

**School of Physics
and Astronomy**



Towards a precision description of precessing
black-hole-binaries

E. Z. Hamilton

Submitted for the degree of Doctor of Philosophy
School of Physics and Astronomy
Cardiff University

28th August 2020

Summary of thesis

There have been a number of gravitational wave events detected by the LIGO and Virgo detectors in recent years. Aside from simply detecting these signals we want to be able to make confident statements about the properties of the systems that emit them— such as the masses and spins of the compact objects that make up the binaries that produce gravitational waves. Inferring these properties enables us to draw conclusions about the population of compact binaries in the universe and their formation mechanisms. To do this we compare the detected signals with theoretical predictions of the signals from sources with known properties.

The aim of this thesis has been to provide a precision description of precessing black-hole-binary systems and the gravitational wave signals they produce. Due to the non-linearity of the Einstein equation, we can only obtain analytical solutions for the gravitational waves emitted by a source while the two black holes are inspiralling towards each other. Obtaining the waveforms during merger and ringdown requires numerical relativity.

In order to explore the phenomenology of precessing signals we produced a catalogue of single spin numerical relativity simulations spanning the precessing parameter space up to mass ratio 8 and spin magnitude 0.8. The waveforms from these simulations can be used for direct comparisons with detected signals or for the development of waveform models. The catalogue presented here provides the basis of the tuned precessing model that forms the bulk of this thesis.

We also further developed a method to reliably specify the orientation of a precessing binary at a given point in the waveform. For waveforms extracted from numerical relativity simulations this method also gives us a way of estimating the time shift between the waveform and dynamics data, thus giving us a way to specify the spins at the same point in the waveform. This is useful both for performing a direct comparison between a numerical waveform and a detected signal and for developing a tuned model of precessing systems based on numerical waveforms.

It is too computationally expensive to perform the number of numerical simulations required to densely sample the parameter space of precessing binaries. We therefore produced a phenomenological model of the signal from precessing binaries. This model is based on the idea that it is possible to “twist up” a non-precessing waveform in order to get a precessing waveform. The model presented here focusses on modelling the precession effects rather than the non-precessing waveform. We used a model from post-Newtonian theory for the precession effects during inspiral and produced a phenomenological model for the effects during merger and ringdown. This phenomenological model has been tuned to the catalogue of numerical simulations described above.

Collaborative work

Parts of this thesis are the result of collaborative work.

- Chapter 2 presents a catalogue of 40 new numerical relativity simulations in addition to a further nine simulations at differing resolutions. I produced approximately one-fifth of these simulations and performed the accuracy analysis presented in this chapter. This will be published as part of a forthcoming paper presenting the catalogue.
- Chapter 3 is based on work published in [82] in collaboration with Mark Hannam. I was the lead author on this paper.
- Chapter 4 presents a tuned waveform model of precessing systems based on the catalogue presented in Chapter 2. I led the development and assessment of this model. This work should be published presently.

Acknowledgements

First and foremost I would like to thank my supervisor Mark Hannam. Thank you for your patience, your advice and your guidance. Thank you for the time and effort you put into your students– and for your enthusiasm and sense of humour while doing it. I consider myself very fortunate to have had such an excellent supervisor.

I would also like to thank the people who have joined me staring into the abyss of precessing systems and their many subtleties and complications. Lionel, Jonathan, Sebastian and Ed– thank you for listening to my frustrations and for your insights. Thank you as well to the other students and post docs who make up the gravity group, who have made the last four years both fun and memorable– from the chair challenge to board games night. Thanks in particular to Iain, Dave and Charlie for listening to me rant and making me laugh.

I would of course like to thank my family. My parents for their encouragement of my decisions, for never putting pressure on me to choose a particular path and for never stopping trying to understand this research. My sister, Catja, for her companionship– partner in mischief and silliness, provider of never-ending support. All of you for putting up with my unexpected presence over the last few months as a result of being locked down together during the pandemic.

