 4.1) with a separate eccentricity-reduction procedure. The phase agreement between these two runs (including an overall time shift and phase shift) is better than 0.01 radians for a total accumulated phase of ~ 100 radians of the 8-orbit run, i.e., better than one part in 10^4 . Run 24b-1 has a similar outer-boundary location as run 30c-3, and indeed both of these runs show similar phase differences from 30c-1.

Finally, the third comparison shown in Fig. 4.9 is the phase difference between the two highest resolutions of the run 30c-1 when a time shift is applied. For $t \gtrsim 1000M$ the agreement is much better than without the time shift (see upper panel), indicating that the dominant error is a small difference in the overall evolution time. For the post-Newtonian comparisons we perform in the second part of this paper, waveforms are always aligned at specific frequencies by applying time and phase shifts. Therefore, the time-shifted phase difference as displayed in the lower panel is the most appropriate measure of numerical-truncation error for these PN comparisons. This difference is less than 0.003 radians after $t = 1000M$ but is larger, about 0.02 radians, at early times where the waveforms

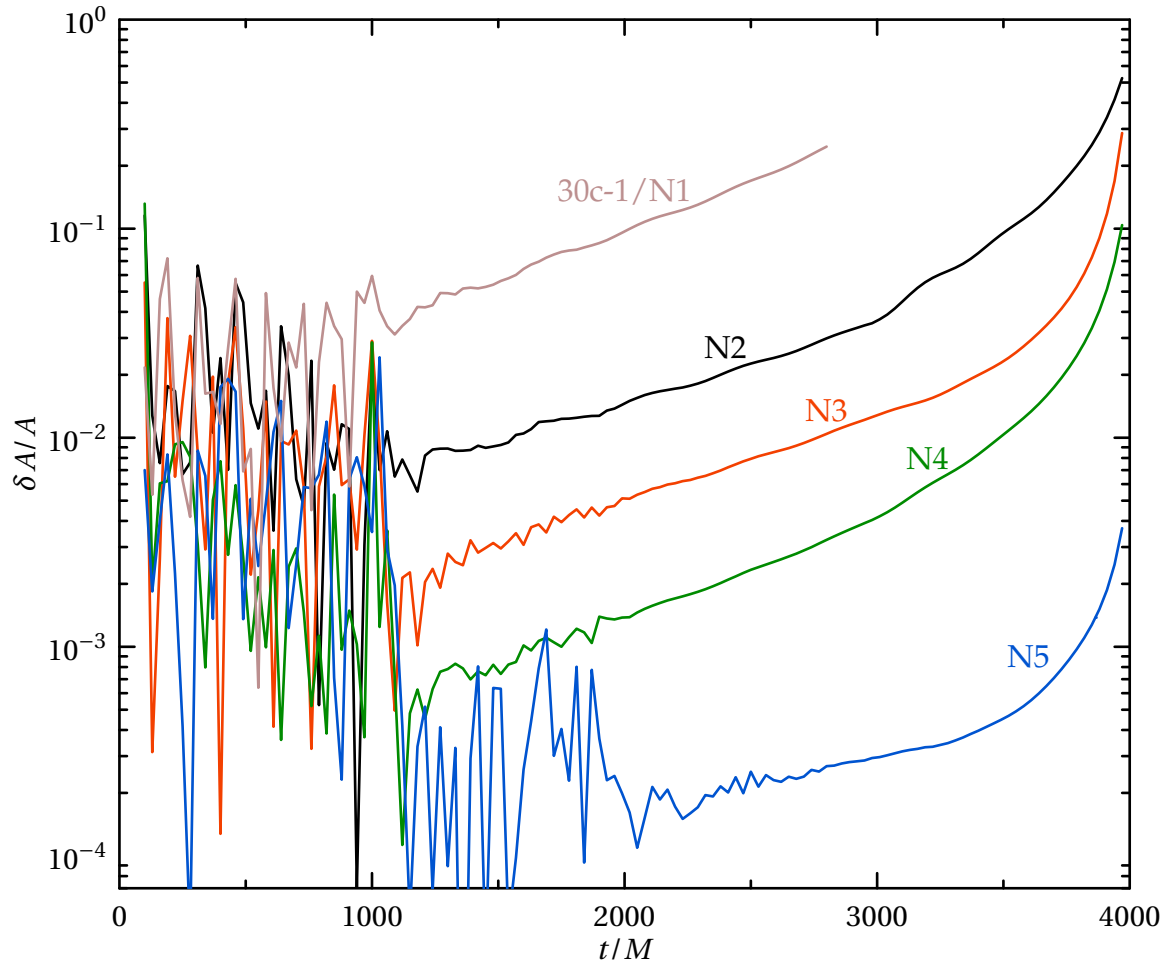


Figure 4.10: **Convergence of the gravitational-wave amplitude**

Data is extracted (not extrapolated) at the radius $R = 77M$. All lines show differences with respect to our highest-resolution run, 30c-1/N6. This plot corresponds to Fig. 4.8, except that relative amplitude differences are shown.

are noisy because of junk radiation.

We also compare the gravitational-wave amplitudes of different runs in the same manner as we compared the gravitational-wave phases. Figures 4.10 and 4.11 present convergence data for the amplitude of the gravitational waves for the same




Figure 4.11: **Convergence of the gravitational-wave amplitude**

All lines show differences with respect to our highest-resolution run, 30c-1/N6. Data is extracted (not extrapolated) at the radius $R = 77M$. This plot corresponds to Fig. 4.9, except that relative amplitude differences are shown. The thin vertical line indicates the time at which $M\omega = -0.1$ for 30c-1/N6.

runs as shown in Figs. 4.8 and 4.9. Spatial-truncation error for the amplitude is less than 0.1 percent for $t/M > 1000$, and earlier than this it is limited by residual noise from the junk radiation. Differences (including time shifts) between runs of different lengths are shown in Fig. 4.10. These differences are even smaller, but because of their small size, they are dominated by noise for about the first half of the run. The oscillations apparent in the comparison to 24b-1 are caused by the larger orbital eccentricity of 24b-1 (see Table 4.1).

4.2.6 Extrapolation to infinity

The quantity of interest to gravitational-wave detectors is the gravitational waveform as seen by an observer effectively infinitely far from the source. Our numerical simulations, in contrast, cover only a region of finite volume around the source, and our numerical waveforms are extracted at a finite radius. Waveforms extracted at a finite radius can differ from those extracted at infinity because of effects discussed in Sec. 4.2.4; these effects can lead to phase errors of several tenths of a radian and relative amplitude errors of several percent. To avoid such errors we extrapolate to infinite extraction radius as follows:

We extract data for Ψ_4 on coordinate spheres of coordinate radii $r/M = 75, 80, 85, \dots, 240$, as described in Sec. 4.2.4. These extracted waveforms are shifted in time relative to one another because of the finite light-travel time between these extraction surfaces. We correct for this by shifting each waveform by the tortoise-coordinate radius at that extraction point [138]

$$r^* = r_{\text{areal}} + 2M_{\text{ADM}} \ln \left(\frac{r_{\text{areal}}}{2M_{\text{ADM}}} - 1 \right). \quad (4.18)$$




Figure 4.12: **Difference between areal radius r_{areal} and coordinate radius r of selected extraction surfaces**

The areal radius r_{areal} remains constant to within $0.01M$ during the evolution. The horizontal dashed line indicates M_{ADM}/M of the initial data. (Note that data for $r = 400M$ ends early because of infall of the outer boundary, as explained below.)

Here M_{ADM} is the ADM mass of the initial data, and $r_{\text{areal}} = \sqrt{A/4\pi}$, where A is the area of the extraction sphere. This is not the only possible choice for the retarded time. For example, the waveforms could be shifted so that the maxima of the amplitude align [157]. It has also been suggested [182] that the time shift should change with the amount of radiated energy—essentially, that the factor of M_{ADM} should be replaced by the amount of mass interior to the extraction radius at each time. We leave investigation of other choices of retarded time for future work.

Figure 4.12 presents the areal radius during the evolution at several typical extraction radii. The areal radius of these extraction surfaces is constant to within about $0.01M$, and to the same precision, $r_{\text{areal}} = r + M_{\text{ADM}}$. This relationship is not surprising, because the initial data is conformally flat, so that for coordinate spheres $r_{\text{areal}} = r + M_{\text{ADM}} + \mathcal{O}(M_{\text{ADM}}/r)$. For convenience, we simply set $r_{\text{areal}} = r + M_{\text{ADM}}$ in Eq. (4.18), rather than explicitly integrating to find the area of each extraction sphere.

After the time shift, each waveform is a function of retarded time, $t - r^*$. At a given value of retarded time, we have a series of data points—one for each extraction radius. We fit phase and amplitude of these data separately to a polynomial in $1/r$:

$$\phi(t - r^*, r) = \phi_{(0)}(t - r^*) + \sum_{k=1}^N \frac{\phi_{(k)}(t - r^*)}{r^k}, \quad (4.19)$$

$$rA(t - r^*, r) = A_{(0)}(t - r^*) + \sum_{k=1}^N \frac{A_{(k)}(t - r^*)}{r^k}. \quad (4.20)$$

The leading-order term of each polynomial, as a function of retarded time, is then




Figure 4.13: **Convergence of phase extrapolation with extrapolating-polynomial order**

Plotted are absolute differences between extrapolation with order N and $N + 1$. Increasing the order of the polynomial increases accuracy, but also amplifies noise. See Eq. (4.19) and surrounding text for discussion.

the desired asymptotic waveform:

$$\phi(t - r^*) = \phi_{(0)}(t - r^*) , \quad (4.21)$$

$$rA(t - r^*) = A_{(0)}(t - r^*) . \quad (4.22)$$

We find good convergence of this method as we increase the order N of the extrapolating polynomial. Figure 4.13 shows the difference in phase between waveforms extrapolated using successively higher-order polynomials. We see a broad improvement in the accuracy of the phase with increasing order, but unfortunately, higher-order extrapolations tend to amplify the noise. Our preferred choice is $N = 3$ extrapolation, resulting in extrapolation errors of $\lesssim 0.003$ radians for $t - r^* \gtrsim 1000M$.

Figure 4.14 is analogous to Fig. 4.13, except that it shows relative differences in the extrapolated amplitudes. The basic picture agrees with the phase extrapolation: Higher order extrapolation reduces the errors, but amplifies noise. Our preferred choice $N = 3$ gives a relative amplitude error of $\lesssim 0.002$ for $t - r^* \gtrsim 1000M$, dropping to less than 0.001 for $t - r^* \gtrsim 2000M$.

Phase and amplitude extrapolation become increasingly accurate at late times. The main obstacle to accuracy seems to be near-zone effects scaling with powers of (λ/r) , where λ is the wavelength of the gravitational wave. The wavelength is quite large at the beginning of the simulation ($\lambda \approx 180M$, see Fig. 4.5), but becomes shorter during the evolution, so that even low-order extrapolation is quite accurate at late times. Alternatively, near-zone effects can be mitigated by using data extracted at large values of r . It is precisely because of these near-zone effects that we have chosen to ignore data extracted at $r < 75M$ when we




Figure 4.14: **Convergence of extrapolated amplitude with extrapolating-polynomial order**

Plotted are relative amplitude differences between extrapolation with orders N and $N+1$. See Eq. (4.20) and surrounding text for discussion.




Figure 4.15: **Effect of choice of wave-extraction radii on extrapolated phase**

Each curve represents the difference from our preferred wave extrapolation using $r \in [75M, 240M]$. The three solid curves represent extrapolation from different intervals of extraction radii. The curves labeled $240M$ and $90M$ represent differences from waves extracted at these two radii, without any extrapolation, for two cases: time and phase shifted so that ϕ and $\dot{\phi}$ match at $M\omega = -0.1$ (dashed), and without these shifts (dotted).

extrapolate to infinity.

In Figs. 4.15 and 4.16, we show the effects of extrapolation using different ranges of extracted data. Using data extracted every $5M$ in the range $r =$




Figure 4.16: **Effect of choice of wave-extraction radii on extrapolated amplitude**

Each curve represents the (relative) amplitude difference to our preferred wave extrapolation using $r \in [75M, 240M]$. The three solid curves represent extrapolation from different intervals of extraction radii. The curves labeled $240M$ and $90M$ represent differences from waves extracted at these two radii, without any extrapolation, for two cases: time and phase shifted so that ϕ and $\dot{\phi}$ match at $M\omega = -0.1$ (dashed), and without these shifts (dotted).

$50M$ – $90M$ results in noticeable differences early in the run—though it is adequate later in the run. For ranges at higher radii (e.g., $[75M, 150M]$ or $[150M, 240M]$), the accuracy is not highly variable, though we find that noise is increased when using data from such a smaller range of extraction radii.

To estimate the errors generated by not extrapolating waveforms to infinity at all, Fig. 4.15 contains also the phase difference between wave extraction at two finite radii ($90M$ and $240M$) and our preferred extrapolated phase at infinity. The dotted lines show such phase differences when only a time shift by the tortoise-coordinate radius of the extraction sphere is applied. The errors are dramatic, tenths of radians or more, even very late in the run. When matching to post-Newtonian waveforms, we are free to add an overall time and phase shift (see Sec. 5.6.1). Therefore, the dashed lines in Fig. 4.15 show phase differences with the same unextrapolated waveforms as shown by the dotted lines, except that a phase and time shift has been applied so that the ϕ and $\dot{\phi}$ agree with those of the extrapolated waveform late in the run (where $M\omega = -0.1$), where the wavelengths are shortest and wave extraction is expected to work best. Even with such an adjustment, the gravitational-wave phase extracted at $r = 90M$ differs by about 0.1 radian at $t \sim 1000M$ before coalescence, with this difference growing to 0.3 radians at the start of our simulation.

Figure 4.16 makes the same comparison for the gravitational-wave amplitude. Wave extraction at $r = 90M$ results in relative amplitude errors of up to 8 percent, and of about 2 percent even in the last $1000M$ of our simulation. We also point out that the errors due to finite extraction radius decay approximately as the inverse of the extraction radius: For waves extracted at $r = 240M$ the errors are

smaller than for waves extracted at $r = 90M$ by about a factor of three, as can be seen in Figs. 4.15 and 4.16; for wave extraction at $r = 45M$, the errors would be approximately twice as large as the $r = 90M$ case. The errors introduced by using a finite extraction radius are significantly larger than our truncation error (even at extraction radius $240M$). Therefore extrapolation to infinity is essential to realize the full accuracy of our simulations.

4.2.7 Estimated time of merger

Since we have not yet been successful with simulating the merger, we do not precisely know when merger occurs. However, by comparing the orbital and gravitational-wave frequencies to already published results, we can nevertheless estimate the time of merger.

The simulation presented in Fig. 4.3 stops at time $t = 3929M$ when the horizons of the black holes become too distorted just before merger. At that point, the proper separation between the horizons is $\sim 4.0M$, and the orbital frequency has reached $M\omega_{\text{orbit}} = 0.125$; comparison with [75] suggests this is about $15M$ before formation of a common apparent horizon, i.e., the common horizon should form in our simulations at $t_{\text{CAH}} \approx 3945M$.

The waveform extrapolated to infinity ends at $t - r^* = 3897M$ at a gravitational-wave frequency of $M\omega \approx -0.16$. This places the end of the waveform at about $50M$ (or ~ 1.5 cycles) before formation of a common apparent horizon⁷ (judged by comparison with [75]). Thus, we estimate the formation of a common horizon

⁷The waveform ends somewhat further from merger than the orbital trajectory, because the artificial boundary is placed initially at a radius $\sim 15M$, and then moves outward somewhat faster than the speed of light, thus overtaking the very last part of the waveform as it travels to the wave-extraction radii.

to correspond to a retarded time of approximately $(t - r^*)_{\text{CAH}} \approx 3950M$.

4.3 Generation of post-Newtonian waveforms

It is not our intention to review all of post-Newtonian (PN) theory, but to summarize the important points that go into the construction of the post-Newtonian waveforms that we will compare to our numerical simulation. For a complete review of post-Newtonian methods applied to inspiralling compact binaries, see the review article by Blanchet [42].

The post-Newtonian approximation is a slow-motion, weak-field approximation to general relativity with an expansion parameter $\epsilon \sim (v/c)^2 \sim (GM/dc^2)$. For a binary system of two point masses M_1 and M_2 , v is the magnitude of the relative velocity, M is the total mass, and r is the separation. In order to produce a post-Newtonian waveform, it is necessary to solve both the post-Newtonian equations of motion describing the binary, and the post-Newtonian equations describing the generation of gravitational waves.

Solving the equations of motion yields explicit expressions for the accelerations of each body in terms of the positions and velocities of the two bodies [174, 175, 117, 119, 48, 49, 118, 46, 171, 170, 169]. The two-body equations of motion can then be reduced to relative equations of motion in the center-of-mass frame in terms of the relative position and velocity [52]. The relative acceleration is currently known through 3.5PN order, where 0PN order for the equations of motion corresponds to Newtonian gravity. The effects of radiation reaction (due to the emission of gravitational waves) enters the relative acceleration starting at 2.5PN order. The relativistic corrections to the relative acceleration at 1PN, 2PN,

and 3PN order (ignoring the radiation reaction terms at 2.5PN and 3.5PN order) admit a conserved center of mass binding energy through 3PN order [125]. There is no 2.5PN or 3.5PN order contribution to the energy.

Solving the post-Newtonian wave generation problem yields expressions for the gravitational waveform h_{ij} and gravitational-wave flux \mathcal{L} in terms of radiative multipole moments [242]. These radiative multipole moments are in turn related to the source multipole moments, which can be given in terms of the relative position and relative velocity of the binary [39]. For the gravitational-wave generation problem, PN orders are named with respect to the leading order waveform and flux, which are given by the quadrupole formalism. Thus, for example, 1.5PN order in the wave-generation problem represents terms of order $(v/c)^3$ beyond quadrupole. Higher order effects enter both through post-Newtonian corrections to the mass quadrupole, as well as effects due to higher multipole moments. Starting at 1.5PN order the radiative multipole moments include non-linear and noninstantaneous (i.e., they depend upon the past history of the binary) interactions among the source multipole moments (e.g., gravitational-wave tails) [39, 45, 40, 38].

It was recognized early that simply plugging in the orbital evolution predicted by the equations of motion into the expressions for the waveform would not generate templates accurate enough for matched filtering in detecting gravitational waves [99]. This is because radiation reaction enters the equations of motion only at the 2.5PN order; hence computing a waveform to k PN order beyond the quadrupole formalism would require $(2.5 + k)$ PN orders in the equations of motion. In order to obtain as accurate a post-Newtonian waveform as possible it is

thus necessary to introduce the assumption of an adiabatic inspiral of a quasi-circular orbit, as well as the assumption of energy balance between the orbital binding energy and the energy emitted by the gravitational waves.

4.3.1 Adiabatic inspiral of quasi-circular orbits

The emission of gravitational radiation causes the orbits of an isolated binary system to circularize [208]. Thus it is a reasonable assumption to model the orbital evolution of the binary as a slow adiabatic inspiral of a quasi-circular orbit. With this assumption, post-Newtonian expressions for the orbital energy E and gravitational energy flux \mathcal{L} are currently known through 3.5PN order [54, 50, 47, 53, 51]. These expressions can be given in terms of a parameter related to either the harmonic coordinate separation r , or to the orbital frequency Ω . We choose to use the expressions given in terms of a frequency-related parameter

$$x := \left(\frac{GM\Omega}{c^3} \right)^{2/3}, \quad (4.23)$$

rather than a coordinate-related parameter, because the coordinate relationship between the numerical simulation and the harmonic coordinates used in post-Newtonian approximations is unknown. The orbital energy for an equal-mass system is given by [42]

$$E = -\frac{Mc^2}{8} x \left[1 - \frac{37}{48} x - \frac{1069}{384} x^2 + \left(\frac{1427365}{331776} - \frac{205}{384} \pi^2 \right) x^3 \right], \quad (4.24)$$

and the gravitational-wave flux for an equal-mass system is given by [42]

$$\mathcal{L} = \frac{2c^5}{5G} x^5 \left\{ 1 - \frac{373}{84}x + 4\pi x^{3/2} - \frac{59}{567}x^2 - \frac{767}{42}\pi x^{5/2} \right. \\ \left. + \left[\frac{18608019757}{209563200} + \frac{355}{64}\pi^2 - \frac{1712}{105}\gamma - \frac{856}{105}\ln(16x) \right] x^3 + \frac{16655}{6048}\pi x^{7/2} \right\}, \quad (4.25)$$

where $\gamma = 0.577216\dots$ is Euler's constant.

4.3.2 Polarization Waveforms

The gravitational polarization waveforms for a quasi-circular orbit in the x - y plane, as measured by an observer at spherical coordinates (R, ϑ, φ) , are given by

$$h_+ = \frac{2G\mu}{c^2 R} x \{ -(1 + \cos \vartheta) \cos 2(\Phi - \varphi) + \dots \} \quad (4.26)$$

$$h_\times = \frac{2G\mu}{c^2 R} x \{ -2 \cos \vartheta \sin 2(\Phi - \varphi) + \dots \}, \quad (4.27)$$

where Φ is the orbital phase (measured from the x axis) and $\mu = M_1 M_2 / M$ is the reduced mass. The polarization waveforms are currently known through 2.5PN order [10, 179].

Optimally oriented observer

For an equal-mass binary the polarization waveforms along the z axis (i.e., the optimally oriented observer along the normal to the orbital plane) are given

by [10, 179]

$$h_+^{(z)} = \frac{GM}{2c^2R} x \left(\cos 2\Phi \left\{ -2 + \frac{17}{4}x - 4\pi x^{3/2} + \frac{15917}{2880}x^2 + 9\pi x^{5/2} \right\} + \sin 2\Phi \left\{ -12 \ln \left(\frac{x}{x_0} \right) x^{3/2} + \left[\frac{59}{5} + 27 \ln \left(\frac{x}{x_0} \right) \right] x^{5/2} \right\} \right), \quad (4.28)$$

$$h_\times^{(z)} = \frac{GM}{2c^2R} x \left(\sin 2\Phi \left\{ -2 + \frac{17}{4}x - 4\pi x^{3/2} + \frac{15917}{2880}x^2 + 9\pi x^{5/2} \right\} + \cos 2\Phi \left\{ 12 \ln \left(\frac{x}{x_0} \right) x^{3/2} - \left[\frac{59}{5} + 27 \ln \left(\frac{x}{x_0} \right) \right] x^{5/2} \right\} \right), \quad (4.29)$$

where

$$\ln x_0 := \frac{11}{18} - \frac{2}{3}\gamma + \frac{2}{3} \ln \left(\frac{GM}{4bc^3} \right) \quad (4.30)$$

is a constant frequency scale that depends upon the constant time scale b entering the gravitational-wave tail contribution to the polarization waveforms [251, 56]. The freely-specifiable constant b corresponds to a choice of the origin of radiative time T with respect to harmonic time t , and enters the relation between the retarded time $T_R = T - R/c$ in radiative coordinates (the coordinates in which the waveform is given) and the retarded time $t - r/c$ in harmonic coordinates (the coordinates in which the equations of motion are given) [251, 56]:

$$T_R = t - \frac{r}{c} - \frac{2GM_{\text{ADM}}}{c^3} \ln \left(\frac{r}{bc} \right). \quad (4.31)$$

Here M_{ADM} is the ADM mass (mass monopole) of the binary system.

The (2,2) mode

When comparing a post-Newtonian waveform with data from a physical gravitational-wave detector, it is necessary to compare waves emitted in a certain direction (ϑ, φ) with respect to the source. However, comparing waveforms between PN

and numerical simulations can be done in all directions simultaneously by decomposing the waveforms in terms of spherical harmonics and then comparing different spherical harmonic modes. Since the power in each spherical harmonic mode decreases rapidly with spherical harmonic index, with the (2,2) mode dominating (for an equal-mass nonspinning binary), it is possible to do a very accurate comparison that is valid for all angles by using only a few modes. In addition, as pointed out by Kidder [178], the dominant (2,2) mode can be computed to 3PN order. For an equal-mass binary, the (2,2) mode is

$$h_{(2,2)} = -2\sqrt{\frac{\pi}{5}} \frac{GM}{c^2 R} e^{-2i\Phi} x \sum_{k=0}^6 H_k x^{k/2}, \quad (4.32a)$$

where the coefficients are given by

$$H_0 = 1, \quad (4.32b)$$

$$H_1 = 0, \quad (4.32c)$$

$$H_2 = -\frac{373}{168}, \quad (4.32d)$$

$$H_3 = \left[2\pi + 6i \ln\left(\frac{x}{x_0}\right) \right], \quad (4.32e)$$

$$H_4 = -\frac{62653}{24192}, \quad (4.32f)$$

$$H_5 = -\left[\frac{197}{42}\pi + \frac{197i}{14} \ln\left(\frac{x}{x_0}\right) + 6i \right], \quad (4.32g)$$

$$H_6 = \left\{ \frac{43876092677}{1117670400} + \frac{99}{128}\pi^2 - \frac{428}{105} \ln x - \frac{856}{105}\gamma \right. \\ \left. - \frac{1712}{105} \ln 2 - 18 \left[\ln\left(\frac{x}{x_0}\right) \right]^2 + \frac{428}{105} i\pi + 12i\pi \ln\left(\frac{x}{x_0}\right) \right\}. \quad (4.32h)$$

Since the (2,2) mode of the numerical waveforms is less noisy than the waveform measured along the z axis, and since we have access to the 3PN amplitude correction of the (2,2) mode, we will use the (2,2) waveforms rather than the

z -axis waveforms for our comparisons between NR and PN in Sec. 4.6. We have verified (for all comparisons using post-Newtonian waveforms of ≤ 2.5 PN order in amplitude) that our results do not change significantly when we use z -axis waveforms instead of (2,2) waveforms.

4.3.3 Absorbing amplitude terms into a redefinition of the phase

The logarithms of the orbital frequency parameter x (as well as the constant frequency scale x_0) that appear in the amplitude expressions (4.28), (4.29), and (4.32) can be absorbed into a redefinition of the phase by introducing an auxiliary phase variable $\Psi = \Phi + \delta$. Noting that the $\ln x$ terms first enter at 1.5 PN order, it is straightforward to show that choosing [36, 10, 178]

$$\delta = -3 \frac{M_{\text{ADM}}}{M} x^{3/2} \ln \left(\frac{x}{x_0} \right), \quad (4.33)$$

where $M_{\text{ADM}}/M = 1 - x/8 + \mathcal{O}(x^2)$ for an equal-mass system, will eliminate the $\ln x$ terms from both the (2,2) mode as well as for the polarization waveforms. This follows from

$$h_{(2,2)} = Ae^{-2i\Psi} \quad (4.34)$$

$$= Ae^{-2i\Phi} e^{-2i\delta} \quad (4.35)$$

$$= Ae^{-2i\Phi} (1 - 2i\delta - 2\delta^2 + \mathcal{O}(x^{9/2})) , \quad (4.36)$$

and similarly for the polarization waveforms. Furthermore, since the orbital phase as a function of frequency goes as $x^{-5/2}$ at leading order (see Eq. (4.43a) below), the $\ln x$ terms, which were 1.5PN, 2.5PN, and 3PN order in the original amplitude

expressions, now appear as phase corrections at relative order 4PN, 5PN, and 5.5PN. As these terms are beyond the order to which the orbital phase evolution is known (3.5PN order), it can be argued that these terms can be ignored. Note that the choices of x_0 in Eq. (4.30) and δ in Eq. (4.33) are not unique; they were made to gather all logarithmic terms into one term, as well as to simplify the waveform [36].

4.3.4 Energy balance

The second assumption that goes into making a post-Newtonian waveform as accurate as possible is that of energy balance. It is assumed that the energy carried away by the emission of gravitational waves is balanced by the change in the orbital binding energy of the binary,

$$\frac{dE}{dt} = -\mathcal{L} . \quad (4.37)$$

While this is extremely plausible, it has only been confirmed through 1.5 PN order [37].

Given the above expressions for the energy, flux, and waveform amplitude, there is still a set of choices that must be made in order to produce a post-Newtonian waveform that can be compared to our numerical waveform. These include

1. The PN order through which terms in the orbital energy and luminosity are retained.
2. The procedure by which the energy balance equation is used to obtain $x(t)$ and $\Phi(t)$.

3. The PN order through which terms in the waveform amplitude are kept.
4. The treatment of the $\ln x$ terms. These terms can be included in the amplitude or included in the orbital phase via the auxiliary phase $\Psi := \Phi + \delta$. If the latter is chosen, these terms can be retained or ignored; ignoring them can be justified because they occur at higher order than all known terms in the orbital phase.

We always expand energy and luminosity to the same order, which may be different from the order of the amplitude expansion; both of these expansion orders are indicated explicitly in each of our comparisons. We ignore the $\ln(x/x_0)$ terms in the amplitude by absorbing them into the phase and dropping them because of their high PN order. In the next section we describe several choices for obtaining $x(t)$ and $\Phi(t)$ from the energy balance equation.

4.3.5 Taylor approximants: Computing $\Phi(t)$

In this section we describe how to obtain the orbital phase as a function of time, $\Phi(t)$, using the energy balance equation (4.37). Different methods of doing this exist; here we follow the naming convention of [108]. These methods, and variations of them, are called Taylor approximants, and all formally agree to a given PN order but differ in how higher-order terms are truncated. We discuss four time-domain approximants here, but more can be defined.

TaylorT1

The TaylorT1 approximant is obtained by numerically integrating the ODEs

$$\frac{dx}{dt} = -\frac{\mathcal{L}}{(dE/dx)}, \quad (4.38)$$

$$\frac{d\Phi}{dt} = \frac{c^3}{GM} x^{3/2}, \quad (4.39)$$

to produce $\Phi(t)$. The fraction on the right side of Eq (4.38) is retained as a ratio of post-Newtonian expansions, and is not expanded further before numerical integration. This is the approximant used in the NR–PN comparisons in [157, 206].

TaylorT2

The TaylorT2 approximant is obtained by starting with the parametric solution of the energy balance equation:

$$t(x) = t_0 + \int_x^{x_0} dx \frac{(dE/dx)}{\mathcal{L}} \quad (4.40)$$

$$\Phi(x) = \Phi_0 + \int_x^{x_0} dx \frac{x^{3/2} c^3 (dE/dx)}{GM \mathcal{L}}. \quad (4.41)$$

The integrand of each expression is re-expanded as a single post-Newtonian expansion in x and truncated at the appropriate PN-order; these integrals are then evaluated analytically to obtain for an equal-mass binary [108, 110]. The time is given by

$$t = t_0 - \frac{5GM}{64c^3} x^{-4} \sum_{k=0}^7 T_k x^{k/2}, \quad (4.42a)$$

where the coefficients T_k are

$$T_0 = 1 , \quad (4.42b)$$

$$T_1 = 0 , \quad (4.42c)$$

$$T_2 = \frac{487}{126} , \quad (4.42d)$$

$$T_3 = -\frac{32}{5}\pi , \quad (4.42e)$$

$$T_4 = \frac{2349439}{254016} , \quad (4.42f)$$

$$T_5 = -\frac{1864}{63}\pi , \quad (4.42g)$$

$$T_6 = \left[-\frac{999777207379}{5867769600} + \frac{1597}{48}\pi^2 + \frac{6848}{105}\gamma + \frac{3424}{105}\ln(16x) \right] , \quad (4.42h)$$

$$T_7 = -\frac{571496}{3969}\pi . \quad (4.42i)$$

The orbital phase may be expressed as

$$\Phi = \Phi_0 - \frac{1}{8}x^{-5/2} \sum_{k=0}^7 F_k x^{k/2} , \quad (4.43a)$$

with coefficients F_k given by

$$F_0 = 1 , \quad (4.43b)$$

$$F_1 = 0 , \quad (4.43c)$$

$$F_2 = \frac{2435}{504} , \quad (4.43d)$$

$$F_3 = -10\pi , \quad (4.43e)$$

$$F_4 = \frac{11747195}{508032} , \quad (4.43f)$$

$$F_5 = \frac{1165}{42}\pi \ln x , \quad (4.43g)$$

$$F_6 = \left[\frac{1573812724819}{4694215680} - \frac{7985}{192}\pi^2 - \frac{1712}{21}\gamma - \frac{856}{21}\ln(16x) \right] , \quad (4.43h)$$

$$F_7 = \frac{357185}{7938}\pi . \quad (4.43i)$$

TaylorT3

The TaylorT3 approximant is closely related to TaylorT2. It is obtained by introducing the dimensionless time variable

$$\tau := \frac{\nu c^3}{5GM} (t_0 - t) , \quad (4.44)$$

where $\nu = M_1 M_2 / M^2$ and $\tau^{-1/4} = \mathcal{O}(\epsilon)$. The TaylorT2 expression $t(x)$ is inverted to obtain $x(\tau)$, and truncated at the desired PN order. Then $x(\tau)$ is integrated to obtain

$$\Phi(\tau) = \Phi_0 - \int_{\tau_0}^{\tau} d\tau \frac{5x^{3/2}}{\nu} . \quad (4.45)$$

This procedure yields for an equal-mass binary [42]:

$$x = \frac{1}{4} \tau^{-1/4} \left\{ 1 + \frac{487}{2016} \tau^{-1/4} - \frac{1}{5} \pi \tau^{-3/8} + \frac{1875101}{16257024} \tau^{-1/2} - \frac{1391}{6720} \pi \tau^{-5/8} + \left[-\frac{999777207379}{1502149017600} + \frac{1597}{12288} \pi^2 + \frac{107}{420} \gamma - \frac{107}{3360} \ln\left(\frac{\tau}{256}\right) \right] \tau^{-3/4} - \frac{88451}{282240} \pi \tau^{-7/8} \right\} \quad (4.46)$$

$$\Phi = \Phi_0 - 4\tau^{5/8} \left\{ 1 + \frac{2435}{4032} \tau^{-1/4} - \frac{3}{4} \pi \tau^{-3/8} + \frac{1760225}{1806336} \tau^{-1/2} - \frac{1165}{5376} \pi \tau^{-5/8} \ln \tau + \left[\frac{24523613019127}{3605157642240} - \frac{42997}{40960} \pi^2 - \frac{107}{56} \gamma + \frac{107}{448} \ln\left(\frac{\tau}{256}\right) \right] \tau^{-3/4} + \frac{28325105}{21676032} \pi \tau^{-7/8} \right\} . \quad (4.47)$$

This is the post-Newtonian approximant used in visual comparisons by [75] and in the PN–NR comparisons in [157] at 3PN order in phase.

TaylorT4

In addition to simply numerically integrating the flux-energy equation (4.38), as is done for TaylorT1, one may instead re-expand the right side of (4.38) as a single series and truncate at the appropriate PN order before doing the integration. The phase evolution $\Phi(t)$ can thus be obtained by numerically integrating the ODEs

$$\frac{dx}{dt} = \frac{16c^3}{5GM} x^5 \left\{ 1 - \frac{487}{168} x + 4\pi x^{3/2} + \frac{274229}{72576} x^2 - \frac{254}{21} \pi x^{5/2} + \left[\frac{178384023737}{3353011200} + \frac{1475}{192} \pi^2 - \frac{1712}{105} \gamma - \frac{856}{105} \ln(16x) \right] x^3 + \frac{3310}{189} \pi x^{7/2} \right\}, \quad (4.48)$$

$$\frac{d\Phi}{dt} = \frac{x^{3/2} c^3}{GM}. \quad (4.49)$$

This approximant was not considered in [108], however for consistency with their notation, we call it TaylorT4. TaylorT4 is the primary approximant used in the PN–NR comparisons in [24, 23], and one of the several approximants considered in the PN–NR comparisons in [206]. Reference [75] pointed out that TaylorT4 at 3.5PN order in phase is close to TaylorT3 at 3PN order in phase, and therefore should give similar agreement with numerical results.

4.4 PN–NR Comparison Procedure

4.4.1 What to compare?

There are many ways to compare numerical relativity and post-Newtonian results. For example, the post-Newtonian orbital phase $\Phi(t)$ could be compared with the coordinate phase of the black hole trajectories. However, this and many other

comparisons are difficult to make in a coordinate-independent manner without expending significant effort to understand the relationship between the gauge choices used in post-Newtonian theory and in the NR simulations. Therefore, in order to obtain the most meaningful comparison possible, we attempt to minimize gauge effects by comparing gravitational waveforms as seen by an observer at infinity. The waveform quantity most easily obtained from the numerical relativity code is the Newman–Penrose quantity Ψ_4 , and we will compare its (2,2) component (see Eq. (4.13)), split into phase ϕ and amplitude A according to Eq. (4.14) and extrapolated to infinite extraction radius.

The post-Newtonian formulas in Sec. 4.3 yield the metric perturbation components h_+ and h_\times , which—for a gravitational wave at infinity—are related to Ψ_4 by

$$\Psi_4(t) = \frac{\partial^2}{\partial t^2} (h_+(t) - ih_\times(t)) . \quad (4.50)$$

We numerically differentiate the post-Newtonian expressions for $h_+(t)$ and $h_\times(t)$ twice before computing amplitude and phase using Eq. (4.14). Note that $\phi(t)$ will differ slightly from the phase computed from the metric perturbation directly, as $\tan^{-1}(h_\times/h_+)$, because both the amplitude and phase of the metric perturbation are time dependent. For the same reason, $\phi(t)$ is not precisely equal to twice the orbital phase.

As in Ref. [157], we compare Ψ_4 rather than $h_{+,\times}$ to avoid difficulties arising with fixing the integration constants when integrating the numerically obtained Ψ_4 (see [209] for more details). Both Ψ_4 and $h_{+,\times}$ contain the same information, so differences between both procedures should be minimal.

4.4.2 Matching procedure

Each of the post-Newtonian waveforms has an arbitrary time offset t_0 and an arbitrary phase offset ϕ_0 . These constants can be thought of as representing the absolute time of merger and the orientation of the binary at merger, and we are free to adjust them in order to match NR and PN waveforms. Following [24, 157], we choose these constants by demanding that the PN and NR gravitational-wave phase and gravitational-wave frequency agree at some fiducial frequency ω_m . Specifically, we proceed as follows: We start with a NR waveform $\Psi_4^{\text{NR}}(t)$ and an unshifted PN waveform $\Psi_4^{\text{PN}'}(t)$ that has an arbitrary time and phase shift. After selecting the matching frequency ω_m , we can find (to essentially unlimited accuracy) the time t_c such that the derivative of the PN phase satisfies $\dot{\phi}_{\text{PN}'}(t_c) = \omega_m$, where $\phi_{\text{PN}'}(t)$ is the phase associated with $\Psi_4^{\text{PN}'}(t)$. Similarly, we find the time t_m such that $\dot{\phi}_{\text{NR}}(t_m) = \omega_m$. The time t_m cannot be found to unlimited accuracy, and the uncertainty in t_m is due mainly to residual eccentricity of the NR waveform, as discussed in Sec. 4.5.5. Once we have t_m and t_c , we leave the NR waveform untouched, but we construct a new, shifted, PN waveform

$$\Psi_4^{\text{PN}}(t) = \Psi_4^{\text{PN}'}(t + t_c - t_m) e^{i(\phi_{\text{NR}}(t_c) - \phi_{\text{PN}'}(t_m))} . \quad (4.51)$$

The phase of this new PN waveform is therefore

$$\phi_{\text{PN}}(t) = \phi_{\text{PN}'}(t + t_c - t_m) - \phi_{\text{PN}'}(t_c) + \phi_{\text{NR}}(t_m) , \quad (4.52)$$

which satisfies $\phi_{\text{PN}}(t_m) = \phi_{\text{NR}}(t_m)$ and $\dot{\phi}_{\text{PN}}(t_m) = \omega_m$, as desired. All our comparisons are then made using the new shifted waveform $\Psi_4^{\text{PN}}(t)$ rather than the unshifted waveform $\Psi_4^{\text{PN}'}(t)$.

4.4.3 Choice of Masses

The post-Newtonian expressions as written in Sec. 4.3 involve the total mass M , which corresponds to the the sum of the bare masses of the point particles in post-Newtonian theory. When comparing PN to NR, the question then arises as to which of the many definitions of the mass of a numerically generated binary black hole solution should correspond to the post-Newtonian parameter M . For nonspinning black holes at very large separation, M reduces to the sum of the irreducible masses of the two holes. Neglecting tidal heating, the irreducible masses should be conserved during the inspiral, so that we identify M with the sum of the irreducible masses of the initial data 30c. As discussed in Sec. 4.5 the black hole spins are sufficiently small so that there is no discernible difference between irreducible mass of the black holes and the Christodoulou mass, Eq. (4.3). Of course, the latter would be more appropriate for spinning black holes.

4.5 Estimation of uncertainties

To make precise statements about agreement or disagreement between numerical and post-Newtonian waveforms, it is essential to know the size of the uncertainties in this comparison. When discussing these uncertainties, we must strive to include all effects that may cause our numerical waveform to differ from the post-Newtonian waveforms we compare to. For instance, in addition to considering effects such as numerical truncation error, we also account for the fact that NR and PN waveforms correspond to slightly different physical scenarios: The PN waveforms have identically zero spin and eccentricity, whereas the numerical sim-

ulations have some small residual spin and eccentricity. Table 4.3 lists all effects we have considered; we discuss these in detail below starting in Sec. 4.5.1. All uncertainties are quoted in terms of phase and amplitude differences, and apply to waveform comparisons via matching at a fixed ω_m according to the procedure in Sec. 5.6.1.

Most of the effects responsible for our uncertainties are time dependent, so that it is difficult to arrive at a single number describing each effect. For simplicity, the error bounds in Table 4.3 ignore the junk-radiation noise that occurs in the numerical waveform for $t - r^* \lesssim 1000M$. The extent to which this noise affects the PN–NR comparisons presented below in Secs. 4.6.1 and 4.6.2 will be evident from the noise in the graphs in these sections. Note that all four matching frequencies ω_m occur after the noise disappears at $t - r^* \sim 1000M$. Furthermore, the post-Newtonian waveforms end at different times depending on the PN order and on which particular post-Newtonian approximant is used. Therefore, in order to produce a single number for each effect listed in Table 4.3, we consider only the part of the waveform prior to some cutoff time, which we choose to be the time at which the numerical waveform reaches gravitational-wave frequency $M\omega = -0.1$.

4.5.1 Errors in numerical approximations

The first three error sources listed in Table 4.3 have already been discussed in detail in Sec. 4.2. We estimate numerical truncation error using the difference between the two highest resolution runs after the waveforms have been shifted to agree at some matching frequency ω_m . For $M\omega_m = -0.1$ this difference is shown as the curves labeled ‘30c-1/N5’ in Figs. 4.9 and 4.11, and corresponds to a phase

Table 4.3: **Summary of uncertainties in the comparison between numerical relativity and post-Newtonian expansions**

Effect	$\delta\phi$ (radians)	$\delta A/A$
Numerical truncation error	0.003	0.001
Finite outer boundary	0.005	0.002
Extrapolation $r \rightarrow \infty$	0.005	0.002
Wave extraction at $r_{\text{areal}}=\text{const?}$	0.002	10^{-4}
Drift of mass M	0.002	10^{-4}
Coordinate time = proper time?	0.002	10^{-4}
Lapse spherically symmetric?	0.01	4×10^{-4}
Residual eccentricity	0.02	0.004
Residual spins	0.03	0.001
root-mean-square sum	0.04	0.005

Quoted error estimates ignore the junk-radiation noise at $t \lesssim 1000M$ and apply to times before the numerical waveform reaches gravitational-wave frequency $M\omega = -0.1$. Uncertainties apply to waveform comparisons via matching at a fixed ω_m according to the procedure in Sec. 5.6.1, and represent the maximum values for all four different matching frequencies ω_m that we consider, unless noted otherwise. For the case of matching at $M\omega_m = -0.04$, the phase uncertainty due to residual eccentricity increases to 0.05 radians, thus increasing the root-mean-square sum to 0.06 radians.

difference of 0.003 radians and a relative amplitude difference of 0.001. For other values of ω_m the differences are similar. The effect of the outer boundary is estimated by the difference between the runs 30c-1/N6 and 30c-2/N6, which for $M\omega_m = -0.1$ is shown as the curves labeled ‘30c-2/N6’ in Figs. 4.9 and 4.11, and amount to phase differences of 0.005 radians and relative amplitude differences of 0.002. Errors associated with extrapolation to infinity have been discussed in detail in Figs. 4.13 and 4.15. Specifically, Fig. 4.13 shows that increasing the

extrapolation order between 3 and 4 changes the extrapolated phase by less than 0.005 radians, and Fig. 4.15 confirms that the extrapolated result is robust under changes of extraction radii.

4.5.2 Constancy of extraction radii

If the physical locations of the coordinate-stationary extraction radii happen to change during the evolution, then the extracted gravitational waves will accrue a timing error equal to the light-travel time between the original location and the final location. From Fig. 4.12, we see that the drift in areal radius is less than $0.02M$, resulting in a time uncertainty of $\delta t = 0.02M$. This time uncertainty translates into a phase uncertainty via

$$\delta\phi = M\omega \times (\delta t/M), \quad (4.53)$$

which yields $\delta\phi \approx 0.002$, when $M\omega = -0.1$ (the value at the end of the PN comparison) was used.

To estimate the effect of this time uncertainty on the amplitude, we first note that to lowest order in the post-Newtonian parameter x (defined in Eq. (4.23)), the wave amplitude of Ψ_4 scales like x^4 . Also, from Eq. (4.48), we have $dx/dt = 16/(5M)x^5$. Therefore,

$$\frac{\delta A}{A} \sim \frac{d \ln A}{dx} \frac{dx}{dt} \delta t \sim \frac{64}{5} (-M\omega/2)^{8/3} \frac{\delta t}{M}, \quad (4.54)$$

where we have used the fact that the gravitational-wave frequency ω is approximately twice the orbital frequency. For a time uncertainty $\delta t = 0.02M$, Eq. (4.54) gives $\delta A/A \approx 10^{-4}$ for $M\omega = -0.1$.

4.5.3 Constancy of mass

Our comparisons with post-Newtonian formulas assume a constant post-Newtonian mass parameter M , which we set equal to the total irreducible mass of the black holes in the numerical simulation. If the total mass of the numerical simulation is not constant, this will lead to errors in the comparison. For example, changes in t/M caused by a changing mass will lead to phase differences. Figure 4.4 demonstrates that the irreducible mass is conserved to a fractional accuracy of about $\delta M/M \approx 5 \times 10^{-6}$.

This change in irreducible mass could be caused by numerical errors, or by a physical increase of the mass of each black hole through tidal heating. For our simulations, $M(t)$ *decreases* during the run (this is not apparent from Fig. 4.4 which plots absolute values), thus contradicting the second law of black hole thermodynamics. Moreover, the increase in $M(t)$ through tidal heating is much smaller than the observed variations in $M(t)$ (see, e.g. [216]). Therefore, the variations in $M(t)$ are numerical errors, and we need to bound the influence of these errors on the comparison to post-Newtonian expansions.

Over an evolution time of $t/M = 4000$, the observed mass uncertainty of $\delta M/M \approx 5 \times 10^{-6}$ results in an uncertainty in the overall time interval of $\delta(t/M) = (t/M) \times (\delta M/M) \approx 0.02$. This time uncertainty translates into a phase uncertainty of $\delta\phi \approx 0.002$, using Eq. (4.53) for $M\omega = -0.1$. Note that the effect of the black-hole spins on the mass is negligible relative to the numerical drift of 5×10^{-6} . This is because the spins of the holes are bounded by $S/M_{\text{irr}} < 2 \times 10^{-4}$ and the spin enters quadratically into the Christodoulou formula (4.3). The error in the gravitational-wave amplitude caused by time uncertainties due to varying mass

is $\delta A/A \approx 10^{-4}$ using Eq. (4.54) for $M\omega = -0.1$. An error in the mass will affect the amplitude not only via a time offset, but also because the amplitude is proportional to $(-M\omega/2)^{8/3}$ (to lowest PN order). However, this additional error is very small, $\delta A/A \approx (8/3)\delta M/M \approx 10^{-5}$.

4.5.4 Time-coordinate ambiguity

We now turn to two possible sources of error that have not yet been discussed, both of which are related to ambiguity in the time coordinate. The basic issue is that the time variable t in post-Newtonian expansions corresponds to proper time in the asymptotically flat region, but the time t in numerical simulations is coordinate time. These two quantities agree only if the lapse function N approaches unity at large distances. To verify this, we decompose N in spherical harmonics centered on the center of mass of the system,

$$N(r, \vartheta, \varphi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l N_{l,m}(r) Y_{l,m}(\vartheta, \varphi). \quad (4.55)$$

The angular average of the lapse function, $\langle N \rangle(r) := \sqrt{4\pi} N_{0,0}(r)$ should then approach unity for $r \rightarrow \infty$, and all other modes $N_{l,m}(r)$ should decay to zero. Fig. 4.17 plots $\langle N \rangle(r) - 1$ versus M/r for three different evolution times. Fitting $\langle N \rangle(r) - 1$ for $r > 100M$ to a polynomial in M/r gives a y intercept of magnitude less than 5×10^{-6} for all three times, and for polynomial orders of two through five. Therefore, the coordinate time of the evolution agrees with proper time at infinity to better than $\delta t/M = t/M \times 5 \times 10^{-6} \approx 0.02$, which induces a phase error of at most $\delta\phi \approx 0.002$ and an amplitude error of $\delta A/A \approx 10^{-4}$ (see Eqs. (4.53) and (4.54)).




Figure 4.17: **Asymptotic behavior of the average lapse at large radii**

This figure displays the angular average of the lapse as a function of (inverse) radius at $t/M = 0, 1900$, and 3800 . The lines all lie on top of each other, but end at different radii, because of the infall of the outer boundary. Note that—at all three times—the lapse is approaching 1 roughly as $1/r$. The horizontal axis is evenly spaced in $1/r$.




Figure 4.18: **Asymptotic behavior of higher angular moments of the lapse at large radii**

This figure shows the dominant higher multipole moments of the lapse at $t = 0, 1900M$, and $3800M$, as a function of (inverse) radius. Despite the obvious oscillations, most components do seem to decay to 0 with increasing radius roughly as fast as $1/r$, at all times. The horizontal axis is evenly spaced in $1/r$.

The second source of error related to the lapse is shown in Fig. 4.18, which presents the three largest higher-order moments $N_{l,m}(r)$. All modes decay to zero as $r \rightarrow \infty$ except, perhaps, the real part of the $N_{2,2}$ mode at $t/M = 3800$. This mode seems to approach a value of about 5×10^{-5} . At $t = 1900M$, this mode still decays nicely to zero, hence the maximum time uncertainty introduced by this effect at late times is $\delta t = 1900M \times 5 \times 10^{-5} \approx 0.1M$, resulting in a potential phase uncertainty of $\delta\phi \approx 0.01$ and a potential amplitude uncertainty of $\delta A/A \approx 4 \times 10^{-4}$.

4.5.5 Eccentricity

We estimated the eccentricity during the numerical simulation with several of the methods described in [75, 209, 167], and have consistently found $e \lesssim 6 \times 10^{-5}$. This eccentricity can affect our comparison to a post-Newtonian waveform of a quasi-circular (i.e., zero eccentricity) inspiral in three ways.

Change in rate of inspiral

The first effect arises because an eccentric binary has a different inspiral rate than a noneccentric binary; physically, this can be understood by noting that the gravitational flux and orbital energy depend upon the eccentricity, and therefore modify the rate at which the orbital frequency evolves assuming energy balance. Reference [184] has derived the first-order correction in the phase of the gravitational wave due to this effect. Converting their result to our notation and restricting to the equal-mass case yields

$$\frac{1}{(dx/dt)} = \frac{5GM}{16c^3x^5} \left(1 - \frac{157}{24} e_i^2 \left(\frac{x_i}{x} \right)^{19/6} \right), \quad (4.56)$$

where e_i is the initial eccentricity and x_i is the initial value of the orbital-frequency parameter. Substituting this into Eq. (4.41) yields

$$\Phi = \Phi_0 - \frac{1}{8}x^{-5/2} + \frac{785}{2176}e_i^2 x_i^{19/6} x^{-17/3} . \quad (4.57)$$

Using $e_i = 6 \times 10^{-5}$ and integrating over the frequency range from the start of our simulation to the matching frequency of $M\omega = -0.1$ yields a phase shift of $\sim 2 \times 10^{-6}$, which is dwarfed by many other error sources, such as the uncertainty in the numerical mass M . (See Sec. 4.5.3.)

Uncertainty in matching time

The second way in which eccentricity affects our comparison is by introducing errors in our procedure for matching the PN and NR waveforms. Recall that our matching procedure involves determining the time t_m at which the gravitational-wave frequency ω takes a certain value $M\omega_m$; eccentricity modulates the instantaneous gravitational-wave frequency $\omega(t)$ via

$$\omega(t) = \langle \omega(t) \rangle [1 + 2e \cos(\Omega_r t)] , \quad (4.58)$$

where $\langle \omega(t) \rangle$ represents the averaged “non-eccentric” evolution of the gravitational-wave frequency, and Ω_r is the frequency of radial oscillations, which is approximately equal to the orbital frequency. We see that ω can differ from $\langle \omega \rangle$ by as much as $2e \langle \omega \rangle \approx 2e\omega$. This could induce an error in the determination of t_m by as much as

$$|\delta t_m| = \frac{|\delta \omega|}{\dot{\omega}} \approx \frac{2e\omega}{\dot{\omega}} . \quad (4.59)$$

We can simplify this expression by using Eq. (4.48) to lowest order, and by noting that the gravitational-wave frequency is approximately twice the orbital fre-

quency. We find

$$|\delta t_m| \leq e \frac{5M}{12} \left(-\frac{M\omega}{2} \right)^{-8/3} . \quad (4.60)$$

This uncertainty is largest at *small* frequencies, because the frequency changes much more slowly. For $M\omega = -0.04$, we find $|\delta t_m| \lesssim 0.9M$, and for $M\omega = -0.1$, we find $|\delta t_m| \lesssim 0.1M$.

To determine how uncertainties in t_m translate into phase differences, recall that in the matching procedure described in Sec. 5.6.1, t_m enters into the phase of the shifted PN waveform according to Eq. (4.52). Therefore the phase difference that we compute between the PN and NR waveforms is

$$\begin{aligned} \Delta\phi(t) &= \phi_{\text{PN}}(t) - \phi_{\text{NR}}(t) \\ &= \phi_{\text{PN}'}(t + t_c - t_m) - \phi_{\text{NR}}(t) + \phi_{\text{NR}}(t_m) - \phi_{\text{PN}'}(t_c) . \end{aligned} \quad (4.61)$$

Then the error in $\Delta\phi$ is found by Taylor-expanding Eq. (4.61):

$$\begin{aligned} \delta\phi := \delta(\Delta\phi(t)) &= (\dot{\phi}_{\text{PN}'}(t + t_c - t_m) - \dot{\phi}_{\text{NR}}(t_m)) \delta t_m \\ &= (\dot{\phi}_{\text{PN}}(t) - \omega_m) \delta t_m . \end{aligned} \quad (4.62)$$

Our simulations (and therefore the comparisons to post-Newtonian theory) start at $M\omega \approx -0.033$, so that the maximum error $\delta\phi$ within our comparison at times *before* the matching frequency will be

$$|\delta\phi_{\text{before}}| \leq |-0.033 - \omega_m| |\delta t_m| . \quad (4.63)$$

Combining Eqs. (4.63) and (4.59), and using $e \approx 6 \times 10^{-5}$, we find that $\delta\phi_{\text{before}} < -0.01$ radians for all four of our matching frequencies $-M\omega_m = 0.04, 0.05, 0.063$, and 0.1. The maximum error $\delta\phi$ within our comparison at times *after* the matching frequency is

$$|\delta\phi_{\text{after}}| \leq |0.1 - \omega_m| |\delta t_m| , \quad (4.64)$$

because we end our comparisons to post-Newtonian theory at $M\omega = -0.1$. Eq. (4.64) evaluates to 0.05 radians for $M\omega_m = -0.04$, and is less than about 0.02 radians for the three higher matching frequencies.

The error in the gravitational-wave amplitude caused by an error in t_m can be estimated by Eq. (4.54). A conservative estimate using $\delta t = 0.9M$ still gives a small error, $\delta A/A \approx 0.004$.

Note that the bounds on $\delta\phi_{\text{before}}$ and $\delta\phi_{\text{after}}$ are proportional to the eccentricity of the numerical simulation. Even with eccentricity as low as 6×10^{-5} , this effect is one of our largest sources of error for the PN–NR comparison. (See Table 4.3). This is the reason why the simpler eccentricity removal procedure of Husa et al. [167] (resulting in $e = 0.0016$) is not adequate for our purposes.

Periodic modulation of phase and amplitude

The third effect of orbital eccentricity is a periodic modulation of the gravitational-wave phase and amplitude. If we assume that $\langle\omega(t)\rangle$ varies on much longer time scales than $1/\Omega_r$ (which is true at large separation) then integration of Eq. (4.58) in time yields

$$\phi(t) = \langle\phi(t)\rangle + 2e \frac{\langle\omega\rangle}{\Omega_r} \sin(\Omega_r t) . \quad (4.65)$$

Because $\Omega_r \approx \Omega \approx -\langle\omega\rangle/2$, we therefore find that the gravitational-wave phase consists of the sum of the desired “circular” phase, $\langle\phi(t)\rangle$, plus an oscillatory component with amplitude $4e \approx 2 \times 10^{-4}$. This oscillatory component, however, is much smaller than other uncertainties of the comparison, for instance the uncertainty in determination of t_m .

Residual eccentricity will also cause a modulation of the gravitational-wave

amplitude in a manner similar to that of the phase. This is because eccentricity explicitly enters the post-Newtonian amplitude formula at 0PN order [243]. This term is proportional to e , and since $e \lesssim 6 \times 10^{-5}$ its contribution to the amplitude error is small compared to the effect due to uncertainty in t_m .

While oscillations in phase and amplitude due to eccentricity are tiny and dwarfed by other uncertainties in the PN–NR comparison, their characteristic oscillatory behavior makes them nevertheless visible on some of the graphs we present below—for instance, Figs. 4.26 and 4.27.

4.5.6 Spin

We now turn our attention to effects of the small residual spins of the black holes. References [135] and [43] compute spin–orbit coupling up to 2.5 post-Newtonian order, and find that the orbital phase, Eq. (4.43a), acquires the following spin contributions

$$\Phi_S(x) = -\frac{1}{32\nu} \sum_{i=1,2} \chi_i \left\{ \left(\frac{565}{24} \frac{M_i^2}{M^2} + \frac{125\nu}{8} \right) x^{-1} - \left[\left(\frac{681145}{4032} + \frac{965\nu}{28} \right) \frac{M_i^2}{M^2} + \frac{37265\nu}{448} + \frac{1735\nu^2}{56} \right] \ln x \right\}, \quad (4.66)$$

where $\chi_i = \vec{S}_i \cdot \hat{L} / M_i^2$ is the projection of the dimensionless spin of the i -th hole onto the orbital angular momentum. For equal-mass binaries with spins $\chi_1 = \chi_2 := \chi$, this reduces to

$$\Phi_S(x) = -\chi \left(\frac{235}{96} x^{-1} - \frac{270625}{16128} \ln x \right). \quad (4.67)$$

Our comparisons to post-Newtonian theory are performed over the orbital-frequency range of $0.0167 \leq M\Omega \leq 0.05$, corresponding to $0.065 \leq x \leq 0.136$. The gravitational-

wave phase is approximately twice the orbital phase, so that the spin-orbit coupling contributes

$$\delta\phi_S = 2 [\Phi_S(0.065) - \Phi_S(0.136)] \approx -64\chi \quad (4.68)$$

to the gravitational-wave phase. In Sec. 4.2.2 we estimated $|\vec{S}|/M_{\text{irr}}^2 < 5 \times 10^{-4}$, where M_{irr} is the irreducible mass of either black hole. Because $\chi \leq |\vec{S}|/M_{\text{irr}}^2 \approx 5 \times 10^{-4}$, the residual black-hole spins contribute less than 0.03 radians to the overall gravitational-wave phase.

We now turn to errors in the amplitude comparison caused by residual spin. We can compute the error in orbital frequency as $\delta\Omega = \dot{\Phi}_S$ using Eq. (4.67):

$$\begin{aligned} \delta\Omega &= \chi \frac{\dot{x}}{x} \left(\frac{235}{96} x^{-1} + \frac{270625}{16128} \right) \\ &= \chi x^4 \frac{16}{5M} \left(\frac{235}{96} x^{-1} + \frac{270625}{16128} \right), \end{aligned} \quad (4.69)$$

where we have used Eq. (4.48). Because the amplitude of Ψ_4 scales like $\Omega^{8/3}$, we arrive at

$$\frac{\delta A}{A} = \frac{8}{3} \frac{\delta\Omega}{\Omega} = \chi x^{5/2} \frac{128}{15} \left(\frac{235}{96} x^{-1} + \frac{270625}{16128} \right), \quad (4.70)$$

which for $M\omega_m = -0.1$ (i.e., $x = 0.136$) gives $\delta A/A = 2.0\chi \sim 1.0 \times 10^{-3}$.

Spin-orbit terms also contribute directly to the amplitude [177, 249]. The leading-order contribution (for an equal-mass binary with equal spins) contributes a term $\delta A/A \sim (4/3)\chi x^{3/2}$, which is the same order of magnitude as the previous error, 10^{-3} .




Figure 4.19: **Comparison of numerical simulation with TaylorT1 3.5/2.5 waveforms—phase difference**

Plotted are comparisons for four values of ω_m . The filled dot on each curve shows the point at which $\dot{\phi} = \omega_m$. Also shown is the difference between the numerical and restricted (i.e., 3.5PN phase, 0PN amplitude) TaylorT1 matched at $M\omega_m = -0.1$, shown as the dashed curve.




Figure 4.20: **Comparison of numerical simulation with TaylorT1 3.5/2.5 waveforms—relative amplitude difference**

Plotted are comparisons for four values of ω_m . The filled dot on each curve shows the point at which $\dot{\phi} = \omega_m$. Also shown is the difference between the numerical and restricted (i.e., 3.5PN phase, 0PN amplitude) TaylorT1 matched at $M\omega_m = -0.1$, shown as the dashed curve.

4.6 Results

4.6.1 Comparison with individual post-Newtonian approximants

We compare our simulations with four different post-Newtonian approximants: the TaylorT1, TaylorT2, TaylorT3, and TaylorT4 waveforms. These four waveforms agree with each other up to their respective post-Newtonian expansion orders, but they differ in the way that the uncontrolled higher-order terms enter. We start with the comparison to TaylorT1.

TaylorT1 (3.5PN phase, 2.5PN amplitude)

Figures 4.19 and 4.20 compare the numerical simulation to TaylorT1 3.5/2.5 waveforms (i.e., expansion order 3.5PN in phase and 2.5PN in amplitude, the highest expansion orders currently available for generic direction; see Sec. 4.3.2). The first of these shows the phase difference, where we find differences of more than a radian for all four matching frequencies we consider: $-M\omega_m = 0.04, 0.05, 0.063,$ and 0.1 .

For our largest matching frequency, $M\omega_m = -0.1$, the phase differences are small toward the end of the run by construction. Nevertheless, a phase difference of more than 0.5 radians builds up in the ~ 1.5 cycles after the matching point, before the TaylorT1 template generation fails. Recall that $M\omega_m = -0.1$ occurs about 2.2 gravitational-wave cycles before our simulations fail, which is still about 1.5 cycles before merger. However, the largest phase disagreement for $M\omega_m = -0.1$ builds up at early times, reaching 1.5 radians at the beginning of




Figure 4.21: **Numerical and TaylorT1 3.5/2.5 waveforms**

This figure shows a detailed view of the dephasing between the numerical and post-Newtonian waveforms during the last 10 gravitational-wave cycles. The waveforms are matched at $M\omega_m = -0.04$.

our simulation, about 28 cycles before the matching (~ 30 cycles before the end of the simulation), and still showing no sign of flattening even at the start of our simulation.

To achieve phase coherence with the early inspiral waveform, it is therefore




Figure 4.22: **Comparison of numerical simulation with TaylorT2 3.5/2.5 waveforms—phase difference**

Plotted are comparisons for four values of $M\omega_m$. The filled dot on each curve shows the point at which $\dot{\phi} = \omega_m$.




Figure 4.23: **Comparison of numerical simulation with TaylorT2 3.5/2.5 waveforms—relative amplitude difference**

Plotted are comparisons for four values of $M\omega_m$. The filled dot on each curve shows the point at which $\dot{\phi} = \omega_m$.

necessary to match earlier than $M\omega_m = -0.1$. Obviously, phase differences at early times become smaller when the matching point itself is moved to earlier time. For instance, $M\omega_m = -0.063$ (about eight gravitational-wave cycles before the end of our simulation), results in phase differences less than 0.5 radians during the 22 earlier cycles of our evolution. However, the phase difference $\phi_{\text{PN}} - \phi_{\text{NR}}$ does not level off at early times within the length of our simulation, so it seems quite possible that the phase difference may grow to a full radian or more if the numerical simulations could cover many more cycles. We thus estimate that for TaylorT1, to achieve 1-radian phase coherence with the early inspiral may require matching more than 10 cycles before merger. To achieve more stringent error bounds in phase coherence will require matching even earlier: for instance it appears one needs to use $M\omega_m = -0.04$ (about 20 cycles before the end of our simulation) for a phase error of less than $\lesssim 0.1$ radians.

While matching at small ω_m yields good phase coherence early in the run, it produces much larger phase differences late in the run. For example, matching at $M\omega_m = -0.04$ results in a phase difference of almost 2 radians at frequency $M\omega = -0.1$. This rather dramatic disagreement is illustrated in Fig. 4.21, which plots both the numerical and the TaylorT1 waveform, matched at $M\omega_m = -0.04$.

Figure 4.19 also includes a comparison to the so-called restricted TaylorT1 template, where only the leading order amplitude terms are used (i.e., 0PN in amplitude). The reason that higher-order amplitude terms affect the phase differences at all is because we are plotting gravitational-wave phase, not orbital phase. However, we see that the effect of these higher-order amplitude terms on the phase difference is small.

We now turn our attention to comparing the amplitudes of the post-Newtonian and numerical waveforms. Figure 4.20 shows relative amplitude differences between TaylorT1 3.5/2.5 and the numerical waveforms. At early times, the amplitudes agree to within 2 or 3 percent, the agreement being somewhat better when the matching is performed at early times. At late times, the amplitudes disagree dramatically; a large fraction of this disagreement lies probably in the fact the post-Newtonian point of merger (i.e., the point at which the amplitude diverges) occurs at a different time than the numerical point of merger. We also plot the amplitude of the restricted TaylorT1 template. The disagreement between restricted TaylorT1 and the numerical result is much larger, about 5 percent.

Hannam et al. [157] performed a similar comparison, matching their waveforms with a restricted TaylorT1 waveform (i.e., 3.5/0.0) generated using the LIGO Algorithm Library (LAL) [94]. The phase difference they observe for waveforms matched at $M\omega_m = -0.1$ is consistent with our results within numerical errors. When matching TaylorT1 3.5/0.0 early in their simulation (at $M\omega_m = -0.0455$), however, Hannam et al. find a cumulative phase difference of 0.6 radians at $M\omega = -0.1$. From Fig. 4.19 we find a quite different value of 1.5 radians for our simulation. This disagreement might be caused by the use of the finite extraction radius $R = 90M$ for the gravitational-wave phase in Hannam et al.: Figure 4.13 shows that extracting at a finite radius leads to a systematic phase error, which will induce a systematic error in determination of the matching time of Hannam et al. This error is amplified by the increasing gravitational-wave frequency toward merger.




Figure 4.24: **Comparison of numerical simulation with TaylorT3 3.5/2.5 waveforms—phase difference**

Plotted are comparisons for three values $M\omega_m$. The filled dot on each curve shows the point at which $\dot{\phi} = \omega_m$. The lines end when the frequency of the TaylorT3 waveform reaches its maximum, which happens before $M\omega = -0.1$, so that matching frequency is absent. However, TaylorT3 3.0/3.0 does achieve that frequency, so we show this curve, matched at $M\omega_m = -0.1$.




Figure 4.25: **Comparison of numerical simulation with TaylorT3 3.5/2.5 waveforms—relative amplitude difference**

Plotted are comparisons for three values $M\omega_m$. The filled dot on each curve shows the point at which $\dot{\phi} = \omega_m$. The lines end when the frequency of the TaylorT3 waveform reaches its maximum, which happens before $M\omega = -0.1$, so that matching frequency is absent.

TaylorT2 (3.5PN phase, 2.5PN amplitude)

Figures 4.22 and 4.23 present the comparison between the numerical waveform and the TaylorT2 approximant. The overall trends are very similar to the TaylorT1 comparison of Figs. 4.19 and 4.20, though the phase differences are smaller by about a factor of 2 when matching at $M\omega_m = -0.1$, and smaller by a factor of 3 to 4 when matching earlier. To our knowledge TaylorT2 has never been compared to a numerical simulation; we include it here mainly for completeness.

TaylorT3 (3.5PN and 3.0PN phase, 2.5PN amplitude)

Figures 4.24 and 4.25 are the same as Figs. 4.19 and 4.20 except that they compare numerical simulations to the TaylorT3 family of waveforms. Two differences between TaylorT1 and TaylorT3 are readily apparent from comparing these two figures. The first is that we do not match TaylorT3 3.5/2.5 waveforms at $M\omega_m = -0.1$. This is because the frequency of TaylorT3 3.5/2.5 waveforms reaches a maximum shortly before the formal coalescence time of the post-Newtonian template, and then *decreases*. The maximal frequency is less than 0.1, so that matching at $M\omega_m = -0.1$ is not possible. For this reason, we have also shown in Figs. 4.24 and 4.25 a comparison with a TaylorT3 3.0/3.0 waveform matched at $M\omega_m = -0.1$. The other major difference between the TaylorT3 3.5/2.5 and TaylorT1 3.5/2.5 comparison is that the phase difference, $\phi_{\text{PN}} - \phi_{\text{NR}}$, has a different sign. While TaylorT1 3.5/2.5 spirals in *more rapidly* than the numerical simulation, TaylorT3 3.5/2.5 *lags behind*. Interestingly, the phase differences from the numerical simulation for both TaylorT1 3.5/2.5 and TaylorT3 3.5/2.5 are of about equal magnitude (but opposite sign). The TaylorT3 3.0/3.0 comparison

matched at $M\omega_m = -0.1$ has smaller phase differences than does the TaylorT3 3.5/2.5 comparison, but the slope of the 3.0/3.0 curve in Fig. 4.24 is nonzero at early times, so it appears that TaylorT3 3.0/3.0 will accumulate significant phase differences at even earlier times, prior to the start of our simulation. In Fig. 4.30 it can be seen that matching TaylorT3 3.0/3.0 at $M\omega_m = -0.04$ leads to a good match early, but leads to a phase difference of 0.6 radians by $M\omega = -0.1$.

Hannam et al. [157] match a TaylorT3 3.0/0.0 waveform at $M\omega_m = -0.1$ and $M\omega_m = -0.0455$. Matching at $M\omega_m = -0.1$ again gives phase differences consistent with our results within numerical errors. Matching at $M\omega_m = -0.0455$, Hannam et al. find a phase difference of 0.9 radians, while we find a smaller value of 0.5 radians. Again, this difference could be due to the finite extraction radius used by Hannam et al.

TaylorT4 (3.5PN phase, 2.5PN amplitude)

Figures 4.26 and 4.27 are the same as Figs. 4.19 and 4.20 except that they compare numerical simulations to the TaylorT4 PN waveforms. The agreement between TaylorT4 waveforms and the numerical results is astonishingly good, far better than the agreement between NR and either TaylorT1 or TaylorT3. The gravitational-wave phase difference lies within our error bounds for the entire comparison region $M\omega \geq -0.1$, agreeing to 0.05 radians or better over 29 of 30 gravitational-wave cycles. Ref. [24] found agreement between TaylorT4 and their numerical simulation to the level of their numerical accuracy (~ 2 radians), agreeing to roughly 0.5 radians in the gravitational frequency range of




Figure 4.26: **Comparison of numerical simulation with TaylorT4 3.5/2.5 waveforms—phase difference**

Plotted are comparisons for four values of $M\omega_m$. The filled dot on each curve shows the point at which $\dot{\phi} = \omega_m$. This plot also includes two phase comparisons with expansions of different PN order in amplitude, as labeled, for $M\omega_m = -0.1$.




Figure 4.27: **Comparison of numerical simulation with TaylorT4 3.5/2.5 waveforms—relative amplitude difference**

Plotted are comparisons for four values of $M\omega_m$. The filled dot on each curve shows the point at which $\dot{\phi} = \omega_m$.

$-0.054 \geq M\omega \geq -0.1$. Ref. [206] found that NR agrees better with TaylorT4 than with TaylorT1, but the larger systematic and numerical errors of the numerical waveforms used in these studies did not allow them to see the surprising degree to which NR and TaylorT4 agree. The gravitational-wave amplitude of TaylorT4 agrees with the NR waveform to about 1–2 percent at early times, and 8 percent at late times. In Fig. 4.28 we plot the NR and TaylorT4 waveforms; the two waveforms are visually indistinguishable on the plot, except for small amplitude differences in the final cycles.

In Fig. 4.26 we also show phase comparisons using PN waveforms computed to 3.5PN order in phase but to 0PN and 3.0PN orders in amplitude, for the case $M\omega_m = -0.1$. The PN order of the amplitude expansion affects the phase comparison because we are plotting differences in gravitational-wave phase and not orbital phase. The differences between using 0PN, 2.5PN, and 3.0PN amplitude expansions are evident on the scale of the graph, but because these differences are smaller than our estimated uncertainties (see Table 4.3), we cannot reliably conclude which of these most closely agrees with the true waveform.

Figure 4.29 presents amplitude differences between NR and TaylorT4 as the post-Newtonian order of the amplitude expansion is varied, but the phase expansion remains at 3.5PN. The 2.5PN amplitude curve was already included in Fig. 4.27. We see clearly that higher order amplitude corrections generally result in smaller differences. The 3PN amplitude correction to the (2,2) mode recently derived by Kidder [178] improves agreement dramatically over the widely known 2.5PN amplitude formulas. Unfortunately, the 3PN amplitude correction to the




Figure 4.28: Numerical and TaylorT4 3.5/3.0 waveforms

The PN waveform is matched to the numerical one at $M\omega_m = -0.04$, indicated by the small circle. The lower panel shows a detailed view of the end of the waveform.




Figure 4.29: **TaylorT4 amplitude comparison for different PN orders**

Shown is the relative difference in gravitational-wave amplitude between TaylorT4 and numerical $\Psi_4^{2,2}$ waveforms as a function of time. Matching is performed at $M\omega_m = -0.04$. All curves use 3.5PN order in phase but different PN orders (as labeled) in the amplitude expansion.

entire waveform, including all (l, m) modes, is not known.⁸

4.6.2 Comparing different post-Newtonian approximants

The previous section presented detailed comparisons of our numerical waveforms with four different post-Newtonian approximants. We now turn our attention to some comparisons between these approximants. In this section we also explore further how the post-Newtonian order influences agreement between numerical and post-Newtonian waveforms.

Figure 4.30 presents phase differences as a function of time for all four PN approximants we consider here and for different PN orders. The post-Newtonian and numerical waveforms are matched at $M\omega_m = -0.04$, about 9 cycles after the beginning of the numerical waveform, and about 21 cycles before its end. We find that some PN approximants at some particular orders agree exceedingly well with the numerical results. The best match is easily TaylorT4 at 3.5PN order, and the next best match is TaylorT4 at 2.0PN order. Some approximants behave significantly worse, such as the TaylorT1 and TaylorT4 waveforms at 2.5PN order. The 2.5PN and 3PN TaylorT3 waveforms agree very well with the numerical waveform at early times, but at late times they accumulate a large phase difference; the 2.5PN TaylorT3 waveform ends even before the numerical waveform reaches $M\omega = -0.1$ (the rightmost vertical brown line in Fig. 4.30).

We also find that all four PN approximants, when computed to 3PN order

⁸To get the complete waveform to 3PN order, only the $(2,2)$ mode must be known to 3PN order; other modes must be known to smaller PN orders. For an equal-mass, nonspinning binary, all modes except the $(3,2)$ mode are currently known to sufficient order to get a complete 3PN waveform [178].




Figure 4.30: **Phase comparison for various PN approximants**

Shown is the difference in gravitational-wave phase between each post-Newtonian approximant and the numerical $\Psi_4^{2,2}$ waveforms as a function of time. The two vertical brown lines indicate when the numerical simulation reaches $M\omega = -0.063$ and -0.1 , respectively; the labels along the top horizontal axes give the number of gravitational-wave cycles before $M\omega = -0.1$.




Figure 4.31: **Late-time phase comparison for various PN approximants**

Same as Fig. 4.30, but showing only the last stage of the inspiral. The horizontal axis ends at the estimated time of merger, $(t-r^*)_{\text{CAH}} = 3950M$. See Sec 4.2.7 for details. Note that the top and bottom panels use different vertical scales.

or higher, match the numerical waveform (and each other) quite closely at early times, when all PN approximants are expected to be accurate. However, at late times, $t - r^* > 2500M$, the four PN approximants begin to diverge, indicating that PN is beginning to break down.

Figure 4.31 is an enlargement of Fig. 4.30 for the last 10 gravitational wave cycles before merger. This figure shows in more detail how the different PN approximants behave near merger.

Figures 4.32 and 4.33 present similar results in a different format. We compute the phase differences between the numerical waveform and the various post-Newtonian approximants at the times when the numerical waveform reaches gravitational-wave frequencies $M\omega = -0.063$ and $M\omega = -0.1$ (the times corresponding to these frequencies are also indicated by brown lines in Fig. 4.30). We then plot these phase differences as a function of the post-Newtonian order (using equal order in phase and amplitude, except for 3.5PN order, where we use 3.0PN in amplitude). Three PN approximants end before the time at which $M\omega = -0.1$: TaylorT1 2.0/2.0, TaylorT3 2.5/2.5, TaylorT3 3.5/3.0. These data points therefore cannot be included in Fig. 4.33.

The general trend seen in Figs. 4.32 and 4.33 is that the phase difference decreases with increasing PN order. However, this convergence is not monotonic, and the scatter in the figures can be larger than the phase differences themselves. For example, the 0PN waveforms are about as good as the 2.5PN waveforms for TaylorT1 and TaylorT4, and the 2PN TaylorT4 waveform agrees with the numerical results much better than do either the 2.5PN or 3PN TaylorT4 waveforms.




Figure 4.32: **Phase differences between numerical and post-Newtonian waveforms at $t = t_{(M\omega = -0.063)}$**

Waveforms are matched at $M\omega_m = -0.04$, and phase differences are computed at the time when the numerical simulation reaches $M\omega = -0.063$. Differences are plotted versus PN order (equal order in phase and amplitude, except the '3.5 PN' points are 3.5/3.0). The thin gray bands indicate upper bounds on the uncertainty of the comparison as discussed in Sec. 4.5.1.




Figure 4.33: **Phase differences between numerical and post-Newtonian waveforms at $t = t_{(M\omega=-0.1)}$**

Waveforms are matched at $M\omega_m = -0.04$, and phase differences are computed at the time when the numerical simulation reaches $M\omega = -0.1$. Differences are plotted versus PN order (equal order in phase and amplitude, except the ‘3.5 PN’ points are 3.5/3.0). The 1PN data points are off scale, clustering at -15 radians. The thin gray bands indicate upper bounds on the uncertainty of the comparison as discussed in Sec. 4.5.1.

Considering these figures, it seems difficult to make statements about the convergence with PN order for any particular PN approximant, or statements about which PN orders are generally “good”, in some sense. Given that at fixed PN order the different approximants differ merely by the treatment of uncontrolled higher-order terms, the scatter in these results represents—in some sense—the truncation error at each PN order. While some PN approximants at certain orders may show better agreement with the numerical simulation, we are not aware of any means to predict this besides direct comparisons to numerical simulations (as is done here). In particular, these figures suggest that the remarkable agreement between our numerical results and the 3.5PN TaylorT4 approximant may be simply due to luck; clearly, more PN–NR comparisons are needed, with different mass ratios and spins, to see if this is the case.

4.7 Conclusions

We have described numerical simulations of an equal-mass, nonspinning binary black hole spacetime covering 15 orbits of inspiral just prior to the merger of the two black holes. Using a multi-domain pseudospectral method we are able to extract the gravitational-wave content measured by a distant observer with a phase accuracy of better than 0.02 radians over the roughly 30 cycles of gravitational radiation observed. We demonstrate that in order to achieve this accuracy it is necessary to accurately extrapolate the waveform from data obtained at extraction surfaces sufficiently far from the center of mass of the system. When comparing to zero-spin, zero-eccentricity PN formulas, our phase uncertainty increases to 0.05 radians because the numerical simulation has a small but nonzero orbital

eccentricity and small but nonzero spins on the holes.

Judging from the case in which we match at $M\omega_m = -0.04$, our numerical simulations are consistent (within our estimated phase uncertainty) with all PN approximants (at the highest PN order) from the beginning of our inspiral until about 15 gravitational-wave cycles prior to the merger of the binary. This agreement provides an important validation of our numerical simulation. It also establishes a regime in which the 3.5-th order post-Newtonian waveforms are accurate to this level, at least for an equal-mass, nonspinning black hole binary. After this point, the various PN approximants begin to diverge, suggesting that the approximation is beginning to break down. Since there are many different PN approximants (including Padé [108] and effective-one-body [76, 103, 105, 78], which were not discussed in this paper) it may be possible to find a clever way to push the PN expansion beyond its breaking point.

Indeed, we find that one approximant, TaylorT4 at 3.5PN in phase, works astonishingly well, agreeing with our numerical waveforms for almost the entire 30-cycle length of our runs. Given the wide scatter plot of predictions by various PN approximants, it is likely that TaylorT4 3.5/3.0 simply got lucky for the equal-mass nonspinning black hole binary. In fact, the assumption of adiabaticity (i.e., circular orbits) is known to lead to much larger phase differences relative to a non-adiabatic inspiral (see Fig. 4 of [77] and [194]) than the phase differences between NR and TaylorT4 we find in Fig. 4.26. Thus it seems that the uncontrolled higher order terms of TaylorT4 3.5/3.0 balance the error introduced by the adiabaticity assumption to a remarkable degree. It remains to be determined whether similar cancellations occur when the black hole masses are unequal or when the holes

have nonzero spin.

Regardless of the robustness of TaylorT4, it seems evident that numerical simulations are needed in order to know which, if any, PN approximant yields the correct waveform after the various approximants begin to diverge. For there is no *a priori* reason why TaylorT4 should be a better choice than TaylorT1 as they differ only in whether the ratio of gravitational-wave flux to the derivative of the orbital energy with respect to frequency is left as a ratio of post-Newtonian expansions or re-expanded as a single post-Newtonian expansion.

The surprising accuracy of TaylorT4 3.5/3.0 in the gravitational frequency range from $M\omega = -0.035$ through $M\omega = -0.15$, for the equal-mass, nonspinning inspiral of two black holes, in principle could form a basis for evaluating the errors of numerical simulations. Instead of worrying about errors due to different formulations, initial data, boundary conditions, extraction methods, etc., perhaps a long inspiral simulation could be compared with TaylorT4 3.5/3.0 in order to get a direct estimate of the phase error. Similarly, because of its good agreement, TaylorT4 3.5/3.0 could also be used to address questions that require much longer waveforms than currently available, for instance the question of when lower order post-Newtonian waveforms become unreliable.

We find that the 3PN contributions to the amplitude of the (2,2) modes improve their accuracy with respect to the numerical waveforms. This suggests that for accurate parameter estimation, it may be desirable to compute the full 3PN amplitude for the polarization waveforms. Despite the formidable nature of the calculation required, it would also be interesting to see how the inclusion of 4PN order corrections to the phasing would affect our comparisons.

Much work still needs to be done to improve the comparison between NR and PN. Our primary goal is to push our simulations through merger and ringdown so that we may compare various resummed PN approximants and the effective-one-body approximants during the last cycle of inspiral and merger, as well as test TaylorT4 3.5/3.0 closer to merger. We also intend to do long inspirals with arbitrary masses and spins in order to test the robustness of PN over a range of these parameters.

Furthermore we wish to improve our initial data. There is a large amount of “junk radiation” present in the initial data that limits how early we can match PN and NR waveforms. Reduction of this junk radiation [189] would improve the accuracy of our simulations as well.

Finally, we have done just a simple comparison between NR and PN, without including any treatment of effects that are important for real gravitational-wave detectors such as limited bandwidth and detector noise. In order to more directly address the suitability of PN formulas for analyzing data from gravitational-wave detectors, it will be necessary to fold in the properties of the detector, to consider specific values for the total mass of the binary, and to fit for the mass from the waveforms rather than assuming that the PN and NR waveforms correspond to the same mass. We leave this for future work.

**HIGH-ACCURACY NUMERICAL SIMULATION OF BLACK-HOLE
BINARIES: COMPUTATION OF THE GRAVITATIONAL-WAVE
ENERGY FLUX AND COMPARISONS WITH POST-NEWTONIAN
APPROXIMANTS¹**

Expressions for the gravitational wave (GW) energy flux and center-of-mass energy of a compact binary are integral building blocks of post-Newtonian (PN) waveforms. In this paper, we compute the GW energy flux and GW frequency derivative from a highly accurate numerical simulation of an equal-mass, nonspinning black hole binary. We also estimate (the derivative of) the center-of-mass energy from the simulation by assuming energy balance. We compare these quantities with the predictions of various PN approximants

¹This chapter is extracted with minor revisions from Ref. [61], which was written in collaboration with Alessandra Buonanno, Lawrence E. Kidder, Abdul H. Mroué, Yi Pan, Harald P. Pfeiffer, and Mark A. Scheel. I extrapolated the waveforms used in this paper, and shared with Harald the determination of uncertainty estimates. I worked to clarify the distinction between the various types of frequencies used, ensuring that we never compared apples with oranges. I also suggested the technique used to match the PN and NR waveforms, and wrote the code to produce the matches, and shared in writing the text.

(adiabatic Taylor and Padé models; non-adiabatic effective-one-body (EOB) models). We find that Padé summation of the energy flux does not accelerate the convergence of the flux series; nevertheless, the Padé flux is markedly closer to the numerical result for the whole range of the simulation (about 30 GW cycles). Taylor and Padé models overestimate the increase in flux and frequency derivative close to merger, whereas EOB models reproduce more faithfully the shape of and are closer to the numerical flux, frequency derivative, and derivative of energy. We also compare the GW phase of the numerical simulation with Padé and EOB models. Matching numerical and untuned 3.5 PN order waveforms, we find that the phase difference accumulated until $M\omega = 0.1$ is -0.12 radians for Padé approximants, and 0.50 (0.45) radians for an EOB approximant with Keplerian (non-Keplerian) flux. We fit free parameters within the EOB models to minimize the phase difference, and confirm the presence of degeneracies among these parameters. By tuning the pseudo-4PN-order coefficients in the radial potential or in the flux, or, if present, the location of the pole in the flux, we find that the accumulated phase difference at $M\omega = 0.1$ can be reduced—if desired—to much less than the estimated numerical phase error (0.02 radians).

5.1 Introduction

The first-generation interferometric gravitational wave (GW) detectors, such as LIGO [26, 245], GEO600 [162], and Virgo [4, 3], are now operating at or near their design sensitivities. One of the most promising sources for these detectors is the inspiral and merger of binary black holes (BBHs) with masses $m_1 \sim m_2 \sim 10-$

$20 M_{\odot}$ [140, 107]. A detailed and accurate understanding of the gravitational waves radiated as the black holes spiral towards each other will be crucial not only for the initial detection of such sources, but also for maximizing the information that can be obtained from signals once they are observed. Both the detection and subsequent analysis of gravitational waves from compact binaries depends crucially on our ability to build an accurate bank of templates, where each template is a theoretical model that accurately represents the gravitational waveform from a binary that has a certain set of parameters (e.g., masses and spins). For detection, the technique of matched filtering is applied to noisy data to extract any signals that match members of the template bank. For analysis, the best-fit parameters are determined, most likely by an iterative process that involves constructing further templates to zero in on the best fit.

When the black holes are far apart and moving slowly, the gravitational waveform (i.e., the template) can be accurately computed using a post-Newtonian (PN) expansion. As the holes approach each other and their velocities increase, the post-Newtonian expansion is expected to become less and less reliable. However, until recently there has been no independent way to determine how close comparable-mass holes must be before PN methods become inaccurate. This has changed with recent advances in numerical relativity (NR), which make it possible for the first time to quantify the disagreement between PN predictions [42] and the true waveform [75, 24, 157, 59, 151, 156]. In a previous paper [59], some of us described numerical simulations of 15 orbits of an equal-mass, nonspinning binary black hole system. Gravitational waveforms from these simulations covering more than 30 GW cycles and ending about 1.5 GW cycles before merger,

were compared with those from quasi-circular PN formulas for several time-domain Taylor approximants computed in the so-called *adiabatic* approximation. We found that there was excellent agreement (within 0.05 radians) in the GW phase between the numerical results and the PN waveforms over the first ~ 15 cycles, thus validating the numerical simulation and establishing a regime where PN theory is accurate. In the last 15 cycles to merger, however, *generic* time-domain Taylor approximants build up phase differences of several radians. But, apparently by coincidence, one specific PN approximant—Taylor T4 at 3.5PN order—agreed much better with the numerical simulations, with accumulated phase differences of less than 0.05 radians over the 30-cycle waveform. Simulations by Hannam et al. [156] for equal-mass, nonprecessing spinning binaries confirm that this agreement in the nonspinning case is a coincidence: they find the phase disagreement between Taylor T4 and the numerical waveform can be a radian or more as the spins of the black holes are increased.