Many thanks go to my grandfather, Gordon, for instilling in me such a passion for study and the pursuit of knowledge. You always supported my every decision and interest– even when so different from your own. (I will never forget our conversation where you proudly declared that of course you knew the elements– earth, air, fire and water). Your support for and interest in my studies have meant everything to me– along with your (to me almost inexplicable) delight that they should take place in Cardiff.

I would finally like to thank the friends who have helped me pull this together– many of whom have been mentioned above. Special thanks go to Mariana who has gone above and beyond what I could have expected.

Contents

1	Introduction	1
1.1	Gravitational Waves	2
1.1.1	Linearised theory	2
1.1.2	Quadrupole formula	3
1.1.3	Post-Newtonian theory	6
1.1.4	Spherical harmonics	7
1.2	Numerical Relativity	7
1.2.1	3+1 decomposition	7
1.2.2	Generalised harmonic co-ordinates	8
1.2.3	Waveform extraction	9
1.2.4	Numerical codes	9
1.3	Gravitational wave sources	10
1.3.1	Black hole binaries	11
1.3.2	Ground based interferometers	12
1.4	Precession Angles	13
1.4.1	Defining a co-precessing frame	14
1.4.2	The minimal-rotation frame	18
1.4.3	Physics in the minimal-rotation frame	20
1.5	Models	21
1.5.1	Waveform families	21
1.5.2	PhenomD	22
1.5.3	PhenomP	24
1.6	Matches	31
1.6.1	Non-precessing matches	31
1.6.2	Precessing matches	32
1.6.3	Power weighted matches	36
2	BAM catalogue	38
2.1	Introduction	38
2.2	Catalogue	39
2.3	Accuracy assessment waveforms	42
2.4	Quantifying Waveform Accuracy	43
2.4.1	Richardson Extrapolation	43
2.4.2	Convergence	44
2.4.3	Combining mismatches	46
2.5	Waveform errors	46
2.6	Matches	49
2.6.1	Convergence order	52
2.6.2	Extrapolation	54

2.7	Conclusion	57
3	Orbital dynamics	59
3.1	Introduction	59
3.2	Frame conventions	61
3.3	Determining $\hat{\mathbf{L}}_{\text{N}}$ and $\hat{\mathbf{n}}$	64
3.3.1	Determining $\hat{\mathbf{L}}$ from the waveform	65
3.3.2	Determining orbital phase and $\hat{\mathbf{n}}$ from the waveform	66
3.3.3	Determining the co-precessing orbital phase	67
3.3.4	Code Conventions	67
3.4	Co-ordinate ambiguities	69
3.4.1	Retarded time	69
3.4.2	Orbital plane nutation	71
3.5	Numerical Comparisons	75
3.6	Conclusion	78
4	Tuned precessing model	80
4.1	Introduction	80
4.2	Model Outline	81
4.2.1	PhenomP: Underlying Principles	81
4.2.2	Tuned angle model: A Summary	83
4.3	Calculating NR precession angles	85
4.3.1	Data processing	86
4.3.2	Conventions for feature alignment of NR angles	86
4.4	Inspiral angles	88
4.4.1	Higher-order PN corrections to β	89
4.4.2	Two spin β	92
4.5	Merger-ringdown angle model	94
4.5.1	Alpha	94
4.5.2	Beta	97
4.6	The phenomenological co-efficients	97
4.6.1	The fitting procedure	97
4.6.2	Key features of the merger-ringdown model	100
4.7	Connecting inspiral angles to merger-ringdown expressions	105
4.7.1	Connection method	105
4.7.2	Full IMR expressions	107
4.8	Model validation: qualitative behaviour	110
4.8.1	Hierarchy in the turnover frequency of the $\ell = 2$ modes	110
4.8.2	Time domain comparisons	111
4.8.3	Behaviour beyond calibration region	111
4.9	Model validation: Matches	113
4.9.1	Verification waveforms	113
4.9.2	Matches: Accuracy of the angle model	113
4.9.3	Matches: Accuracy of the full model	119
4.10	Conclusion	123
5	Conclusion	124
	Appendices	130