To build a template bank to be used by ground-based GW detectors, one possibility would be to run a separate numerical simulation for each template. This is not currently possible, however, due to the large computational cost per numerical waveform (on the order of a week for a single waveform) and the large number of templates needed to cover the parameter space, especially when spins are present. A more realistic possibility is to perform a small number of simulations and develop an analytic template family (i.e., a fitting formula) which interpolates the parameter space between the simulations [206, 78, 5, 122, 123, 124].

Before the NR breakthrough several analytic prescriptions were proposed to

address the loss of accuracy of the adiabatic Taylor approximants. Damour, Iyer, and Sathyaprakash [106] introduced the Padé summation of the PN center-of-mass energy and gravitational energy flux in order to produce a series of Padé approximants for the waveforms in the adiabatic. Buonanno and Damour [76, 77, 114, 103] introduced the effective-one-body (EOB) approach which gives an analytic description of the motion and radiation beyond the adiabatic approximation of the binary system through inspiral, merger, and ringdown. The EOB approach also employs the Padé summation of the energy flux and of some crucial ingredients, such as the radial potential entering the conservative dynamics. So far, the EOB waveforms have been compared with several numerical waveforms of nonspinning binary black holes [75, 206, 78, 122, 123, 124]. Buonanno et al. [78] showed that by using three quasi-normal modes [75] and by tuning the pseudo 4PN order coefficient [105] in the EOB radial potential to a specific value, the phase difference accumulated by the end of the ringdown phase can be reduced to $\sim 3\text{--}8 \times 10^{-2}$ of a GW cycle, depending on the mass ratio and the number of multipole moments included in the waveform. Those results were obtained using waveforms with 5–16 GW cycles and mass ratios 1:4, 1:2, 2:3, and 1:1. In Refs. [122, 123, 124] the authors introduced other improvements in the EOB approach, in part obtained by tuning the test-mass limit results [121]—for example Padé summation of the PN amplitude corrections in the inspiral waveform; ringdown matching over an interval instead of a point; inclusion of noncircular terms in the tangential damping force; use of five quasi-normal modes. They found that the phase differences accumulated by the end of the inspiral (ringdown) can be reduced to $\pm 2 \times 10^{-4}$ ($\pm 5 \times 10^{-3}$) of a GW cycle for equal-mass binaries [122, 123]

and to $\pm 8 \times 10^{-3}$ of a GW cycle for binaries with mass ratio 1:2 [124]. Note that these phase differences are smaller than the numerical errors in the simulations.

The energy flux and the center-of-mass energy are two fundamental quantities of the binary dynamics and crucial ingredients in building GW templates. In this paper we extract these quantities, and compare the results from our numerical inspiral simulation [59] with PN results in both their Taylor-expanded and summed (Padé and EOB) forms. The agreement between the numerical and analytical results for the energy flux and the center-of-mass energy is a further validation of the numerical simulation. It also allows us to study whether or not the agreement of the phase evolution of PN and numerical waveforms is accidental. In addition, we compute waveforms based on adiabatic Padé and non-adiabatic EOB approximants in their *untuned* form (i.e., without introducing fitting coefficients) and study their agreement with our numerical simulations.

We try to understand whether these approximants can reproduce features of the numerical simulations that can be exploited to develop a faithful analytic template family. By introducing unknown higher-order PN coefficients into the dynamics and tuning them to the numerical data, we investigate how to improve the agreement with the numerical results. Although our study only examines nonspinning, equal-mass binary black holes, by combining it with other studies [206, 78, 5, 122, 123, 124] one can already pinpoint which parameters are degenerate and which have the largest effect on the waveforms. This is particularly relevant during the last stages of inspiral and plunge. The overall methodology can be extended to a larger region of the parameter space. We will defer to a future paper a complete study of the flexibility of the EOB approach with the

extension of our numerical waveform through merger and ringdown.

This paper is organized as follows: Section 5.2 gives a quick review of the numerical simulations presented in [59], and then presents the computation of the GW energy flux from the simulation. In Sec. 5.3 we summarize the PN approximants that will be compared to the numerical simulation. In Sec. 5.4, we compare the GW energy flux for the various PN approximants with numerical results and explore the possibility of improving the agreement with the numerical flux by adding phenomenological parameters [206, 78, 122, 123, 124]. In Sec. 5.5, we examine the evolution of the center-of-mass energy for the various PN approximants and compare to the numerical results assuming balance between the change in the center-of-mass energy and the energy carried from the system by the gravitational waves. In Sec. 5.6 we compare waveforms constructed from the Padé and EOB approximants with our numerical results, and study how to improve the agreement by exploiting the flexibility of the EOB model (i.e., by fitting free parameters of the EOB model). Finally, we present some concluding remarks in Sec. 5.7. We also include an appendix (Sec. 5.8) reviewing the performance of the Padé summation of the Taylor series of the energy flux in the test particle limit.




Figure 5.1: **Some aspects of the numerical simulation**

From top panel to bottom: the leading mode \dot{h}_{22} ; the two next largest modes, \dot{h}_{44} and \dot{h}_{32} (smallest); the frequency of \dot{h}_{22} (see Eq. (5.5)).

5.2 Computation of the numerical gravitational-wave energy flux

5.2.1 Overview and Definitions

The data used in this paper is the same as that described in Sec. II of Boyle et al. [59]. The simulation is a 16-orbit inspiral, with very low spin and eccentricity. Figure 5.1 presents a view of some relevant quantities of that simulation.

The Newman-Penrose scalar Ψ_4 , defined using a coordinate-based tetrad, is extracted from the simulation at several extraction radii and expanded in spin-weighted spherical harmonics,

$$\Psi_4(t, r, \vartheta, \varphi) = \sum_{l, m} \Psi_4^{l, m}(t, r) {}_{-2}Y_{l, m}(\vartheta, \varphi) . \quad (5.1)$$

Then $\Psi_4^{l, m}(t, r)$ is extrapolated to infinite extraction radius using an n -th order polynomial in $1/r$, where typically $n = 3$. This results in the asymptotic field $r\Psi_4^{l, m}(t - r_*)$ as function of retarded time² $t - r_*$.

Gravitational radiation may also be expressed via the standard metric-perturbation quantities h_+ and h_\times , which we similarly write in terms of spin-weighted spherical harmonic components,

$$h := h_+ - ih_\times = \sum_{l, m} h_{l, m} {}_{-2}Y_{l, m} . \quad (5.2)$$

For linear perturbations around Minkowski space, $\Psi_4^{l, m}(t - r_*) = -\ddot{h}_{l, m}(t - r_*)$. In particular, this relation should be true for the waveforms we have extrapolated to infinity.

²See Sec. II F of Ref. [59] for a precise definition of r_* and a description of the extrapolation.

However, to compute the energy flux we do not need to determine h ; we need only its time derivative \dot{h} . The energy flux depends on the spin-weighted spherical harmonic coefficients of the time derivative \dot{h} via

$$F = \frac{1}{16\pi} \sum_{l=2}^{\infty} \sum_{m=-l}^l |r \dot{h}_{l,m}|^2 . \quad (5.3)$$

We obtain $\dot{h}_{l,m}$ by time-integration of $\Psi_4^{l,m}$, as discussed in detail below.

Finally, we define gravitational wave phase and frequency in two ways—one based on Ψ_4^{22} , and one based on \dot{h}_{22} :

$$\phi = \arg(\Psi_4^{22}) , \quad \omega = \frac{d}{dt}\phi , \quad (5.4)$$

$$\varphi = \arg(\dot{h}_{22}) , \quad \varpi = \frac{d}{dt}\varphi . \quad (5.5)$$

In both cases, we define the \arg function to be the usual function, with discontinuities of 2π removed. Many PN formulas (see Sec. 5.3) involve yet another frequency and phase: the *orbital* phase Φ and *orbital* frequency Ω . Although the three frequencies satisfy $\omega \approx \varpi \approx 2\Omega$, the slight differences between different frequencies are significant at the level of precision of our comparison (see Fig. 5.6 below), so it is important to distinguish carefully between them.

When discussing our numerical solution, we write all dimensionful quantities in terms of the mass scale M , which we choose to be the sum of the irreducible masses of the two black holes.³

³This quantity was denoted by m in Ref. [59].

5.2.2 Calculation of \dot{h}

The energy flux depends on the spin-weighted spherical harmonic coefficients of \dot{h} via Eq. (5.3). We therefore need to perform one time integration on $\Psi_4^{l,m}$:

$$\dot{h}_{l,m}(t) = - \int_{t_0}^t \Psi_4^{l,m}(t') dt' + H_{l,m}. \quad (5.6)$$

This integration is performed for each mode (l, m) separately and requires the choice of two integration constants, which are contained in the complex number $H_{l,m}$. Ideally, $H_{l,m}$ should be chosen such that $\dot{h}_{l,m} \rightarrow 0$ for $t \rightarrow -\infty$. Because our numerical simulations do not extend into the distant past, this prescription cannot be implemented. Rather, we make use of the approximation that the real and imaginary parts of $\dot{h}_{l,m}$ should oscillate symmetrically around zero.

Let us consider a pure sine/cosine wave, with constant amplitude and phase:

$$\Psi_4^{\text{ex}} = -A[\cos(\omega t) + i \sin(\omega t)], \quad (5.7)$$

$$\dot{h}^{\text{ex}} = \frac{A}{\omega}[\sin(\omega t) - i \cos(\omega t)] + H^{\text{ex}}, \quad (5.8)$$

where the superscript ‘ex’ stands for example. The amplitude is given by

$$|\dot{h}^{\text{ex}}|^2 = \frac{A^2}{\omega^2} + 2\frac{A}{\omega}[\Re H^{\text{ex}} \sin(\omega t) - \Im H^{\text{ex}} \cos(\omega t)] + |H^{\text{ex}}|^2. \quad (5.9)$$

Only for the correct choice of integration constants, $H^{\text{ex}} = 0$, is the amplitude $|\dot{h}^{\text{ex}}|$ constant.

Therefore, we propose to determine the integration constants $H_{l,m}$ in Eq. (5.6) by minimizing the time derivative of the amplitude over the entire waveform. In particular we minimize

$$\mathcal{I}_{l,m} := \int_{t_1}^{t_2} \left(\frac{d}{dt} |\dot{h}_{l,m}|^2 \right)^2 dt. \quad (5.10)$$

From this minimization principle it follows that $H_{l,m}$ is determined by the linear system

$$\Re H \int (\Re \Psi_4)^2 dt + \Im H \int \Re \Psi_4 \Im \Psi_4 dt = - \int [(\Re \Psi_4)^2 \Re \dot{h}_0 + \Re \Psi_4 \Im \Psi_4 \Im \dot{h}_0] dt , \quad (5.11a)$$

$$\Re H \int \Re \Psi_4 \Im \Psi_4 dt + \Im H \int (\Im \Psi_4)^2 dt = - \int [(\Im \Psi_4)^2 \Im \dot{h}_0 + \Re \Psi_4 \Im \Psi_4 \Re \dot{h}_0] dt . \quad (5.11b)$$

Here, we have suppressed the indices l, m for clarity, all integrals are definite integrals from t_1 to t_2 , and $\dot{h}_0(t) := -\int_{t_0}^t \Psi_4(t') dt'$. For a given integration interval $[t_1, t_2]$, Eqs. (5.11) provide a deterministic procedure to determine the integration constants $H_{l,m}$. We note that there have been several earlier proposals to fix integration constants [209, 32, 217, 124, 232]. While we have not tested those proposals, we point out that Eqs. (5.11) allow for very accurate determination of the integration constants and one can easily obtain an error estimate, as we discuss in the next subsection.

5.2.3 Uncertainties in numerical quantities

Because the amplitude and frequency of the waveform are not constant, this procedure is imperfect, and the result depends somewhat on the chosen values of t_1 and t_2 . To estimate the residual uncertainty in H due to this choice, we select nine different values for t_1 and eleven values for t_2 : $t_1 = 200M, 220M, \dots, 360M$; $t_2 = 2000M, 2100M, \dots, 3000M$. The values of t_1 vary over roughly one GW cycle and test the sensitivity to the GW phase at the start of the integration interval; the values of t_2 are designed to test the dependence on the amplitude at the end of the integration interval. For $t_2 > 3000M$ we find that the errors in our procedure

rapidly increase for several reasons: (a) the minimization principle is based on the approximation that the amplitude is constant; this approximation becomes worse toward merger; (b) $\mathcal{I}_{l,m}$ in Eq. (5.10) weights absolute changes in $|\dot{h}|$, not relative ones; close to merger, the amplitude becomes so large that it dominates $\mathcal{I}_{l,m}$; and (c) the integration constants shift the waveform $\dot{h}_{l,m}$ vertically, and we are trying to determine the particular vertical shift such that $\dot{h}_{l,m}$ is centered around zero. Determination of such an offset is most accurate in a regime where the oscillations are *small*, i.e., at early times.

For each of these 99 integration intervals, we compute integration constants using Eqs. (5.11) for the three dominant modes, \dot{h}_{22} , \dot{h}_{44} , and \dot{h}_{32} , and we compute $F(t)$ from Eq. (5.3) using only these modes and we compute $\varpi(t)$. (We will show below that the contributions of other modes are far below our numerical errors on the flux.) We average the 99 functions $F(t)$ and $\varpi(t)$ and then use a parametric plot of $F(t)$ versus $\varpi(t)$ in our comparisons presented below. The variation in these 99 values yields an uncertainty in F due to the choice of integration constants.

The lower panel of Fig. 5.2 shows the variation in flux from the 99 different integration intervals. We find that the *maximum* deviation can be well approximated by $\max|\delta F|/F = 1.5 \times 10^{-5}(-M\varpi)^{-3/2}$ (see the solid line in lower panel of Fig. 5.2). The *average* F computed from all 99 intervals $[t_1, t_2]$ will have a smaller error. Inspection of the lower panel of Fig. 5.2 reveals that the $\delta F/F$ curves fall into 11 groups, corresponding to the 11 values of t_2 . Assuming that δF between these groups is randomly distributed, the error of the average will be reduced by a factor $\sqrt{11}$, i.e., $\delta F/F = 5 \times 10^{-6}(-M\varpi)^{-3/2}$. This error is indicated as the gray




Figure 5.2: Accuracy of the numerical flux

Lower panel: Relative difference between flux $F(\omega)$ computed with 99 different intervals $[t_1, t_2]$ and the average of these. **Upper panel:** Relative change in the flux $F(\omega)$ under various changes to the numerical simulation. The gray area in the upper panel indicates the uncertainty due to the choice of integration constants, which is always dominated by numerical error. The dashed line in the upper panel is our final error estimate, which we plot in later figures.

shaded area in the upper panel of Fig. 5.2.

The upper panel of Fig. 5.2 plots the relative change in $F(\omega)$ for several changes in our numerical simulation: (a) Computing the flux from a run with lower resolution (0030c/N5 in the language of Boyle et al. [59]); (b) using a different set of extraction radii for the extraction of the gravitational wave; (c) increasing the polynomial order of extrapolation of Ψ_4 to infinite extraction radius from $n = 3$ to $n = 4$; and (d) computing the flux from a separate evolution with a different outer boundary radius (0030c-2/N6). At low frequencies, the error is dominated by extrapolation to infinite radius and is a few tenths of a percent; at intermediate frequencies, $0.055 \lesssim -M\omega < 0.083$, all errors are smaller than 0.1 percent. At frequency $-M\omega \approx 0.084$ we change the gauge conditions in the evolutions to allow wave-escorting; this introduces high-frequency features, which are small when extrapolation order $n = 3$ is used, but which dominate for $n = 4$ extrapolation. The numerical data we use in the PN comparisons below is extrapolated with $n = 3$, for which the features due to change of gauge are small, but nevertheless we will use conservative error bars encompassing the $n = 4$ extrapolation as indicated in Fig. 5.2, i.e., a relative error of 0.2 percent for $-M\omega > 0.083$. We find that the uncertainty in the flux due to numerical error in determining Ψ_4 is always larger than the uncertainty due to the choice of integration constants.

The contributions of the various (l, m) modes to the total flux [see Eq. (5.3)] are plotted in Fig. 5.3. The top panel plots the flux as a function of time; the lower panel as a function of frequency $M\omega$. The dashed line in the lower panel corresponds to the error estimate of Fig. 5.2. Because the modes (5,4), (6,6), and





Figure 5.3: Contributions of various (l, m) modes to the total numerical gravitational wave flux

Upper panel: plotted as a function of time. **Lower panel:** Plotted as a function of frequency $M\omega$. The lower panel also contains the error estimate derived in Fig. 5.2.

Figure 5.4: Accuracy of numerical $\dot{\omega}$

Lower panel: Difference between frequency derivative $\dot{\omega}$ computed with 99 different intervals $[t_1, t_2]$ and the average of these. **Upper panel:** Change in the frequency derivative $\dot{\omega}$ under various changes to the numerical simulation. The gray area in the upper panel indicates the uncertainty due to choice of integration constants, which dominates the overall uncertainty for low frequencies. The dashed line in the upper panel is our final error estimate, which we plot in later figures.

(8,8) are significantly smaller than our error estimate, we do not include them in the present analysis.

To estimate the uncertainty in $\dot{\omega}$, we proceed in a similar fashion. Each one of the 99 different integration intervals yields an \dot{h}_{22} from which we determine $\dot{\omega}$. We average these to obtain the final $\dot{\omega}$ to be used in the post-Newtonian comparisons. The lower panel of Fig. 5.4 shows the variation in $\dot{\omega}$ between the 99 different integration intervals. We find that the *maximum* deviation can be well approximated by $\max|M^2\delta\dot{\omega}| = 5 \times 10^{-6}(-M\omega)^{-0.3}$ (see the solid line in lower panel of Fig. 5.4). The *average* $\dot{\omega}$ computed from all 99 intervals $[t_1, t_2]$ will have a smaller error. Inspection of the lower panel of Fig. 5.4 reveals that the $\delta\dot{\omega}$ curves fall into 11 groups, corresponding to the 11 values of t_2 . As for the case of δF , if we assume that $\delta\dot{\omega}$ between these groups is randomly distributed, then the error of the average will be reduced by a factor $\sqrt{11}$, i.e., $M^2\delta\dot{\omega} = 1.5 \times 10^{-6}(-M\omega)^{-0.3}$. This error is indicated as the gray shaded area in the upper panel of Fig. 5.4.

The upper panel of Fig. 5.4 plots also the change in $\dot{\omega}(\omega)$ for the same changes in our numerical simulation already discussed above. We find that at $-M\omega < 0.083$, the uncertainty in $\dot{\omega}$ is dominated by the choice of integration constants, whereas at higher frequencies the uncertainty is dominated by the numerical errors in the calculation of Ψ_4 . As discussed above, at frequency $-M\omega \approx 0.084$ we change the gauge conditions in the evolutions to allow wave-escorting; this introduces high-frequency features leading to more conservative error estimates.

Note that $\dot{\omega}$ is a very steep function of ω . While the absolute errors in $\dot{\omega}$ are roughly constant for our simulation, the *relative* errors change significantly: $\delta\dot{\omega}/\dot{\omega}$ drops from about 10 percent early in the run to about 0.2 percent at late times.

We also point out that the first $1000M$ of our simulation are contaminated by noise due to a pulse of “junk-radiation” at the start of the simulation. While this contamination is not apparent on a plot of the waveform as in Fig. 5.1, it nevertheless limits accurate PN-NR comparisons to the region, $t - r_* \gtrsim 1000M$, i.e., $-M\omega \gtrsim 0.037$.

5.3 Post-Newtonian approximants

In this paper we will compare the numerical simulation to various approximants based on the PN expansion. The PN expansion is a slow-motion, weak-field approximation to general relativity with an expansion parameter $\epsilon \sim (v/c)^2 \sim (GM/rc^2)$. For a binary system of two point masses m_1 and m_2 , v is the magnitude of the relative velocity, M is the total mass, and r is the separation. For a review of the PN expansion applied to gravitational radiation from inspiralling compact binaries, see Ref. [42].

In Table 5.1 we summarize the PN approximants that we use and our notation. We shall use the PN approximants in the so-called adiabatic approximation, both in the standard Taylor-expanded form (reviewed in Sec. 5.3.1) and in a form based on Padé summation (reviewed in Sec. 5.3.2). In addition we shall use the non-adiabatic EOB model (reviewed in Sec. 5.3.3) in its original form [76, 77, 114], as well as several variations that differ in the form of the radiation-reaction force [73, 72, 104]. After summarizing the various PN approximants in Secs. 5.3.1, 5.3.2, and 5.3.3, we describe how we construct the waveform for these approximants in Sec. 5.3.4.

Table 5.1: Summary of PN approximants

Approximant	Notation	See Eqs.	Adiabatic	Keplerian
Taylor (T-)	F_n/E_p	(5.19), (5.14)	yes	yes
Padé (P-)	F_n^m/E_p^q	(5.39), (5.33)	yes	yes
EOB (E-)	F_n^m/H_p	(5.64), (5.44)	no	yes
EOB (E-)	${}^{\text{nK}}F_n^m/H_p$	(5.65), (5.44)	no	no
EOB (E-)	F_n/H_p	(5.69), (5.44)	no	yes
EOB (E-)	${}^{\text{nK}}F_n/H_p$	(5.70), (5.44)	no	no

The T-approximants are always Taylor T4 [59] except in Fig. 5.16. The P-approximant in the second row was introduced in Refs. [106, 114, 73] and the *original* E-approximant in third row was introduced in Refs. [76, 77, 114]. The last three rows refer to three possible variations of E-approximants introduced in Refs. [73, 72]. In a few tests aimed at improving the closeness between numerical data and E-approximants, we vary v_{pole} and treat the logarithms as constants when Padé summation to the flux is applied [122]. We shall denote this flux by \bar{F}_n^m . Finally, when using tuned PN approximants with pseudo 4PN order terms in the flux, energy, or Hamiltonian, we denote the latter as pF , pE , and pH . Note that if known test-mass limit coefficients in the flux are used, the latter is still denoted as F even at PN orders larger than 3.5PN. Finally, the values of v_{pole} and v_{iso} used in the P-approximants F_n^m and ${}^{\text{nK}}F_n^m$ are $v_{\text{pole}}^{2\text{PN}} = 0.6907$ and $v_{\text{iso}}^{2\text{PN}} = 0.4456$.

In the adiabatic approximation the inspiral is modeled as a quasi-stationary sequence of circular orbits. The evolution of the inspiral (and in particular of the orbital phase Φ) is completely determined by the *energy-balance equation* [42]

$$\frac{dE(v_\Omega)}{dt} = -F(v_\Omega) . \quad (5.12)$$

This equation relates the time derivative of the center-of-mass energy $E(v_\Omega)$ (which is conserved in absence of radiation reaction) to the gravitational wave energy flux $F(v_\Omega)$. Both functions are known for quasi-circular orbits as a PN

expansion in the invariantly defined velocity

$$v_\Omega = (M\Omega)^{1/3} , \quad (5.13)$$

where $\Omega = \dot{\Phi}$ is the orbital frequency (in units such that $G = c = 1$).⁴ We will denote the Taylor-expanded flux (energy) by F_k (E_k) where k denotes the maximum power of v_Ω retained in the series. (Recall that $k = 2N$ for an N th order PN expansion.) We will denote the Padé-expanded flux (energy) by F_n^m (E_n^m) where $m + n = k$, with m and n denoting the order of the polynomial in the numerator and denominator, respectively.

5.3.1 Adiabatic Taylor approximants

For generic values of the symmetric mass ratio $\nu = m_1 m_2 / M^2$, the center-of-mass energy is known through 3PN order [174, 125, 115, 49, 119]. For circular orbits the Taylor PN approximants (henceforth, T-approximants) to the energy are given by

$$E_{2k}(v_\Omega) = -\frac{M\nu}{2} v_\Omega^2 \sum_{i=0}^k \mathcal{E}_{2i}(\nu) v_\Omega^{2i} , \quad (5.14)$$

where the known coefficients are

$$\mathcal{E}_0(\nu) = 1 , \quad (5.15)$$

$$\mathcal{E}_2(\nu) = -\frac{3}{4} - \frac{\nu}{12} , \quad (5.16)$$

$$\mathcal{E}_4(\nu) = -\frac{27}{8} + \frac{19}{8}\nu - \frac{1}{24}\nu^2 , \quad (5.17)$$

$$\begin{aligned} \mathcal{E}_6(\nu) = & -\frac{675}{64} + \left(\frac{34445}{576} - \frac{205}{96}\pi^2 \right) \nu - \frac{155}{96}\nu^2 \\ & - \frac{35}{5184}\nu^3 . \end{aligned} \quad (5.18)$$

⁴In Ref. [59] we used $x = v_\Omega^2$ as the expansion parameter.

The GW energy flux for arbitrary masses has been computed through 3.5PN order [50, 46]:

$$F_k(v_\Omega) = \frac{32}{5} v^2 v_\Omega^{10} \sum_{i=0}^k \mathcal{F}_i(v) v_\Omega^i, \quad (5.19)$$

where

$$\mathcal{F}_0(v) = 1, \quad (5.20)$$

$$\mathcal{F}_1(v) = 0, \quad (5.21)$$

$$\mathcal{F}_2(v) = -\frac{1247}{336} - \frac{35}{12}v, \quad (5.22)$$

$$\mathcal{F}_3(v) = 4\pi, \quad (5.23)$$

$$\mathcal{F}_4(v) = -\frac{44711}{9072} + \frac{9271}{504}v + \frac{65}{18}v^2, \quad (5.24)$$

$$\mathcal{F}_5(v) = -\left(\frac{8191}{672} + \frac{583}{24}v\right)\pi, \quad (5.25)$$

$$\begin{aligned} \mathcal{F}_6(v) = & \frac{6643739519}{69854400} + \frac{16}{3}\pi^2 - \frac{1712}{105}\gamma_E \\ & - \frac{856}{105}\log(16v_\Omega^2) + \left(-\frac{134543}{7776} + \frac{41}{48}\pi^2\right)v \\ & - \frac{94403}{3024}v^2 - \frac{775}{324}v^3, \end{aligned} \quad (5.26)$$

$$\mathcal{F}_7(v) = \left(-\frac{16285}{504} + \frac{214745}{1728}v + \frac{193385}{3024}v^2\right)\pi, \quad (5.27)$$

where γ_E is Euler's constant. Notice that starting at 3PN order ($k = 6$) logarithms enter the flux.

5.3.2 Adiabatic Padé approximants

Center-of-mass energy

Damour, Iyer, and Sathyaprakash [106] (henceforth DIS) proposed a new class of approximate waveforms constructed by introducing new energy and flux func-

tions and by applying Padé summation [31] to build successive approximants to these two functions (henceforth P-approximants). Their motivation for introducing these new functions and using their P-approximants came from an examination of the behavior of the standard PN-expansion and the new P-approximants in the test-mass limit in which the exact gravitational energy flux is known numerically [215], the PN expansion of the flux is known through 5.5PN order [241], and the center-of-mass energy is known analytically as

$$\frac{E(v_\Omega; \nu = 0)}{\mu} = \frac{1 - 2v_\Omega^2}{\sqrt{1 - 3v_\Omega^2}} - 1, \quad (5.28)$$

where $\mu = M\nu$ is the reduced mass.

DIS first observed that in the quantum two-body problem the symmetric quantity

$$\epsilon := \frac{E_{\text{tot}}^2 - m_1^2 - m_2^2}{2m_1 m_2}, \quad (5.29)$$

(where the total relativistic energy $E_{\text{tot}} = E + M$), is the best energy function when treating the two-body problem as an effective one-body problem in an external field. Because in the test-mass limit

$$\epsilon(v_\Omega; \nu = 0) = \frac{1 - 2v_\Omega^2}{\sqrt{1 - 3v_\Omega^2}}, \quad (5.30)$$

DIS defined the new energy function as

$$e(v_\Omega) := \epsilon^2 - 1, \quad (5.31)$$

as this function has a simple pole singularity on the real axis in the test-mass limit, and DIS conjectured that such a pole would continue to exist in the comparable mass case.⁵ The energy function $E(v_\Omega)$ entering the balance equation (5.12) can

⁵A motivation for having used Eq. (5.31) instead of Eq. (5.29) as a basic quantity is that the former (unlike the latter) is amenable to Padé summation in the test mass limit.

be expressed in terms of $e(v_\Omega)$ as

$$E(v_\Omega) = \left\{ M^2 + 2vM^2 \left[\sqrt{1 + e(v_\Omega)} - 1 \right] \right\}^{1/2} - M . \quad (5.32)$$

by combining Eqs. (5.29) and (5.31). [Note that the map between the adiabatic functions e and E given by Eq. (5.32) is the same map found in the EOB model between the effective Hamiltonian H^{eff} and the real Hamiltonian H^{real} , as given by Eq. (5.44).]

Finally, DIS proposed as approximants to the energy function $e(v_\Omega)$ the diagonal or subdiagonal P-approximants, depending on whether the PN order is even or odd.⁶ Investigating the behavior of the P-approximants under variations of an (at the time) unknown coefficient in the 3PN center-of-mass energy, Damour, Jaranowski, and Schäfer [114] found it more robust to use the superdiagonal P-approximant instead of the subdiagonal P-approximant at 3PN order.⁷ This suggestion was also adopted in Ref. [73] and will be used here; that is, we use subdiagonal P-approximants for 1PN, diagonal for 2PN, and superdiagonal for 3PN.

The P-approximants for the center-of-mass energy are defined as

$$E_p^q(v_\Omega) = \left\{ M^2 + 2vM^2 \left[\sqrt{1 + e_p^q(v_\Omega)} - 1 \right] \right\}^{1/2} - M , \quad (5.33)$$

where at 2PN order [106]

$$e_2^2(v_\Omega) = -v_\Omega^2 \frac{1 + \frac{1}{3}v - \left(4 - \frac{9}{4}v + \frac{1}{9}v^2\right) v_\Omega^2}{1 + \frac{1}{3}v - \left(3 - \frac{35}{12}v\right) v_\Omega^2} , \quad (5.34)$$

⁶As the energy is only a function of even powers of v_Ω , the choice of using diagonal or subdiagonal (superdiagonal) is based on the order of v_Ω^2 that is retained. For notational consistency, the indices on all approximants will refer to the power of v_Ω . Other references define the indices on the energy approximants with respect to v_Ω^2 .

⁷Subdiagonal P-approximants were extended to 3PN order in Ref. [110], and LAL [94] software uses those P-approximants for the energy function.

and at 3PN order [114]

$$e_2^4(v_\Omega) = -v_\Omega^2 \frac{1}{1 - w_3(v) v_\Omega^2} \left[1 - \left(1 + \frac{1}{3}v + w_3(v) \right) v_\Omega^2 - \left(3 - \frac{35}{12}v - \left(1 + \frac{1}{3}v \right) w_3(v) \right) v_\Omega^4 \right], \quad (5.35)$$

where

$$w_3(v) = \frac{40}{36 - 35v} \left[\frac{27}{10} + \frac{1}{16} \left(\frac{41}{4} \pi^2 - \frac{4309}{15} \right) v + \frac{103}{120} v^2 - \frac{1}{270} v^3 \right]. \quad (5.36)$$

Gravitational-wave energy flux

As originally pointed out in Refs. [214, 101], the flux function in the test-mass limit has a simple pole at the light-ring position (i.e., the last unstable circular orbit of a photon). Motivated by this, DIS introduced a new flux-type function

$$f_k(v_\Omega) = \left(1 - \frac{v_\Omega}{v_{\text{pole}}(v)} \right) F_k(v_\Omega; v), \quad (5.37)$$

with the suggestion that v_{pole} be chosen to be at the light ring (pole singularity) of the new energy function.

In order to construct well-behaved approximants, DIS proposed to normalize the velocity v_Ω entering the logarithms in Eq. (5.26) to some relevant scale which they chose to be $v_{\text{lso}}(v)$, where the last stable orbit (LSO) is defined as the minimum of the energy. Also, they factored out the logarithms yielding

$$f_k(v_\Omega) = \frac{32}{5} v^2 v_\Omega^{10} \left[1 + \log \frac{v_\Omega}{v_{\text{lso}}(v)} \left(\sum_{i \geq 6}^k \ell_i v_\Omega^i \right) \right] \times \left(1 - \frac{v_\Omega}{v_{\text{pole}}(v)} \right) \sum_{i=0}^k \mathcal{F}_i^{\text{log-fac}} v_\Omega^i, \quad (5.38)$$

where ℓ_i and $\mathcal{F}_i^{\log\text{-fac}}$ are functions of \mathcal{F}_i . Through 3.5PN order, $\ell_6 = -1712/105$, $\ell_7 = 0$, and $\mathcal{F}_i^{\log\text{-fac}} = \mathcal{F}_i$ with the replacement of $\nu_\Omega \rightarrow \nu_{\text{iso}}$ in \mathcal{F}_6 [see Eq. (5.26)].

Finally, DIS proposed to define the P-approximant of the GW energy flux as

$$F_n^m(\nu_\Omega) = \frac{1}{1 - \nu_\Omega/\nu_{\text{pole}}(\nu)} f_n^m(\nu_\Omega) . \quad (5.39)$$

where

$$f_n^m(\nu_\Omega) = \frac{32}{5} \nu^2 \nu_\Omega^{10} \left[1 + \log \frac{\nu_\Omega}{\nu_{\text{iso}}(\nu)} \left(\sum_{i \geq 6}^k \ell_i \nu_\Omega^i \right) \right] \\ \times P_n^m \left[\left(1 - \frac{\nu_\Omega}{\nu_{\text{pole}}(\nu)} \right) \sum_{i=0}^k \mathcal{F}_i^{\log\text{-fac}} \nu_\Omega^i \right] , \quad (5.40)$$

where $P_n^m[x]$ denotes Padé summation of the series x . DIS proposed to use the diagonal or subdiagonal P-approximants, depending on whether $k = n + m$ is even or odd. Furthermore, DIS proposed to use $\nu_{\text{iso}}(\nu)$ and $\nu_{\text{pole}}(\nu)$ as the minimum and pole of the center-of-mass energy P-approximant of the same PN order. At 2PN (the order to which the PN expansion was known by DIS) ν_{pole} is determined from the pole of the Padé energy function e_2^2 , yielding

$$\nu_{\text{pole}}^{2\text{PN}}(\nu) = \sqrt{\frac{1}{3} \frac{1 + \frac{1}{3}\nu}{1 - \frac{35}{36}\nu}} . \quad (5.41)$$

When the PN expansion was extended to 3PN order, it was found that none of the 3PN P-approximants have a physical pole. Therefore, somewhat arbitrarily, we will follow previous analyses and use the value (5.41) also at 3PN order. We denote the P-approximants defined by Eqs. (5.39) and (5.33) as F_n^m/E_p^q .

The denominator in the Padé summation of the GW energy flux can have zeros. They are called *extraneous poles* of the P-approximant [31]. It is desirable that these poles be located at high frequency (i.e., beyond the transition from inspiral

to plunge). We shall see that depending on the PN order and also the mass ratio, extraneous poles can be present at low frequencies. This could indicate poor convergence of the Padé summation.

In Secs. 5.4.2, 5.6.2, and 5.6.3 we shall investigate how to improve the closeness of the PN approximants to the numerical data by varying a_5 [105, 78, 122], ν_{pole} [105, 122] and also by introducing higher-order PN coefficients in the flux function. When varying ν_{pole} in the P-approximant at 3.5PN order, extraneous poles appear at low values of ν_{Ω} . Therefore, in order to push these poles to very high frequency, we follow the suggestion of Ref. [122], and use P-approximants at 4PN order, where the 4PN coefficient is set to its known value in the test-mass limit. This cure may fail for different mass ratios if new extraneous poles appear at low frequency. Furthermore the logarithm in the flux is not factored out as in Eq. (5.38), but treated as a constant when Padé summation is done. In this case the flux function is denoted \bar{F}_n^m .

We notice that DIS motivated the introduction of the P-approximants first in the test-mass limit case by observing much faster and monotonic convergence of the Padé energy, flux, and waveforms with respect to Taylor energy, flux, and waveforms. Quantitative tests of the convergence were done only for the Padé waveforms (see e.g., Tables III and IV in Ref. [106]), while for the flux and the energy conclusions were drawn qualitatively from Figs. 3 and 4 of Ref. [106]. DIS then conjectured that the comparable mass case is a smooth deformation of the test-mass limit case, and proposed to use close-to-diagonal P-approximants for the flux and the energy when $\nu \neq 0$. In the appendix we perform a few convergence tests of the P-approximants of the flux function in the test-mass limit case, and

conclude that whereas the P-approximants provide a better fit to the numerical flux at 5.5PN order, they do not accelerate the convergence of the Taylor series expansion of the energy flux.

5.3.3 Non-adiabatic effective-one-body approximants

The EOB model goes beyond the adiabatic approximation and can incorporate deviations from the Keplerian law when the radial separation becomes smaller than the last stable circular orbit.

Here we briefly review the main equations defining the EOB dynamics and refer the reader to previous papers for more details [77, 76, 114, 72, 206, 78, 122, 123]. The nonspinning EOB effective Hamiltonian is [76, 114]:

$$\begin{aligned} H^{\text{eff}}(\vec{r}, \vec{p}) &= \mu \widehat{H}^{\text{eff}}(\vec{r}, \vec{p}) \\ &= \mu \left\{ A(r) \left[1 + \vec{p}^2 + \left(\frac{A(r)}{D(r)} - 1 \right) (\vec{n} \cdot \vec{p})^2 \right. \right. \\ &\quad \left. \left. + \frac{1}{r^2} 2(4 - 3\nu) \nu (\vec{n} \cdot \vec{p})^4 \right] \right\}^{1/2}, \end{aligned} \quad (5.42)$$

with \vec{r} and \vec{p} being the reduced dimensionless variables; $\vec{n} = \vec{r}/r$ where we set $r = |\vec{r}|$. In absence of spins the motion is constrained to a plane. Introducing polar coordinates (r, Φ, p_r, p_Φ) , the EOB effective metric reads

$$\begin{aligned} ds_{\text{eff}}^2 &:= g_{\mu\nu}^{\text{eff}} dx^\mu dx^\nu \\ &= -A(r) c^2 dt^2 + \frac{D(r)}{A(r)} dr^2 + r^2 (d\vartheta^2 + \sin^2 \vartheta d\varphi^2). \end{aligned} \quad (5.43)$$

The EOB real Hamiltonian is

$$H^{\text{real}} = M \sqrt{1 + 2\nu \left(\frac{H^{\text{eff}} - \mu}{\mu} \right)} - M, \quad (5.44)$$

and we define $\hat{H}^{\text{real}} = H^{\text{real}}/\mu$. The T-approximants to the coefficients $A(r)$ and $D(r)$ in Eqs. (5.42) and (5.43) read [76, 114]

$$A_k(r) = \sum_{i=0}^{k+1} \frac{a_i}{r^i}, \quad (5.45)$$

$$D_k(r) = \sum_{i=0}^k \frac{d_i}{r^i}, \quad (5.46)$$

where

$$a_0 = 1, \quad a_1 = 2, \quad a_2 = 0, \quad a_3(\nu) = 2\nu, \quad a_4(\nu) = \left(\frac{94}{3} - \frac{41}{32}\pi^2 \right) \nu, \quad (5.47)$$

$$d_0 = 1, \quad d_1 = 0, \quad d_2(\nu) = 6\nu, \quad d_3(\nu) = 2(3\nu - 26)\nu. \quad (5.48)$$

In Sec. 5.6.3, we will explore the flexibility of the EOB model by tuning the pseudo-4PN-order coefficients $a_5(\nu)$ which we will take to have the following functional form⁸

$$a_5(\nu) = a_5 \nu. \quad (5.49)$$

In order to assure the presence of a horizon in the effective metric, we need to factor out a zero of $A(r)$. This is obtained by applying the Padé summation [114]. Thus, the coefficients $A_k(r)$ and $D_k(r)$ are replaced by the Padé approximants [114]

$$A_2^1(r) = \frac{r(-4 + 2r + \nu)}{2r^2 + 2\nu + r\nu}, \quad (5.50)$$

at 2PN order, and

$$A_3^1(r) = \frac{\text{Num}(A_3^1)}{\text{Den}(A_3^1)}, \quad (5.51)$$

with

$$\text{Num}(A_3^1) = r^2 [(a_4(\nu) + 8\nu - 16) + r(8 - 2\nu)], \quad (5.52)$$

⁸Note that what we denote a_5 in this paper was denoted λ in Ref. [78].

and

$$\text{Den}(A_3^1) = r^3 (8 - 2\nu) + r^2 [a_4(\nu) + 4\nu] + r [2a_4(\nu) + 8\nu] + 4[\nu^2 + a_4(\nu)] , \quad (5.53)$$

at 3PN order. When exploring the flexibility of the EOB model, we use the following Padé approximant at 4 PN order [105, 78]:

$$A_4^1(r) = \frac{\text{Num}(A_4^1)}{\text{Den}(A_4^1)} , \quad (5.54)$$

with

$$\text{Num}(A_4^1) = r^3 [32 - 24\nu - 4a_4(\nu) - a_5(\nu)] + r^4 [a_4(\nu) - 16 + 8\nu] , \quad (5.55)$$

and

$$\begin{aligned} \text{Den}(A_4^1) = & -a_4^2(\nu) - 8a_5(\nu) - 8a_4(\nu)\nu + 2a_5(\nu)\nu - 16\nu^2 \\ & + r [-8a_4(\nu) - 4a_5(\nu) - 2a_4(\nu)\nu - 16\nu^2] \\ & + r^2 [-4a_4(\nu) - 2a_5(\nu) - 16\nu] \\ & + r^3 [-2a_4(\nu) - a_5(\nu) - 8\nu] \\ & + r^4 [-16 + a_4(\nu) + 8\nu] . \end{aligned} \quad (5.56)$$

For the coefficient $D(r)$, the P-approximant used at 2PN, 3PN, and 4PN order respectively are [114, 105, 78]:

$$D_2^0(r) = 1 - \frac{6\nu}{r^2} , \quad (5.57)$$

$$D_3^0(r) = \frac{r^3}{r^3 + 6\nu r + 2\nu(26 - 3\nu)} , \quad (5.58)$$

$$D_4^0(r) = \frac{r^4}{r^4 + 6\nu r^2 + 2\nu(26 - 3\nu)r - d_4(\nu) + 36\nu^2} , \quad (5.59)$$

and we choose somewhat arbitrarily $d_4(\nu) = 36\nu^2$, so that $D_4^0 = D_3^0$. (We note that the value of d_4 does not affect much the EOB evolution [78].) The EOB

Hamilton equations written in terms of the reduced quantities $\widehat{H}^{\text{real}}$ and $\widehat{t} = t/M$, $\widehat{\Omega} = \Omega M$ [77], are

$$\frac{dr}{d\widehat{t}} = \frac{\partial \widehat{H}^{\text{real}}}{\partial p_r}(r, p_r, p_\Phi) , \quad (5.60)$$

$$\frac{d\Phi}{d\widehat{t}} := \widehat{\Omega} = \frac{\partial \widehat{H}^{\text{real}}}{\partial p_\Phi}(r, p_r, p_\Phi) , \quad (5.61)$$

$$\frac{dp_r}{d\widehat{t}} = - \frac{\partial \widehat{H}^{\text{real}}}{\partial r}(r, p_r, p_\Phi) , \quad (5.62)$$

$$\frac{dp_\Phi}{d\widehat{t}} = \widehat{\mathcal{F}}[\widehat{\Omega}(r, p_r, p_\Phi)] , \quad (5.63)$$

where for the Φ component of the radiation-reaction force a few approximants are available. Originally, Ref. [77] suggested the following Keplerian P-approximants to the flux

$${}^{\text{K}}\widehat{\mathcal{F}}_n^m := - \frac{1}{v v_\Omega^3} F_n^m(v_\Omega; v, v_{\text{pole}}) , \quad (5.64)$$

where F_n^m is given by the Padé flux in Eqs. (5.39) and (5.40). Here by Keplerian we mean that in the flux the tangential velocity $V_\Phi = \dot{\Phi} r$ is set to $V_\Phi := v_\Omega = \dot{\Phi}^{1/3}$, having assumed the Keplerian relation $\dot{\Phi}^2 r^3 = 1$. It was then pointed out in Ref. [104] that the Keplerian relation becomes less and less accurate once the binary passes through the last stable orbit. A more appropriate approximant to the flux would be

$${}^{\text{nK}}\widehat{\mathcal{F}}_n^m := - \frac{v_\Omega^3}{v V_\Phi^6} F_n^m(V_\Phi; v, v_{\text{pole}}) , \quad (5.65)$$

where $V_\Phi := \dot{\Phi} r_\Omega$. Notice that because the EOB Hamiltonian is a deformation of the Schwarzschild Hamiltonian, the exact Keplerian relation is $\dot{\Phi}^2 r_\Omega^3 = 1$ with $r_\Omega := r[\psi(r, p_\Phi)]^{1/3}$ and ψ is defined following the argument presented around

Eq. (19) to (22) in Ref. [104]:

$$\begin{aligned} \frac{1}{\psi r^3} &:= \omega_{\text{circ}}^2 = \left(\frac{\partial \mathcal{H}(r, p_r = 0, p_\phi)}{\partial p_\phi} \right)^2 \\ &= \frac{1}{r^3} \frac{p_\phi^2 A(r)}{\left(1 + \frac{p_\phi^2}{r^2}\right) r (1 + 2\eta(\sqrt{w(r, p_\phi)} - 1))} \end{aligned} \quad (5.66)$$

where $w(r, p_\phi) = A(r) \left(1 + \frac{p_\phi^2}{r^2}\right)$. The value of p_ϕ of circular orbits are obtained by minimizing with respect to r the circular orbit Hamiltonian $\mathcal{H}(r, p_r = 0, p_\phi)$ and it yields the following relation between r and p_ϕ

$$\frac{2p_\phi^2 A(r)}{r^3} = \left(1 + \frac{p_\phi^2}{r^2}\right) \frac{dA(r)}{dr} . \quad (5.67)$$

By inserting Eq. (5.67) in the definition of ψ , and replacing all p_ϕ except the one which implicitly appears in $w(r, p_\phi)$ we obtain

$$\psi = \frac{1 + 2\eta(\sqrt{w(r, p_\phi)} - 1)}{r^2 dA(r)/dr/2} . \quad (5.68)$$

Finally, Refs. [73, 72] introduced another possible variation of the EOB flux approximants which use T-approximants for the flux given by Eq. (5.19), in either the Keplerian or non-Keplerian form, i.e.,

$${}^{\text{K}}\widehat{\mathcal{F}}_n = - \frac{1}{v v_\Omega^3} F_n(v_\Omega) , \quad (5.69)$$

and

$${}^{\text{nK}}\widehat{\mathcal{F}}_n = - \frac{v_\Omega^3}{v V_\Phi^6} F_n(V_\Phi) . \quad (5.70)$$

Note that the flux for the non-Keplerian EOB models are not simply functions of the orbital frequency Ω . We denote the original E-approximants [76, 77, 114] which use the Padé flux (5.40) as F_n^m/H_p where H_p is H^{real} computed from A_p^1 and D_p^0 . Other E-approximants used in this paper are summarized in Table 5.1. The initial conditions for Eqs. (5.60)–(5.63) are obtained following Ref. [77] and starting the evolution far apart to reduce the eccentricity to negligible values.

5.3.4 Waveforms

The PN waveforms are obtained by substituting the orbital phase and frequency into the spherical harmonic mode (2,2) with amplitude corrections through 3PN order [178, 10]

$$\begin{aligned}
h_{22} = & -8\sqrt{\frac{\pi}{5}} \frac{\nu M}{R} e^{-2i\Phi} \nu_\Omega^2 \left\{ 1 - \nu_\Omega^2 \left(\frac{107}{42} - \frac{55}{42} \nu \right) \right. \\
& + 2\pi \nu_\Omega^3 - \nu_\Omega^4 \left(\frac{2173}{1512} + \frac{1069}{216} \nu - \frac{2047}{1512} \nu^2 \right) \\
& - \nu_\Omega^5 \left[\left(\frac{107}{21} - \frac{34}{21} \nu \right) \pi + 24i\nu \right] \\
& + \nu_\Omega^6 \left[\frac{27027409}{646800} - \frac{856}{105} \gamma_E + \frac{2}{3} \pi^2 - \frac{1712}{105} \ln 2 \right. \\
& - \frac{856}{105} \ln \nu_\Omega - \left(\frac{278185}{33264} - \frac{41}{96} \pi^2 \right) \nu - \frac{20261}{2772} \nu^2 \\
& \left. + \frac{114635}{99792} \nu^3 + \frac{428i}{105} \pi \right] + \mathcal{O}(\epsilon^{7/2}) \left. \right\}. \tag{5.71}
\end{aligned}$$

For the adiabatic models, the orbital phase is obtained by rewriting the energy balance equation (5.12) as

$$\frac{d\Omega}{dt} = -\frac{F}{dE/d\Omega}. \tag{5.72}$$

and integrating this equation along with $d\Phi/dt = \Omega$. The Taylor approximants are formed first by substituting $F = F_n$ and $E = E_n$ into Eq. (5.72). The P-approximant

waveform is formed similarly by substituting $F = F_n^m$ and $E = E_n^m$ into Eq. (5.72). The Taylor T1 and Padé approximants then numerically integrate Eq. (5.72). The Taylor T4 approximant is formed by first re-expanding the right side of Eq. (5.72) as a single Taylor expansion truncated at the appropriate order, and then numerically integrating the resulting equation. The Taylor T2 and Taylor T3 approximants perform the integration analytically. The various Taylor approximants are reviewed in Sec. III E of Ref. [59].

For the non-adiabatic EOB models, the orbital phase is determined by solving Hamilton's equations (5.60)–(5.63). After computing h_{22} , the appropriate time derivatives are taken to form \dot{h}_{22} and Ψ_4^{22} .

5.4 Comparison with post-Newtonian approximants: Energy flux

We now compare the numerical GW energy flux with predictions from PN theory. In Sec. 5.4.1 we present comparisons with T-, P-, and E-approximants, and in Sec. 5.4.2 we explore ways of fitting the numerical flux by introducing higher-order PN coefficients and varying the value of ν_{pole} away from $\nu_{\text{pole}}^{2\text{PN}}$ [Eq. (5.41)].

The PN flux is derived as a function of frequency, so it is natural to perform this comparison as a function of frequency. One alternative, comparison as a function of time, would require computation of the PN phase as a function of time. This depends on the PN energy, so that a comparison with respect to time would mix effects due to flux and energy. Furthermore, comparisons with respect to time are sensitive to (and likely dominated by) secularly accumulating phase

differences [23].

The PN flux is given in terms of the *orbital* frequency Ω —see Eqs. (5.19) and (5.13)—so at first glance, it might seem natural to compare PN and NR energy fluxes at particular values of Ω . However, the orbital frequency is gauge-dependent, and there is no simple relation between the NR orbital frequency and the PN orbital frequency. Nor is there a simple relation between the NR orbital frequency and any quantity measured at infinity (where the energy flux is defined). In particular, it is very difficult to determine the NR orbital frequency as a function of retarded time. In contrast, the frequency ϖ (see Eq. (5.5)) of the GWs at infinity is an observable quantity, and is easily obtained from both PN formulas and from the NR simulation. Therefore, to achieve a meaningful comparison, we compare the PN and NR energy flux at particular values of ϖ .

In order to compute the PN flux as a function of ϖ , we need to find the mapping $\varpi_{\text{PN}}: \Omega \rightarrow \varpi$. In order to find this mapping, we must build a PN waveform as a function of Ω and compute ϖ as defined by Eq. (5.5). We construct the waveforms as described in Sec. 5.3.4. For the T-approximant of the flux, we will use the Taylor T4 waveform. In Fig. 5.5 we plot both GW frequencies (defined in Eqs. (5.4) and (5.5)). We then invert the mapping to obtain $\Omega_{\text{PN}} = \varpi_{\text{PN}}^{-1}: \varpi \rightarrow \Omega$. So, given the PN flux $F(\Omega)$ from Sec. 5.3, the flux as a function of the GW frequency is given by $F(\varpi) = F(\Omega_{\text{PN}}(\varpi))$. The relation $\Omega_{\text{PN}}(\varpi)$ depends on the instantaneous evolution of the PN model around frequency Ω , and is therefore (unfortunately) dependent on the PN model, in particular the choice of PN energy. This dependence, however, is local and will not lead to secularly accumulating differences.




Figure 5.5: Ratio of GW frequencies ω and $\bar{\omega}$ to orbital frequency Ω

The data are shown as functions of (twice) the orbital frequency, for different PN models. The GW frequencies ω and $\bar{\omega}$ are defined in Eqs. (5.4) and (5.5). Solid lines correspond to 3.5PN, dashed and dotted lines to 3PN and 2.5PN, respectively.




Figure 5.6: Effect of choice of frequency type

Shown are the PN fluxes for two representative PN approximants, plotted (correctly) as function of ω and (incorrectly) as function of 2Ω . Plotting as a function of 2Ω changes the PN fluxes significantly relative to the numerical flux F_{NR} .

Notice from Fig. 5.5 that the orbital frequency and the GW frequency differ by $\sim 1\%$ – 3% at large frequencies, depending on the PN model and the PN order, and the difference in ω between different PN models is about 5%. Because the energy flux is roughly proportional to $|\omega|^{10/3}$ (more precisely, $d \log F / d \log(M\omega)$ increases to ~ 3.6 at $-M\omega = 0.15$), the difference in the flux caused by using GW frequency from different PN models is about three to four times the difference in GW frequencies. Fig. 5.6 illustrates this effect by intentionally plotting the PN flux versus the incorrect frequency Ω . Because changing the PN model has a significant effect on the flux, we consider flux comparisons for several different

PN models below.

Note that for the flux comparison (and the comparisons of the derivative of the energy in Sec. 5.5), the PN waveforms are used only to define the mapping between Ω and ω . The PN flux is taken directly from the PN flux expressions—e.g., Eq. (5.19)—and *not* computed by applying Eq. (5.3) to PN waveforms $h(t)$. Equation (5.3) is used only to compute the numerical flux.

5.4.1 Flux comparison

Figure 5.7 plots the NR flux and the fluxes for the T-, P-, and E-approximants at 3.5PN order as a function of the GW frequency ω computed from \dot{h}_{22} . The T-approximant is Taylor T4 [59]. Along the top of this figure (as in several figures below) we indicate the number of gravitational wave cycles up to merger, where we define “merger” as the maximum of $|\Psi_4^{22}|$. Figure 5.8 zooms over the first 15 GW cycles. We notice that during the first 15 GW cycles the numerical data are fit best by the P- and E-approximants at 3PN and 3.5PN order. At these low frequencies the NR flux is best matched by the Keplerian and non-Keplerian EOB models and the Padé model.

To more clearly show the behavior of the PN approximants, we plot in Fig. 5.9 the energy flux normalized by the Newtonian flux. The normalized flux is computed as

$$\frac{F(\omega)}{F_{\text{Newt}}(\omega)} := \frac{F(\omega)}{\frac{32}{5}v^2 \left(\frac{M|\omega|}{2}\right)^{10/3}}, \quad (5.73)$$

where for the same reason mentioned above, the Newtonian flux is expressed in




Figure 5.7: **Comparison of NR and PN energy flux**

Here we compare the numerical energy flux and several PN approximants at 3.5PN order versus GW frequency ω extracted from \dot{h}_{22} in the equal-mass case.




Figure 5.8: Early-time comparison of NR and PN energy flux

Here we compare PN approximants versus GW frequency ω extracted from \dot{h}_{22} in the equal-mass case at early times. We show the relative difference between numerical flux and PN flux, as well as the estimated error of the numerical flux (blue bars, see Fig. 5.2). Solid lines represent 3.5PN models and NR; dashed and dotted lines correspond to 3PN and 2.5PN models, respectively. For notation see Table 5.1 and its caption.




Figure 5.9: **Comparison of normalized energy flux F/F_{Newt} for equal-mass systems**

See Eq. (5.73) for the definition of the normalized F_{Newt} . Solid lines represent 3.5PN models and NR; dashed and dotted lines correspond to 3PN and 2.5PN models, respectively. For notation see Table 5.1 and its caption.

Table 5.2: Normalized energy flux F/F_{Newt} for the T- and P-approximants at various PN orders and velocities v_{Ω}

PN order ($n+m$)/2	$v_{\Omega} = 0.1$		$v_{\Omega} = 0.25$		$v_{\Omega} = 0.3$		$v_{\Omega} = 0.35$		$v_{\Omega} = 0.4$	
	$2M\Omega = 0.002$	$2M\Omega = 0.031$	$2M\Omega = 0.054$	$2M\Omega = 0.086$	$2M\Omega = 0.128$	F_{n+m}/F_{Newt}	F_n^m/F_{Newt}	F_{n+m}/F_{Newt}	F_n^m/F_{Newt}	F_{n+m}/F_{Newt}
0.0	1.0000000	1.1692906	1.0000	1.5673	1.000	1.7678	1.000	2.027	1.000	2.376
0.5	1.0000000	1.0214102	1.0000	1.1507	1.000	1.2325	1.000	1.345	1.000	1.505
1.0	0.9555952	0.9251084	0.7225	-0.8648	0.939	-7.8434	0.456	16.01	1.091	8.443
1.5	0.9681616	0.9686094	0.9188	0.9074	0.940	0.9069	0.995	0.924	1.094	0.967
2.0	0.9681512	0.9676191	0.9184	0.8850	0.939	0.8671	0.993	0.860	1.091	0.867
2.5	0.9675775	0.9676981	0.8624	0.8890	0.799	0.8754	0.692	0.875	0.504	0.893
3.0	0.9677265	0.9677247	0.8951	0.8914	0.895	0.8804	0.928	0.883	1.022	0.903
3.5	0.9677274	0.9677233	0.8957	0.8912	0.897	0.8798	0.934	0.882	1.036	0.900

$v_{\Omega} = 0.25$ corresponds to the start of the numerical simulation. The P-approximant flux is given by Eq. (5.39). Note that the P-approximant has an extraneous pole at 1PN order at $v_{\Omega} = 0.326$. We use $v_{\text{iso}} = v_{\text{iso}}^{2\text{PN}} = 0.4456$ and $v_{\text{pole}} = v_{\text{pole}}^{2\text{PN}} = 0.6907$. We use boldface to indicate the range of significant figures that do not change with increasing PN order.

terms of the GW frequency. Notice that the P-approximants and some of the E-approximants use the same Padé flux, but they start differing at $-M\omega \sim 0.12$ due to their different GW frequencies (obtained from an adiabatic and nonadiabatic evolution, respectively). The E-approximants with Keplerian and non-Keplerian flux increase less abruptly at high frequency than the P- and T-approximants. This is a consequence of nonadiabatic effects captured by the EOB model. Quite remarkably, the E-approximants with non-Keplerian fluxes are rather close to the NR result for the entire range of frequency spanned by the simulation.⁹ We observe that somewhat accidentally the PN approximants at 2.5PN order are also close to the numerical flux.

The normalized NR flux starts to decrease at $-M\omega \sim 0.13$. We notice that this behavior is rather different from the behavior of the normalized flux in the test-mass limit (see Figs. 5.19 and 5.20 in Sec. 5.8). The E-approximants with non-Keplerian Padé or Taylor flux show a similar decreasing behavior at high frequency.

Both Figs. 5.8 and 5.9 show that in the equal-mass case P-approximants fit the numerical results better than T-approximants. In numerical analysis, however, Padé summation is often used as a technique to accelerate the convergence of a slowly converging Taylor series (e.g., see Tables 8.9 and 8.12 in Ref. [31]); hence it is natural to ask in the PN case whether Padé summation indeed accelerates the convergence of the series. In Table 5.2 we list the T- and P-approximants of F/F_{Newt} computed at subsequent PN orders and for several values of ν_Ω

⁹We notice that whereas the Keplerian Padé-based (or Taylor-based) approximants to the flux differ from each other only when expressed in terms of the GW frequency, the non-Keplerian Padé-based (or Taylor-based) approximants to the flux differs from the others because their functional dependence on the frequency is different (e.g., compare Eq. (5.65) with Eq. (5.64)).




Figure 5.10: Cauchy convergence test of F/F_{Newt} for T- and P-approximants

We plot $\Delta F_{n+m} := F_{n+m+1} - F_{n+m}$, and $\Delta F_n^m := F_{n+1}^m - F_n^m$ for different values of ν_Ω . The T- and P-approximants are given by Eqs. (5.19) and (5.39), respectively. Note that the P-approximant has an extraneous pole at 1PN order at $\nu_\Omega = 0.326$. We use $v_{\text{iso}} = v_{\text{iso}}^{2\text{PN}}$, and $v_{\text{pole}} = v_{\text{pole}}^{2\text{PN}}$.

[from left to right $\nu_\Omega = 0.1, 0.25$ (i.e., beginning of the numerical simulation), 0.3, 0.35, and 0.4]. In Fig. 5.10 we perform a Cauchy convergence test and compute the difference between T- and P-approximants at subsequent PN orders. The figures do not suggest an acceleration of the convergence. We notice that in the equal-mass case P-approximants are converging more systematically than T-approximants. However, this fact seems to depend on the mass ratio, as can be seen by comparing Fig. 5.10 with Table 5.4 and Fig. 5.22 in the appendix which are obtained in the test-mass limit.

5.4.2 On the fitting of the numerical relativity energy flux

In view of building accurate analytical templates that can interpolate the NR waveforms during inspiral, merger, and ringdown, we explore here the possibility of improving the PN approximants to the energy flux by introducing *phenomenological* higher-order PN coefficients and/or by varying the value of ν_{pole} . This study should be considered a first exploration of the problem, demonstrating only the *flexibility* of the PN models. None of the *quantities* derived here should be used as the basis for further work.

We will minimize the difference between the PN flux and the numerical flux by varying particular coefficients in the PN model. Ideally, the PN and numerical fluxes should be expressed as functions of ϖ before taking this difference, so that the fluxes are compared in a physically meaningful way. Unfortunately, the calculation of ϖ for the PN models is time-consuming, because for each trial value of the phenomenological coefficient it is necessary to compute a full waveform to determine the mapping between ϖ and Ω . So instead, in this section we




Figure 5.11: Fitting several PN approximants to the numerical flux

The x axis denotes the orbital frequency Ω . Because the numerical flux is computed as function of the GW frequency, we use for the numerical flux $\Omega := -\omega/2$. The blue bars indicate estimated errors on the numerical flux, see Fig. 5.2. For notation see Table 5.1 and its caption.

simply compare PN and numerical fluxes as functions of Ω , where we define the numerical orbital frequency as $\Omega := -\varpi/2$. In Fig. 5.6, we can see that the error introduced by the discrepancy between Ω and $\varpi/2$ will be significant. As we will show in Sec. 5.6.2, the waveforms produced using these “tuned” flux functions will improve agreement with the numerical waveform at a significant level. Nevertheless, the values derived in this section may not be optimal. Thus, we emphasize that the results of this section constitute merely an exercise demonstrating the feasibility of adjusting the PN parameters to optimize the agreement of the PN flux function with numerical data.

The least-squares fits are done on $F(\varpi)/F_{\text{Newt}}(\varpi)$ [see Eq. (5.73)]. In the case of T-approximants, we fit for the unknown 4PN-order coefficient in Eq. (5.19) for the equal-mass case. We perform a least-squares fit of the 4PN-order function $\mathcal{F}_8(\nu = 1/4) = A_8 + B_8 \log \nu_\Omega$ over the orbital-frequency range $M\Omega = 0.02\text{--}0.08$ which starts after the first 9 GW cycles. We obtain $A_8 = -141, B_8 = 102$. We notice that when we perform the fit over the first 15 (or 20) GW cycles, spanning the frequency region $M\Omega = 0.0168\text{--}0.0235$ ($M\Omega = 0.0168\text{--}0.0283$), the agreement becomes worse. The resulting flux is shown in Fig. 5.11. The relative difference with the numerical flux is at most $\sim 0.8\%$.

We repeat this analysis in the case of P-approximants. Because the latter also depend upon ν_{pole} , we perform two least-squares fits. In the first fit, we fix ν_{pole} to the value given by Eq. (5.41) and apply the least-squares fit to $\mathcal{F}_8(\nu = 1/4)$ obtaining $A_8 = -1382, B_8 = 197$.

In the second fit, we vary ν_{pole} . When varying ν_{pole} in the P-approximant at 3.5PN order, extraneous poles appear at low values of ν_Ω . Therefore, in order to

push these poles to very high frequency, we follow the suggestion of Ref. [122], and use P-approximants at 4PN order, where the 4PN coefficient is set to its known value in the test-mass limit. Furthermore the logarithm in the flux is not factored out, but treated as a constant when Padé summation is done. This cure may fail for different mass ratios if new extraneous poles appear at low frequency. The least-squares fit gives $\nu_{\text{pole}} = 0.74$. All the results for the P-approximants are displayed in Fig. 5.11, where we also show the T- and P-approximants at 3.5PN order without any fit.

Figure 5.11 might suggest that by introducing higher-order PN coefficients in the flux, the numerical flux can be fit better by T-approximants than by P-approximants. However, this result can depend on the use of orbital frequency instead of GW frequency. In Sec. 5.6.3 (see Fig. 5.18) we employ the fit values obtained in this study and show phase differences between NR and tuned EOB models.

Finally, we attempted to extract PN coefficients higher than 3.5PN order from the numerical flux, as was done at 2PN, 2.5PN, and 3PN order in Ref. [101] in the test-mass limit. Unfortunately, the differences between numerical flux and T-approximants are so large—even at the beginning of the numerical waveform—that we were not able to extract even known PN coefficients, like the ones at 3PN and 3.5PN order. Thus, to fit unknown PN coefficients would require a numerical simulation with more cycles starting at lower frequency.




Figure 5.12: GW frequency derivative $\dot{\omega}$ for the numerical relativity simulation and various PN approximants at 3.5PN order

For notation see Table 5.1 and its caption.

5.5 Estimation of (the derivative of) the center-of-mass energy

In the previous section, we analyzed and compared PN and numerical energy fluxes. The energy of the binary is the second fundamental ingredient in the construction of adiabatic PN approximants. Unfortunately, there is no way to extract the energy for the numerical simulation as a function of a gauge-invariant quantity such as the GW frequency, so that it is impossible to compare PN and NR energies directly. The frequency derivative, $\dot{\omega}$, however, is easily accessible in the numerical data, and, in the adiabatic approximation is intimately related to the energy, as can be seen by rewriting the energy balance, Eq. (5.12), in the form

$$\frac{d\omega}{dt} = -\frac{F}{dE/d\omega}. \quad (5.74)$$

Therefore, we begin this section with a comparison between numerical $\dot{\omega}$ and the predictions of various PN approximants. For the PN approximants, we compute h_{22} as usual (i.e., using energy balance to compute the orbital frequency derivative $\dot{\Omega}$), and take a time derivative to obtain \dot{h}_{22} and extract $\dot{\omega}$ from it. The waveform h_{22} for the E-approximants is computed using Eqs. (5.42), (5.44), (5.45), and (5.46) in Sec. 5.3.3. Figure 5.12 plots the numerical $\dot{\omega}$ and its value for T-, P-, and E-approximants at 3.5PN order.

In order to emphasize differences between the different $\dot{\omega}$, we normalize the data in Fig. 5.12 by the Newtonian value of $\dot{\omega}$,

$$\frac{\dot{\omega}}{\dot{\omega}_{\text{Newt}}} := -\frac{\dot{\omega}}{\frac{192}{5} \frac{v}{M^2} \left(\frac{M|\omega|}{2}\right)^{11/3}}. \quad (5.75)$$




Figure 5.13: Comparison of $\dot{\omega}$ for the numerical results and various PN approximants

Dotted, dashed, and solid lines correspond to 1.5PN, 2.5PN, and 3.5PN models, respectively. For notation see Table 5.1 and its caption.

The normalization is used only to eliminate the leading-order behavior of the various curves in Fig. 5.12; therefore, to compute the denominator of Eq. (5.75) we have simply substituted $-\omega/2$ rather than Ω into the Newtonian formula for the frequency derivative.

The normalized frequency derivatives are shown in Fig. 5.13. At low frequencies, $\dot{\omega}$ is very challenging to compute in numerical simulations, resulting in comparatively large numerical uncertainties. Therefore, for frequencies $-M\omega \lesssim 0.045$ we can merely conclude that PN and NR are consistent with each other (i.e., are within the numerical error bars of about 10 percent).

The 3.5PN Taylor T4 model (labeled F_7/E_6T4) agrees very well with the numerical simulation up to $-M\omega \approx 0.1$; this observation is consistent with the excellent agreement between Taylor T4 (3.5PN) and the numerical simulation observed in Boyle et al. [59], who compared up to this frequency. Beyond $-M\omega = 0.1$, however, $\dot{\omega}/\dot{\omega}_{\text{Newt}}$ for Taylor T4 continues to increase (as for all other Taylor and Padé models considered here), whereas for the numerical simulation, $\dot{\omega}/\dot{\omega}_{\text{Newt}}$ flattens (this behavior was also observed in Ref. [122].) Only the E-approximants at 3.5PN order reproduce the flattening of $\dot{\omega}/\dot{\omega}_{\text{Newt}}$ at high frequencies, with the closest being the one which uses the non-Keplerian Padé flux (${}^{\text{nk}}F_4^3$). Because the frequency derivative is the relevant quantity that determines the phase evolution, the turning over of $\dot{\omega}/\dot{\omega}_{\text{Newt}}$ for the nonadiabatic models in Fig. 5.13 suggests that, at high frequency, nonadiabatic analytical models might be superior to adiabatic models.

If sufficient smoothing is applied to the numerical $\dot{\omega}$ it becomes a smooth curve even at low frequencies. Figure 5.14 presents a comparison between such a




Figure 5.14: Comparison of PN $\dot{\omega}$ with a heavily smoothed version of the numerical $\dot{\omega}$

Solid lines represent 3.5PN models and NR; dashed and dotted lines correspond to 3PN and 2.5PN models, respectively. For notation see Table 5.1 and its caption.

heavily smoothed numerical curve and the PN approximants. As already pointed out, all PN approximants are consistent to within our estimated numerical errors at low frequencies. However, the NR result in Fig. 5.14 is notably closer to the 3.5PN approximants than to lower order PN approximants. This good agreement provides a further validation of the numerical code used in Boyle et al. [59]. It also indicates that our error analysis in Sec. 5.2 may be overly conservative.

Our comparisons of $\dot{\omega}$ reveal a lot of information about the PN approximants. However, $\dot{\omega}$ depends on both flux and energy (see Eq. (5.74)), and so these comparisons do not yield information about flux or energy separately. To isolate effects due to the PN energy, we rearrange Eq. (5.74) further, such that it yields in the adiabatic approximation the derivative of the center-of-mass energy for the numerical simulation:

$$\left[\frac{dE}{d\omega} \right]_{\text{NR}} = - \frac{F_{\text{NR}}}{[d\omega/dt]_{\text{NR}}} . \quad (5.76)$$

The relative error in $[dE/d\omega]_{\text{NR}}$ is obtained as the root-square-sum of the relative errors of flux and frequency derivative (see Figs. 5.2 and 5.4). In Fig. 5.15 we compare the latter with T-, P-, and E-approximants. For adiabatic T4 and Padé models, we compute $dE/d\omega$ by taking derivatives of $E(\Omega)$ in Eq. (5.14) with respect to Ω and then expressing the derivative in terms of $\omega(\Omega)$. For nonadiabatic EOB models, we compute $dE/d\omega$ from the ratio of F_{PN} and $[d\omega/dt]_{\text{PN}}$ as obtained from Figs. 5.7 and 5.12. The closeness between the numerical result and adiabatic PN approximants is expected only in the range of frequencies over which the balance equation and the adiabatic approximation are valid. The upper panel of Fig. 5.15 shows the Taylor and Padé adiabatic models. The plot suggests that around $-M\dot{\omega} \sim 0.08$ non-adiabatic effects are no longer negligible.

At lower frequencies, both 3.5PN-order adiabatic approximants (Padé and Taylor T4) match the numerical result very well. Taylor T4 at 2.5PN matches well, too, although its frequency derivative $\dot{\omega}$ and flux differ significantly from NR (see Figs. 5.13 and 5.9). The T-approximant at 3.5PN order is closest to the numerical result. The lower panel of Fig. 5.15 shows the nonadiabatic E-approximants. We notice that the nonadiabatic models, especially at 3.5PN order, follow quite nicely the behavior of the numerical derivative of the center of mass energy. The E-approximant with non-Keplerian flux is closest to the numerical result. This analysis emphasizes again the relevance of including nonadiabatic effects in the analytical model [77].

5.6 Comparing waveforms

Here we compare the numerical waveform to various PN waveforms, basically extending the analysis of Boyle et al. [59] to include Padé and EOB waveforms. Because the (2,2) mode dominates the waveform for an equal-mass, nonspinning binary, we restrict the comparison to only this mode. As in [59], we use Ψ_4^{22} and the GW phase and frequency ω defined by Eq. (5.4) when comparing waveforms.

For the comparisons presented in this section, the uncertainty in the phase of the numerical waveform is roughly 0.02 radians. This number includes numerical errors (e.g., due to convergence and extrapolation of the waveform to infinite extraction radius), as well as modeling errors due to slightly nonzero eccentricity and spin of the numerical simulation; see Ref. [59] Sec. V for details. We note that the modeling errors have decreased since the analysis in Ref. [59] because




Figure 5.15: **Comparison of NR and PN $dE/d\omega$ versus GW frequency ω**

(See Eq. (5.76) for our definition of $dE/d\omega$). Solid lines represent 3.5PN models and NR; dashed and dotted lines correspond to 3PN and 2.5PN models, respectively. For notation see Table 5.1 and its caption.

the new matching procedure reduces the impact of eccentricity, and because the more sophisticated spin-diagnostics presented in Ref. [190]) resulted in a smaller bound on the residual spin.

5.6.1 Matching procedure

Each PN waveform has an arbitrary time offset, t_0 , and phase offset, ϕ_0 with respect to the NR waveform. The procedure used by Boyle et al. [59]—as well as in various other papers before it, such as [24, 157]—sets these constants by ensuring that the GW phase and frequency match at a fiducial time. Unfortunately, when matching at low frequency this method is sensitive to noise and to residual eccentricity in the numerical waveform, and does not easily translate into a robust and automatic algorithm. Since we want to match as early as possible (where we expect the PN approximants to be valid), we propose to use, instead, a matching procedure which achieves the same goal, but extends over a range of data. This procedure is similar to the one proposed by Ajith et al. [5], but whereas we match only the GW phase, Ajith et al. match the entire gravitational waveform—including the amplitude—and include an overall amplitude scaling. This method can be easily implemented as a fairly automatic algorithm, robust against noise and residual eccentricity.

Using the phase of the numerical and PN waveforms, we define the quantity

$$\Xi(\Delta t, \Delta\phi) = \int_{t_1}^{t_2} [\phi_{\text{NR}}(t) - \phi_{\text{PN}}(t - \Delta t) - \Delta\phi]^2 dt . \quad (5.77)$$

Here, t_1 and t_2 represent the chosen range over which to compare. Minimizing this quantity by varying the time and phase offsets Δt and $\Delta\phi$ produces the optimal values for these quantities in a least-squares sense. Then to compare PN

and NR waveforms, we compare the (unchanged) NR waveform with an offset PN waveform defined by

$$\Psi_{4,\text{PN}}(t) = A_{\text{PN}}(t + \Delta t) e^{-i[\phi_{\text{PN}}(t + \Delta t) + \Delta\phi]} . \quad (5.78)$$

With reasonable first guesses for Δt and $\Delta\phi$, the function Ξ is quite nicely paraboloidal. Thus, even simple minimization routines work well. However, in cases where speed is an issue, the problem can be reduced to one dimension. For a given value of Δt , the optimization over $\Delta\phi$ may be done analytically by setting

$$\Delta\phi(\Delta t) = \frac{\int_{t_1}^{t_2} [\phi_{\text{NR}}(t) - \phi_{\text{PN}}(t - \Delta t)] dt}{t_2 - t_1} . \quad (5.79)$$

Using this value of $\Delta\phi$ for a given value of Δt decreases the number of function evaluations needed to find the minimum. This can be very useful for large data sets, or situations where many such matches need to be done.

The choice of t_1 and t_2 involves some degree of judgment. Preferably, t_1 should be as early as possible, while not being contaminated by junk radiation. We choose $t_1 = 1100M$, corresponding to $-M\omega = 0.037$. Similarly, t_2 should be as early as possible, but far enough from t_1 so that the integration averages over the noise. In addition, the effects of the small but nonzero orbital eccentricity show up as oscillations in the phase, as can be seen, for example, in the range $t \in [1100, 1900]M$ in Fig. 5.17. We would like t_2 to be large enough so that the integration averages over several cycles of this oscillation, thus resulting in less bias due to eccentricity. Here we use $t_2 = 1900M$, corresponding to $-M\omega = 0.042$. We have checked that changing the values of t_1 and t_2 by $\pm 100M$ changes the resulting phases by less than a few thousandths of a radian through the end of the numerical waveform.

This method is quite similar to the one suggested in Ref. [5]. However, here we consider only the phase and not the amplitude of the waveform. Because we restrict the analysis only to the (2,2) waveform mode of an equal-mass binary and compare only the phase and not the amplitude, we think it is reasonable to have neglected the amplitude in the matching procedure.

5.6.2 Padé waveforms

In Fig. 5.16 we plot the phase difference between the numerical, T- and P-approximants [106, 114, 73] at the times when the numerical waveform reaches GW frequencies $-M\omega = 0.063$ and $-M\omega = 0.1$. The phase differences are plotted versus the PN order. The phase difference at $-M\omega = 0.1$ of the P-approximant at 3.5PN order is -0.12 radians. When comparing with generic Taylor approximants, we notice that the phase differences of the P-approximants are less scattered as the PN order is increased. This might be due to the fact that P-approximants of the energy flux are closer to the NR flux, especially for lower ν_Ω where the phase accumulates the most. Figure 5.16 could be contrasted with Tables III and IV of Ref. [106] which show the overlaps between the numerical waveform and P-approximants at subsequent PN orders, in the test-mass limit case. The behavior of the P-approximants in Fig. 5.16 are consistent with the behavior of $\dot{\omega}$ seen in Fig. 5.13: At 1.5PN, Padé has larger $\dot{\omega}$ than the numerical simulation, at 2.5PN, Padé has smaller $\dot{\omega}$. Consequently, $\Phi_{\text{PN}} - \Phi_{\text{NR}}$ is negative at 1.5PN order and positive at 2.5PN order. For 3.5PN order, the P-approximant in Fig. 5.13 agrees very well with the numerical simulation (at least for $-M\omega \lesssim 0.1$), which translates into excellent agreement in Fig. 5.16.




Figure 5.16: Phase differences between the numerical waveform, and untuned, original EOB, untuned Padé, and Taylor waveforms, at two selected times close to merger

The E-approximants are F_n^m/H_p , while the P-approximants are F_n^m/E_p^q (see Table 5.1 and its caption). Waveforms are matched with the procedure described in Sec. 5.6.1 and phase differences are computed at the time when the numerical simulation reaches $-M\omega = 0.063$ (left panel) and $-M\omega = 0.1$ (right panel). Differences are plotted versus PN order. Note that at 1PN order the Padé flux has an extraneous pole at $\nu = 0.326$ causing a very large phase difference. The thick black line indicates the uncertainty of the comparison as discussed in Sec. 5.6, $|\Phi_{\text{PN}} - \Phi_{\text{NR}}| \leq 0.02$ radians.

In Fig. 5.17 we explore the possibility of reducing the phase differences between the numerical waveform and P-approximants: By (i) varying ν_{pole} or (ii) introducing the pseudo-4PN-order coefficient $\mathcal{F}_8(\nu = 1/4) = A_8 + B_8 \log \nu_\Omega$ in the energy flux. We tune the coefficients by minimizing the sum of the squares of the phase difference at $t_{0.063}$ and $t_{0.1}$. We find that if $\nu_{\text{pole}} = 0.633$, the P-approximant F_4^4/E_2^4 has a maximum phase difference before $-M\omega = 0.1$ smaller than the numerical error in the simulation. A similar result is obtained for the the P-approximant pF_4^4/E_2^4 if we use $\nu_{\text{pole}} = \nu_{\text{pole}}^{2\text{PN}} = 0.6907$, and tune $A_8 = -493$, $B_8 = 330$.

5.6.3 Effective-one-body waveforms

In Fig. 5.16 we also plot the phase differences between the numerical and the untuned, original E-approximants [76, 77, 114] F_n^m/H_p . At 3.5PN order the phase difference at $-M\omega = 0.1$ is 0.50 radians. We also computed the phase differences at $-M\omega = 0.1$ of the E-approximants ${}^{\text{nK}}F_4^3/H_7$, ${}^{\text{nK}}F_7/H_7$, and F_7/H_7 and found 0.45, 2.56, and 2.7 radians, respectively. Thus, for untuned EOB models it is crucial to have introduced the Padé flux. When contrasting the original E-approximants with generic Taylor approximants, we find that the phase differences are less scattered as the PN order is increased. However, despite the fact that the Padé-based EOB flux is closer to the numerical flux (see Figs. 5.8 and 5.9), untuned, original E-approximants accumulate more phase difference than P-approximants. This could be a consequence of the fact that independently of the flux and the energy functions, what seems to matter is the way the equations of motions are solved to get the phasing.




Figure 5.17: **Phase differences between untuned and tuned P-approximants and NR waveforms**

The untuned P-approximant is F_4^3/E_2^4 ($v_{\text{iso}} = v_{\text{iso}}^{2\text{PN}}$, $v_{\text{pole}} = v_{\text{pole}}^{2\text{PN}}$). The tuned P-approximants are F_4^4/E_2^4 and tunable v_{pole} ($v_{\text{iso}} = v_{\text{iso}}^{2\text{PN}}$) and pF_4^4/E_2^4 ($v_{\text{iso}} = v_{\text{iso}}^{2\text{PN}}$, $v_{\text{pole}} = v_{\text{pole}}^{2\text{PN}}$) with tunable A_8 and B_8 . In all cases, waveforms are matched over $t - r_* \in [1100, 1900]M$.

Because of the reduction of the dynamics to *a few* crucial functions determining the inspiral evolution [76, 77, 103], notably A , D and \mathcal{F} , and because of the rather simple procedure to match the inspiral(-plunge) waveform to the ringdown waveform, the EOB model turned out to be particularly suitable for matching the full numerical waveforms [75, 78, 121, 122, 124]. In view of a future study which will include merger and ringdown, we start here exploring the possibility of improving the agreement with numerical waveforms by tuning the pseudo-4PN-order coefficients a_5 , A_8 , and B_8 and/or, if present, the pole location ν_{pole} . In the lower panel of Fig. 5.18, using different ν_{pole} values, we show the phase differences computed at $t_{0.063}$ and $t_{0.1}$ as functions of the unknown PN-expansion coefficient a_5 [see Eq. (5.49)]. As first pointed out and discussed in Ref. [122] (see, e.g., Fig. 3 in that paper), we find that there is a strong degeneracy between a_5 and ν_{pole} . In fact, for different ν_{pole} values, the curves in Fig. 5.18 are almost identical except for a shift in a_5 . Although in this test we use the E-approximant $F_4^4/pH_8(\nu_{\text{lso}} = \nu_{\text{lso}}^{2\text{PN}})$, we find that this degeneracy appears in all E-approximants considered.

To obtain the optimal a_5 and ν_{pole} that minimize phase differences during the entire numerical simulation, we first choose an arbitrary ν_{pole} in the range of degeneracy. Then, we determine the a_5 value by minimizing the sum of the squares of the phase difference at $t_{0.063}$ and $t_{0.1}$. In the upper panel of Fig. 5.18, we show phase differences in time and GW frequency for several E-approximants using those optimal a_5 and ν_{pole} values, which are given in Table 5.3. In Fig. 5.18, we also show phase differences for E-approximants with pseudo 4PN order coefficients determined by the flux fit of Sec. 5.4.2 (see Fig. 5.11) and tunable a_5 .

Table 5.3: **Optimal a_5 and ν_{pole} that minimize phase differences between tuned EOB models and the numerical simulation**

EOB model and fixed parameters	a_5	ν_{pole}
${}^{\text{nK}}\bar{F}_4^4/pH_8$	—	29.78
F_4^4/pH_8	$\nu_{\text{iso}} = \nu_{\text{iso}}^{\text{2PN}}$	39.35
pF_8/pH_8	$A_8 = -141, B_8 = 102$	5.32
pF_4^4/pH_8	$A_8 = -1382, B_8 = 197,$ $\nu_{\text{iso}} = \nu_{\text{iso}}^{\text{2PN}}, \nu_{\text{pole}} = \nu_{\text{pole}}^{\text{2PN}}$	-3.10

The optimal a_5 values are shown in Table 5.3. The smaller phase differences along the entire inspiral are obtained with the E-approximants with Padé flux F_4^4/pH_8 ($\nu_{\text{iso}} = \nu_{\text{iso}}^{\text{2PN}}$) and tunable ν_{pole}, a_5 and Taylor flux pF_8/pH_8 with tunable A_8, B_8, a_5 . We notice that for $t > t_{0.1}$ the phase difference increases more abruptly for the latter model. In the best case, the absolute phase difference during the entire numerical simulation is within the numerical error, i.e., within 0.02 radians. The choice of the best tuned E-approximant [206, 78, 122, 123, 124] will be determined once merger and ringdown are included, and when long and accurate comparisons with numerical simulations are extended to BBH with mass ratio different from one.

Finally, in Ref. [122], Damour and Nagar extracted the data of the numerical simulation used in the present paper from one of the figures of Ref. [59] and compared those data with the EOB approach. They found for their “non-tuned” EOB model phase differences of ± 0.05 radians. This phase difference is smaller

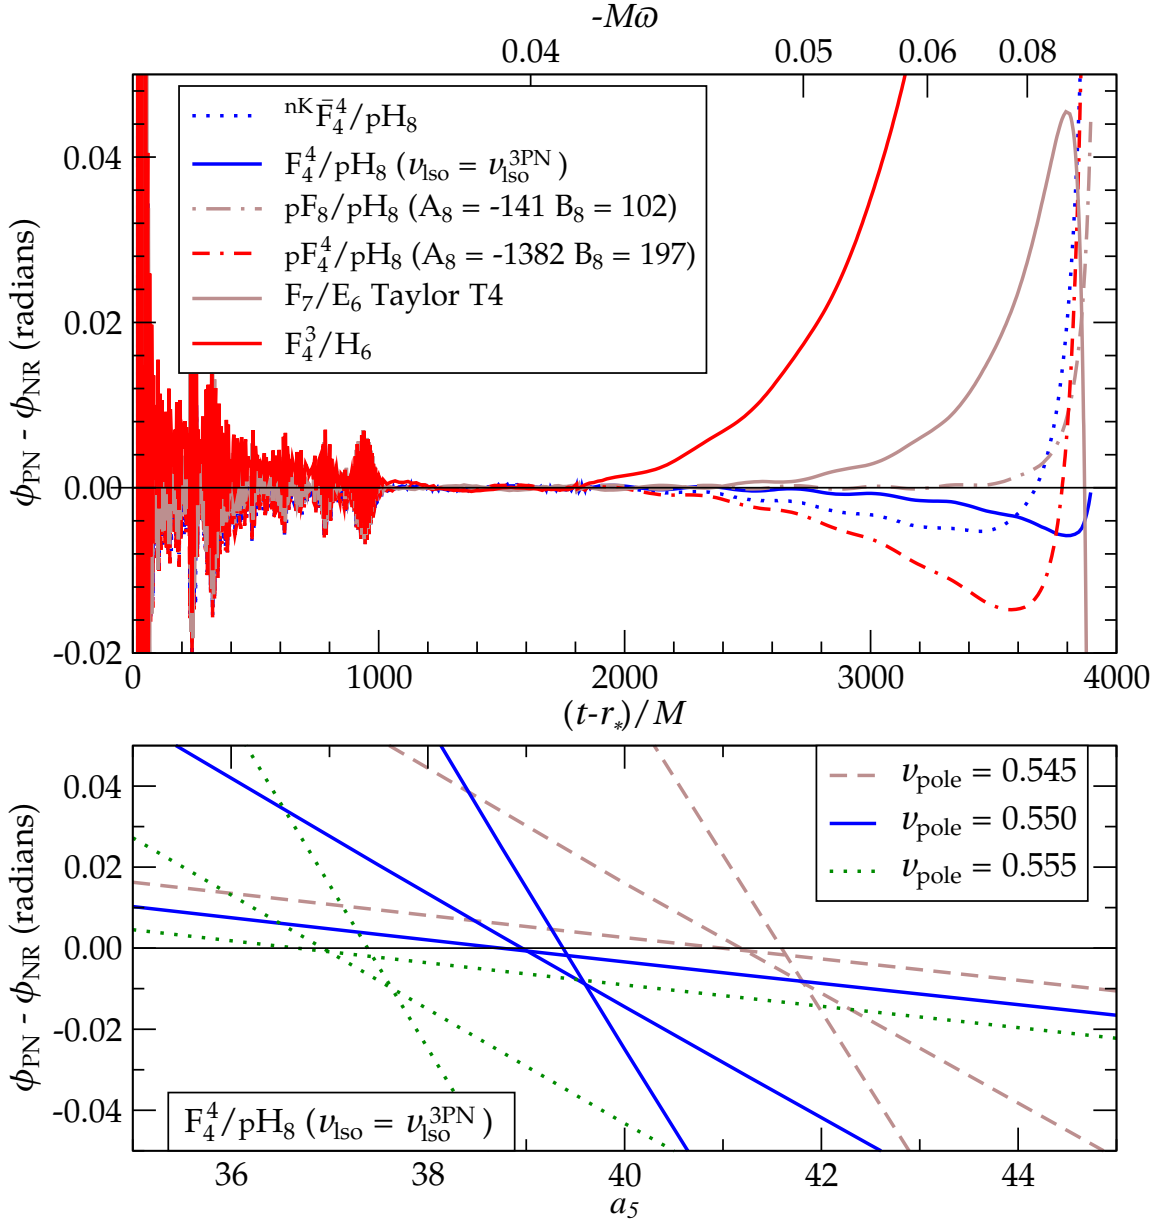


Figure 5.18: Phase accuracy of various PN approximants

The **upper panel** shows phase differences versus time (lower x axis) and versus GW frequency (upper x axis) for several tuned and untuned E-approximants. For the tuned models, the optimal a_5 and u_{pole} values are displayed in Table 5.3. In the **lower panel** we show phase differences between numerical and E-approximants computed at $t_{0.063}$, $t_{0.1}$, and the end of the numerical simulation $t_{0.16}$, as functions of a_5 . For the same color and style, the curve with the steepest slope corresponds to $t_{0.16}$ and the curve with the smallest slope corresponds to $t_{0.063}$. For notation see Table 5.1 and its caption.

than the phase differences we discuss in this paper for untuned EOB models (see Fig. 5.16 and discussion around it). However, we notice that ± 0.05 radians in Ref. [122] refers to *half* the maximum phase difference accumulated over the entire evolution when matching the numerical and EOB phases at $-M\omega = 0.1$. By contrast, in this paper, and in particular in Fig. 5.16, we match numerical and EOB phases in a time interval and compute the phase differences at $-M\omega = 0.1$.