List of Figures

1.1	Idealised sensitivity curve for aLIGO at design sensitivity.	12
1.2	The three Euler angles $\{\alpha, \beta, \gamma\}$ which make up the precession angles which describe the transformation from an inertial frame into the co-precessing frame. There are different choices for the definition of \mathbf{V} ; the maximum emission direction, the Newtonian orbital angular momentum and varying orders of the post-Newtonian orbital-angular momentum.	14
2.1	The value δ as given by equation 2.15 as a function of convergence order. The left panel shows the results for waveforms of differing resolution. The right panel shows the results for waveforms at differing extraction radii. The configuration $q = 8, \chi = 0.8, \theta = 150^\circ$ was used in this analysis.	47
2.2	Variation of the absolute error in the time domain co-precessing phase with simulation resolution. The phases have been aligned at merger and are measured in radians.	48
2.3	Variation of the absolute error in the time domain co-precessing phase with extraction radius. The phases have been aligned at merger and are measured in radians.	48
2.4	Mismatch between waveforms at varying resolution against the high resolution waveform as a function of the total mass of the system.	50
2.5	Mismatch between waveforms at varying extraction radii and the waveform extracted at $R_{\text{ext}} = 90\text{M}$ as a function of the total mass of the system.	51
2.6	Mismatches demonstrating fourth order convergence of the BAM code with respect to resolution. The solid lines show the calculated mismatch while the dotted red lines show the predicted mismatch for varying convergence order. This is for the $q = 8, \chi = 0.8, \theta = 150^\circ$ configuration.	52
2.7	Mismatches demonstrating first order convergence of the BAM code with respect to extraction radius. The solid lines show the calculated mismatch, while the dotted red line shows the predicted mismatch for the pair of waveforms shown in purple based on the mismatch in blue, assuming first order convergence.	53

2.8	Variation of the mismatch with resolution for a $q = 8$, $\chi = 0.8$, $\theta = 150^\circ$ system with a total mass of 100M. The mismatches shown are with respect to the high resolution run and so the mismatch is 0 at $N = 120$. The line shows the relation in equation 2.16 with κ_{res} assuming fourth order convergence calculated using equation 2.17 and is consistent with all resolutions except the lowest one at $N = 80$	55
2.9	Projected mismatch between a waveform extracted at a resolution of $N = 96$ and one that is infinitely well resolved.	56
2.10	Variation of the mismatch with extraction radius for a system with a total mass of 100M. The mismatches shown are with respect to the waveform extracted at $R_{\text{ext}} = 90\text{M}$. The line shows the fit based on equation 2.15 assuming first order convergence. It is consistent with all extraction radii except at $R_{\text{ext}} = 50\text{M}$	56
2.11	Projected mismatch between a waveform extracted at a radius of $R_{\text{ext}} = 90\text{M}$ and one extracted infinitely far away.	57
3.1	The binary properties in the NR simulation frame (black) at a time t_{ref} . The binary is then rotated to the LAL source frame (blue) where $\hat{\mathbf{z}}$ is parallel to the (Newtonian) orbital angular momentum $\hat{\mathbf{L}}_{\text{N}}$ at time t_{ref} and $\hat{\mathbf{x}}$ is aligned along $\hat{\mathbf{n}}$	62
3.2	Simulation SXSBBH0152 ($q = 1$, $M\omega_{22}^{\text{start}} = 0.0297$, $\chi_1 = (0, 0, 0.6) = \chi_2$). The left hand plot shows the difference between the orbital phase estimate from the gravitational wave signal, Φ , and the orbital phase ϕ_{orb} for each of the time shifts shown in the right hand plot.	70
3.3	The same quantities as in figure 3.2, but for the BAM q8 simulation ($q = 8$, $M\omega_{22}^{\text{start}} = 0.0625$, $\chi_1 = (0, 0, 0.756)$, $\chi_2 = (0, 0, 0.0945)$).	70
3.4	θ_L calculated from the Newtonian orbital angular momentum (i.e. the normal to the orbital plane), the post-Newtonian orbital angular momentum, and from ψ_4 and the gravitational wave strain. (See text for discussion).	72
3.5	Difference between the calculation of θ_L and φ_L from ψ_4 , and that calculated from $\hat{\mathbf{L}}_{\text{N}}$ (black line), $\hat{\mathbf{L}}_{\text{PN}}$ (red line) and $\hat{\mathbf{L}}_{\text{h}}$ (dashed blue) for a post-Newtonian waveform with $q = 3$, $\chi = 0.75$ on the larger black hole, on average in the orbital plane.	72
3.6	PN waveform ($q = 3$, $\chi = 0.75$ on the larger black hole, on average in the orbital plane). A comparison of half the co-precessing gravitational wave phase Φ and the orbital phase ϕ_{orb}	74
3.7	BAM q1 ($q = 1$, $M\omega_{22}^{\text{start}} = 0.0354$, $\chi_1 = (0, -0.8, 0)$, $\chi_2 = 0$). On the left is shown the evolution of θ_L and φ_L as calculated for the Newtonian orbital angular momentum direction from the dynamics and for the direction of maximum emission from the waveform. The right hand side shows the difference between the quantities calculated from the dynamics and those calculated using ψ_4 (red) and those calculated using strain (blue). The time shift used for this comparison is that obtained by aligning the orbital and signal phases.	75
3.8	A comparison of the orbital phase Φ estimated from the gravitational wave signal, and the orbital phase ϕ_{orb} calculated from the dynamics, for the four precessing waveforms listed in table 3.1. The comparison was made using the time shifts provided by the group that produced each NR simulation.	77

3.9	A comparison of $\hat{\mathbf{n}}$ calculated from the dynamics (solid black) and using the waveform phase (dashed red) for the four precessing waveforms listed in table 3.1. This comparison was made using the time shifts provided by the group that produced each NR simulation.	77
3.10	A comparison of the difference between θ and φ (shown in figure 3.9) calculated from the dynamics and from the waveform for the four precessing waveforms listed in table 3.1. This comparison was made using the time shifts provided by the group that produced each NR simulation.	78
4.1	Comparison of the post-Newtonian expressions for each of the precession angles (blue line) with the NR data (black dots) for the case with $(q, \chi, \theta) = (8, 0.8, 60^\circ)$	89
4.2	Inclination angles for the $(q = 8, \chi = 0.8, \theta = 60^\circ)$ configuration. Shown is the NR inclination angle of the QA frame, β (black dots); the PN inclination angle of the orbital plane, ι (solid blue) and the approximate QA angle β as a function of ι (dashed magenta). See text for details.	90
4.3	Various options for the PN expression for the opening angle. The left-hand panel shows the PN value of ι for a two spin system (blue) and for the equivalent single spin system (green) calculated using the expressions used in PhenomPv3. In light blue is shown the effect of tapering the two-spin oscillations to the single-spin value at the connection frequency f_c , shown as a grey vertical line. In the right-hand panel the value for β used in the model (pink) is compared with the NR value of β found for this case. The configuration shown is Case 24 in table 4.1.	94
4.4	Comparison of the phenomenological ansatz presented in equation 4.33 with the NR data over the frequency range to which the coefficients in the ansatz were tuned for a selection of the cases which comprise the NR catalogue.	95
4.5	Comparison of the phenomenological ansatz presented in equation 4.34 with the NR data over the frequency range to which the coefficients in the ansatz were tuned for a selection of the cases which comprise the NR catalogue.	96
4.6	Comparison of the fits for each of the co-efficients for the ansatz for α given in equation 4.33 with the co-efficients found from the data as described in section 4.6.1. The fits are shown as two-dimensional surfaces covering the parameter space described by η and $\cos\theta$. On the left in blue are the fits for the simulations with $\chi = 0.4$ and on the right in red are the fits for $\chi = 0.8$. Above each of these surfaces are shown the residuals.	98
4.7	Comparison of the fits for each of the co-efficients for the ansatz for β given in equation 4.34 with the co-efficients found from the data as described in section 4.6.1. The fits are shown as two-dimensional surfaces covering the parameter space described by η and $\cos\theta$. On the left in blue are the fits for the simulations with $\chi = 0.4$ and on the right in red are the fits for $\chi = 0.8$. Above each of these surfaces are shown the residuals.	99