Moreover, we observe that their “non-tuned” EOB model is not really untuned, because it uses the Padé summation of the radial potential at 4PN order and *then* sets $a_5 = 0$. This is not equivalent to using the radial potential at 3.5PN order with $a_5 = 0$. In fact, to recover the 3.5PN-order Padé radial potential from the 4PN-order Padé potential one should use $a_5 = -17.16$. They also use the non-Keplerian flux at 4PN order ${}^{\text{nK}}\overline{F}_4^4$ which is different from the 3.5PN order ${}^{\text{nK}}F_4^3$. For our untuned EOB model at 3.5PN order which uses ${}^{\text{nK}}F_4^3$ and the EOB dynamics at 3PN order, if we apply Ref. [122] procedure and compute *half* the maximum phase difference when matching the numerical and EOB phases at $-M\omega = 0.1$, we find a phase difference of ± 0.18 radians

5.7 Conclusions

In this paper, using a highly accurate and long numerical simulation [59] of a non-spinning equal-mass black hole binary, we compute the gravitational waveform, GW energy flux, and GW frequency derivative. Imposing the balance equation, we also estimate the (derivative of) center-of-mass energy. We compare these quantities to those computed using adiabatic Taylor T4 and Padé [106, 114, 73], and nonadiabatic EOB PN approximants [76, 77, 114].

We find that for the first 15 GW cycles, the 3.5PN-order T-approximant and the 3.5PN-order untuned P- and E-approximants (see Table 5.1) reproduce the numerical results for energy flux, GW frequency derivative and (derivative of) center-of-mass energy quite well (see Figs. 5.8, 5.9, 5.13, 5.14, and 5.15), but with interesting differences.

We attempted to study the convergence of the PN expansion for the energy flux.¹⁰ We find that Padé approximants to the flux introduced in Ref. [106] do not accelerate the convergence of the Taylor series, but are closer to the numerical flux than are the T-approximants. In particular, the Taylor flux at all orders through 3.5 PN is outside the numerical flux error bars even ~ 25 GW cycles before merger (see Fig. 5.8). We find that the nonadiabatic non-Keplerian E-approximants to the flux at 3.5PN order are within $\sim 2\%$ of the numerical flux over the entire frequency range we consider (see Fig. 5.9).

Quite interestingly, in the equal-mass case the numerical normalized energy flux F/F_{Newt} starts decreasing at high frequency during the late part of the inspiral and blurred plunge (see Fig. 5.9). This differs from the behavior of F/F_{Newt} in the test-mass limit (see Figs. 5.19 and 5.20). Both the Taylor and Padé-based E-approximants with non-Keplerian flux [104] show a similar decreasing behavior at high frequency. This fact suggests that if a pole is present in the energy flux of equal-mass binaries, it is located at a larger frequency than that at which the common apparent horizon forms. As seen in Sec. 5.4.2, when fitting for ν_{pole} we obtain $\nu_{\text{pole}}(\nu = 1/4) = 0.74$, which is to be contrasted with the test-mass case

¹⁰We also tried to apply the criterion suggested in Ref. [254] to assess the region of validity of the PN series for the flux in the equal-mass case. Unfortunately, the numerical simulation starts at too high a frequency, when the Taylor series at 3.5PN order seems to already be outside the region of validity.

$\nu_{\text{pole}}(\nu = 0) = 1/\sqrt{3} \approx 0.58$. These values of ν_{pole} correspond to orbital frequencies $M\Omega = 0.405$ and $M\Omega = 0.192$, respectively.

For the GW frequency derivative $\dot{\omega}$, we find that at low frequency the Taylor, Padé, and EOB models at 3.5PN order are within the numerical error (see Fig. 5.13). At high frequency, as already observed in Ref. [122], only the non-adiabatic E-approximant has a GW frequency derivative that flattens out, as does the numerical result. The non-Keplerian E-approximant at 3.5PN order is closest to the numerical data (see Fig. 5.14).

When estimating the derivative of center-of-mass energy $dE/d\omega$, we expect the numerical result and adiabatic PN approximants to be close only in the range of frequencies over which the balance equation and the adiabatic approximation are valid. We find that this range of frequencies is $-M\omega \gtrsim 0.08$ (see Fig. 5.15) for the 2.5PN T-approximant and all the 3.5PN approximants.¹¹ At higher frequency, the 3.5PN-order nonadiabatic E-approximants are closer to the numerical $dE/d\omega$ than are the adiabatic approximants, and the non-Keplerian E-approximant is the closest.

Applying a new matching procedure, we compared the numerical waveforms with Taylor T4, Padé, and EOB waveforms. We find that the accumulated phase difference from the numerical solution at $-M\omega = 0.1$ is 0.12 radians for the untuned 3.5PN P-approximant [106, 114, 73], -0.50 radians for the untuned, original 3.5PN E-approximant [76, 77, 114], and 0.45 radians for the untuned non-Keplerian [104] 3.5PN E-approximant (see Fig. 5.16). Although those phase differences are larger than for 3.5PN Taylor T4 (0.04 radians), the phase differences

¹¹It is not clear whether the failure of the adiabatic models is a result of the assumption of adiabaticity, or if the accuracy of those models would continue to improve if terms at order higher than 3.5PN were known.

for the P-approximants are less scattered as a function of PN order than are the phase differences for generic Taylor approximants.

The analyses of the flux, GW frequency derivative, and (derivative of the) center-of-mass energy emphasize again the importance of including nonadiabatic effects during the last stages of inspiral [77]. Roughly, we can say that non-adiabatic effects are no longer negligible starting from a frequency $-M\omega \sim 0.08\text{--}0.12$, as can be seen in Figs. 5.9, 5.13, and 5.15. As seen in these figures, nonadiabatic E-approximants can capture some of the relevant features of the late time evolution. We expect that by further improving these models by fitting higher-order PN coefficients to the numerical data, they will become excellent candidates for developing an analytic template bank of coalescing BBHs [75, 78, 121, 122, 124].

In this paper we started to explore the possibility of reducing the phase differences between numerical and E-approximant waveforms by fitting the unknown parameters a_5 , \mathcal{F}_8 , and ν_{pole} (see Fig. 5.18). As a first step, for several E-approximants we searched for a local minimal phase difference by varying a_5 , \mathcal{F}_8 , and ν_{pole} . We found that we were able to reduce phase differences to below the numerical uncertainty. In a future work which will include merger and ring-down, we plan to determine the *region* of the parameter space $(a_5, \mathcal{F}_8, \nu_{\text{pole}})$ in which the phase difference is within the numerical uncertainty of the simulation.

Acknowledgments

We thank Emanuele Berti, Lee Lindblom, Etienne Racine, Bangalore Sathyaprakash, Saul Teukolsky, and Kip Thorne for informative discussions. We also thank Emanuele Berti and Eric Poisson for providing us the numerical data of the GW

flux in the test-mass limit case. We thank Thibault Damour and Alessandro Nagar for clarifications on the “nontuned” EOB model used in Ref. [122].

A.B. and Y.P. acknowledge support from NSF grant PHY-0603762, and A.B. also acknowledges support from the Alfred P Sloan Foundation. M.B., L.K., A.M., H.P., and M.S. are supported in part by grants from the Sherman Fairchild Foundation to Caltech and Cornell, and from the Brinson Foundation to Caltech; by NSF grants PHY-0601459, PHY-0652995, DMS-0553302 and NASA grant NNG05GG52G at Caltech; by NSF grants PHY-0652952, DMS-0553677, PHY-0652929, and NASA grant NNG05GG51G at Cornell.

5.8 Appendix: Padé approximants to the energy flux in the test particle limit

In the test-mass-limit case the GW energy flux is known through 5.5PN order [241]. The explicit coefficients entering Eq. (5.19) for $i \geq 8$ and $\nu = 0$ can be read from Eqs. (4.1) and (4.2) of Ref. [106].

In Figs. 5.19 and 5.20 we compare the normalized energy flux function [215] F/F_{Newt} to the T- and P-approximants. To easily compare these figures with the other figures in the paper, we plot quantities as functions of the approximate GW frequency defined by $2M\Omega$. As noticed in Ref. [106], the P-approximants approach the numerical data more systematically. The differences between different PN orders are difficult to see in these figures. To obtain a clearer view,




Figure 5.19: **Low-order normalized energy flux F/F_{Newt} versus GW frequency 2Ω in the test-mass limit**

For notation see Table 5.1 and its caption. For comparison, both panels also include the result of the numerical calculation of Poisson [215], labeled with 'NR'.




Figure 5.20: **High-order normalized energy flux F/F_{Newt} versus GW frequency 2Ω in the test-mass limit**

For notation see Table 5.1 and its caption. For comparison, both panels also include the result of the numerical calculation of Poisson [215], labeled with ‘NR’.




Figure 5.21: **Convergence of the PN approximants in the test-mass limit**

Plotted are differences of F/F_{Newt} from the numerical result. The P-approximants do not converge faster than the Taylor series.




Figure 5.22: Cauchy convergence test of F/F_{Newt} in the test-mass limit for the T- and P-approximants

We plot $\Delta F_{n+m} := F_{n+m+1} - F_{n+m}$, and $\Delta F_n^m := F_{n+1}^m - F_n^m$ at three different frequencies. At high frequencies, the 4.5 and 5 PN Padé approximants are contaminated by the extraneous pole of the 5PN Padé series; for low frequencies ($\nu_{\Omega} = 0.1$), the pole is apparently irrelevant.

Table 5.4: Normalized energy flux F/F_{Newt} in the test-mass limit for the T- and P-approximants at different PN orders and at three different frequencies

PN order ($n+m$)/2	$\nu_{\Omega} = 0.1; 2M\Omega = 0.002$	$\nu_{\Omega} = 0.3; 2M\Omega = 0.054$	$\nu_{\Omega} = 0.4; 2M\Omega = 0.128$
	F_{n+m}/F_{Newt}	F_n^m/F_{Newt}	F_{n+m}/F_{Newt}
0.0	1.000000000	1.20948977	1.000
0.5	1.000000000	1.03092783	1.000
1.0	0.9628869047	0.94287089	0.406
1.5	0.9754532753	0.97587569	1.210
2.0	0.9749604292	0.97462770	1.084
2.5	0.9745775009	0.97469475	0.692
3.0	0.9747307757	0.97471937	1.227
3.5	0.9747206248	0.97471854	1.061
4.0	0.9747182352	0.97471874	0.952
4.5	0.9747194262	0.97471859	1.190
5.0	0.9747192776	0.97471930	1.051
5.5	0.9747192763	0.97471928	1.073

We use boldface to indicate the range of significant figures that do not change with increasing PN order.

Fig. 5.21 plots the differences between PN flux and numerical flux at four fixed frequencies. Fig. 5.21 shows this somewhat better behavior of Padé; however, the Padé-approximants show little improvement between PN orders 3.5 and 4.5, and at order 5 there occurs an extraneous pole. At frequency $2M\Omega = 0.04$, P-approximants with order ≥ 2.5 are within 0.5 percent of the numerical data, as are T-approximants with order ≥ 3.5 . Good agreement at low frequency is rather important because that is where the majority of the waveform phasing accumulates.

Table 5.4 and Fig. 5.22 test the internal convergence of T- and P-approximants *without* referring to a numerical result. Table 5.4 displays the flux at all known PN-orders at select frequencies, with boldface highlighting the digits that have already converged. Although the Padé summation does not accelerate the convergence, the P-approximant at 5.5PN order is closest to the numerical data (see Fig. 5.21).

Comparing Table 5.4 with Table 5.2, and Fig. 5.22 with Fig. 5.10 we observe that the P-approximants converge more systematically in the equal-mass case than in the test-mass limit. This is also evident by comparing Fig. 5.21 with Fig. 5.8: We see that P-approximants at 3PN and 3.5PN orders are inside the numerical flux error whereas T-approximants at all orders through 3.5 PN are outside the numerical flux error bars even ~ 25 GW cycles before merger. However, as the Padé approximant does not converge faster, it is not immediately clear whether similar superior behavior of Padé can be expected for more generic binary black holes.

HYBRID WAVEFORMS IN GRAVITATIONAL-WAVE-DETECTOR DATA ANALYSIS¹

We study the detection efficiency of stationary-phase approximated post-Newtonian waveforms currently used by ground-based gravitational-wave detectors to search for the coalescence of binary black holes by comparing them to an accurate waveform obtained from numerical simulation of an equal-mass non-spinning binary black hole inspiral, merger and ringdown. We perform this study for the Initial- and Advanced-LIGO detectors. We find that detection efficiency can be improved by integrating the match filter to higher frequencies than used currently. We propose simple analytic frequency cutoffs for both Initial and Advanced LIGO, which achieve nearly optimal match, and can easily be extended to unequal-mass, spinning systems. We also find that templates that include terms in the phase evolution up to 3.5 pN order are nearly always better, and rarely significantly worse, than 2.0 pN templates

¹This chapter is extracted from a paper to be published in collaboration with Duncan Brown, Larne Pekowsky, Harald Pfeiffer, and Mark Scheel [60]. I extrapolated the numerical waveforms, produced the long post-Newtonian waveform, joined the two, and produced the error estimates. I verified the matches produced by the searches of parameter space, which were run by Larne. I also introduced the weighted-average frequency cutoff, and shared in writing the text.

currently in use. For Initial LIGO we recommend a strategy using templates that include a recently introduced pseudo-4.0 pN term in the low-mass ($M \leq 30 M_{\odot}$) region, and 3.5 pN templates allowing unphysical values of the symmetric reduced mass η above this. This strategy always achieves overlaps within 0.3% of the optimum, for the data used here. For Advanced LIGO we recommend a strategy using 3.5 pN templates up to $M = 12 M_{\odot}$, 2.0 pN templates up to $M = 21 M_{\odot}$, pseudo-4.0 pN templates for masses in the range 21–65 M_{\odot} , and 3.5 pN templates with unphysical η for higher masses. This strategy always achieves overlaps within 0.7% of the optimum for Advanced LIGO.

6.1 Introduction

The coalescence of binary black holes is one of the most promising sources of gravitational waves for interferometric gravitational-wave detectors, such as LIGO, Virgo, and GEO600. The first-generation LIGO detectors have achieved their design sensitivity and recorded over one year of coincident data. These data, together with data from the Virgo detector, are currently being searched for gravitational waves from compact binary coalescence. Upgrades to improve the sensitivity of these detectors by a factor of two, and ultimately ten, are underway. Optimal searches using the enhanced detectors in 2009 will be sensitive to black-hole coalescence out to hundreds of megaparsecs. The advanced detectors, operational next decade, could detect black-hole binaries at distances of over one gigaparsec.

Optimal searches for gravitational waves use matched filtering, which requires

accurate knowledge of the waveform. Previous searches in LIGO data have used post-Newtonian and phenomenological templates to search for the coalescence of black-hole binaries. Over the last several years numerical relativity has made remarkable breakthroughs in simulating the late inspiral, merger, and ringdown of black-hole binaries. The computational cost of these simulations is high, however, making it impractical to use them directly as template waveforms for use in a matched-filter search. It has been shown [75, 24, 206, 78, 157, 59, 152, 156, 61, 197, 163] that there is good agreement between the waveforms generated by numerical relativity with analytic post-Newtonian waveforms to within just a few orbits of merger.

This paper uses the high-accuracy Caltech–Cornell numerical-relativity waveforms to suggest improvements to the analytic waveforms currently used in gravitational-wave searches by LIGO and Virgo. A similar study has been performed by Pan et al. using numerical data from Pretorius and the Goddard groups [206]. Our main results are in agreement with their conclusion that a simple extension of the existing stationary phase approximation to the adiabatic post-Newtonian waveforms (called *TaylorF2* in Ref. [109]) yields high overlaps with numerical waveforms.

In Sec. 6.2, we review the current techniques used for searching for gravitational waves in gravitational-wave detector data. We discuss the construction of the waveform—a pN–NR hybrid—in Sec. 6.3. In Sec. 6.4 we employ the hybrid waveform in a comparison of the detection efficiency of gravitational-wave templates that may be used in upcoming searches of LIGO and Virgo data. Finally, in Sec. 6.5, we discuss improvements that may be made to the current data-analysis

techniques to optimize overlaps.

Throughout this paper, we use only the $(l, m) = (2, 2)$ component of the waveform $\Psi_4^{2,2}$ (as defined, e.g., in [61]). For convenience, we drop the superscript. Whenever possible, we use dimensionless quantities, like $r M |\Psi_4|$, where r is the areal radius of the observation sphere, and M is the total apparent-horizon mass of the holes in the initial data. However, for any calculation involving the LIGO noise curve, we have a physical scale, and thus use standard mks units, where

$$G = 6.67259 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2} , \quad (6.1)$$

$$c = 299792458 \text{ m s}^{-1} , \quad (6.2)$$

$$M_\odot = 1.98892 \times 10^{30} \text{ kg} , \quad (6.3)$$

$$1 \text{ Mpc} = 3.08568025 \times 10^{22} \text{ m} . \quad (6.4)$$

6.2 Searches for gravitational waves from black-hole binaries

6.2.1 Matched filtering

Current searches for gravitational waves from binary black hole coalescence use matched filtering to search for a waveform buried in noise. The matched filter is the optimal filter for detecting a signal in stationary Gaussian noise. Suppose that $n(t)$ is a stationary Gaussian noise process with one-sided power spectral density $S_n(f)$ given by $\langle \tilde{n}(f) \tilde{n}^*(f') \rangle = \frac{1}{2} S_n(|f|) \delta(f - f')$. For long integration times, the data stream $s(t)$ output by the detector will always be dominated by the noise. Thus, we can simply approximate $n \approx s$ to calculate $S_n(f)$.

Using this power spectral density (PSD), we can define the inner product between two real-valued signals—the data stream s and the filter template h —by

$$(s|h) := 2 \Re \int_{-\infty}^{\infty} \frac{\tilde{s}(f) \tilde{h}^*(f)}{S_n(|f|)} df \quad (6.5a)$$

$$= 4 \Re \int_0^{\infty} \frac{\tilde{s}(f) \tilde{h}^*(f)}{S_n(f)} df . \quad (6.5b)$$

For simplicity of presentation, we will assume that the signals are normalized, so that $(s|s) = 1$ and $(h|h) = 1$. Obviously, this condition can always be imposed, if necessary, by taking

$$s \rightarrow \frac{s}{\sqrt{(s|s)}} \quad \text{and} \quad h \rightarrow \frac{h}{\sqrt{(h|h)}} . \quad (6.6)$$

The filter template includes arbitrary time and phase offsets, encoded by the arrival time and phase, t_a and ϕ_a . Under a change of these quantities, the Fourier transform behaves as

$$\tilde{h}(f) \rightarrow \tilde{h}(f) e^{-2\pi i f t_a - i \phi_a} . \quad (6.7)$$

The matched-filter output—the *overlap* between the two waveforms—is then defined as the (normalized) inner product of the signals, maximized over these two variables:

$$\langle s|h \rangle := \max_{t_a, \phi_a} (s|h) \quad (6.8)$$

$$= \max_{t_a, \phi_a} 4 \Re \int_0^{\infty} \frac{\tilde{s}(f) \tilde{h}^*(f)}{S_n(f)} e^{2\pi i f t_a + i \phi_a} df \quad (6.9)$$

$$= 4 \max_{t_a} \left| \int_0^{\infty} \frac{\tilde{s}(f) \tilde{h}^*(f)}{S_n(f)} e^{2\pi i f t_a} df \right| . \quad (6.10)$$

Note that this integral is just the (inverse) Fourier transform of the quantity $\tilde{s}(f) \tilde{h}^*(f)/S_n(f)$ evaluated at t_a . Thus finding the maximum over t_a involves

taking the Fourier transform and selecting the largest element of the finite set that results from discretization.

6.2.2 Post-Newtonian template

Searches for gravitational waves in LIGO and Virgo use a post-Newtonian waveform known as *TaylorF2*. This is a frequency-domain waveform obtained via the stationary-phase approximation [100], which assumes that the frequency-domain amplitude is simply proportional to $f^{-7/6}$ (the lowest-order behavior), while its phasing is given by the phase of the time-domain waveform, as a function of frequency. For a binary consisting of masses m_1 and m_2 , located at an “effective” distance D_{eff} , we have

$$\tilde{h}(f; M, \eta, f_c) = \Theta(f_c - f) \left(\frac{1 \text{ Mpc}}{D_{\text{eff}}} \right) \mathcal{A}_{1 \text{ Mpc}}(M, \eta) f^{-7/6} e^{i\Psi(f; M, \eta)} , \quad (6.11a)$$

where

$$\mathcal{A}_{1 \text{ Mpc}}(M, \eta) := \left(\frac{5\pi}{24} \right)^{1/2} \left(\frac{GM_\odot/c^2}{1 \text{ Mpc}} \right) \left(\frac{\pi GM_\odot}{c^3} \right)^{-1/6} \left(\frac{\eta}{M_\odot} \right)^{1/2} \left(\frac{M}{M_\odot} \right)^{1/3} , \quad (6.11b)$$

$$M := m_1 + m_2 , \quad (6.11c)$$

$$\eta := \frac{m_1 m_2}{(m_1 + m_2)^2} , \quad (6.11d)$$

$$v := \left(\frac{GM}{c^3} \pi f \right)^{1/3} , \quad (6.11e)$$

and the phasing Ψ of the frequency-domain waveform is given to 3.5pN accuracy by the formula

$$\begin{aligned}
\Psi(f; M, \eta) = & 2\pi f t_0 - 2\phi_0 - \pi/4 \\
& + \frac{3}{128\eta} \left[\nu^{-5} + \left(\frac{3715}{756} + \frac{55}{9}\eta \right) \nu^{-3} - 16\pi \nu^{-2} \right. \\
& + \left(\frac{15293365}{508032} + \frac{27145}{504}\eta + \frac{3085}{72}\eta^2 \right) \nu^{-1} \\
& + \pi \left[\frac{38645}{756} - \frac{65}{9}\eta \right] \left[1 + 3\ln\left(\frac{\nu}{\nu_0}\right) \right] \\
& + \left[\frac{11583231236531}{4694215680} - \frac{640}{3}\pi^2 - \frac{6848}{21}\gamma \right] \nu \\
& + \left[\left(-\frac{15335597827}{3048192} + \frac{2255}{12}\pi^2 - \frac{47324.0}{63.0} - \frac{7948}{9} \right) \eta \right. \\
& \quad \left. + \frac{76055}{1728}\eta^2 - \frac{127825}{1296}\eta^3 \right] \nu \\
& \left. + \pi \left[\frac{77096675}{254016} + \frac{378515}{1512}\eta - \frac{74045}{756}\eta^2 \right] \nu^2 \right]. \quad (6.11f)
\end{aligned}$$

The overall frequency scale is set by the total mass M , as can be seen by observing that each occurrence of f is accompanied by a factor of M .² Thus, going to a higher-mass system shifts the waveform to lower frequencies. On the other hand, to first order, the timescale for the rate of change of the frequency is given by the *chirp mass*:

$$\mathcal{M} := \left(\frac{m_1^3 m_2^3}{m_1 + m_2} \right)^{1/5} = M\eta^{3/5}. \quad (6.12)$$

Clearly, the total mass and chirp mass give us two very different handles on the behavior of the waveform. These two handles will be important when we try to match the template to our waveform in regions where the post-Newtonian and

²The term $2\pi f t_0$ might be rewritten as $2\pi M f \times t_0/M$. This term accounts for a time offset altering the phase of the Fourier transform.

stationary-phase approximations are poor. This is typically the case for high-mass systems, which only enter the detector band late in the inspiral. In this case, we can still obtain a high match, at the cost of using templates with the wrong values of M and η .

We also note that physical binary systems are restricted to $0 < \eta \leq 1/4$. However, for higher values of η , the formulas shown above still give plausible waveforms; in fact, in some cases these templates match the true waveform better than any physical template. We will explore the implications of allowing unphysical values for η in searches over the templates in Sec. 6.4.2.

Note the Heaviside function in Eq. (6.11a). This contains a cutoff frequency f_c which is used to ensure that the template does not extend to frequencies much greater than the frequencies contained in the expected signal. This is essentially a third parameter for the template waveform, and will be searched over. See Sec. 6.4.1 for a discussion of strategies for optimizing detection by changing this cutoff.

It has previously been shown that second-order post-Newtonian stationary-phase waveforms will provide acceptable detection templates for binary neutron stars and sub-solar mass black holes [129], but not necessarily for higher-mass black holes. This is an issue we will investigate below by testing 2pN and 3.5pN templates.

We also use the *TaylorT4* waveform [59] to create a hybrid waveform used for evaluation of the templates. This hybrid is described in Sec. 6.3.3.

6.2.3 Discretization

The direct and inverse Fourier transforms are defined (using the standard LIGO conventions [8]) as

$$\tilde{s}(f) := \int_{-\infty}^{\infty} s(t) e^{-2\pi i f t} dt , \quad (6.13)$$

$$s(t) = \int_{-\infty}^{\infty} \tilde{s}(f) e^{2\pi i f t} df . \quad (6.14)$$

In transferring these and the continuum expressions of preceding sections to computer, we need to introduce two changes.

First, the ranges of integration must be restricted to finite intervals. We need to assume that the physical signal contains nothing of interest at frequencies higher than f_{Ny} or, considering negative frequencies, lower than $-f_{\text{Ny}}$. For notational simplicity, we define the discretized Fourier transform to be periodic, with period $2f_{\text{Ny}}$. Similarly, we will assume that the signal s is periodic, with period T . Thus, we can restrict each of the integrals given above to one period of the relevant quantity.

Second, the quantities must be given on a discrete grid. We will assume that the signal s is sampled at N uniform intervals of $\Delta t = T/N$. This will give rise to a frequency discretization of $\Delta f = 1/T = 1/N\Delta t$. We define the quantities

$$t_j = j\Delta t \quad \text{and} \quad f_k = k\Delta f , \quad (6.15)$$

for *all* integers j and k . It is not hard to see that the highest frequency that can be represented on this discrete set is bounded by the Nyquist frequency $f_{\text{Ny}} = 1/2\Delta t$.

The combined operation of discretizing and restricting to finite range will be

denoted by \rightsquigarrow , so

$$\tilde{s}(f_k) \rightsquigarrow \sum_{t_j > -N\Delta t/2}^{N\Delta t/2} s(t_j) e^{-2\pi i f_k t_j} \Delta t \quad (6.16)$$

$$\rightsquigarrow \Delta t \sum_{j=0}^{N-1} s(t_j) e^{-2\pi i j k / N} , \quad (6.17)$$

$$s(t_k) \rightsquigarrow \sum_{f_j > -f_{Ny}}^{f_{Ny}} \tilde{s}(f_j) e^{2\pi i f_k t_j} \Delta f \quad (6.18)$$

$$\rightsquigarrow \Delta f \sum_{j=0}^{N-1} \tilde{s}(f_k) e^{2\pi i j k / N} . \quad (6.19)$$

Note that in the second step of each of these expressions, we have used the periodic character of s to re-express negative times as positive, and the periodic character of \tilde{s} to re-express negative frequencies as positive. We have also used the relation $f_k t_j = j k / N$.

A notational subtlety arises when using the frequency-domain quantity. The symbol \tilde{s}_k is defined as the sum in Eq. (6.17), without the factor of Δt [65]. In particular, we have $\tilde{s}(f_k) \rightsquigarrow \Delta t \tilde{s}_k$. The expressions given above for $\tilde{s}(f_k)$ and $s(t_j)$ should not depend strongly on the fineness of the discretization (for sufficiently fine discretizations). Clearly, then, \tilde{s}_k will depend strongly on the discretization. Though the factor of Δt should drop out for calculations using normalized quantities, for calculations of the signal-to-noise ratio, or demonstrations of $\tilde{s}(f_k)$, it is an important distinction that needs to be kept in mind. For example, the FFTW and MATLAB software packages use \tilde{s}_k as their standard frequency-domain quantity. Throughout the remainder of this paper, we will use $\tilde{s}(f_k)$ exclusively.

For completeness, we include the expressions

$$\langle s|h \rangle \rightsquigarrow 2 \Re \sum_{k=0}^{N-1} \frac{\tilde{s}(f_k) \tilde{h}^*(f_k)}{S_n(|f_k|)} \Delta f \quad (6.20)$$

$$\rightsquigarrow 4 \Re \sum_{k=0}^{\lfloor N/2 \rfloor} \frac{\tilde{s}(f_k) \tilde{h}^*(f_k)}{S_n(f_k)} \Delta f , \quad (6.21)$$

$$\langle s|h \rangle \rightsquigarrow 2 \max_{t_a} \left| \sum_{k=0}^{N-1} \frac{\tilde{s}(f_k) \tilde{h}^*(f_k)}{S_n(|f_k|)} e^{2\pi i f_k t_a} \Delta f \right| \quad (6.22)$$

$$\rightsquigarrow 4 \max_{t_a} \left| \sum_{k=0}^{\lfloor N/2 \rfloor} \frac{\tilde{s}(f_k) \tilde{h}^*(f_k)}{S_n(f_k)} e^{2\pi i f_k t_a} \Delta f \right| . \quad (6.23)$$

The notation $\lfloor N/2 \rfloor$ denotes the greatest integer less than or equal to $N/2$.

6.3 PN–NR hybrid waveform

We need to construct a “true” black-hole binary waveform, which we might expect to observe with detectors. A numerical simulation will provide the data for the crucial nonlinear merger phase. We carefully extract the data and extrapolate it to large radius, and investigate the effects of numerical error on the final result. Because this waveform is very computationally expensive to produce, it covers only about 32 cycles, which is not sufficient for a thorough investigation of the possibility of detecting it in searches of data from gravitational-wave detectors. Thus, we match the numerical waveform to a post-Newtonian waveform, producing a hybrid which extends for many thousands of cycles, covering the entire band of interest.

6.3.1 Numerical simulation, extraction, and extrapolation

The numerical simulation is the same as that described in Refs. [59, 229]: an equal-mass, nonspinning, black-hole binary with reduced eccentricity [209], beginning roughly 16 orbits before merger, continuing through merger and ringdown [229]. It is performed with the Caltech–Cornell pseudospectral code, using boundary conditions designed to prevent constraint violations and gravitational radiation from entering the domain [164, 187].

Data is extracted from the simulation in the form of the Newman–Penrose scalar

$$\Psi_4 = -C_{\alpha\beta\gamma\delta} l^\alpha \bar{m}^\beta l^\gamma \bar{m}^\delta, \quad (6.24)$$

where l^α and the complex vector \bar{m}^β are constructed with reference to the coordinate basis. Along the positive z axis, we have

$$l^\alpha = \frac{1}{\sqrt{2}} \left(\frac{\partial}{\partial t} - \frac{\partial}{\partial z} \right)^\alpha, \quad (6.25)$$

$$\bar{m}^\beta = \frac{1}{\sqrt{2}} \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right)^\beta. \quad (6.26)$$

This quantity is extracted as a function of time, at various radii along the positive z axis. This is then extrapolated to large radii, as described in Ref. [59], and in greater detail in Ref. [63].

The measured (instantaneous) frequency at the beginning of the simulation is

$$f_{\text{initial}} = (1.08 \pm 0.01) \times 10^3 \text{ Hz} \frac{M_\odot}{M}. \quad (6.27)$$

The measured ringdown frequency is

$$f_{\text{ringdown}} = (1.78 \pm 0.02) \times 10^4 \text{ Hz} \frac{M_\odot}{M}. \quad (6.28)$$

The measured *Christodoulou* mass and spin of the final black hole are

$$M_{\chi, \text{ final}} = (0.95162 \pm 0.00002) M_{\chi, \text{ initial}} , \quad (6.29)$$

$$S_{\text{ final}} = (0.68646 \pm 0.00004) M_{\chi, \text{ final}}^2 . \quad (6.30)$$

Using this value for the spin, a quasi-analytic formula due to Echeverria [130] predicts a value of $1.77 \times 10^4 \text{ Hz } \frac{M_{\odot}}{M}$, for the ringdown frequency, in close agreement with the measured frequency.

6.3.2 Accuracy of the numerical simulation

The numerical waveform will be the standard against which we will judge the SPA waveforms used in LIGO data analysis. To understand how precisely we should trust our final results, we need to understand the accuracy of the waveform itself. The most obvious measure of the error in this fiducial waveform is its convergence with increasing numerical resolution. Fig. 6.1 shows the normalized inner product (Eq. (6.5)) and overlap (Eq. (6.8)) between waveforms computed at different resolutions.

Because of the short extent of the numerical waveforms, we need to be careful when using their Fourier transforms. The signal can be corrupted easily by the nonperiodicity of the waveforms, and the discontinuous jumps that result. We can mitigate this problem by increasing the sampling frequency of the input data, and restricting the Fourier transform to frequencies corresponding to instantaneous frequencies contained in the data. The input data can easily be upsampled in the time domain by interpolating the phase and amplitude of the complex data to a finer time grid. We then perform the transform, and explicitly set the data to zero

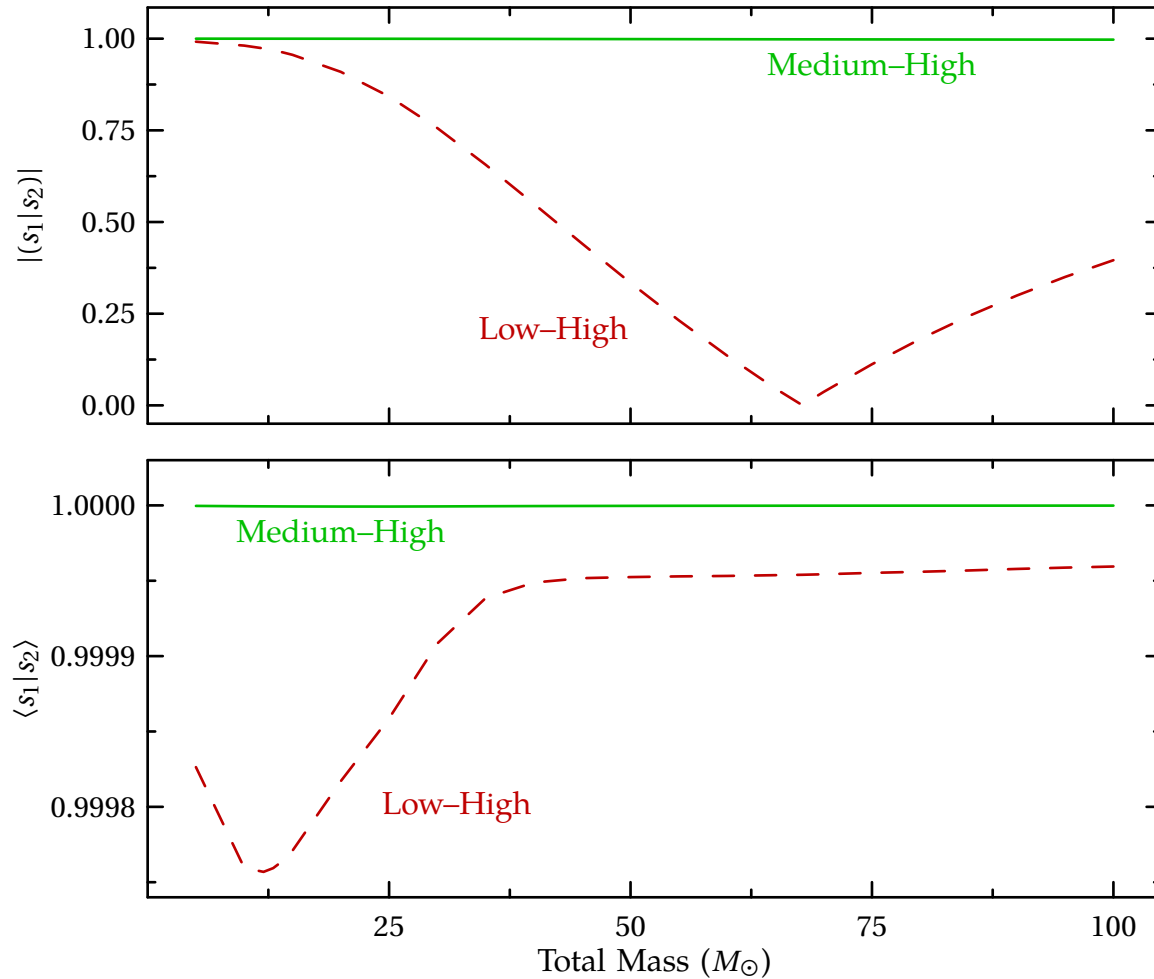


Figure 6.1: **Convergence testing for numerical waveforms from a data-analysis perspective**

We plot the match between waveforms computed at different numerical resolutions. The waveforms are scaled to various masses, and the Initial-LIGO noise curve is used in the calculation of the match. The upper panel shows the simple inner product (that is, there is no maximization over arrival time and phase); the lower panel shows the overlap (after maximization). In each panel, the lower (dashed red) line compares the lowest- and highest-resolution simulations, while the upper (solid green) line compares the medium- and highest-resolution simulations. Note that this plot uses only numerical data, with no post-Newtonian contribution.

at frequencies below f_{initial} and above f_{ringdown} , as given in Eqs. 6.27 and 6.28. While the results do depend on whether or not we impose these cutoffs, they do not depend sensitively on the actual cutoff frequencies.

The inner product between the lowest- and highest-resolution simulations (dashed red lines) actually passes through zero, as shown in the upper panel. Presumably, this is because of loss of phase accuracy over the course of the simulation. All three simulations begin with the same initial data, so the waveforms are most similar at the beginning. Masses for which this is the most important segment (the lowest masses) will naturally have the highest inner product between resolutions. As the simulation progresses, numerical error accumulates— notably in the phase—so the inner product decreases with masses for which later segments dominate the inner product (higher masses). When the inner product is optimized over arrival time and phase, we can see that the overlap becomes much better, as shown in the lower panel, indicating sufficient accuracy within any frequency band for which phase coherence is required.

In either case, the medium- and highest-resolutions are much more nearly the same. Without optimization, their inner product is within a few tenths of a percent of 1; after optimization, the overlap is within 10^{-6} of 1.

In the rest of our analysis, we use the highest-resolution waveform. Because we always optimize over arrival time and phase, the lower panel of Fig. 6.1 is the most relevant, and shows that the waveform has converged to very high accuracy. The overlaps we quote below will only be given to three decimal places at most, because this is roughly the accuracy of the single-precision numerical methods used in the rest of the paper. This accuracy is also sufficient for searches of

gravitational-wave data. Thus, the truncation error of the simulated waveform is irrelevant for those purposes.

Other sources of error include residual eccentricity and spin, the influence of the outer boundary of the simulation, extrapolation errors, and coordinate effects, as discussed in Ref. [59]. The eccentricity had a disproportionately large effect on the error quoted in that paper because of the matching technique, which is not used here. Restricting attention to the other effects of eccentricity, the uncertainty falls below that due to numerical error. Similarly, using the techniques of Ref. [190], the initial spins of the black holes have been measured more reliably, and found to be more than an order of magnitude smaller than previously determined, allowing us to reduce the estimate for that error to less than the numerical truncation error. The various coordinate effects were all estimated to be of roughly the same magnitude as the numerical error.

With the numerical error being many times more accurate than needed for this analysis, and the other sources of uncertainty being of roughly the same size, these considerations indicate that the overall error in our fiducial waveform is substantially less than the precision needed for this analysis.

6.3.3 Hybrid waveform

Numerical simulations cannot simulate a very large portion of the inspiral of a black-hole binary system. Indeed, the longest such simulation currently in the literature is the one used here—which extends over just 32 GW cycles before merger. Fortunately, this is the only stage in which simulations are needed. It has been shown previously [59] that the *TaylorT4* waveform with 3.5-pN phase and

3.0-pN amplitude matches the early part of this simulation to very high accuracy. We generate a *TaylorT4* waveform of over 8000 GW cycles ($t \sim 1.2 \times 10^6 M$, starting at $Mf = 0.004$), and transition between the two to create a hybrid. This long waveform is sufficient to ensure that—even for the lowest-mass systems we will consider—the waveform begins well before it enters the frequency band of interest to LIGO.

Following Ref. [61], we match the numerical waveform to the pN waveform by adjusting the time and phase offsets of the pN waveform to minimize the quantity

$$\Xi(\Delta t, \Delta\phi) = \int_{t_1}^{t_2} [\phi_{\text{NR}}(t) - \phi_{\text{PN}}(t - \Delta t) - \Delta\phi]^2 dt . \quad (6.31)$$

Here, we choose $t_1 = 900 M$ and $t_2 = 1730 M$, which is closer to the beginning of the waveform than in the previous paper. This particular interval is chosen to begin and end at troughs of the small oscillations due to the residual eccentricity $e \sim 5 \times 10^{-5}$ in our numerical waveform. Taking a range from trough to trough or peak to peak—rather than node to node, for example—of the eccentricity effects minimizes their influence on the matching. The eccentricity oscillations can be seen more easily after low-pass filtering the waveform, though we find filtering to be unnecessary for this paper. The junk radiation apparent in the waveform as shown here has no effect on the resulting match—as we have verified by filtering, and redoing the match. Because the final waveform will incorporate no numerical data before t_1 and very little immediately thereafter (as explained below), the junk radiation will have no effect on any of our results—as we have also explicitly verified. In particular, by integrating Ψ_4 to obtain h , we effectively smooth the junk radiation.




Figure 6.2: **Amplitude and phase differences between the numerical and post-Newtonian waveforms blended to create the hybrid waveform**

The (red) vertical lines at $900M$ and $1730M$ denote the region over which matching and hybridization occur. Note that the agreement is well within the numerical accuracy of the simulation, represented by the horizontal bands, throughout the matching region. Also note that the phase difference is fairly flat for a significant period of time after the matching range, which indicates that the match is not sensitive to the particular range chosen for matching.

In Fig. 6.2 we compare the phase of the numerical and pN waveforms. The quantities plotted are

$$\delta\phi := \phi_{\text{PN}} - \phi_{\text{NR}} , \quad (6.32)$$

$$\frac{\delta A}{A} := \frac{A_{\text{PN}} - A_{\text{NR}}}{A_{\text{NR}}} , \quad (6.33)$$

shown over the interval on which both data sets exist. The vertical bars denote the matching region. Note that the phase difference is well within the accuracy of the simulation (about 0.01 radians, represented by the horizontal band) over a range extending later than the matching region. Also, the difference between the two is fairly flat, which implies that the match is not very sensitive to the region chosen for matching. Because of this, we expect that the phase coherence between the early pN data and the late NR data will be physically accurate to high precision.