4.8	Variation in the location of the centre of the Lorentzian dip seen in α across the three-dimensional parameter space over which the modelling was performed.	100
4.9	Variation in the height of the Lorentzian dip seen in α across the three-dimensional parameter space over which the modelling was performed.	101
4.10	Variation in the width of the Lorentzian dip seen in α across the three-dimensional parameter space over which the modelling was performed.	101
4.11	Variation in the amplitude of the merger-ringdown part of β across the three-dimensional parameter space over which the modelling was performed.	102
4.12	Variation in the location of the inflection point in the merger-ringdown part of β across the three-dimensional parameter space over which the modelling was performed.	103
4.13	Possible morphologies of the ansatz given by equation 4.34 depending on the values taken by the co-efficients in different regions of the parameter space. From left to right the panels show systems with $(q, \chi, \theta) = (8, 0.2, 155)$, $(2.5, 0.4, 90)$ and $(5, 0.8, 160)$. The red dots mark the extrema, the green crosses show the inflection points and the blue dot indicates the inflection point chosen as described in section 4.6.2. The points of maximum curvature around this inflection point are shown by the black lines, which give a measure of the width of the turnover.	103
4.14	Variation in the width of the turnover of β across the three-dimensional parameter space over which the modelling was performed.	104
4.15	Variation in the final value towards which β tends during ringdown across the three-dimensional parameter space over which the modelling was performed.	105
4.16	Comparison of the complete model for each of the precession angles (red line) with the NR data (black dots) for the case with $(q, \chi, \theta) = (1, 0.4, 30)$. The vertical black lines show the connection frequencies for α and β	109
4.17	Comparison of the complete model for each of the precession angles (red line) with the NR data (black dots) for the case with $(q, \chi, \theta) = (8, 0.8, 60)$. The vertical black lines show the connection frequencies for α and β	109
4.18	A comparison of the $\ell = 2$ modes between the NR data (shown in black on both panels) with PhenomPv3 (top panel) and the new model (bottom panel). PhenomPv3 does not show the hierarchy in the turnover frequency if each of the $\ell = 2$ modes, which is now seen in our new model. The system plotted here has $q = 4$, $\chi = 0.4$, $\theta = 90^\circ$ with a total mass of 100 solar masses at a distance of 100Mpc.	110
4.19	Comparison of the time domain precession angles from both the PhenomPNR and PhenomPv3 models with the NR data. These angles are for the case with $(q, \chi, \theta) = (4, 0.4, 60)$	112
4.20	Comparison of the time domain precession angles from both the PhenomPNR and PhenomPv3 models with the NR data. These angles are for the case with $(q, \chi, \theta) = (8, 0.4, 30)$	112

4.21	Comparison of the value of α from the PN expression with the value calculated from the NR waveform. In order to see a region over which the two values agree well we would need a longer NR waveform. In order to have a complete IMR model for α in this case we would also require a model of the ‘intermediate’ region which was tuned to the NR data.	115
4.22	Precessing matches for each of the 40 NR waveforms used to calibrate the model, averaged over inclination. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (blue) and twisted up with the angle model used by PhenomPv3 (green). The plots show cases with $\chi = 0.4$ on the left and $\chi = 0.8$ on the right. Starting from the top, each of the rows show $q = 1, 2, 4$ and 8 respectively.	116
4.23	Precessing matches for four of the 40 NR waveforms used to calibrate the model. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (blue) and twisted up with the angle model used by PhenomPv3 (green).	117
4.24	Precessing matches for the single spin verification NR waveforms, averaged over inclination. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (blue) and twisted up with the angle model used by PhenomPv3 (green).	118
4.25	Precessing matches for the two spin verification NR waveforms, averaged over inclination. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (blue) and twisted up with the angle model used by PhenomPv3 (green).	118
4.26	Precessing matches for each of the 40 NR waveforms used to calibrate the model, averaged over inclination. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and complete precessing model– PhenomPNR is shown in blue and PhenomPv3 is shown in green.	120
4.27	Precessing matches for each of the 40 NR waveforms used to calibrate the model, averaged over inclination. These matches are between the unsymmetrised NR waveforms in the \mathbf{J} -aligned frame and complete precessing model– PhenomPNR is shown in blue and PhenomPv3 is shown in green.	121
4.28	Non-precessing matches for each of the 40 NR waveforms used to calibrate the model. These matches are between the symmetrised co-precessing NR waveforms in the \mathbf{J} -aligned frame and modified PhenomD which provides the co-precessing model employed by PhenomP versions 1-3.	122