The hybrid waveform is then constructed by blending the two matched waveforms together according to

$$A_{\text{hyb}}(t) = \tau(t) A_{\text{NR}} + [1 - \tau(t)] A_{\text{PN}}(t) , \quad (6.34)$$

$$\phi_{\text{hyb}}(t) = \tau(t) \phi_{\text{NR}} + [1 - \tau(t)] \phi_{\text{PN}}(t) . \quad (6.35)$$

The blending function τ is defined by

$$\tau(t) = \begin{cases} 0 & \text{if } t < t_1 , \\ \frac{t-t_1}{t_2-t_1} & \text{if } t_1 \leq t < t_2 , \\ 1 & \text{if } t_2 \leq t . \end{cases} \quad (6.36)$$

The values of t_1 and t_2 are the same as those used for the matching. The amplitude discrepancy between the pN waveform and the NR waveform over this




Figure 6.3: **The last $t = 5000M_\odot$ of the hybrid waveform used in this analysis**

This is the h_+ waveform seen by an observer on the positive z axis. The vertical lines denote the matching and hybridization region. The 0 on the time axis corresponds to the beginning of the numerical simulation.

interval is within numerical uncertainty—roughly 0.4%. As with the matching technique (Eq. (6.31)), this method is similar to that of Ref. [5], but distinct, in that we blend the phase and amplitude, rather than the real and imaginary parts. This leads to a smoothly blended waveform, shown in Fig 6.3.

Up to this point, the waveform has been Ψ_4 data. With the long waveform

in hand, we numerically integrate twice to obtain h , and set the four integration constants so that the final waveform has zero average and first moment [209]. Because of the very long duration of the waveforms, this gives a reasonable result, which is only incorrect at very low frequencies—lower than any frequency of interest to us. We have also checked that our results do not change when we effectively integrate in the frequency domain by taking

$$\tilde{h} = \frac{\tilde{\Psi}_4}{4\pi f^2}, \quad (6.37)$$

which is the frequency-domain analog of the equation $\Psi_4 = -\ddot{h}$.

6.4 Detection efficiency of gravitational-wave templates

In this section we study the efficiency of restricted, stationary phase *TaylorF2* post-Newtonian templates with terms up to order 2.0, order 3.5, and a “pseudo-4.0 pN order” term recommended in Ref. [206]. Overlaps are calculated using the techniques of Sec. 6.2.1, with the signal s being the hybrid waveform described in Sec. 6.3 scaled to a range of masses. We consider both the Initial- and Advanced-LIGO noise curves.

Plots of the hybrid waveforms in comparison to the Initial-LIGO noise curve are shown in Fig. 6.4. The masses are chosen so that various frequencies of interest (the final stitching frequency, the ISCO, and the ringdown) occur at the seismic wall for Initial LIGO, 40 Hz. The waveforms \tilde{s} are scaled to depict the




Figure 6.4: **Hybrid Caltech–Cornell waveform scaled to various total masses shown against the Initial- and Advanced-LIGO noise curves**

Sources are optimally oriented and placed at 100 Mpc. Markers are placed along the lines at frequencies corresponding to various instantaneous frequencies of the waveforms. The triangles represent the beginning and end of the blending region; the circle represents the ISCO frequency; the square the light-ring; and the diamond the measured ringdown frequency. See the text for discussion of the normalization. The values given for ρ use the Initial-LIGO noise curve, with sources at a distance of 100 Mpc.

detectability of the signal, typically quantified by the signal-to-noise ratio (SNR)

$$\rho^2 := \int_0^\infty \frac{4 \tilde{s}(f) \tilde{s}^*(f)}{S_n(f)} df \quad (6.38)$$

$$= \int_0^\infty \frac{|2 \tilde{s}(f) \sqrt{f}|^2}{S_n(f)} d \ln f . \quad (6.39)$$

Note that the SNR is just the un-normalized inner product of s with itself. In the final expression, the numerator and denominator have the same units, and are directly comparable. Because the square root of the denominator is familiar, we plot that along with the square root of the numerator. Plotting these two quantities together gives a graphical impression of the detectability of the waveform, and the relative importance of each part of the waveform, by its height above the noise curve. In Ref. [64], Brady and Creighton define a slightly different quantity, the characteristic strain:

$$h_{\text{char}} := f |\tilde{s}(f)| . \quad (6.40)$$

The relative factor of \sqrt{f} they use is present so that they can plot h_{char} against $\sqrt{f S_n(f)}$. Cutler and Thorne [102] define still another quantity, the signal strength $\tilde{h}_s(f)$, which is related to the Fourier transform by

$$\tilde{h}(f) = \sqrt{5} \frac{T}{N} \tilde{h}(s) . \quad (6.41)$$

The factor of $\sqrt{5}$ comes from averaging over the sky which we do not do, and T/N is the ratio of the threshold to the rms noise at the endpoint of signal processing.

For each template family we initially optimize over signal mass M , symmetric mass ratio $\eta = m_1 m_2 / (m_1 + m_2)^2$, and upper cutoff frequency f_c . The optimization is performed using a Nelder–Mead (“amoeba”) algorithm [218]. The amoeba starts with a simplex in the parameter space, and proceeds through a series of




Figure 6.5: **Histogram of overlaps found by 300 instances of the Amoeba algorithm**

The amoeba optimizes the overlap against a given waveform over M, η, f_c , with randomized initial conditions. Note the logarithmic scale on the vertical axis. The majority of instances produced a lower overlap than the optimum. We interpret this as pointing to the existence of a broad local maximum which did not coincide with the global maximum.

steps, each of which will improve the value of the function at at least one vertex. The algorithm terminates when all vertices have converged to the same point to within a specified tolerance. This process is deterministic, and amounts to an enhanced steepest-ascent algorithm. It is therefore only guaranteed to find a local maximum. We supplement the basic amoeba by running 300 instances with random starting values, and taking the best match obtained over all instances.

In all cases the point representing the optimal set of parameters is not the one that is found by the most amoebas. We interpret this as being due to a large region in parameter space containing a local maximum and a relatively smaller region containing the global maximum. However, in repeated runs the same optimal parameters were found by at least some of the amoebas. A sample histogram displaying this behavior is shown in Fig. 6.5. The horizontal axis shows the overlap found, while the vertical axis shows the number of instances which found that overlap.

The results of optimizing over all parameters for selected masses for Initial LIGO are given in Table 6.1 and summarized in Fig. 6.6. The same plot for Advanced LIGO is shown in Fig. 6.7. We see that 3.5 pN TaylorF2 outperforms 2.0 pN templates in most cases. For masses between 10 and $50 M_{\odot}$, with the Advanced-LIGO noise curve, 2.0 pN does actually achieve slightly higher overlaps. However, in most of that region, pseudo-4.0 pN templates outperform both.

We see from Tables 6.1 and 6.2 that the parameters of the optimal templates are often far from the parameters of the physical waveform, especially for high-mass systems, which emphasize portions of the waveform for which the pN and SPA assumptions are poor. For Initial LIGO, pseudo-4.0 pN templates generally

Table 6.1: **Maximum overlaps between Caltech–Cornell hybrid waveforms and restricted SPA pN templates using the Initial-LIGO noise curve**

	$(10 + 10)M_{\odot}$	$(20 + 20)M_{\odot}$	$(30 + 30)M_{\odot}$	$(50 + 50)M_{\odot}$
$\langle s^{\text{NR-CC}} h^{\text{SPA}_c^{\text{ext}(2.0)}} \rangle$	0.99	0.98	0.97	0.96
M/M_{\odot}	$23.27^{+0.13}_{-0.12}$	$25.99^{+0.61}_{-0.56}$	$35.2^{+1.8}_{-1.9}$	$47.5^{+6.9}_{-4.7}$
η	$0.199^{+0.003}_{-0.003}$	$0.771^{+0.049}_{-0.042}$	$1.000_{-0.139}$	$1.000_{-0.249}$
f_{cut} (Hz)	500^{+520}_{-150}	430^{+360}_{-80}	296^{+53}_{-31}	191^{+20}_{-14}
$\langle s^{\text{NR-CC}} h^{\text{SPA}_c^{\text{ext}(3.5)}} \rangle$	0.98	0.99	0.99	0.99
M/M_{\odot}	$18.75^{+0.1}_{-0.1}$	$31.88^{+0.77}_{-0.71}$	$47.2^{+4.4}_{-3.3}$	260_{-190}
η	$0.290^{+0.004}_{-0.004}$	$0.493^{+0.053}_{-0.041}$	$0.76^{+0.24}_{-0.23}$	$0.95^{+0.05}_{-0.21}$
f_{cut} (Hz)	507^{+520}_{-160}	450^{+580}_{-80}	330^{+150}_{-40}	197^{+24}_{-16}
$\langle s^{\text{NR-CC}} h^{\text{SPA}_c^{\text{ext}(4)}} \rangle$	0.99	0.96	0.95	0.96
M/M_{\odot}	$23.64^{+0.13}_{-0.12}$	$47.9^{+1.3}_{-1.1}$	$61.8^{+8.7}_{-6.2}$	90^{+20}_{-17}
η	$0.182^{+0.003}_{-0.003}$	$0.181^{+0.016}_{-0.014}$	$0.52^{+0.43}_{-0.18}$	$0.53^{+0.47}_{-0.31}$
f_{cut} (Hz)	510^{+650}_{-150}	352^{+73}_{-61}	310^{+72}_{-47}	196^{+21}_{-15}

The first number in each block is the overlap; subsequent numbers are the template parameters that achieve this overlap. Parameter values within the specified ranges keep the overlap within 1% of the maximum by varying that parameter, while leaving others fixed. We restrict the search to $0 \leq \eta \leq 1.000$, so the upper error bounds when $\eta \sim 1.000$ may be artificially small.

achieve higher accuracy in the optimal parameters than the other two, despite the lower overlaps. Conversely, for Advanced LIGO, 3.5 pN templates achieve higher accuracy.

Table 6.2: Maximum overlaps between Caltech–Cornell hybrid waveforms and restricted SPA pN templates using the Advanced-LIGO noise curve

	$(10 + 10)M_{\odot}$	$(20 + 20)M_{\odot}$	$(30 + 30)M_{\odot}$	$(50 + 50)M_{\odot}$
$\langle s^{\text{NR-CC}} h^{\text{SPA}_c^{\text{ext}(2.0)}} \rangle$	0.98	0.92	0.91	0.94
M/M_{\odot}	$25.15^{+0.02}_{-0.02}$	$47.73^{+0.12}_{-0.11}$	$54.39^{+0.51}_{-0.43}$	$60.2^{+1.6}_{-1.3}$
η	$0.170^{+0.001}_{-0.001}$	$0.188^{+0.001}_{-0.001}$	$0.335^{+0.008}_{-0.007}$	$0.891^{+0.066}_{-0.049}$
f_{cut} (Hz)	440^{+130}_{-120}	267^{+48}_{-50}	262^{+34}_{-36}	182^{+24}_{-18}
$\langle s^{\text{NR-CC}} h^{\text{SPA}_c^{\text{ext}(3.5)}} \rangle$	0.97	0.92	0.92	0.96
M/M_{\odot}	$20.27^{+0.02}_{-0.02}$	$38.11^{+0.11}_{-0.09}$	$50.09^{+0.49}_{-0.42}$	$78.1^{+1.9}_{-1.5}$
η	$0.245^{+0.001}_{-0.001}$	$0.277^{+0.002}_{-0.002}$	$0.386^{+0.013}_{-0.010}$	$0.494^{+0.076}_{-0.033}$
f_{cut} (Hz)	356^{+97}_{-88}	263^{+47}_{-48}	281^{+41}_{-37}	186^{+30}_{-19}
$\langle s^{\text{NR-CC}} h^{\text{SPA}_c^{\text{ext}(4)}} \rangle$	0.97	0.96	0.94	0.90
M/M_{\odot}	$22.24^{+0.02}_{-0.02}$	$46.57^{+0.11}_{-0.11}$	$72.06^{+0.35}_{-0.35}$	$118.5^{+2.0}_{-1.6}$
η	$0.208^{+0.001}_{-0.001}$	$0.190^{+0.001}_{-0.001}$	$0.177^{+0.002}_{-0.003}$	$0.186^{+0.010}_{-0.007}$
f_{cut} (Hz)	470^{+550}_{-140}	353^{+73}_{-69}	242^{+37}_{-36}	152^{+19}_{-19}

The first number in each block is the overlap; subsequent numbers are the template parameters that achieve this overlap. Parameter values within the specified ranges keep the overlap within 1% of the maximum by varying that parameter, while leaving others fixed. We restrict the search to $0 \leq \eta \leq 1.000$, so the upper error bounds when $\eta \sim 1.000$ may be artificially small.




Figure 6.6: **Overlaps between Caltech–Cornell hybrid waveforms and restricted stationary-phase pN waveforms for the Initial-LIGO PSD**

We scale the numerical waveforms to various masses and optimize the pN waveform over M and η , allowing η to range over $0 < \eta \leq 1$. The cutoff frequency f_c is prescribed by the weighted average described below.

This leads us to consider the range of possible template parameters which may give high overlaps. We next consider the reduction in overlap as the parameters f_c and η are independently varied from the optimal value.




Figure 6.7: **Overlaps between Caltech–Cornell hybrid waveforms and restricted stationary-phase pN waveforms for the Advanced-LIGO PSD**

We scale the numerical waveforms to various masses and optimize the pN waveform over M and η , allowing η to range over $0 < \eta \leq 1$. The cutoff frequency f_c is prescribed by the weighted average described below.

6.4.1 Effect of upper frequency cutoff

The optimal cutoff frequency for a template used in a search depends crucially on the total mass of the binary. This dependence comes from the need to cut off the integral in the overlap, Eq. (6.5). We can get a feeling for the relative contributions to the overlap from different parts of a waveform by considering the overlap of a template with itself. This quantity is shown in Fig. 6.8 for the Initial-LIGO noise curve. The template used here is the SPA waveform for an equal-mass $10 M_{\odot}$ binary, but only the overall amplitude changes as we change the masses of the binary, not the shape or position of the curve.

If we ignore the relative phasing of a waveform and a well-matched template, the integrand of the overlap will look approximately like this. The only difference is that a physical binary waveform would end at some frequency—roughly the ringdown frequency of the final black hole. At higher frequencies, the curve would drop steeply. Thus, there is no improvement to the overlap to be gained from extending the template to higher frequencies.

On the contrary, there may be something to be lost by pushing the template to higher frequencies. Note the normalization condition of Eq. (6.6). If the template extends to higher frequencies than the physical waveform can match, the normalization factor will simply tend to decrease the overall match.

For systems with ringdown frequencies well above the peak of the integrand in Fig. 6.8, this is not important; there is so little contribution to the integral at a high frequencies that the change in normalization will be small. For example, binaries of total mass roughly $40 M_{\odot}$ have ringdown frequencies at roughly 450 Hz. We can see from the plot that very little of the overlap is found at higher frequencies.




Figure 6.8: **Integrand of the inner product for a *TaylorF2* 3.5 pN waveform**

This equal-mass waveform is scaled to $M = 10$ and placed at a distance of 100 Mpc. Note, however, that the shape of this curve does not change as we change M and η ; only the vertical scale changes.

Thus, we expect that systems with lower masses should not suffer great loss in overlap if the cutoff frequency is extended higher. This is indeed what we find, as shown by a representative example in Fig. 6.9. For this $40 M_{\odot}$ system, using the Initial-LIGO noise curve, the optimal cutoff frequency is around 450 Hz—roughly the ringdown frequency. Decreasing the cutoff quickly decreases the overlap. The cutoff may be increased almost indefinitely, however, with only 0.5% loss in overlap. This, of course, changes when using the Advanced-LIGO noise curve. We revisit this issue in Sec. 6.5.

6.4.2 Unrestricted η

The physical symmetric mass ratio is restricted to the range $0 < \eta \leq 0.25$, values above this imply complex-valued masses. However the pN waveforms are well-behaved for $0 < \eta < 1.0$, and as seen from Tables 6.1 and 6.2, the highest overlaps are often obtained at unphysical values of η . In Fig. 6.10 we show the effect of limiting the optimization to physical η . At high masses, the limitation reduces the optimal overlap by up to 12%. This is a result of the fact that the integral of the overlap, Eq. (6.5) is effectively cut off by merger for high-mass systems. That is, the part of the physical waveform where we expect the pN and SPA approximations to be valid is buried in detector noise. Thus, if the SPA waveform is to match the physical waveform, the match will be effectively spurious. Nonetheless, this spurious match can be quite high, as seen in the figure.




Figure 6.9: **Overlap between Caltech–Cornell waveform and restricted *TaylorF2*, 3.5 pN waveform as a function of cutoff frequency f_c**

The numerical waveform is scaled to $M = 40 M_{\odot}$, and the overlap is calculated using the Initial-LIGO noise curve. The vertical bars delineate 1% loss. Note that the upper bound extends to higher frequencies indefinitely. That is, the cutoff frequency may, in this case, be extended to high frequency with little cost.




Figure 6.10: **Maximum overlaps obtained by allowing η to range over unphysical values, compared to those obtained by restricting the range of η**

These overlaps are generated using 3.5 pN TaylorF2 templates, searching over values of the total mass and mass ratio. We see that extending to unphysical values of η improves the match by up to 11% for masses near $60 M_{\odot}$.

6.5 Recommendations for improvements

Based on the analysis of the previous sections we propose a series of adjustments to searches using SPA template waveforms to enhance the efficiency of those searches.

First, as seen in Fig. 6.6 for Initial LIGO, adding terms up to 3.5 pN order does as well as or better than the current 2.0 pN templates over most of the mass range, and does significantly better ($> 3\%$) above $50 M_{\odot}$. While the pseudo-4.0 pN templates recommended in Ref. [206] do slightly better at masses near $20 M_{\odot}$, the difference is less than about 1%. Thus, we recommend 3.5 pN templates for all masses, when using an Initial-LIGO noise curve, resulting in overlaps of more than 0.98 for all masses.

The improvement due to 3.5 pN templates over 2.0 pN generally holds for Advanced LIGO as well, as shown in Fig. 6.7. The 3.5 pN templates do better than 2.0 pN templates above $50 M_{\odot}$ without a significant loss (within 1%) at lower masses. However, there is a large region for which the pseudo-4.0 pN term does significantly better. When using an Advanced-LIGO noise curve, we recommend 3.5 pN templates generally, being replaced by pseudo-4.0 pN templates for masses in the range $20\text{--}60 M_{\odot}$.

As a second improvement, we note from Fig. 6.10 that allowing η to range over unphysical values significantly improves matches above $30 M_{\odot}$. In preliminary studies we have found that extending to $\eta \leq 1$ roughly doubles the size of the template bank, and the advantages must therefore be weighed against the increase in false alarm rate.

Our third recommendation involves the cutoff frequency used for the tem-

plate waveform. Optimization over the cutoff frequency is too computationally intensive to be done in searches. Currently, the cutoff frequency is typically taken to be the Schwarzschild ISCO frequency. To examine the effect of this choice we repeat the variation done in Fig. 6.9 for all masses in our range, and plot the regions within which the overlap drops by less than 1% (dark gray) and 3% (light gray) of the optimal value. The results are shown in Figs. 6.11 and 6.12, along with several other physically motivated cutoff frequencies. We find that none of these cutoffs stays within 1%. In particular, the ISCO is a poor choice for both Initial and Advanced LIGO except at very low masses, where the precise value of the cutoff is not critical.

The ISCO is often pointed to—somewhat arbitrarily—as a good estimate of the breakdown of post-Newtonian approximations [42]. So, for instance, if we were to match a pN template to a physical waveform, beginning at some point in the distant past, we might expect them to separate quite badly near the ISCO. Of course, for realistic black-hole binaries, the gravitational waves will only enter the LIGO band late in the inspiral—just before the ISCO for low-mass systems, or after the ISCO for high-mass systems. We can see from Fig. 6.4 that, for masses below about $30 M_{\odot}$, the ISCO is high enough that lower-frequency parts of the waveform contribute the most to the SNR.

For very high masses, however, this basically cuts the waveform down to nothing. In Initial LIGO, the ISCO is completely buried in seismic noise for masses above about $100 M_{\odot}$. Thus, we must move the cutoff frequency up. We cannot push the cutoff far above ringdown, because the physical waveform simply




Figure 6.11: **Candidate f_c values for 3.5 pN templates with Initial LIGO**

The dark gray band contains cutoff frequencies with matches within 1% of the value at which the best overlap was obtained. The light gray band contains frequencies with matches within 3%.




Figure 6.12: Candidate f_c values for 3.5 pN templates with Advanced LIGO

The dark gray band contains cutoff frequencies with matches within 1% of the value at which the best overlap was obtained. The light gray band contains frequencies with matches within 3%. Note that the weighted-average cutoff extends past the 1% error bars for $12 < M/M_{\odot} < 40$. However, in that same region, the 3.5 pN templates do poorly overall, and we recommend pseudo 4.0 pN templates. The optimal cutoff frequency for pseudo 4.0 pN templates is much closer to the weighted-average cutoff in this mass range.

ceases to exist (see Fig. 6.4). It has been suggested that an “effective ringdown” (ERD) frequency

$$f_{\text{ERD}} := 1.07 f_{\text{Ringdown}} \quad (6.42)$$

is a useful upper limit [206].

For intermediate masses, we would like to interpolate somehow between these two extremes of ISCO and ERD. We suggest setting the cutoff frequency to a weighted average of the two, where the weights are the contributions to the SNR below the given frequency. If we assume coherent phasing between the template and the physical waveform, we can simply take the amplitudes of the two waveforms. Also, note that the restricted SPA approximation for the amplitude is reasonable. Thus, define

$$\rho_{\text{ISCO}}^2 := \int_0^{f_{\text{ISCO}}} \frac{f^{-7/3}}{S_n(f)} df, \quad (6.43)$$

$$\rho_{\text{ERD}}^2 := \int_{f_{\text{ISCO}}}^{f_{\text{ERD}}} \frac{f^{-7/3}}{S_n(f)} df, \quad (6.44)$$

$$\rho_{\text{tot}}^2 := \int_0^{f_{\text{ERD}}} \frac{f^{-7/3}}{S_n(f)} df, \quad (6.45)$$

$$f_{\text{cut}} := \frac{f_{\text{ISCO}} \rho_{\text{ISCO}} + f_{\text{ERD}} \rho_{\text{ERD}}}{\rho_{\text{tot}}}. \quad (6.46)$$

We have already dropped constant factors in the expressions for ρ that will cancel out.

Note that these expressions only depend on the total mass by way of the limits of integrations, which are very simple known functions of the mass. Thus, these integrals could be done just once for a given noise curve, storing the intermediate values. When the cutoff needs to be calculated, the integral could be evaluated at the given ISCO and ringdown frequencies. Thus, it would be a fast way of

calculating the cutoff, with no need to do the integrals each time the cutoff is needed.

We can test this recommended frequency by comparing it to the optimal cutoff frequency found by the amoeba search described in Sec. 6.4. For 3.5 pN templates in Initial LIGO, we find that it is an excellent match to the optimal frequency. Fig. 6.11 shows these two values, along with dark and light bands showing the regions in which changing f_c results in a loss of overlap of 1% and 3%, respectively. Of course, the same figure shows that using the ERD recommendation would stay within the 1% error bounds. Nonetheless, the close match between this recommendation and the true optimum suggests that it is sound. While our analysis has been restricted to equal-mass systems, the cutoff frequency we've defined here could be applied to unequal-mass systems as well. It will be interesting to see how this cutoff fares in those situations. Thus, our final recommendation is to use the weighted-average frequency cutoff throughout the entire mass range.

Similar results hold for Advanced LIGO, when using our recommended template for each mass. That is, in regions where 3.5 pN templates do poorly (see Fig. 6.7), the weighted average is a poor predictor of the optimal cutoff frequency using those templates, as shown in Fig. 6.12. However, in those same regions—where pseudo-4.0 pN templates do well—the weighted average is a good predictor of the optimal cutoff frequency for 4.0 pN templates. Thus, again, we recommend using the weighted-average frequency cutoff throughout the entire mass range with Advanced LIGO.

We make these recommendations to reduce the burden of searching. By prescribing a cutoff frequency, the search no longer needs to extend over that pa-

parameter. Similarly, by prescribing a post-Newtonian order, we need use only one template for a given total mass. On the other hand, if these recommendations decrease the overlap found when using them compared to the overlap found by an unconstrained search, it may be better to search the larger parameter space. We can evaluate the loss in overlap by comparing the results found using our recommendations to the results found when searching over the set of all three template families, and all masses, mass ratios, and cutoff frequencies. The difference is shown in Fig. 6.13. We see that the loss in overlap when using our recommendations is always less than 0.0025, and less than 0.007 for Advanced LIGO.





Figure 6.13: **Loss in overlap when using our recommendations, compared to results searching over all template families, masses, mass ratios, and cutoff frequencies**

Our recommendations prescribe the template family for a given total mass and the cutoff frequency to be used. In that case, the search is performed for the optimal mass and mass ratio of the template. For Initial LIGO, the loss in overlap when using our recommendations is always less than 0.0025; for Advanced LIGO the loss is always less than 0.007.



NOTATION AND CONVENTIONS

A.1 Conventions in this thesis

$:=$ imposed absolute equality (definition),

\equiv derived absolute equality,

\bar{z} complex conjugate of z ,

$\Re z$ real part of z ,

$\Im z$ imaginary part of z ,

$\lceil x \rceil$ least integer greater than x ,

$\lfloor x \rfloor$ greatest integer less than x ,

$\mathbf{T}, \mathbf{v}; T^{\mu\nu}, v^\alpha$ 4-tensor or 4-vector,

$\vec{v}; v^j$ 3-vector.

A.1.1 Units

Throughout most of this paper, geometric units with $G = 1$ and $c = 1$ are used. Where the link to physical measurements is more immediately relevant, mks units are used. In those cases, masses will often be quoted in units of the solar mass, and distances in megaparsecs.

$$G = 6.67259 \times 10^{-11} \frac{\text{m}^3}{\text{kg sec}^2} \quad (\text{A.1})$$

$$c = 2.99792458 \times 10^8 \frac{\text{m}}{\text{sec}} \quad (\text{A.2})$$

$$1 M_{\odot} = 1.98892 \times 10^{22} \text{ kg} \quad (\text{A.3})$$

$$1 \text{ Mpc} = 3.08568025 \times 10^{22} \text{ m} \quad (\text{A.4})$$

A.1.2 Metric

We use a metric signature $(-+++)$. This choice is compared with the one made in several other references in Table A.1. The symbol $g_{\mu\nu}$ denotes the full spacetime metric, whereas $\underline{g}_{\mu\nu}$ denotes a background metric. Greek indices run over the values 0–3. Latin indices *typically* run over the values 1–3. However, when discussing Regge–Wheeler–Zerilli methods, lowercase Latin indices (e.g., g_{ab}) refer to the t – r sector, whereas uppercase Latin indices (e.g., g_{AB}) refer to the ϑ – φ sector. These can also be combined, as in g_{Ab} .

A.1.3 Metric perturbation

Let us define the symbol

$$h^{\alpha\beta} := \eta^{\alpha\beta} - \sqrt{-g} g^{\alpha\beta}, \quad (\text{A.5})$$

where

$$g := \det(g_{\alpha\beta}) , \quad (\text{A.6})$$

and $\eta^{\alpha\beta}$ represents an auxiliary Minkowski metric. We will distinguish this quantity from the metric perturbation

$$\delta g_{\alpha\beta} := g_{\alpha\beta} - \eta_{\alpha\beta} . \quad (\text{A.7})$$

For linear perturbations of Minkowski space, $h^{\alpha\beta}$ is the trace-reversed metric perturbation:

$$h^{\alpha\beta} \rightarrow \delta g^{\alpha\beta} - \frac{1}{2} \eta^{\alpha\beta} \delta g^\gamma{}_\gamma . \quad (\text{A.8})$$

Note that various references (e.g. [195, 242, 57]) use the symbol $\bar{h}^{\alpha\beta}$ for (the linearized version of) our quantity $h^{\alpha\beta}$, and $h^{\alpha\beta}$ for our symbol $\delta g^{\alpha\beta}$. In transverse, traceless gauge, of course, the distinction is irrelevant. The choice here is made for agreement (up to a sign) with Blanchet [42], among others. Also, the overbar notation will be reserved for complex conjugation in this work.

In the linear regime, the metric perturbation is encoded in the transverse, traceless components h_+ and h_\times . To construct these quantities, we will define the polarization tensors

$$\boldsymbol{\varepsilon}_+ := \boldsymbol{\vartheta} \otimes \boldsymbol{\vartheta} - \boldsymbol{\varphi} \otimes \boldsymbol{\varphi} , \quad (\text{A.9a})$$

$$\boldsymbol{\varepsilon}_\times := \boldsymbol{\vartheta} \otimes \boldsymbol{\varphi} + \boldsymbol{\varphi} \otimes \boldsymbol{\vartheta} . \quad (\text{A.9b})$$

(Along the z axis, we will simply use the limit of these expressions as ϑ goes to 0 or π , along the line $\varphi = 0$.) Thus, we define

$$h_+ := \frac{1}{2} \varepsilon_{+ij} [h^{ij}]^{\text{TT}} , \quad (\text{A.10a})$$

$$h_\times := \frac{1}{2} \varepsilon_{\times ij} [h^{ij}]^{\text{TT}} , \quad (\text{A.10b})$$

where TT denotes the transverse, traceless part. Note that each polarization tensor defined in Eq. (A.9) differs by a sign from the one used by Blanchet. On the other hand, so does our definition of $h^{\alpha\beta}$, so h_x and h_+ should agree in sign. Finally, we define the complex combination of these components:

$$h := h_+ - ih_x . \quad (\text{A.11})$$

The choice of relative sign in this definition is motivated by the linear relation $\Psi_4 \rightarrow -\ddot{h}$, as discussed in Sec. 3.3.2, so that no complex conjugation is necessary.

A.1.4 Curvature

We define the Riemann tensor by its action on a dual-vector field ω_α as

$$\nabla_\alpha \nabla_\beta \omega_\gamma - \nabla_\beta \nabla_\alpha \omega_\gamma = R_{\alpha\beta\gamma}{}^\delta \omega_\delta . \quad (\text{A.12})$$

The Ricci tensor is then defined as the contraction

$$R_{\alpha\gamma} := R_{\alpha\beta\gamma}{}^\beta . \quad (\text{A.13})$$

This choice of sign is compared with those of other references in Table A.1.

A.1.5 Weyl scalars

We define the tetrad

$$l^\alpha := \frac{1}{\sqrt{2}} (t^\alpha + r^\alpha) , \quad (\text{A.14a})$$

$$n^\alpha := \frac{1}{\sqrt{2}} (t^\alpha - r^\alpha) , \quad (\text{A.14b})$$

$$m^\alpha := e^{i\xi} \frac{1}{\sqrt{2}} (\vartheta^\alpha + i\varphi^\alpha) , \quad (\text{A.14c})$$

$$\bar{m}^\alpha := e^{-i\xi} \frac{1}{\sqrt{2}} (\vartheta^\alpha - i\varphi^\alpha) , \quad (\text{A.14d})$$

where the vectors $\mathbf{t}, \mathbf{r}, \mathbf{\vartheta}, \mathbf{\varphi}$ are simply the usual coordinate vectors. Note that, in general, this tetrad does not satisfy any orthonormality conditions. Also note the bookkeeping parameter ξ , taking inspiration from Dray [127]. We would like to set this to zero, but keep it throughout to ease comparison with results using other conventions. Using this tetrad, we define the Weyl scalars:

$$\Psi_0 := C_{\alpha\beta\gamma\delta} l^\alpha m^\beta l^\gamma m^\delta e^{i(\zeta-2\xi)} ; \quad (\text{A.15a})$$

$$\Psi_1 := C_{\alpha\beta\gamma\delta} l^\alpha m^\beta l^\gamma n^\delta e^{i(\zeta-\xi)} ; \quad (\text{A.15b})$$

$$\Psi_2 := C_{\alpha\beta\gamma\delta} l^\alpha m^\beta \bar{m}^\gamma n^\delta e^{i\zeta} ; \quad (\text{A.15c})$$

$$\Psi_3 := C_{\alpha\beta\gamma\delta} l^\alpha n^\beta \bar{m}^\gamma n^\delta e^{i(\zeta+\xi)} ; \quad (\text{A.15d})$$

$$\Psi_4 := C_{\alpha\beta\gamma\delta} n^\alpha \bar{m}^\beta n^\gamma \bar{m}^\delta e^{i(\zeta+2\xi)} . \quad (\text{A.15e})$$

Again, we define the parameter ζ to ease comparisons with results using other conventions. The phases of ξ are included here to cancel those introduced in the tetrad. With $\xi = \zeta = 0$, these choices coincide with those of Brown et al. [66]. Comparisons for the choices of ξ and ζ with other references are shown in Table A.1.

A.1.6 Angles and frequencies

There are several distinct—though sometimes related—angles we need to discriminate between:

$$\Phi \quad \in (-\infty, \infty) \quad \text{orbital phase;} \quad (\text{A.16})$$

$$\phi \quad \in (-\infty, \infty) \quad \text{gravitational-wave phase;} \quad (\text{A.17})$$

$$\varphi \quad \in [0, 2\pi) \quad \text{azimuthal angular coordinate;} \quad (\text{A.18})$$

$$\vartheta \quad \in [0, \pi] \quad \text{zenith angular coordinate (colatitude);} \quad (\text{A.19})$$

$$\alpha \quad \in [0, 2\pi) \quad \text{right ascension;} \quad (\text{A.20})$$

$$\gamma \quad \in [0, \pi] \quad \text{codeclination;} \quad (\text{A.21})$$

$$\Psi \quad \in [0, \pi) \quad \text{polarization angle.} \quad (\text{A.22})$$

Note that these angles fall naturally into three groups: the first two being continuous phases taking arbitrary values; the next two being coordinates on the sphere; and the last three describing the orientation of an observer, or the position of a source relative to the observer.

The first two angles above are assumed to be continuous, and can be made so by adding appropriate multiples of 2π at each discontinuity. The associated frequencies are then defined by

$$\Omega := \frac{d}{dt}\Phi, \quad (\text{A.23})$$

$$\omega := \frac{d}{dt}\phi, \quad (\text{A.24})$$

where t will generally be the local coordinate time. Note that ϕ may refer to various components of a gravitational wave—e.g., the (2,2) mode or the z -axis

signal. This will be made clear in the context, by a subscript if necessary. We may also define the frequency $f := \omega/2\pi$.

The third and fourth angles— ϑ and φ —refer to coordinates on the sphere. Typically, we will think of the sphere as being centered on a binary system. We will associate to these coordinates the vector fields $\boldsymbol{\vartheta}$ and $\boldsymbol{\varphi}$ in the usual way.

The last three angles— α , γ , and Ψ —denote the relative orientation of the observer and the binary system. For example, the angles α and γ may be considered the (nearly) standard coordinates for the binary on the celestial sphere. We, however, denote by γ the “codeclination,” which is defined with respect to the more usual declination δ by

$$\gamma := \pi/2 - \delta . \tag{A.25}$$

Thus, we have $\gamma \in [0, \pi]$. The polarization angle Ψ is needed to represent the rotation of the observer in a plane transverse to the gravitational waves. Because plane-propagating gravitational waves have spin-weight $(\pm)2$, the range of this angle only needs to run from 0 to π .

A.1.7 Condon–Shortley phase

There is an ambiguity in the definition of the standard spherical harmonic functions, amounting to a disagreement in sign of $(-1)^m$. This quantity is called the “Condon–Shortley phase”—after the authors who introduced it in [95] to simplify the ladder operators of quantum mechanics, when the factor is not included in the definition of the Legendre polynomials. The factor carries through directly into the definition of the spin-weighted spherical harmonics. We choose

the Condon–Shortley convention, so that

$$Y_{1,1}(\vartheta, \varphi) = -\sqrt{\frac{3}{8\pi}} \sin(\vartheta) e^{i\varphi} , \quad (\text{A.26})$$

for example. This choice agrees with the conventions of Brown et al. [66] and MATHEMATICA, among others.

A.1.8 Fourier transforms

We relate a continuous time signal $s(t)$ to its continuous Fourier transform $\tilde{s}(f)$ by the formulas

$$\tilde{s}(f) := \int_{-\infty}^{\infty} s(t) e^{-2\pi i f t} dt , \quad (\text{A.27})$$

$$s(t) \equiv \int_{-\infty}^{\infty} \tilde{s}(f) e^{2\pi i f t} df . \quad (\text{A.28})$$

This normalization is chosen so that the transform of a pure signal $s(t) = e^{2\pi i f t}$ will simply be a Dirac δ function, with no additional factor. Note that the entire range of f (including negative frequencies) is necessary for general complex functions. In particular, if $s(t)$ is purely real, we need

$$\tilde{s}(-f) = [\tilde{s}(f)]^* . \quad (\text{A.29})$$

The notation for discretized signals involves subtleties. Suppose the continuous time signal $s(t)$ is sampled at N uniform intervals, beginning at $t_0 = 0$ and separated by Δt . This will give rise to a frequency discretization interval of $\Delta f = 1/N\Delta t$. Then, we define the quantities

$$f_j := j\Delta f , \quad (\text{A.30})$$

$$t_k := k\Delta t , \quad (\text{A.31})$$

for integers j and k . We will need to assume that the time signal is periodic, over a time $N\Delta t$. Under restriction of the range of integration to one period and discretization (represented by \rightsquigarrow), the formulas for the Fourier transform become

$$\tilde{s}(f_j) \rightsquigarrow \sum_{k=-\lfloor(N-1)/2\rfloor}^{\lfloor N/2\rfloor} s(t_k) e^{-2\pi i f_j t_k} \Delta t \equiv \Delta t \sum_{k=0}^{N-1} s(t_k) e^{-2\pi i j k / N}, \quad (\text{A.32})$$

$$s(t_k) \rightsquigarrow \sum_{j=-\lfloor(N-1)/2\rfloor}^{\lfloor N/2\rfloor} \tilde{s}(f_j) e^{-2\pi i f_j t_k} \Delta f \equiv \Delta f \sum_{j=0}^{N-1} \tilde{s}(f_j) e^{-2\pi i j k / N}. \quad (\text{A.33})$$

Note that we have used the assumption that $s(t + N\Delta t) = s(t)$, as well as the fact that $\tilde{s}(f_j) \rightsquigarrow \tilde{s}(f_{j+N})$, which results from our restriction to a finite range.

Perhaps because $\tilde{s}(f_j)$ as given by the expression in (A.32) does not equal $\tilde{s}(f)$ as given by Eq. (A.27) when $f = f_j$, there is a subtle notational distinction commonly employed [65]. We define $s_k := s(t_k)$, quite naturally. Using this, we write

$$\tilde{s}_j := \sum_{k=0}^{N-1} s_k e^{-2\pi i j k / N}, \quad (\text{A.34})$$

$$s_k \equiv \frac{1}{N} \sum_{j=0}^{N-1} \tilde{s}_j e^{2\pi i j k / N}. \quad (\text{A.35})$$

Note that $\tilde{s}(f_j) \rightsquigarrow \Delta t \tilde{s}_j$; in particular, $\tilde{s}(f_j) \neq \tilde{s}_j$. This normalization results in the circumstance that the transform of a pure signal will be N for the relevant frequency and zero for other values of j in $0 \leq j \leq N-1$ —which is roughly the discrete version of the δ function. Of course, this means that \tilde{s}_j depends on the sampling frequency; for a given signal, \tilde{s}_j will typically scale as N . This notation and choice of normalization agree with the standard conventions of LIGO [8] and the FFTW and MATLAB software packages.

The assumption that the time-domain signal is periodic, and the restriction of the Fourier transform—Eq. (A.27)—to a finite range make the frequency-domain

signal periodic on the scale of the sampling frequency

$$f_s := \frac{1}{\Delta t} . \quad (\text{A.36})$$

Any frequency $f \geq f_s$ will be *aliased* to a lower frequency, meaning that power will appear at a frequency $f - f_s$. In fact, any continuous signal at frequency f will be represented in the discretized case by a Fourier transform with power at $f + n f_s$, for *all* integers n . That is, what was a δ function in the continuous case has been made into a frequency “comb” (with infinitely many finitely large teeth) in the discrete case.

More stringent than the cutoff of the sampling frequency, there is an upper limit to the frequencies that can be unambiguously represented on a given discrete grid: the Nyquist frequency

$$f_{\text{Ny}} := \frac{1}{2\Delta t} = \frac{f_s}{2} . \quad (\text{A.37})$$

The “comb” effect combines with the negative-frequency effect to produce this limit. For example, if we take

$$s(t) = \sin(2\pi f_{\text{Ny}} t) = -\frac{i}{2} \left(e^{2\pi i f_{\text{Ny}} t} - e^{-2\pi i f_{\text{Ny}} t} \right) , \quad (\text{A.38})$$

it is easy to see that its discrete Fourier transform will be identically zero.

A.2 Comparison with other references

We present a comparison of the conventions used in this paper with those of other references. (See also the insert before the half-title page of MTW [195].) Each row presents the equivalent symbol in each reference. Thus, for example,

$$g_{\alpha\beta}^{(\text{MTW})} = -g_{\alpha\beta}^{(\text{Newman-Penrose})} . \quad (\text{A.39})$$

A blank space indicates that the quantity does not appear in that reference, so no choice is made. A question mark indicates ambiguity: the quantity is used, but not defined clearly enough.

Table A.1: Comparison of sign conventions for geometric quantities.

This reference	$g_{\alpha\beta}$	$h^{\alpha\beta}$	$\delta g^{\alpha\beta}$	h_+	h_\times	h	$R_{\alpha\beta\gamma\delta}$	Ψ_n
[92] Chandrasekhar (1992) ^a	$-g_{ij}$			$-h_{11}$	h_{31}		R_{ijkl}	$-\Psi_n$
[195] MTW (1973) ^a	$g_{\alpha\beta}$	$\bar{h}^{\alpha\beta}$	$h^{\alpha\beta}$	A_+	A_\times		$R_{\alpha\beta\gamma\delta}$	
[236] Stephani et al. (2003) ^b	g_{ab}						R_{abcd}	Ψ_n
[244] Wald (1984) ^a	$g_{\alpha\beta}$	$\bar{\gamma}^{\alpha\beta}$	$\gamma^{\alpha\beta}$				$R_{\alpha\beta\gamma\delta}$	$-\Psi_n$
[66] Brown et al. (2007) ^c	$g_{\alpha\beta}$			$\pm h_+$	$\pm h_\times$	$\pm \frac{t}{M} H$	$\pm R_{\alpha\beta\gamma\delta}$	$\pm \Psi_n$
[42] Blanchet (2006)	$g_{\alpha\beta}$	$-h^{\alpha\beta}$		h_+	h_\times			
[201] Newman–Penrose (1962)	$-g_{\alpha\beta}$						$-R_{\alpha\beta\gamma\delta}$	Ψ_n

^aNote that Chandrasekhar, MTW, and Wald only deal with linearized perturbations. The expressions given here for the metric perturbations are true only in that limit.

^bNote that Pirani et al. use a signature of (+, +, +, -), and tetrad vectors k and l in place of l and n , respectively.

^cBrown et al. have an overall sign ambiguity. This is not crucial, of course, because the waveforms are shifted in phase by an arbitrary amount.

SPIN-WEIGHTED SPHERICAL HARMONICS

B.1 Spin-weighted functions

We defined the complex vector \mathbf{m} in Eq. (A.14c) with respect to the coordinate vectors $\boldsymbol{\vartheta}, \boldsymbol{\varphi}$. If we change the coordinates, the vector *as defined* will also change as a result. In particular, choose a point p on the sphere. If we rotate the coordinates on the sphere about this point by an angle η , we will alter the vector \mathbf{m} at p as

$$\mathbf{m} \rightarrow \mathbf{m} e^{i\eta} . \tag{B.1}$$

In this sense, we may think of \mathbf{m} defined in this way not as a vector, but as a *spin-weighted* vector [148].

Any scalar formed by contraction of this vector with any other tensor would

also change. In particular, for the Weyl scalars, we have

$$\Psi_0 \rightarrow \Psi_0 e^{2i\eta} , \quad (\text{B.2})$$

$$\Psi_1 \rightarrow \Psi_1 e^{i\eta} , \quad (\text{B.3})$$

$$\Psi_2 \rightarrow \Psi_2 , \quad (\text{B.4})$$

$$\Psi_3 \rightarrow \Psi_3 e^{-i\eta} , \quad (\text{B.5})$$

$$\Psi_4 \rightarrow \Psi_4 e^{-2i\eta} . \quad (\text{B.6})$$

In analogy with the vector, we call these *spin-weighted scalars*. The coefficient of $i\eta$ in the above is the *spin weight* of the object.

For a given choice of the coordinates (ϑ, φ) , each of the Weyl scalars is a continuous, complex-valued function on the sphere. Thus, we may hope to decompose them into spherical harmonics. Unfortunately, if we use the standard spherical harmonics, we obtain undesirable features. For a simple, equal-mass, noninspiraling Keplerian binary, we might expect a pure quadrupolar radiation pattern. However, when Ψ_4 (or h) is decomposed in standard spherical harmonics, we find a nonzero $m = 2$ component for *all* values of l . In particular, there is a great deal of power in the higher l values. For example, the magnitude of the $(l, m) = (8, 2)$ mode is over 20% of the magnitude of the $(2, 2)$ mode. For large l , the components go roughly as $l^{-3/2}$. To obtain high accuracy, we would clearly need to use many l modes, in addition to losing any familiar physical intuition.

This failure can be traced to the spin weight of the functions, and the fact that the standard coordinates (ϑ, φ) are singular at the poles of the sphere. We might hope to decompose the field in terms of harmonics with the same spin weight. It turns out that we can obtain the desired behavior by using spin-weighted spherical harmonics (SWSHs) to decompose the field. These objects may

be constructed from the standard spherical harmonics through the use of spin-ladder operators. For a function f of spin weight s , Newman and Penrose [202] define the spin-raising operator^{1,2}

$$\bar{\partial}f = -(\sin\vartheta)^s \left(\frac{\partial}{\partial\vartheta} + \frac{i}{\sin\vartheta} \frac{\partial}{\partial\varphi} \right) [(\sin\vartheta)^{-s}f] . \quad (\text{B.7})$$

It is straightforward, if tedious, to show that $\bar{\partial}f$ has spin weight $s+1$. That is, $\bar{\partial}$ raises the value of the spin weight by one. Similarly, the complex conjugate $\bar{\partial}$ lowers the spin weight by one. We can use this operator to define the SWSH by

$${}_sY_{l,m} = \begin{cases} \sqrt{\frac{(l-s)!}{(l+s)!}} \bar{\partial}^s Y_{l,m} & 0 \leq s \leq l , \\ \sqrt{\frac{(l+s)!}{(l-s)!}} (-1)^s \bar{\partial}^{-s} Y_{l,m} & -l \leq s \leq 0 , \\ 0 & l < |s| . \end{cases} \quad (\text{B.8})$$

This function has the property of transforming as a spin-weighted function of weight s . Given the closed-form expressions for the spherical harmonics, we can explicitly evaluate this formula, and find³

$${}_sY_{l,m}(\varphi, \vartheta) = (-1)^m \sqrt{\frac{2l+1}{4\pi} \frac{(l+m)!(l-m)!}{(l+s)!(l-s)!}} e^{im\varphi} \\ \times \sum_{\rho} \binom{l-s}{\rho} \binom{l+s}{\rho-m+s} \left(\sin \frac{\vartheta}{2} \right)^{2l-2\rho+m-s} \left(\cos \frac{\vartheta}{2} \right)^{2\rho-m+s} . \quad (\text{B.9})$$

¹The symbol $\bar{\partial}$ is (the serifed version of) a standard character in the Icelandic alphabet, pronounced *th* with the help of the vocal chords—as the *th* sound in *then*. It also represents this sound—the voiced dental fricative—in the International Phonetic Alphabet (IPA) [168]. There, it is contrasted with the voiceless dental fricative θ , pronounced without the vocal chords—as the *th* sound in *thin*. Appearing on its own, as it does here, it is pronounced *eth*—or $\varepsilon\bar{\partial}$, in the notation of the IPA.

²Note that Dray [127] defines this operator with a relative minus sign. As a result, his SWSHs differ from those of Newman–Penrose, those of Goldberg et al., and these by a factor of $(-1)^s$. This, of course, is irrelevant if we restrict ourselves to fields of spin ± 2 .

³This equation differs from that of Goldberg et al. by a factor of $(-1)^m$. This is precisely the Condon–Shortley phase.

Here, the sum over ρ ranges over all integers for which the binomial coefficients are nonzero. Explicitly,

$$\max(0, m - s) \leq \rho \leq \min(l - s, l + m) . \quad (\text{B.10})$$

B.2 Behavior under rotation

Following Goldberg et al. [148], we can also express the spin-weighted spherical harmonics in terms of the Wigner \mathcal{D} matrices, which describe the behavior of angular-momentum eigenfunctions under rotation. This allows us to easily determine the behavior of SWSHs under rotation.

We will think of the rotation operator $\mathcal{R}(\alpha, \beta, \gamma)$ as corresponding to a rotation of the *physical system*, leaving the coordinates fixed. We then define \mathcal{R} to be composed of a rotation about the z axis by an angle γ , followed by a rotation by β about the y axis, followed by another rotation about the z axis by α , leaving the coordinate axes fixed at all times, with all rotations being in the positive sense. This is equivalent to the more usual method of defining the Euler angles [131].

Given the usual eigenvectors $|l, m\rangle$ of the angular-momentum operators, Wigner [248] defines (with different conventions⁴)

$$\mathcal{D}_{m',m}^{(l)}(\alpha, \beta, \gamma) := \langle l, m' | \mathcal{R}(\alpha, \beta, \gamma)^{-1} | l, m \rangle . \quad (\text{B.11})$$

Note that there is no mixing of different l modes under rotation. In language familiar from quantum mechanics, we say that the $|l, m\rangle$ vectors transform among themselves under an irreducible representation of the rotation group. (Of course,

⁴Wigner [248], Edmonds [131], and most other texts define the matrix \mathcal{D} with respect to \mathcal{R} , rather than \mathcal{R}^{-1} . In choosing this convention, we follow Goldberg et al. We can convert to the usual convention by taking $\mathcal{D}_{m',m}^{(l)}(\alpha, \beta, \gamma) \rightarrow \mathcal{D}_{m',m}^{(l)}(-\gamma, -\beta, -\alpha)$, which is just the inverse rotation.

more complicated transformations than rigid rotations *would* lead to mixing of l values.) The components of this matrix are⁵

$$\begin{aligned} \mathcal{D}_{m',m}^{(l)}(\alpha, \beta, \gamma) &= (-1)^{m+m'} \sqrt{\frac{(l+m)!(l-m)!}{(l+m')!(l-m')!}} e^{im\alpha + im'\gamma} \\ &\times \sum_{\rho} \binom{l+m'}{\rho} \binom{l-m'}{\rho-m-m'} \left(\sin \frac{\beta}{2}\right)^{2l-2\rho+m+m'} \left(\cos \frac{\beta}{2}\right)^{2\rho-m-m'}. \end{aligned} \quad (\text{B.12})$$

For $\gamma = 0$, this formula reduces—up to a coefficient—to the formula for the spin-weighted spherical harmonics. Specifically,

$${}_s Y_{l,m}(\vartheta, \varphi) = (-1)^s \sqrt{\frac{2l+1}{4\pi}} \mathcal{D}_{-s,m}^{(l)}(\varphi, \vartheta, 0). \quad (\text{B.13})$$

We can use this relationship to find the behavior of spin-weighted spherical harmonics under rotation.

The standard spherical harmonics are simply the projection of an angular-momentum eigenstate into a position eigenstate:

$$Y_{l,m}(\vartheta, \varphi) := \langle \vartheta, \varphi | l, m \rangle. \quad (\text{B.14})$$

Suppose the rotation operator $\mathcal{R}(\alpha, \beta, \gamma)$ takes the (physical) position eigenstate characterized by coordinates (ϑ, φ) into the eigenstate characterized by coordinates (ϑ', φ') . That is,

$$|\vartheta', \varphi'\rangle = \mathcal{R}(\alpha, \beta, \gamma) |\vartheta, \varphi\rangle. \quad (\text{B.15})$$

⁵Again, the pre-factor of $(-1)^{m+m'}$ is a result of our choice of the Condon–Shortley phase.

Then, we can write

$$Y_{l,m}(\vartheta', \varphi') = \langle \vartheta', \varphi' | l, m \rangle \quad (\text{B.16a})$$

$$= \langle \vartheta, \varphi | \mathcal{R}^\dagger | l, m \rangle \quad (\text{B.16b})$$

$$= \sum_{l', m'} \langle \vartheta, \varphi | l', m' \rangle \langle l', m' | \mathcal{R}^\dagger | l, m \rangle \quad (\text{B.16c})$$

$$= \sum_{m'} \langle \vartheta, \varphi | l, m' \rangle \langle l, m' | \mathcal{R}^{-1} | l, m \rangle \quad (\text{B.16d})$$

$$= \sum_{m'} Y_{l, m'}(\vartheta, \varphi) \mathcal{D}_{m', m}^{(l)}(\alpha, \beta, \gamma) . \quad (\text{B.16e})$$

Obviously, we can compose rotations:

$$Y_{l,m}(\vartheta'', \varphi'') = \sum_{m'} \sum_{m''} Y_{l, m''}(\vartheta, \varphi) \mathcal{D}_{m'', m'}^{(l)}(\alpha', \beta', \gamma') \mathcal{D}_{m', m}^{(l)}(\alpha, \beta, \gamma) . \quad (\text{B.17})$$

On the other hand, there must exist some rotation $\mathcal{R}(\alpha'', \beta'', \gamma'')$ for which

$$Y_{l,m}(\vartheta'', \varphi'') = \sum_{m''} Y_{l, m''}(\vartheta, \varphi) \mathcal{D}_{m'', m}^{(l)}(\alpha'', \beta'', \gamma'') . \quad (\text{B.18})$$

Now, orthonormality of the $Y_{l,m}$ functions allows us to write

$$\mathcal{D}_{m'', m}^{(l)}(\alpha'', \beta'', \gamma'') = \sum_{m'} \mathcal{D}_{m'', m'}^{(l)}(\alpha', \beta', \gamma') \mathcal{D}_{m', m}^{(l)}(\alpha, \beta, \gamma) , \quad (\text{B.19})$$

with

$$\mathcal{R}(\alpha'', \beta'', \gamma'') = \mathcal{R}(\alpha, \beta, \gamma) \mathcal{R}(\alpha', \beta', \gamma') . \quad (\text{B.20})$$

Using this result and Eq. (B.13), we can express the behavior of SWSHs under rotation:

$${}_s Y_{l,m}(\vartheta', \varphi') = \sum_{m'} {}_s Y_{l, m'}(\vartheta, \varphi) \mathcal{D}_{m', m}^{(l)}(\alpha, \beta, \gamma) , \quad (\text{B.21})$$

where the angles are related by

$$\mathcal{R}(\varphi', \vartheta', 0) = \mathcal{R}(\alpha, \beta, \gamma) \mathcal{R}(\varphi, \vartheta, 0) . \quad (\text{B.22})$$

This is the same formula as in the spin-0 case, with spin indices included.

B.3 Multipole decompositions

From Eq. (B.6), we see that Ψ_4 has spin weight -2 . It is not hard to show from the definition of h in Eq. (A.11) that h also has spin weight -2 . (We also have, in the linear approximation, $\Psi_4 \rightarrow -\ddot{h}$.) Thus, we can expect that the asymptotic forms of these fields can both be written uniquely as expansions in spin-weight -2 spherical harmonics:

$$\Psi_4(t, r, \vartheta, \varphi) = \sum_{l=-2}^{\infty} \sum_{m=-l}^l \Psi_4^{l,m}(t, r) {}_{-2}Y_{l,m}(\vartheta, \varphi) , \quad (\text{B.23})$$

$$h(t, r, \vartheta, \varphi) = \sum_{l=-2}^{\infty} \sum_{m=-l}^l h^{l,m}(t, r) {}_{-2}Y_{l,m}(\vartheta, \varphi) . \quad (\text{B.24})$$

Given the functions $\Psi_4(t, r, \vartheta, \varphi)$ and $h(t, r, \vartheta, \varphi)$, we would like to be able to find the components $\Psi_4^{l,m}$ and $h^{l,m}$. This can be done, as usual, using orthogonality properties of the SWSHs.

Using the explicit formula of Eq. (B.9), we can perform the integration

$$\begin{aligned} & \int_0^{2\pi} \int_0^\pi {}_s Y_{l,m}(\vartheta, \varphi) {}_{s'} \bar{Y}_{l',m'}(\vartheta, \varphi) \sin \vartheta \, d\vartheta \, d\varphi \\ &= (-1)^{l+l'+m+m'+s+s'} \sqrt{\frac{2l+1}{4\pi} \frac{(l+m)!(l-m)!}{(l+s)!(l-s)!}} \sqrt{\frac{2l'+1}{4\pi} \frac{(l'+m')!(l'-m')!}{(l'+s')!(l'-s')!}} \\ & \quad \times \frac{2}{(l+l'+1)!} \int_0^{2\pi} e^{i(m-m')\varphi} \, d\varphi \\ & \quad \times \sum_{\rho, \rho'} (-1)^{\rho+\rho'} \binom{l-s}{\rho} \binom{l+s}{\rho-m+s} \binom{l'-s'}{\rho'} \binom{l'+s'}{\rho'-m'+s'} \Gamma[\alpha] \Gamma[l+l'-\alpha] , \end{aligned} \quad (\text{B.25})$$

where Γ is the usual Gamma function, and

$$\alpha = \frac{2\rho - m + s + 2\rho' - m' + s'}{2} . \quad (\text{B.26})$$

While this relation is true for general s and s' , it simplifies greatly for $s = s'$. In that case, we find that the SWSHs form an orthonormal set:

$$\int_0^{2\pi} \int_0^\pi {}_s Y_{l,m}(\vartheta, \varphi) {}_s \bar{Y}_{l',m'}(\vartheta, \varphi) \sin \vartheta \, d\vartheta \, d\varphi = \delta_{l,l'} \delta_{m,m'} . \quad (\text{B.27})$$

Using this orthonormality property, we can find the components of the decompositions:

$$\Psi_4^{l,m}(t, r) = \int_0^{2\pi} \int_0^\pi \Psi_4(t, r, \vartheta, \varphi) {}_{-2} \bar{Y}_{l,m}(\vartheta, \varphi) \sin \vartheta \, d\vartheta \, d\varphi , \quad (\text{B.28})$$

$$h^{l,m}(t, r) = \int_0^{2\pi} \int_0^\pi h(t, r, \vartheta, \varphi) {}_{-2} \bar{Y}_{l,m}(\vartheta, \varphi) \sin \vartheta \, d\vartheta \, d\varphi . \quad (\text{B.29})$$

BIBLIOGRAPHY

- [1] Andrew M. Abrahams and Charles R. Evans, *Reading off gravitational radiation waveforms in numerical relativity calculations: Matching to linearized gravity*, Phys. Rev. D **37** (1988), no. 2, 318–332.
<http://link.aps.org/abstract/PRD/v37/p318>
(Cited on page 62.)
- [2] ———, *Gauge-invariant treatment of gravitational radiation near the source: Analysis and numerical simulations*, Phys. Rev. D **42** (1990), no. 8, 2585–2594.
<http://link.aps.org/abstract/PRD/v42/p2585>
(Cited on pages 62 and 99.)
- [3] Fausto Acernese et al., *The Virgo status*, Class. Quant. Grav. **23** (2006), no. 19, S635–S642.
<http://stacks.iop.org/0264-9381/23/S635>
(Cited on pages 5, 113, and 218.)
- [4] Fausto Acernese et al., *The present status of the VIRGO Central Interferometer*, Class. Quant. Grav. **19** (2002), no. 7, 1421–1428.
<http://stacks.iop.org/0264-9381/19/1421>
(Cited on pages 113 and 218.)
- [5] Parameswaran Ajith et al., *Template bank for gravitational waveforms from coalescing binary black holes: Nonspinning binaries*, Phys. Rev. D **77** (2008), no. 10, 104017.

<http://link.aps.org/abstract/PRD/v77/e104017>

(Cited on pages 220, 222, 273, 275, and 312.)

- [6] Parameswaran Ajith et al., *A phenomenological template family for black-hole coalescence waveforms*, *Class. Quant. Grav.* **24** (2007), no. 19, S689–S699.

<http://stacks.iop.org/0264-9381/24/S689>

(Cited on page 112.)

- [7] Miguel Alcubierre et al., *Dynamical evolution of quasicircular binary black hole data*, *Phys. Rev. D* **72** (2005), no. 4, 044004.

<http://link.aps.org/abstract/PRD/v72/e044004>

(Cited on pages 14, 15, 16, 22, 23, 27, and 52.)

- [8] Stuart Anderson et al., *Conventions for data and software products of LIGO and the LSC*, Tech. report, California Institute of Technology and Massachusetts Institute of Technology, 2001.

<http://www.ligo.caltech.edu/docs/T/T010095-00.pdf>

(Cited on pages 301 and 343.)

- [9] Masaki Ando et al., *Stable operation of a 300-m laser interferometer with sufficient sensitivity to detect gravitational-wave events within our galaxy*, *Phys. Rev. Lett.* **86** (2001), no. 18, 3950–3954.

<http://link.aps.org/abstract/PRL/v86/p3950>

(Cited on page 6.)

- [10] K. G. Arun, Luc Blanchet, Bala R. Iyer, and Moh'd S. S. Qusailah, *The 2.5PN gravitational wave polarizations from inspiralling compact binaries in circular orbits*, *Class. Quant. Grav.* **21** (2004), no. 15, 3771–3801.

<http://stacks.iop.org/0264-9381/21/3771>

(Cited on pages 163, 164, 166, and 249.)

- [11] Abhay Ashtekar, Christopher Beetle, and Jerzy Lewandowski, *Mechanics of rotating isolated horizons*, Phys. Rev. D **64** (2001), no. 4, 044016.

<http://link.aps.org/abstract/PRD/v64/e044016>

(Cited on page 131.)

- [12] Abhay Ashtekar and Badri Krishnan, *Dynamical horizons and their properties*, Phys. Rev. D **68** (2003), no. 10, 104030.

<http://link.aps.org/abstract/PRD/v68/e104030>

(Cited on page 131.)

- [13] Pia Astone et al., *Long-term operation of the Rome “Explorer” cryogenic gravitational wave detector*, Phys. Rev. D **47** (1993), no. 2, 362–375.

<http://link.aps.org/abstract/PRD/v47/p362>

(Cited on page 6.)

- [14] Pia Astone et al., *The gravitational wave detector NAUTILUS operating at $T = 0.1\text{ K}$* , Astropart. Phys. **7** (1997), no. 3, 231–243.

<http://www.sciencedirect.com/science/article/B6TJ1-3SPTTBC-5/2/6d0322e43c18d2453b727c40394306c5>

(Cited on page 6.)

- [15] Maria Babiuc, Bela Szilagyi, Ian Hawke, and Yosef Zlochower, *Gravitational wave extraction based on Cauchy-characteristic extraction and characteristic evolution*, Class. Quant. Grav. **22** (2005), no. 23, 5089–5107.

<http://stacks.iop.org/0264-9381/22/5089>

(Cited on page 65.)

- [16] Maria C. Babiuc, Nigel T. Bishop, Bela Szilagy, and Jeffrey Winicour, *Strategies for the characteristic extraction of gravitational waveforms*, arXiv:0808.0861, August 2008.

<http://arxiv.org/abs/0808.0861>

(Cited on page 65.)

- [17] Maria C. Babiuc, Béla Szilágyi, and Jeffrey Winicour, *Testing numerical evolution with the shifted gauge wave*, *Class. Quant. Grav.* **23** (2006), S319–S342.

<http://arxiv.org/abs/gr-qc/0511154>

(Cited on pages 15, 45, and 55.)

- [18] John G. Baker et al., *Modeling kicks from the merger of nonprecessing black hole binaries*, *Astrophys. J.* **668** (2007), no. 2, 1140–1144.

<http://www.journals.uchicago.edu/doi/abs/10.1086/521330>

(Cited on page 112.)

- [19] John G. Baker, Manuela Campanelli, Frans Pretorius, and Yosef Zlochower, *Comparisons of binary black hole merger waveforms*, *Class. Quant. Grav.* **24** (2007), S25–S31.

<http://stacks.iop.org/0264-9381/24/S25>

(Cited on page 112.)

- [20] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter, *Binary black hole merger dynamics and waveforms*, *Phys. Rev. D* **73** (2006), no. 10, 104002.

<http://link.aps.org/abstract/PRD/v73/e104002>

(Cited on pages 112 and 141.)

- [21] ———, *Gravitational-wave extraction from an inspiraling configuration of merging black holes*, *Phys. Rev. Lett.* **96** (2006), no. 11, 111102.

<http://link.aps.org/abstract/PRL/v96/e111102>

(Cited on pages 14 and 112.)

- [22] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, James R. van Meter, and M. Coleman Miller, *Getting a kick out of numerical relativity*, *Astrophys. J. Lett.* **653** (2006), no. 2, L93–L96.

<http://www.journals.uchicago.edu/doi/abs/10.1086/510448>

(Cited on page 112.)

- [23] John G. Baker et al., *Binary black hole late inspiral: Simulations for gravitational wave observations*, *Phys. Rev. D* **75** (2007), 124024.

<http://arxiv.org/abs/gr-qc/0612117>

(Cited on pages 113, 115, 116, 172, and 251.)

- [24] John G. Baker, James R. van Meter, Sean T. McWilliams, Joan Centrella, and Bernard J. Kelly, *Consistency of post-Newtonian waveforms with numerical relativity*, *Phys. Rev. Lett.* **99** (2007), no. 18, 181101.

<http://link.aps.org/abstract/PRL/v99/e181101>

(Cited on pages 113, 115, 116, 172, 174, 200, 219, 273, and 295.)

- [25] James M. Bardeen and Luisa T. Buchman, *Numerical tests of evolution systems, gauge conditions, and boundary conditions for 1D colliding gravitational plane waves*, *Phys. Rev. D* **65** (2002), no. 6, 064037.

<http://link.aps.org/abstract/PRD/v65/e064037>

(Cited on page 125.)

- [26] Barry C. Barish and Rainer Weiss, *LIGO and the detection of gravitational waves*, *Phys. Today* **52** (1999), no. 10, 44–50.

<http://scitation.aip.org/getabs/servlet/GetabsServlet?prog=normal&id=PHTOAD000052000010000044000001>

(Cited on pages 113 and 218.)

- [27] Thomas W. Baumgarte, Patrick R. Brady, Jolien D. E. Creighton, Luis Lehner, Frans Pretorius, and Ricky DeVoe, *Learning about compact binary merger: The interplay between numerical relativity and gravitational-wave astronomy*, *Phys. Rev. D* **77** (2008), no. 8, 084009.

<http://link.aps.org/abstract/PRD/v77/e084009>

(Cited on page 112.)

- [28] Thomas W. Baumgarte, Gregory B. Cook, Mark A. Scheel, Stuart L. Shapiro, and Saul A. Teukolsky, *Implementing an apparent-horizon finder in three dimensions*, *Phys. Rev. D* **54** (1996), no. 8, 4849–4857.

<http://link.aps.org/abstract/PRD/v54/p4849>

(Cited on page 131.)

- [29] Christopher Beetle, Marco Bruni, Lior M. Burko, and Andrea Nerozzi, *Towards a wave-extraction method for numerical relativity. I. Foundations and initial-value formulation*, *Phys. Rev. D* **72** (2005), no. 2, 024013.

<http://link.aps.org/abstract/PRD/v72/e024013>

(Cited on pages 62 and 140.)

- [30] Matthew J. Benacquista, *Relativistic binaries in globular clusters*, Living Rev. Relativity, 2002.
<http://www.livingreviews.org/lrr-2002-2>
(Cited on page 4.)
- [31] Carl M. Bender and Steven A. Orszag, *Advanced mathematical methods for scientists and engineers*, McGraw-Hill, New York, 1978.
(Cited on pages 239, 242, and 259.)
- [32] Emanuele Berti et al., *Inspiral, merger, and ringdown of unequal mass black hole binaries: A multipolar analysis*, Phys. Rev. D **76** (2007), no. 6, 064034.
<http://link.aps.org/abstract/PRD/v76/e064034>
(Cited on pages 112 and 228.)
- [33] Nigel T. Bishop, Roberto Gómez, Luis Lehner, and Jeffrey Winicour, *Cauchy-characteristic extraction in numerical relativity*, Phys. Rev. D **54** (1996), no. 10, 6153–6165.
<http://link.aps.org/abstract/PRD/v54/p6153>
(Cited on page 65.)
- [34] Morten Bjørhus, *The ODE formulation of hyperbolic PDEs discretized by the spectral collocation method*, SIAM J. Sci. Comput. **16** (1995), no. 3, 542–557.
<http://link.aip.org/link/?SCE/16/542/1>
(Cited on page 125.)
- [35] David G Blair, *Present status of resonant-mass detectors*, Class. Quant. Grav. **18** (2001), no. 19, 4087–4100.

<http://stacks.iop.org/0264-9381/18/4087>

(Cited on page 6.)

- [36] Luc Blanchet, *Energy losses by gravitational radiation in inspiraling compact binaries to 5/2 post-Newtonian order*, Phys. Rev. D **54** (1996), no. 2, 1417–1438.

<http://link.aps.org/abstract/PRD/v54/p1417>

(Cited on pages 166 and 167.)

- [37] ———, *Gravitational radiation reaction and balance equations to post-Newtonian order*, Phys. Rev. D **55** (1997), no. 2, 714–732.

<http://link.aps.org/abstract/PRD/v55/p714>

(Cited on page 167.)

- [38] ———, *Gravitational-wave tails of tails*, Class. Quant. Grav. **15** (1998), no. 1, 113–141, Erratum: [41].

<http://stacks.iop.org/0264-9381/15/113>

(Cited on page 161.)

- [39] ———, *On the multipole expansion of the gravitational field*, Class. Quant. Grav. **15** (1998), no. 7, 1971–1999.

<http://stacks.iop.org/0264-9381/15/1971>

(Cited on page 161.)

- [40] ———, *Quadrupole-quadrupole gravitational waves*, Class. Quant. Grav. **15** (1998), no. 1, 89–111.

<http://stacks.iop.org/0264-9381/15/89>

(Cited on page 161.)

- [41] ———, *Erratum: Gravitational-wave tails of tails*, *Class. Quant. Grav.* **22** (2005), no. 16, 3381.
<http://stacks.iop.org/0264-9381/22/3381>
(Cited on page 362.)
- [42] ———, *Gravitational radiation from post-Newtonian sources and inspiralling compact binaries*, *Living Rev. Relativity*, June 2006.
<http://www.livingreviews.org/lrr-2006-4>
(Cited on pages 8, 76, 160, 162, 163, 171, 219, 235, 236, 328, 337, and 346.)
- [43] Luc Blanchet, Alessandra Buonanno, and Guillaume Faye, *Higher-order spin effects in the dynamics of compact binaries. II. Radiation field*, *Phys. Rev. D* **74** (2006), no. 10, 104034.
<http://link.aps.org/abstract/PRD/v74/e104034>
(Cited on page 187.)
- [44] Luc Blanchet and Thibault Damour, *Radiative gravitational fields in general relativity I. General structure of the field outside the source*, *Proc. R. Soc. A* **320** (1986), no. 1555, 379–430.
<http://www.jstor.org/stable/37878>
(Cited on page 96.)
- [45] ———, *Hereditary effects in gravitational radiation*, *Phys. Rev. D* **46** (1992), no. 10, 4304–4319.
<http://link.aps.org/abstract/PRD/v46/p4304>
(Cited on page 161.)

- [46] Luc Blanchet, Thibault Damour, and Gilles Esposito-Farèse, *Dimensional regularization of the third post-Newtonian dynamics of point particles in harmonic coordinates*, Phys. Rev. D **69** (2004), no. 12, 124007.
<http://link.aps.org/abstract/PRD/v69/e124007>
(Cited on pages 160 and 238.)
- [47] Luc Blanchet, Thibault Damour, Gilles Esposito-Farèse, and Bala R. Iyer, *Gravitational radiation from inspiralling compact binaries completed at the third post-Newtonian order*, Phys. Rev. Lett. **93** (2004), no. 9, 091101.
<http://link.aps.org/abstract/PRL/v93/e091101>
(Cited on page 162.)
- [48] Luc Blanchet and Guillaume Faye, *On the equations of motion of point-particle binaries at the third post-Newtonian order*, Phys. Lett. **271** (2000), 58–64(7).
<http://www.ingentaconnect.com/content/els/03759601/2000/00000271/00000001/art00360>
(Cited on page 160.)
- [49] ———, *General relativistic dynamics of compact binaries at the third post-Newtonian order*, Phys. Rev. D **63** (2001), no. 6, 062005.
<http://link.aps.org/abstract/PRD/v63/e062005>
(Cited on pages 160 and 237.)
- [50] Luc Blanchet, Guillaume Faye, Bala R. Iyer, and Benoit Joguet, *Gravitational-wave inspiral of compact binary systems to 7/2 post-Newtonian order*, Phys. Rev. D **65** (2002), no. 6, 061501, Erratum: [51].

<http://link.aps.org/abstract/PRD/v65/e061501>

(Cited on pages 162 and 238.)

- [51] ———, *Erratum: Gravitational-wave inspiral of compact binary systems to 7/2 post-Newtonian order*, *Phys. Rev. D* **71** (2005), no. 12, 129902.

<http://link.aps.org/abstract/PRD/v71/e129902>

(Cited on pages 162 and 364.)

- [52] Luc Blanchet and Bala R. Iyer, *Third post-Newtonian dynamics of compact binaries: Equations of motion in the centre-of-mass frame*, *Class. Quant. Grav.* **20** (2003), no. 4, 755–776.

<http://stacks.iop.org/0264-9381/20/755>

(Cited on page 160.)

- [53] ———, *Hadamard regularization of the third post-Newtonian gravitational wave generation of two point masses*, *Phys. Rev. D* **71** (2005), no. 2, 024004.

<http://link.aps.org/abstract/PRD/v71/e024004>

(Cited on page 162.)

- [54] Luc Blanchet, Bala R. Iyer, and Benoit Joguet, *Gravitational waves from inspiraling compact binaries: Energy flux to third post-Newtonian order*, *Phys. Rev. D* **65** (2002), no. 6, 064005, Erratum: [55].

<http://link.aps.org/abstract/PRD/v65/e064005>

(Cited on page 162.)

- [55] ———, *Erratum: Gravitational waves from inspiraling compact binaries: Energy flux to third post-Newtonian order*, *Phys. Rev. D* **71** (2005), no. 12, 129903.

<http://link.aps.org/abstract/PRD/v71/e129903>

(Cited on page 365.)

- [56] Luc Blanchet and Gerhard Schäfer, *Gravitational wave tails and binary star systems*, *Class. Quant. Grav.* **10** (1993), no. 12, 2699–2721.

<http://stacks.iop.org/0264-9381/10/2699>

(Cited on page 164.)

- [57] Roger D. Blandford and Kip S. Thorne, *Applications of classical physics*, California Institute of Technology, 2004. Not yet published; web version.

(Cited on pages 100 and 337.)

- [58] John P. Boyd, *Chebyshev and Fourier spectral methods*, second ed., Dover Publications, Inc., 1999.

<http://citeseer.ist.psu.edu/boyd99chebyshev.html>

(Cited on pages 43 and 44.)

- [59] Michael Boyle et al., *High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions*, *Phys. Rev. D* **76** (2007), no. 12, 124038.

<http://link.aps.org/abstract/PRD/v76/e124038>

(Cited on pages 8, 10, 81, 111, 219, 222, 223, 225, 226, 231, 236, 237, 250, 254, 268, 270, 271, 273, 280, 282, 295, 300, 304, and 308.)

- [60] Michael Boyle, Duncan A. Brown, and Larne Pekowsky, *High-accuracy numerical simulations of black-hole binaries: The detection efficiency of post-Newtonian template waveforms for black-hole binary searches in ground-based gravitational-wave detectors*, in preparation, 2008, Chapter 6 of this thesis

represents an early draft.

(Cited on pages 11, 81, and 293.)

- [61] Michael Boyle et al., *High-accuracy numerical simulation of black-hole binaries: Computation of the gravitational-wave energy flux and comparisons with post-Newtonian approximants*, arXiv:0804.4184, 2008.

<http://arxiv.org/abs/0804.4184>

(Cited on pages 10, 81, 217, 295, 296, and 309.)

- [62] Michael Boyle, Lee Lindblom, Harald P. Pfeiffer, Mark A. Scheel, and Lawrence E. Kidder, *Testing the accuracy and stability of spectral methods in numerical relativity*, *Phys. Rev. D* **75** (2007), no. 2, 024006, reprinted here as Chapter 2.

<http://link.aps.org/abstract/PRD/v75/e024006>

(Cited on pages 9 and 13.)

- [63] Michael Boyle and Abdul H. Mroué, *Extrapolating gravitational-wave data from numerical simulations*, in preparation; Chapter 3 of this thesis will be incorporated into this paper, 2008.

(Cited on pages 9 and 304.)

- [64] Patrick R. Brady and Jolien D. Creighton, *Gravitational wave astronomy*, *Encyclopedia of Physical Science and Technology*, vol. 7, third ed., Academic Press, Inc, 2002, 33–48.

(Cited on pages 4 and 315.)

- [65] Duncan A. Brown, *Searching for gravitational radiation from binary black hole MACHOs in the galactic halo*, Ph.D. thesis, University of Wisconsin–

Milwaukee, 2004.

<http://arxiv.org/abs/0705.1514>

(Cited on pages 7, 302, and 343.)

- [66] Duncan A. Brown et al., *Data formats for numerical relativity waves*, arXiv:0709.0093v2, 2007.

<http://arxiv.org/abs/0709.0093>

(Cited on pages 339, 342, and 346.)

- [67] J. David Brown and James W. York, *Quasilocal energy and conserved charges derived from the gravitational action*, Phys. Rev. D **47** (1993), no. 4, 1407–1419.

<http://link.aps.org/abstract/PRD/v47/p1407>

(Cited on page 131.)

- [68] Bernd Brügmann, José A. González, Mark Hannam, Sascha Husa, and Ulrich Sperhake, *Exploring black hole superkicks*, Phys. Rev. D **77** (2008), no. 12, 124047.

<http://link.aps.org/abstract/PRD/v77/e124047>

(Cited on page 112.)

- [69] Bernd Brügmann, José A. González, Mark Hannam, Sascha Husa, Ulrich Sperhake, and Wolfgang Tichy, *Calibration of moving puncture simulations*, Phys. Rev. D **77** (2008), no. 2, 024027.

<http://link.aps.org/abstract/PRD/v77/e024027>

(Cited on pages 112, 140, and 141.)

- [70] Luisa T. Buchman and Olivier C. A. Sarbach, *Towards absorbing outer boundaries in general relativity*, Class. Quant. Grav. **23** (2006), no. 23, 6709–6744.

<http://stacks.iop.org/0264-9381/23/6709>

(Cited on pages 125 and 145.)

- [71] ———, *Improved outer boundary conditions for Einstein's field equations*, *Class. Quant. Grav.* **24** (2007), no. 12, S307–S326.

<http://stacks.iop.org/0264-9381/24/S307>

(Cited on page 74.)

- [72] Alessandra Buonanno, Yanbei Chen, and Thibault Damour, *Transition from inspiral to plunge in precessing binaries of spinning black holes*, *Phys. Rev. D* **74** (2006), no. 10, 104005.

<http://link.aps.org/abstract/PRD/v74/e104005>

(Cited on pages 114, 235, 236, 244, and 248.)

- [73] Alessandra Buonanno, Yanbei Chen, and Michele Vallisneri, *Detection template families for gravitational waves from the final stages of binary–black-hole inspirals: Nonspinning case*, *Phys. Rev. D* **67** (2003), no. 2, 024016, Erratum: [74].

<http://link.aps.org/abstract/PRD/v67/e024016>

(Cited on pages 114, 117, 235, 236, 240, 248, 275, 282, and 284.)

- [74] ———, *Erratum: Detection template families for gravitational waves from the final stages of binary–black-hole inspirals: Nonspinning case*, *Phys. Rev. D* **74** (2006), no. 2, 029903.

<http://link.aps.org/abstract/PRD/v74/e029903>

(Cited on page 369.)

- [75] Alessandra Buonanno, Gregory B. Cook, and Frans Pretorius, *Inspiral, merger, and ring-down of equal-mass black-hole binaries*, Phys. Rev. D **75** (2007), no. 12, 124018.
<http://link.aps.org/abstract/PRD/v75/e124018>
(Cited on pages 112, 115, 116, 140, 159, 171, 172, 183, 219, 221, 279, 285, and 295.)
- [76] Alessandra Buonanno and Thibault Damour, *Effective one-body approach to general relativistic two-body dynamics*, Phys. Rev. D **59** (1999), no. 8, 084006.
<http://link.aps.org/abstract/PRD/v59/e084006>
(Cited on pages 114, 213, 221, 235, 236, 244, 245, 249, 277, 279, 282, and 284.)
- [77] ———, *Transition from inspiral to plunge in binary black hole coalescences*, Phys. Rev. D **62** (2000), no. 6, 064015.
<http://link.aps.org/abstract/PRD/v62/e064015>
(Cited on pages 114, 213, 221, 235, 236, 244, 247, 249, 271, 277, 279, 282, 284, and 285.)
- [78] Alessandra Buonanno et al., *Approaching faithful templates for nonspinning binary black holes using the effective-one-body approach*, Phys. Rev. D **76** (2007), no. 10, 104049.
<http://link.aps.org/abstract/PRD/v76/e104049>
(Cited on pages 112, 213, 220, 221, 222, 223, 243, 244, 245, 246, 279, 280, 285, and 295.)
- [79] Lior M. Burko, Thomas W. Baumgarte, and Christopher Beetle, *Towards a wave-extraction method for numerical relativity. III. Analytical examples for the*

Beetle–Burko radiation scalar, Phys. Rev. D **73** (2006), no. 2, 024002.

<http://link.aps.org/abstract/PRD/v73/e024002>

(Cited on page 140.)

- [80] Gioel Calabrese, Jorge Pullin, Olivier Sarbach, and Manuel Tiglio, *Well posed constraint-preserving boundary conditions for the linearized Einstein equations*, Commun. Math. Phys. **240** (2003), 377–395.

<http://www.springerlink.com/content/eekmn3htjd0uandf/>

(Cited on page 125.)

- [81] Manuela Campanelli, Bernard Kelly, and Carlos O. Lousto, *The Lazarus project. II. Spacelike extraction with the quasi-Kinnersley tetrad*, Phys. Rev. D **73** (2006), no. 6, 064005.

<http://link.aps.org/abstract/PRD/v73/e064005>

(Cited on pages 14 and 140.)

- [82] Manuela Campanelli, Carlos Lousto, Yosef Zlochower, and David Merritt, *Large merger recoils and spin flips from generic black hole binaries*, Astrophys. J. Lett. **659** (2007), no. 1, L5–L8.

<http://www.journals.uchicago.edu/doi/abs/10.1086/516712>

(Cited on page 112.)

- [83] Manuela Campanelli and Carlos O. Lousto, *Second order gauge invariant gravitational perturbations of a Kerr black hole*, Phys. Rev. D **59** (1999), 124022.

<http://link.aps.org/abstract/PRD/v59/e124022>

(Cited on pages 74 and 87.)

- [84] Manuela Campanelli, Carlos O. Lousto, Pedro Marronetti, and Yosef Zlochower, *Accurate evolutions of orbiting black-hole binaries without excision*, *Phys. Rev. Lett.* **96** (2006), no. 11, 111101.
<http://link.aps.org/abstract/PRL/v96/e111101>
(Cited on page 112.)
- [85] Manuela Campanelli, Carlos O. Lousto, and Yosef Zlochower, *Gravitational radiation from spinning-black-hole binaries: The orbital hang up*, *Phys. Rev. D* **74** (2006), no. 4, 041501.
<http://link.aps.org/abstract/PRD/v74/e041501>
(Cited on page 112.)
- [86] ———, *Last orbit of binary black holes*, *Phys. Rev. D* **73** (2006), no. 6, 061501.
<http://link.aps.org/abstract/PRD/v73/e061501>
(Cited on page 112.)
- [87] ———, *Spin-orbit interactions in black-hole binaries*, *Phys. Rev. D* **74** (2006), no. 8, 084023.
<http://link.aps.org/abstract/PRD/v74/e084023>
(Cited on page 112.)
- [88] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, Badri Krishnan, and David Merritt, *Spin flips and precession in black-hole-binary mergers*, *Phys. Rev. D* **75** (2007), no. 6, 064030.
<http://link.aps.org/abstract/PRD/v75/e064030>
(Cited on page 112.)

- [89] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, and David Merritt, *Maximum gravitational recoil*, Phys. Rev. D **98** (2007), no. 23, 231102.
<http://link.aps.org/abstract/PRL/v98/e231102>
(Cited on page 112.)
- [90] Matthew Caudill, Gregory B. Cook, Jason D. Grigsby, and Harald P. Pfeiffer, *Circular orbits and spin in black-hole initial data*, Phys. Rev. D **74** (2006), no. 6, 064011.
<http://link.aps.org/abstract/PRD/v74/e064011>
(Cited on pages 120 and 123.)
- [91] Massimo Cerdonio et al., *The ultracryogenic gravitational-wave detector AURIGA*, Class. Quant. Grav. **14** (1997), no. 6, 1491–1494.
<http://stacks.iop.org/0264-9381/14/1491>
(Cited on page 6.)
- [92] Subrahmanyan Chandrasekhar, *The mathematical theory of black holes*, Oxford University Press, 1992.
(Cited on pages 70, 77, and 346.)
- [93] Dae-Il Choi, Bernard J. Kelly, William D. Boggs, John G. Baker, Joan Centrella, and James van Meter, *Recoiling from a kick in the head-on collision of spinning black holes*, Phys. Rev. D **76** (2007), no. 10, 104026.
<http://link.aps.org/abstract/PRD/v76/e104026>
(Cited on page 112.)
- [94] LIGO Scientific Collaboration, *LSC Algorithm Library software packages LAL, LALWRAPPER, and LALAPPS*.

<http://www.lsc-group.phys.uwm.edu/lal>

(Cited on pages 196 and 240.)

- [95] Edward U. Condon and George H. Shortley, *The theory of atomic spectra*, Cambridge University Press, 1935.

(Cited on page 341.)

- [96] Gregory B. Cook, *Corotating and irrotational binary black holes in quasicircular orbits*, Phys. Rev. D **65** (2002), no. 8, 084003.

<http://link.aps.org/abstract/PRD/v65/e084003>

(Cited on page 120.)

- [97] Gregory B. Cook and Harald P. Pfeiffer, *Excision boundary conditions for black-hole initial data*, Phys. Rev. D **70** (2004), no. 10, 104016.

<http://link.aps.org/abstract/PRD/v70/e104016>

(Cited on page 120.)

- [98] Gregory B. Cook and Bernard F. Whiting, *Approximate Killing vectors on S^2* , Phys. Rev. D **76** (2007), no. 4, 041501.

<http://link.aps.org/abstract/PRD/v76/e041501>

(Cited on page 131.)

- [99] Curt Cutler et al., *The last three minutes: Issues in gravitational-wave measurements of coalescing compact binaries*, Phys. Rev. Lett. **70** (1993), no. 20, 2984–2987.

<http://link.aps.org/abstract/PRL/v70/p2984>

(Cited on page 161.)

- [100] Curt Cutler and Éanna E. Flanagan, *Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform?*, Phys. Rev. D **49** (1994), no. 6, 2658–2697.
<http://link.aps.org/abstract/PRD/v49/p2658>
(Cited on page 298.)
- [101] Curt Cutler, Lee Samuel Finn, Eric Poisson, and Gerald Jay Sussman, *Gravitational radiation from a particle in circular orbit around a black hole. II. Numerical results for the nonrotating case*, Phys. Rev. D **47** (1993), no. 4, 1511–1518.
<http://link.aps.org/abstract/PRD/v47/p1511>
(Cited on pages 241 and 264.)
- [102] Curt Cutler and Kip S. Thorne, *An overview of gravitational-wave sources*, arXiv:gr-qc/0204090v1, 2002.
<http://arxiv.org/abs/gr-qc/0204090v1>
(Cited on pages 5 and 315.)
- [103] Thibault Damour, *Coalescence of two spinning black holes: An effective one-body approach*, Phys. Rev. D **64** (2001), no. 12, 124013.
<http://link.aps.org/abstract/PRD/v64/e124013>
(Cited on pages 114, 213, 221, and 279.)
- [104] Thibault Damour and Achamveedu Gopakumar, *Gravitational recoil during binary black hole coalescence using the effective one body approach*, Phys. Rev. D **73** (2006), no. 12, 124006.
<http://link.aps.org/abstract/PRD/v73/e124006>
(Cited on pages 235, 247, 248, 283, and 284.)

- [105] Thibault Damour, Bala R. Iyer, Piotr Jaranowski, and Bangalore S. Sathyaprakash, *Gravitational waves from black hole binary inspiral and merger: The span of third post-Newtonian effective-one-body templates*, Phys. Rev. D **67** (2003), no. 6, 064028.
<http://link.aps.org/abstract/PRD/v67/e064028>
(Cited on pages 114, 213, 221, 243, and 246.)
- [106] Thibault Damour, Bala R. Iyer, and Bangalore S. Sathyaprakash, *Improved filters for gravitational waves from inspiraling compact binaries*, Phys. Rev. D **57** (1998), no. 2, 885–907.
<http://link.aps.org/abstract/PRD/v57/p885>
(Cited on pages 114, 221, 236, 238, 240, 243, 275, 282, 283, 284, and 286.)
- [107] ———, *Frequency-domain P-approximant filters for time-truncated inspiral gravitational wave signals from compact binaries*, Phys. Rev. D **62** (2000), no. 8, 084036.
<http://link.aps.org/abstract/PRD/v62/e084036>
(Cited on page 219.)
- [108] ———, *Comparison of search templates for gravitational waves from binary inspiral*, Phys. Rev. D **63** (2001), no. 4, 044023, Erratum: [111].
<http://link.aps.org/abstract/PRD/v63/e044023>
(Cited on pages 114, 115, 117, 168, 169, 172, and 213.)
- [109] ———, *Comparison of search templates for gravitational waves from binary inspiral*, Phys. Rev. D **63** (2001), no. 4, 044023, Erratum: [112].

<http://link.aps.org/abstract/PRD/v63/e044023>

(Cited on page 295.)

- [110] ———, *Comparison of search templates for gravitational waves from binary inspiral: 3.5PN update*, Phys. Rev. D **66** (2002), no. 2, 027502, Erratum: [113].

<http://link.aps.org/abstract/PRD/v66/e027502>

(Cited on pages 114, 117, 169, and 240.)

- [111] ———, *Erratum: Comparison of search templates for gravitational waves from binary inspiral*, Phys. Rev. D **72** (2005), no. 2, 029902.

<http://link.aps.org/abstract/PRD/v72/e029902>

(Cited on page 376.)

- [112] ———, *Erratum: Comparison of search templates for gravitational waves from binary inspiral*, Phys. Rev. D **72** (2005), no. 2, 029902.

<http://link.aps.org/abstract/PRD/v72/e029902>

(Cited on page 376.)

- [113] ———, *Erratum: Comparison of search templates for gravitational waves from binary inspiral: 3.5PN update*, Phys. Rev. D **72** (2005), no. 2, 029901.

<http://link.aps.org/abstract/PRD/v72/e029901>

(Cited on page 377.)

- [114] Thibault Damour, Piotr Jaranowski, and Gerhard Schäfer, *Determination of the last stable orbit for circular general relativistic binaries at the third post-Newtonian approximation*, Phys. Rev. D **62** (2000), no. 8, 084011.

<http://link.aps.org/abstract/PRD/v62/e084011>

(Cited on pages 114, 221, 235, 236, 240, 241, 244, 245, 246, 249, 275, 277, 282, and 284.)

[115] ———, *Dynamical invariants for general relativistic two-body systems at the third post-Newtonian approximation*, Phys. Rev. D **62** (2000), no. 4, 044024.

<http://link.aps.org/abstract/PRD/v62/e044024>

(Cited on page 237.)

[116] ———, *Erratum: Poincaré invariance in the ADM Hamiltonian approach to the general relativistic two-body problem*, Phys. Rev. D **63** (2000), no. 2, 029903.

<http://link.aps.org/abstract/PRD/v63/e029903>

(Cited on page 378.)

[117] ———, *Poincaré invariance in the ADM Hamiltonian approach to the general relativistic two-body problem*, Phys. Rev. D **62** (2000), no. 2, 021501, Erratum:

[116].

<http://link.aps.org/abstract/PRD/v62/e021501>

(Cited on page 160.)

[118] ———, *Dimensional regularization of the gravitational interaction of point masses*, Phys. Lett. **513** (2001), 147–155(9).

<http://www.ingentaconnect.com/content/els/03702693/2001/00000513/00000001/art00642>

(Cited on page 160.)

[119] ———, *Equivalence between the ADM-Hamiltonian and the harmonic-coordinates approaches to the third post-Newtonian dynamics of compact binaries*, Phys. Rev. D **63** (2001), no. 4, 044021, Erratum: [120].

<http://link.aps.org/abstract/PRD/v63/e044021>

(Cited on pages 160 and 237.)

- [120] ———, *Erratum: Equivalence between the ADM-Hamiltonian and the harmonic-coordinates approaches to the third post-Newtonian dynamics of compact binaries*, Phys. Rev. D **66** (2002), no. 2, 029901.

<http://link.aps.org/abstract/PRD/v66/e029901>

(Cited on page 378.)

- [121] Thibault Damour and Alessandro Nagar, *Faithful effective-one-body waveforms of small-mass-ratio coalescing black hole binaries*, Phys. Rev. D **76** (2007), no. 6, 064028.

<http://link.aps.org/abstract/PRD/v76/e064028>

(Cited on pages 221, 279, and 285.)

- [122] ———, *Comparing effective-one-body gravitational waveforms to accurate numerical data*, Phys. Rev. D **77** (2008), no. 2, 024043.

<http://link.aps.org/abstract/PRD/v77/e024043>

(Cited on pages 220, 221, 222, 223, 236, 243, 244, 264, 268, 279, 280, 282, 284, 285, and 286.)

- [123] Thibault Damour, Alessandro Nagar, Ernst Nils Dorband, Denis Pollney, and Luciano Rezzolla, *Faithful effective-one-body waveforms of equal-mass coalescing black-hole binaries*, Phys. Rev. D **77** (2008), no. 8, 084017.

<http://link.aps.org/abstract/PRD/v77/e084017>

(Cited on pages 220, 221, 222, 223, 244, and 280.)

- [124] Thibault Damour, Alessandro Nagar, Mark Hannam, Sascha Husa, and Bernd Brügmann, *Accurate effective-one-body waveforms of inspiralling and coalescing black-hole binaries*, *Phys. Rev. D* **78** (2008), no. 4, 044039.
<http://link.aps.org/abstract/PRD/v78/e044039>
(Cited on pages 220, 221, 222, 223, 228, 279, 280, and 285.)
- [125] Vanessa C. de Andrade, Luc Blanchet, and Guillaume Faye, *Third post-Newtonian dynamics of compact binaries: Noetherian conserved quantities and equivalence between the harmonic-coordinate and ADM-Hamiltonian formalisms*, *Class. Quant. Grav.* **18** (2001), no. 5, 753–778.
<http://stacks.iop.org/0264-9381/18/753>
(Cited on pages 161 and 237.)
- [126] Peter Diener et al., *Accurate evolution of orbiting binary black holes*, *Phys. Rev. Lett.* **96** (2006), no. 12, 121101.
<http://link.aps.org/abstract/PRL/v96/e121101>
(Cited on pages 14 and 112.)
- [127] Tevian Dray, *The relationship between monopole harmonics and spin-weighted spherical harmonics*, *J. Math. Phys.* **26** (1985), no. 5, 1030–1033.
<http://link.aip.org/link/?JMP/26/1030/1>
(Cited on pages 339 and 349.)
- [128] Olaf Dreyer, Badri Krishnan, Deirdre Shoemaker, and Erik Schnetter, *Introduction to isolated horizons in numerical relativity*, *Phys. Rev. D* **67** (2003), no. 2, 024018.

<http://link.aps.org/abstract/PRD/v67/e024018>

(Cited on page 131.)

- [129] Serge Droz, Daniel J. Knapp, Eric Poisson, and Benjamin J. Owen, *Gravitational waves from inspiraling compact binaries: Validity of the stationary-phase approximation to the Fourier transform*, *Phys. Rev. D* **59** (1999), no. 12, 124016.

<http://link.aps.org/abstract/PRD/v59/e124016>

(Cited on page 300.)

- [130] Fernando Echeverria, *Gravitational-wave measurements of the mass and angular momentum of a black hole*, *Phys. Rev. D* **40** (1989), no. 10, 3194–3203.

<http://link.aps.org/abstract/PRD/v40/p3194>

(Cited on page 305.)

- [131] A. R. Edmonds, *Angular momentum in quantum mechanics*, Princeton University Press, 1996.

(Cited on page 350.)

- [132] Albert Einstein, *Die grundlage der allgemeinen relativitätstheorie*, *Ann. Phys.* (Leipzig) **49** (1916), no. 7, 769–822, Translated in Ref. [133].

[http://www.alberteinstein.info/gallery/pdf/CP6Doc30_pp284-339.](http://www.alberteinstein.info/gallery/pdf/CP6Doc30_pp284-339.pdf)

pdf

(Cited on page 7.)

- [133] ———, “The Foundation of the General Theory of Relativity”, *The Berlin years: Writings, 1914–1917*, *The Collected Papers of Albert Einstein*, vol. 6, Princeton University Press, 1997, 146–201, (Alfred Engel, trans.).

http://www.alberteinstein.info/gallery/pdf/CP6Doc30_English_

pp146-200.pdf

(Cited on page 381.)

- [134] Zachariah B. Etienne, Joshua A. Faber, Yuk Tung Liu, Stuart L. Shapiro, and Thomas W. Baumgarte, *Filling the holes: Evolving excised binary black hole initial data with puncture techniques*, *Phys. Rev. D* **76** (2007), no. 10, 101503.

<http://link.aps.org/abstract/PRD/v76/e101503>

(Cited on page 112.)

- [135] Guillaume Faye, Luc Blanchet, and Alessandra Buonanno, *Higher-order spin effects in the dynamics of compact binaries. I. Equations of motion*, *Phys. Rev. D* **74** (2006), no. 10, 104033.

<http://link.aps.org/abstract/PRD/v74/e104033>

(Cited on page 187.)

- [136] Lee Samuel Finn, *Detection, measurement, and gravitational radiation*, *Phys. Rev. D* **46** (1992), no. 12, 5236–5249.

<http://link.aps.org/abstract/PRD/v46/p5236>

(Cited on page 61.)

- [137] Lee Samuel Finn and David F. Chernoff, *Observing binary inspiral in gravitational radiation: One interferometer*, *Phys. Rev. D* **47** (1993), no. 6, 2198–2219.

<http://link.aps.org/abstract/PRD/v47/p2198>

(Cited on page 61.)

- [138] David R. Fiske, John G. Baker, James R. van Meter, Dae-Il Choi, and Joan M. Centrella, *Wave zone extraction of gravitational radiation in three-dimensional numerical relativity*, *Phys. Rev. D* **71** (2005), no. 10, 104036.

<http://link.aps.org/abstract/PRD/v71/e104036>

(Cited on pages 140 and 150.)

- [139] Éanna É. Flanagan and Scott A. Hughes, *Measuring gravitational waves from binary black hole coalescences. I. Signal to noise for inspiral, merger, and ringdown*, Phys. Rev. D **57** (1998), no. 8, 4535–4565.

<http://link.aps.org/abstract/PRD/v57/p4535>

(Cited on page 113.)

- [140] ———, *Measuring gravitational waves from binary black hole coalescences. II. The waves' information and its extraction, with and without templates*, Phys. Rev. D **57** (1998), no. 8, 4566–4587.

<http://link.aps.org/abstract/PRD/v57/p4566>

(Cited on page 219.)

- [141] Helmut Friedrich, *On the hyperbolicity of Einstein's and other gauge field equations*, Commun. Math. Phys. **100** (1985), no. 4, 525–543.

<http://www.springerlink.com/content/w602g633428x8365>

(Cited on page 125.)

- [142] Helmut Friedrich and Gabriel Nagy, *The initial boundary value problem for Einstein's vacuum field equation*, Commun. Math. Phys. **201** (1999), no. 3, 619–655.

<http://www.springerlink.com/content/vvtmtkvptv3p58qw>

(Cited on page 125.)

- [143] Peter Fritschel, *Second generation instruments for the Laser Interferometer Gravitational-wave Observatory (LIGO)*, Proc. SPIE (Mike Cruise and Peter

Saulson, eds.), vol. 4856, 2003, 282–291.

<http://arxiv.org/abs/gr-qc/0308090v1>

(Cited on page 113.)

- [144] David Garfinkle, *Harmonic coordinate method for simulating generic singularities*, Phys. Rev. D **65** (2002), no. 4, 044029.

<http://link.aps.org/abstract/PRD/v65/e044029>

(Cited on page 125.)

- [145] Sara J. Gettel, Michael T. Geske, , and Tim A. McKay, *A catalog of 1022 bright contact binary stars*, Astron. J. **131** (2006), no. 1, 621–632.

<http://stacks.iop.org/1538-3881/131/621>

(Cited on page 3.)

- [146] Reinaldo J. Gleiser, Carlos O. Nicasio, Richard H. Price, and Jorge Pullin, *Second-order perturbations of a Schwarzschild black hole*, Class. Quant. Grav. **13** (1996), no. 10, L117–L124.

<http://stacks.iop.org/0264-9381/13/L117>

(Cited on pages 74 and 87.)

- [147] ———, *Gravitational radiation from Schwarzschild black holes: The second-order perturbation formalism*, Phys. Rep. **325** (2000), no. 2, 41–81.

<http://www.sciencedirect.com/science/article/B6TVP-3YJ9Y9H-2/2/9a81c5e5dde85ab58517c382c96f3c77>

[9a81c5e5dde85ab58517c382c96f3c77](http://www.sciencedirect.com/science/article/B6TVP-3YJ9Y9H-2/2/9a81c5e5dde85ab58517c382c96f3c77)

(Cited on pages 74 and 87.)

- [148] Josh N. Goldberg, Alan J. Macfarlane, Ezra T. Newman, Fritz Rohrlich, and Ennackal Chandy George Sudarshan, *Spin- s spherical harmonics and δ* , J.

Math. Phys. **8** (1967), no. 11, 2155–2161.

<http://link.aip.org/link/?JMP/8/2155/1>

(Cited on pages 347 and 350.)

- [149] José A. González, Mark Hannam, Ulrich Sperhake, Bernd Brügmann, and Sascha Husa, *Supermassive recoil velocities for binary black-hole mergers with antialigned spins*, Phys. Rev. Lett. **98** (2007), no. 23, 231101.

<http://link.aps.org/abstract/PRL/v98/e231101>

(Cited on page 112.)

- [150] José A. González, Ulrich Sperhake, Bernd Brügmann, Mark Hannam, and Sascha Husa, *Maximum kick from nonspinning black-hole binary inspiral*, Phys. Rev. Lett. **98** (2007), no. 9, 091101.

<http://link.aps.org/abstract/PRL/v98/e091101>

(Cited on page 112.)

- [151] Achamveedu Gopakumar, Mark Hannam, Sascha Husa, and Bernd Brügmann, *Comparison between numerical relativity and a new class of post-Newtonian gravitational-wave phase evolutions: The non-spinning equal-mass case*, arXiv:0712.3737, 2007.

<http://arxiv.org/abs/0712.3737>

(Cited on page 219.)

- [152] Achamveedu Gopakumar, Mark Hannam, Sascha Husa, and Bernd Brügmann, *Comparison between numerical relativity and a new class of post-Newtonian gravitational-wave phase evolutions: The non-spinning equal-mass case*, arXiv:0712.3737, 2007.

<http://arxiv.org/abs/0712.3737>

(Cited on page 295.)

- [153] David Gottlieb and Jan S. Hesthaven, *Spectral methods for hyperbolic problems*, J. Comput. Appl. Math. **128** (2001), no. 1–2, 83–131.

[http://dx.doi.org/10.1016/S0377-0427\(00\)00510-0](http://dx.doi.org/10.1016/S0377-0427(00)00510-0)

(Cited on page 125.)

- [154] Carsten Gundlach, *Pseudospectral apparent horizon finders: An efficient new algorithm*, Phys. Rev. D **57** (1998), no. 2, 863–875.

<http://link.aps.org/abstract/PRD/v57/p863>

(Cited on page 131.)

- [155] Eric Gustafson, Deirdre Shoemaker, Kenneth Strain, and Rainer Weiss, *LSC white paper on detector research and development*, Tech. report, California Institute of Technology and Massachusetts Institute of Technology, September 1999.

<http://www.ligo.caltech.edu/docs/T/T990080-00.pdf>

(Cited on page 6.)

- [156] Mark Hannam, Sascha Husa, Bernd Brügmann, and Achamveedu Gopakumar, *Comparison between numerical-relativity and post-Newtonian waveforms from spinning binaries: The orbital hang-up case*, arXiv:0712.3787, 2007.

<http://arxiv.org/abs/0712.3787>

(Cited on pages 219, 220, and 295.)

- [157] Mark Hannam, Sascha Husa, José A. González, Ulrich Sperhake, and Bernd Brügmann, *Where post-Newtonian and numerical-relativity waveforms meet*,

Phys. Rev. D **77** (2008), no. 4, 044020.

<http://link.aps.org/abstract/PRD/v77/e044020>

(Cited on pages 113, 116, 152, 169, 171, 173, 174, 196, 200, 219, 273, and 295.)

- [158] Frank Herrmann, Ian Hinder, Deirdre M. Shoemaker, and Pablo Laguna, *Unequal mass binary black hole plunges and gravitational recoil*, Class. Quant. Grav. **24** (2007), no. 12, S33–S42.

<http://stacks.iop.org/0264-9381/24/S33>

(Cited on pages 14 and 112.)

- [159] Frank Herrmann, Ian Hinder, Deirdre M. Shoemaker, Pablo Laguna, and Richard A. Matzner, *Binary black holes: Spin dynamics and gravitational recoil*, Phys. Rev. D **76** (2007), no. 8, 084032.

<http://link.aps.org/abstract/PRD/v76/e084032>

(Cited on page 112.)

- [160] _____, *Gravitational recoil from spinning binary black hole mergers*, Astrophys. J. **661** (2007), no. 1, 430–436.

<http://www.journals.uchicago.edu/doi/abs/10.1086/513603>

(Cited on page 112.)

- [161] Jan S. Hesthaven, *Spectral penalty methods*, Appl. Num. Math. **33** (2000), 23–41(19).

<http://www.ingentaconnect.com/content/els/01689274/2000/00000033/00000001/art00068>

(Cited on page 125.)

- [162] Stefan Hild (for the LIGO Scientific Collaboration), *The status of GEO 600*, *Class. Quant. Grav.* **23** (2006), no. 19, S643–S651.
<http://stacks.iop.org/0264-9381/23/S643>
(Cited on pages 113 and 218.)
- [163] Ian Hinder, Frank Herrmann, Pablo Laguna, and Deirdre Shoemaker, *Comparisons of eccentric binary black hole simulations with post-Newtonian models*, arXiv:0806.1037, 2008.
<http://arxiv.org/abs/0806.1037>
(Cited on page 295.)
- [164] Michael Holst, Lee Lindblom, Robert Owen, Harald P. Pfeiffer, Mark A. Scheel, and Lawrence E. Kidder, *Optimal constraint projection for hyperbolic evolution systems*, *Phys. Rev. D* **70** (2004), no. 8, 084017.
<http://link.aps.org/abstract/PRD/v70/e084017>
(Cited on page 304.)
- [165] Peter Hübner, *From now to timelike infinity on a finite grid*, *Class. Quant. Grav.* **18** (2001), no. 10, 1871–1884.
<http://stacks.iop.org/0264-9381/18/1871>
(Cited on page 65.)
- [166] Sascha Husa, José A. González, Mark Hannam, Bernd Brügmann, and Ulrich Sperhake, *Reducing phase error in long numerical binary black hole evolutions with sixth order finite differencing*, arXiv:0706.0740v1, 2007.
<http://arxiv.org/abs/0706.0740v1>
(Cited on pages 113 and 116.)

- [167] Sascha Husa, Mark Hannam, José A. González, Ulrich Sperhake, and Bernd Brügmann, *Reducing eccentricity in black-hole binary evolutions with initial parameters from post-Newtonian inspiral*, arXiv:0706.0904, 2007.
<http://arxiv.org/abs/0706.0904>
(Cited on pages 116, 123, 183, and 186.)
- [168] International Phonetic Association, *Handbook of the International Phonetic Association: A guide to the use of the International Phonetic Alphabet*, Cambridge University Press, June 1999.
(Cited on page 349.)
- [169] Yousuke Itoh, *Equation of motion for relativistic compact binaries with the strong field point particle limit: Third post-Newtonian order*, Phys. Rev. D **69** (2004), no. 6, 064018.
<http://link.aps.org/abstract/PRD/v69/e064018>
(Cited on page 160.)
- [170] Yousuke Itoh and Toshifumi Futamase, *New derivation of a third post-Newtonian equation of motion for relativistic compact binaries without ambiguity*, Phys. Rev. D **68** (2003), no. 12, 121501.
<http://link.aps.org/abstract/PRD/v68/e121501>
(Cited on page 160.)
- [171] Yousuke Itoh, Toshifumi Futamase, and Hideki Asada, *Equation of motion for relativistic compact binaries with the strong field point particle limit: The second and half post-Newtonian order*, Phys. Rev. D **63** (2001), no. 6, 064038.

<http://link.aps.org/abstract/PRD/v63/e064038>

(Cited on page 160.)

[172] John David Jackson, *Classical electrodynamics*, third ed., John Wiley and Sons, Inc., 1999.

(Cited on page 95.)

[173] Piotr Jaranowski and Andrzej Królak, *Gravitational-wave data analysis. Formalism and sample applications: The Gaussian case*, *Living Rev. Relativity*, 2005.

<http://www.livingreviews.org/lrr-2005-3>

(Cited on page 7.)

[174] Piotr Jaranowski and Gerhard Schäfer, *Third post-Newtonian higher order ADM Hamilton dynamics for two-body point-mass systems*, *Phys. Rev. D* **57** (1998), no. 12, 7274–7291, Erratum: [176].

<http://link.aps.org/abstract/PRD/v57/p7274>

(Cited on pages 160 and 237.)

[175] ———, *Binary black-hole problem at the third post-Newtonian approximation in the orbital motion: Static part*, *Phys. Rev. D* **60** (1999), no. 12, 124003.

<http://link.aps.org/abstract/PRD/v60/e124003>

(Cited on page 160.)

[176] ———, *Erratum: Third post-Newtonian higher order ADM Hamilton dynamics for two-body point-mass systems*, *Phys. Rev. D* **63** (2000), no. 2, 029902.

<http://link.aps.org/abstract/PRD/v63/e029902>

(Cited on page 390.)

- [177] Lawrence E. Kidder, *Coalescing binary systems of compact objects to (post)^{5/2}-Newtonian order. V. Spin effects*, Phys. Rev. D **52** (1995), no. 2, 821–847.
<http://link.aps.org/abstract/PRD/v52/p821>
(Cited on page 188.)
- [178] ———, *Using full information when computing modes of post-Newtonian waveforms from inspiralling compact binaries in circular orbit*, Phys. Rev. D **77** (2008), no. 4, 044016.
<http://link.aps.org/abstract/PRD/v77/e044016>
(Cited on pages 118, 165, 166, 203, 206, and 249.)
- [179] Lawrence E. Kidder, Luc Blanchet, and Bala R. Iyer, *A note on the radiation reaction in the 2.5PN waveform from inspiralling binaries in quasi-circular orbits*, Class. Quant. Grav. **24** (2007), no. 20, 5307–5312.
<http://stacks.iop.org/0264-9381/24/5307>
(Cited on pages 163 and 164.)
- [180] Lawrence E. Kidder, Lee Lindblom, Mark A. Scheel, Luisa T. Buchman, and Harald P. Pfeiffer, *Boundary conditions for the Einstein evolution system*, Phys. Rev. D **71** (2005), no. 6, 064020.
<http://link.aps.org/abstract/PRD/v71/e064020>
(Cited on pages 17 and 125.)
- [181] Lawrence E. Kidder, Mark A. Scheel, and Saul A. Teukolsky, *Extending the lifetime of 3D black hole computations with a new hyperbolic system of evolution equations*, Phys. Rev. D **64** (2001), no. 6, 064017.

<http://link.aps.org/abstract/PRD/v64/e064017>

(Cited on pages 15 and 16.)

- [182] Bence Kocsis and Abraham Loeb, *Distortion of gravitational-wave packets due to their self-gravity*, Phys. Rev. D **76** (2007), no. 8, 084022.

<http://link.aps.org/abstract/PRD/v76/e084022>

(Cited on pages 71 and 152.)

- [183] Michael Koppitz et al., *Recoil velocities from equal-mass binary-black-hole mergers*, Phys. Rev. Lett. **99** (2007), no. 4, 041102.

<http://link.aps.org/abstract/PRL/v99/e041102>

(Cited on page 112.)

- [184] Andrzej Królak, Kostas D. Kokkotas, and Gerhard Schäfer, *Estimation of the post-Newtonian parameters in the gravitational-wave emission of a coalescing binary*, Phys. Rev. D **52** (1995), no. 4, 2089–2111.

<http://link.aps.org/abstract/PRD/v52/p2089>

(Cited on page 183.)

- [185] Luis Lehner and Osvaldo M. Moreschi, *Dealing with delicate issues in waveform calculations*, Phys. Rev. D **76** (2007), no. 12, 124040.

<http://link.aps.org/abstract/PRD/v76/e124040>

(Cited on pages 95 and 140.)

- [186] Lee Lindblom, Benjamin J. Owen, and Duncan A. Brown, *Model waveform accuracy standards for gravitational wave data analysis*, arXiv:0809.3844, 2008.

<http://arxiv.org/abs/0809.3844>

(Cited on page 62.)

- [187] Lee Lindblom, Mark A. Scheel, Lawrence E. Kidder, Robert Owen, and Oliver Rinne, *A new generalized harmonic evolution system*, *Classical and Quantum Gravity* **23** (2006), no. 16, S447–S462.
<http://stacks.iop.org/0264-9381/23/S447>
(Cited on pages 124, 125, 138, and 304.)
- [188] Duncan R. Lorimer, *Binary and millisecond pulsars*, *Living Rev. Relativity*, 2005.
<http://www.livingreviews.org/lrr-2005-7>
(Cited on page 4.)
- [189] Geoffrey Lovelace, *Topics in gravitational wave physics*, Ph.D. thesis, California Institute of Technology, 2007.
<http://etd.caltech.edu/etd/available/etd-05232007-115433>
(Cited on page 215.)
- [190] Geoffrey Lovelace, Robert Owen, Harald P. Pfeiffer, and Tony Chu, *Binary-black-hole initial data with nearly-extremal spins*, arXiv:0805.4192, 2008.
<http://arxiv.org/abs/0805.4192>
(Cited on pages 273 and 308.)
- [191] Pedro Marronetti et al., *Binary black holes on a budget: Simulations using workstations*, *Class. Quant. Grav.* **24** (2007), no. 12, S43–S58.
<http://stacks.iop.org/0264-9381/24/S43>
(Cited on page 112.)
- [192] Evan Mauceli et al., *The Allegro gravitational wave detector: Data acquisition and analysis*, *Phys. Rev. D* **54** (1996), no. 2, 1264–1275.

<http://link.aps.org/abstract/PRD/v54/p1264>

(Cited on page 6.)

- [193] Jeffrey E. McClintock and Ronald A. Remillard, *Black hole binaries*, Compact Stellar X-ray Sources (Walter H. G. Lewin and Michiel van der Klis, eds.), Cambridge University Press, 2006.

<http://arxiv.org/abs/astro-ph/0306213>

(Cited on page 5.)

- [194] Mark Miller, *Circular orbit approximation for binary compact objects in general relativity*, Phys. Rev. D **69** (2004), no. 12, 124013.

<http://link.aps.org/abstract/PRD/v69/e124013>

(Cited on page 213.)

- [195] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler, *Gravitation*, W. H. Freeman, Sept. 1973.

(Cited on pages 337, 344, and 346.)

- [196] Vincent Moncrief, *Gravitational perturbations of spherically symmetric systems. I. The exterior problem*, Ann. Phys. (NY) **88** (1974), no. 2, 323–342.

<http://www.sciencedirect.com/science/article/B6WB1-4DF50YP-17W/2/f1926ff9bd2916c718567373fd27ec7e>

(Cited on page 73.)

- [197] Abdul H. Mroué, Lawrence E. Kidder, and Saul A. Teukolsky, *Ineffectiveness of Padé resummation techniques in post-Newtonian approximations*, arXiv:0805.2390, 2008.

<http://arxiv.org/abs/0805.2390>

(Cited on page 295.)

- [198] Alessandro Nagar and Luciano Rezzolla, *Gauge-invariant non-spherical metric perturbations of Schwarzschild black-hole spacetimes*, *Class. Quant. Grav.* **22** (2005), no. 16, R167–R192.

<http://stacks.iop.org/0264-9381/22/R167>

(Cited on page 74.)

- [199] Andrea Nerozzi, Christopher Beetle, Marco Bruni, Lior M. Burko, and Denis Pollney, *Towards a wave-extraction method for numerical relativity. II. The quasi-Kinnersley frame*, *Phys. Rev. D* **72** (2005), no. 2, 024014.

<http://link.aps.org/abstract/PRD/v72/e024014>

(Cited on page 140.)

- [200] Andrea Nerozzi, Marco Bruni, Lior M. Burko, and Virginia Re, *Towards a novel wave-extraction method for numerical relativity*, *Proceedings of the Albert Einstein Century International Conference, Paris, France, 2005*, 702–707.

<http://scitation.aip.org/getabs/servlet/GetabsServlet?prog=normal&id=APCPCS000861000001000702000001&idtype=cvips&gifs=Yes>

(Cited on page 140.)

- [201] Ezra T. Newman and Roger Penrose, *An approach to gravitational radiation by a method of spin coefficients*, *J. Math. Phys.* **3** (1962), no. 3, 566–578.

<http://link.aip.org/link/?JMP/3/566/1>

(Cited on pages 76 and 346.)

- [202] ———, *Note on the Bondi–Metzner–Sachs group*, *J. Math. Phys.* **7** (1966), no. 5, 863–870.
<http://link.aip.org/link/?JMP/7/863/1>
(Cited on pages 80 and 349.)
- [203] Carlos O. Nicasio, Reinaldo Gleiser, and Jorge Pullin, *Second order perturbations of a Schwarzschild black hole: Inclusion of odd parity perturbations*, *Gen. Relativ. Gravit.* **32** (2000), no. 10, 2021–2042.
<http://www.springerlink.com/content/g18267241584g27k>
(Cited on pages 74 and 87.)
- [204] Robert Owen, *Constraint damping in first-order evolution systems for numerical relativity*, *Phys. Rev. D* **76** (2007), no. 4, 044019.
<http://link.aps.org/abstract/PRD/v76/e044019>
(Cited on pages 18 and 25.)
- [205] ———, *Topics in numerical relativity: The periodic standing-wave approximation, the stability of constraints in free evolution, and the spin of dynamical black holes*, Ph.D. thesis, California Institute of Technology, 2007.
<http://resolver.caltech.edu/CaltechETD:etd-05252007-143511>
(Cited on page 131.)
- [206] Yi Pan et al., *Data-analysis driven comparison of analytic and numerical coalescing binary waveforms: Nonspinning case*, *Phys. Rev. D* **77** (2008), no. 2, 024014.
<http://link.aps.org/abstract/PRD/v77/e024014>
(Cited on pages 112, 116, 169, 172, 203, 220, 221, 222, 223, 244, 280, 295, 313, 327, and 331.)

- [207] Enrique Pazos, Ernst Nils Dorband, Alessandro Nagar, Carlos Palenzuela, Erik Schnetter, and Manuel Tiglio, *How far away is far enough for extracting numerical waveforms, and how much do they depend on the extraction method?*, arXiv:gr-qc/0612149, 2006.
<http://arxiv.org/abs/gr-qc/0612149>
(Cited on page 140.)
- [208] Philip C. Peters, *Gravitational radiation and the motion of two point masses*, Phys. Rev. **136** (1964), no. 4B, B1224–B1232.
<http://link.aps.org/abstract/PR/v136/pB1224>
(Cited on pages 113, 114, and 162.)
- [209] Harald P. Pfeiffer, Duncan A. Brown, Lawrence E. Kidder, Lee Lindblom, Geoffrey Lovelace, and Mark A. Scheel, *Reducing orbital eccentricity in binary black hole simulations*, Class. Quant. Grav. **24** (2007), no. 12, S59–S81.
<http://stacks.iop.org/0264-9381/24/S59>
(Cited on pages 113, 116, 120, 123, 126, 134, 139, 173, 183, 228, 304, and 313.)
- [210] Harald P. Pfeiffer, Lawrence E. Kidder, Marl A. Scheel, and Saul A. Teukolsky, *A multidomain spectral method for solving elliptic equations*, Comput. Phys. Commun. **152** (2003), 253–273(21).
<http://www.ingentaconnect.com/content/els/00104655/2003/00000152/00000003/art00847>
(Cited on page 120.)
- [211] Harald P. Pfeiffer, Saul A. Teukolsky, and Gregory B. Cook, *Quasicircular orbits for spinning binary black holes*, Phys. Rev. D **62** (2000), no. 10, 104018.

<http://link.aps.org/abstract/PRD/v62/e104018>

(Cited on page 131.)

- [212] Harald P. Pfeiffer and James W. York, *Extrinsic curvature and the Einstein constraints*, Phys. Rev. D **67** (2003), no. 4, 044022.

<http://link.aps.org/abstract/PRD/v67/e044022>

(Cited on page 120.)

- [213] Felix A. E. Pirani, *Introduction to gravitational radiation theory*, Lectures in Theoretical Physics (Stanley Deser and Kenneth W. Ford, eds.), Brandeis Summer Institute in Theoretical Physics, vol. 1, Prentice-Hall, 1964, 249–373.

(Cited on pages 96 and 97.)

- [214] Eric Poisson, *Gravitational radiation from a particle in circular orbit around a black hole. I. Analytical results for the nonrotating case*, Phys. Rev. D **47** (1993), no. 4, 1497–1510.

<http://link.aps.org/abstract/PRD/v47/p1497>

(Cited on page 241.)

- [215] ———, *Gravitational radiation from a particle in circular orbit around a black hole. VI. Accuracy of the post-Newtonian expansion*, Phys. Rev. D **52** (1995), no. 10, 5719–5723.

<http://link.aps.org/abstract/PRD/v52/p5719>

(Cited on pages 239, 286, 287, and 288.)

- [216] ———, *Absorption of mass and angular momentum by a black hole: Time-domain formalisms for gravitational perturbations, and the small-hole or slow-motion ap-*

- proximation*, Phys. Rev. D **70** (2004), no. 8, 084044.
<http://link.aps.org/abstract/PRD/v70/e084044>
(Cited on page 179.)
- [217] Denis Pollney et al., *Recoil velocities from equal-mass binary black-hole mergers: A systematic investigation of spin-orbit aligned configurations*, Phys. Rev. D **76** (2007), no. 12, 124002.
<http://link.aps.org/abstract/PRD/v76/e124002>
(Cited on page 228.)
- [218] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery, *Numerical recipes: The art of scientific computing*, third ed., Cambridge University Press, 2007.
(Cited on pages 67, 92, and 315.)
- [219] Frans Pretorius, *Evolution of binary black-hole spacetimes*, Phys. Rev. Lett. **95** (2005), no. 12, 121101.
<http://link.aps.org/abstract/PRL/v95/e121101>
(Cited on pages 14, 112, and 115.)
- [220] ———, *Numerical relativity using a generalized harmonic decomposition*, Class. Quant. Grav. **22** (2005), no. 2, 425–451.
<http://stacks.iop.org/0264-9381/22/425>
(Cited on page 125.)
- [221] ———, *Simulation of binary black hole spacetimes with a harmonic evolution scheme*, Class. Quant. Grav. **23** (2006), no. 16, S529–S552.

<http://stacks.iop.org/0264-9381/23/S529>

(Cited on page 112.)

- [222] Tullio Regge and John A. Wheeler, *Stability of a Schwarzschild singularity*, Phys. Rev. **108** (1957), no. 4, 1063–1069.

<http://link.aps.org/abstract/PR/v108/p1063>

(Cited on page 73.)

- [223] Oliver Rinne, *Stable radiation-controlling boundary conditions for the generalized harmonic Einstein equations*, Class. Quant. Grav. **23** (2006), no. 22, 6275–6300.

<http://stacks.iop.org/0264-9381/23/6275>

(Cited on page 125.)

- [224] Oliver Rinne, Luisa T. Buchman, Mark A. Scheel, and Harald P. Pfeiffer, *Implementation of higher-order absorbing boundary conditions for the Einstein equations*, arXiv:0811.3593, 2008.

<http://arxiv.org/abs/0811.3593>

(Cited on page 76.)

- [225] Oliver Rinne, Lee Lindblom, and Mark A. Scheel, *Testing outer boundary treatments for the Einstein equations*, Class. Quant. Grav. **24** (2007), no. 16, 4053–4078.

<http://stacks.iop.org/0264-9381/24/4053>

(Cited on pages 125, 142, and 145.)

- [226] Milton Ruiz, Miguel Alcubierre, Darío Núñez, and Ryoji Takahashi, *Multipole expansions for energy and momenta carried by gravitational waves*, Gen. Relativ. Gravit. **40** (2008), no. 8, 1705–1729. [Note that the final sign in the

fourth line of Eq. (25) should be positive.]

<http://www.springerlink.com/content/c527258554223613>

(Cited on page 74.)

- [227] Ray K. Sachs, *Gravitational waves in general relativity. VI. The outgoing radiation condition*, Proc. R. Soc. A **264** (1961), no. 1318, 309–338.

<http://www.jstor.org/stable/2414993>

(Cited on pages 77 and 96.)

- [228] Olivier Sarbach and Manuel Tiglio, *Gauge-invariant perturbations of Schwarzschild black holes in horizon-penetrating coordinates*, Phys. Rev. D **64** (2001), no. 8, 084016.

<http://link.aps.org/abstract/PRD/v64/e084016>

(Cited on pages 73 and 76.)

- [229] Mark A. Scheel et al., *High-accuracy waveforms for binary black-hole inspiral, merger, and ringdown*, arXiv:0810.1767, 2008, submitted to Phys. Rev. D.

<http://arxiv.org/abs/0810.1767>

(Cited on pages 81 and 304.)

- [230] 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** (2006), no. 10, 104006.

<http://link.aps.org/abstract/PRD/v74/e104006>

(Cited on pages 14, 112, 124, 125, and 126.)

- [231] Mark A. Scheel, Oliver Rinne, and Luisa T. Buchman, personal communication, 2009.
(Cited on page 76.)
- [232] Jeremy D. Schnittman et al., *Anatomy of the binary black hole recoil: A multipolar analysis*, *Phys. Rev. D* **77** (2008), no. 4, 044031.
<http://link.aps.org/abstract/PRD/v77/e044031>
(Cited on pages 112 and 228.)
- [233] Carlos F. Sopuerta, Nicolas Yunes, and Pablo Laguna, *Gravitational recoil velocities from eccentric binary black hole mergers*, *Astrophys. J. Lett.* **656** (2007), no. 1, L9–L12.
<http://www.journals.uchicago.edu/doi/abs/10.1086/512067>
(Cited on page 112.)
- [234] Ulrich Sperhake, *Binary black-hole evolutions of excision and puncture data*, *Phys. Rev. D* **76** (2007), no. 10, 104015.
<http://link.aps.org/abstract/PRD/v76/e104015>
(Cited on page 112.)
- [235] Hans Stephani, *General relativity: An introduction to the theory of the gravitational field*, second ed., Cambridge University Press, 1990.
(Cited on page 77.)
- [236] Hans Stephani, Dietrich Kramer, Malcolm MacCallum, Cornelius Hoense-laers, and Eduard Herlt, *Exact solutions of Einstein's field equations*, second ed., Cambridge University Press, 2003.
(Cited on page 346.)

- [237] John M. Stewart, *The Cauchy problem and the initial boundary value problem in numerical relativity*, *Class. Quant. Grav.* **15** (1998), no. 9, 2865–2889.
<http://stacks.iop.org/0264-9381/15/2865>
(Cited on page 125.)
- [238] Béla Szilágyi, Denis Pollney, Luciano Rezzolla, Jonathan Thornburg, and Jeffrey Winicour, *An explicit harmonic code for black-hole evolution using excision*, *Class. Quant. Grav.* **24** (2007), no. 12, S275–S293.
<http://stacks.iop.org/0264-9381/24/S275>
(Cited on page 112.)
- [239] Béla Szilágyi, Bernd Schmidt, and Jeffrey Winicour, *Boundary conditions in linearized harmonic gravity*, *Phys. Rev. D* **65** (2002), no. 6, 064015.
<http://link.aps.org/abstract/PRD/v65/e064015>
(Cited on page 125.)
- [240] Béla Szilágyi and Jeffrey Winicour, *Well-posed initial-boundary evolution in general relativity*, *Phys. Rev. D* **68** (2003), no. 4, 041501.
<http://link.aps.org/abstract/PRD/v68/e041501>
(Cited on page 125.)
- [241] Takahiro Tanaka, Hideyuki Tagoshi, and Misao Sasaki, *Gravitational waves by a particle in circular orbits around a Schwarzschild black hole—5.5 post-Newtonian formulation*, *Prog. Theor. Phys.* **96** (1996), no. 6, 1087.
<http://ptp.ipap.jp/link?PTP/96/1087/>
(Cited on pages 239 and 286.)

- [242] Kip S. Thorne, *Multipole expansions of gravitational radiation*, *Rev. Mod. Phys.* **52** (1980), no. 2, 299–339.
<http://link.aps.org/abstract/RMP/v52/p299>
(Cited on pages 95, 96, 97, 161, and 337.)
- [243] Hugo Wahlquist, *The Doppler response to gravitational waves from a binary star source*, *Gen. Relativ. Gravit.* **19** (1987), no. 11, 1101–1113.
<http://www.springerlink.com/content/k472327452285616>
(Cited on page 187.)
- [244] Robert M. Wald, *General relativity*, University of Chicago Press, 1984.
(Cited on page 346.)
- [245] Samuel J. Waldman (for the LIGO Science Collaboration), *Status of LIGO at the start of the fifth science run*, *Class. Quant. Grav.* **23** (2006), no. 19, S653–S660.
<http://stacks.iop.org/0264-9381/23/S653>
(Cited on pages 5, 113, and 218.)
- [246] Roald K. Wangsness, *Electromagnetic fields*, second ed., John Wiley and Sons, Inc., 1986.
(Cited on page 95.)
- [247] Joel M. Weisberg and Joseph H. Taylor, *Relativistic binary pulsar B1913+16: Thirty years of observations and analysis*, *Binary Radio Pulsars* (Fred A. Rasio and Ingrid H. Stairs, eds.), *Astronomical Society of the Pacific Conference Series*, vol. 328, July 2005.

<http://arxiv.org/abs/astro-ph/0407149v1>

(Cited on page 5.)

- [248] Eugene P. Wigner and James J. Griffin, *Group theory and its application to the quantum mechanics of atomic spectra*, Academic Press, Inc., 1959.

(Cited on page 350.)

- [249] Clifford M. Will and Alan G. Wiseman, *Gravitational radiation from compact binary systems: Gravitational waveforms and energy loss to second post-Newtonian order*, *Phys. Rev. D* **54** (1996), no. 8, 4813–4848.

<http://link.aps.org/abstract/PRD/v54/p4813>

(Cited on page 188.)

- [250] Benno Willke et al., *The GEO 600 gravitational wave detector*, *Class. Quant. Grav.* **19** (2002), no. 7, 1377–1387.

<http://stacks.iop.org/0264-9381/19/1377>

(Cited on page 6.)

- [251] Alan G. Wiseman, *Coalescing binary systems of compact objects to (post)^{5/2}-Newtonian order. IV. The gravitational wave tail*, *Phys. Rev. D* **48** (1993), no. 10, 4757–4770.

<http://link.aps.org/abstract/PRD/v48/p4757>

(Cited on page 164.)

- [252] James W. York, Jr., *Kinematics and dynamics of general relativity*, *Sources of Gravitational Radiation* (Larry L. Smarr, ed.), Cambridge University Press, Cambridge, England, 1979, 83–126.

(Cited on pages 16 and 17.)

- [253] ———, *Conformal “thin-sandwich” data for the initial-value problem of general relativity*, *Phys. Rev. Lett.* **82** (1999), no. 7, 1350–1353.
<http://link.aps.org/abstract/PRL/v82/p1350>
(Cited on page 120.)
- [254] Nicolás Yunes and Emanuele Berti, *Accuracy of the post-Newtonian approximation: Optimal asymptotic expansion for quasicircular, extreme-mass ratio inspirals*, *Phys. Rev. D* **77** (2008), no. 12, 124006.
<http://link.aps.org/abstract/PRD/v77/e124006>
(Cited on page 283.)
- [255] Frank J. Zerilli, *Effective potential for even-parity Regge–Wheeler gravitational perturbation equations*, *Phys. Rev. Lett.* **24** (1970), no. 13, 737–738.
<http://link.aps.org/abstract/PRL/v24/p737>
(Cited on page 73.)