List of Tables

2.1	Properties of twenty of the waveforms included in the catalogue with mass ratios 1 and 2.	40
2.2	Properties of twenty of the waveforms included in the catalogue with mass ratios 4 and 8.	41
2.3	Properties of the various resolutions at which the simulations were performed.	43
2.4	Relative error in the waveform quantities compared with the Richardson extrapolated quantities for the $q = 8$, $\chi = 0.8$, $\theta = 150^\circ$ configuration.	49
3.1	List of simulations used in numerical comparisons. $M\omega_{22}^{\text{start}}$ gives the gravitational wave frequency at the start of the waveform. The number of precession cycles is estimated from the number of orbits \mathbf{L}_N completes around $\mathbf{J} = \mathbf{L}_N + \mathbf{S}$	76
4.1	Additional configurations used to verify the accuracy of the model in cases where it had not been tuned.	114
A.1	Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are up to linear order in up to two parameters.	131
A.2	Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are up to second order in one of the parameters.	132
A.3	Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are up to second order in up to two parameters.	133
A.4	Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are up to third order in one of the parameters.	134
A.5	Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are first or third order in each of the parameters.	135
A.6	Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are second or third order in two of the parameters.	136
A.7	Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are fourth order in spin angle and linear in symmetric mass ratio.	137

A.8 Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are fourth order in spin angle and up to third order in symmetric mass ratio. 138

Everything has a meaning, if only we could read it.
Philip Pullman

Chapter 1

Introduction

The General Theory of Relativity [71] was published in 1915 and describes how gravity arises as a geometric property of space-time and relates the curvature of this space-time to the energy and momenta of the matter and radiation contained within it. Gravitational waves are disturbances to space-time as a result of the motion of matter or radiation. One possible source of gravitational waves are accelerating massive bodies; known astrophysical sources are compact objects (such as black holes) orbiting each other and coalescing. The first detection of such a source was made by the LIGO and Virgo collaborations in 2015 [7] and there have since been a number of further detections [6, 9, 10, 13–16].

The detection of gravitational waves is greatly assisted by models of the signals being searched for. Since General Relativity is a highly non-linear theory, it can only be solved analytically for very simple systems and analytic approximations can only be developed under certain simplifying assumptions. In order to accurately determine the signals emitted by complicated sources we therefore require numerical solutions of the Einstein equation, but this is very computationally expensive. It is therefore necessary to develop models of the signals based on available analytic and numerical data. The family of models on which I will focus in this thesis are *phenomenological* models. We need models which describe the signals from all the different types of binary systems which can be observed, including systems in which the black holes are spinning in ways such that the orbital plane of the binary precesses.

The intention of this thesis is to provide a precision description of the gravitational waves emitted by precessing black hole binaries. In Chapter 2 I present and discuss the accuracy of a catalogue of numerical simulations of precessing binaries upon which the rest of the work presented in this thesis is based. In Chapter 3 I present a method of reliably relating the dynamical information extracted from a numerical simulation to the waveform information, enabling the exact configuration of the binary to be stated at a given point in the waveform. In Chapter 4 I present a model of the precessional motion of black hole binary systems.

In the rest of this introduction I will discuss gravitational waves and methods of modelling and detecting them, as well as possible sources. In section 1.1 I will introduce gravitational waves and show how they emerge from General Relativity followed by some discussion of analytic treatments of gravitational waves. In section 1.2 I will then discuss how the nonlinear equations of General Relativity can be solved numerically and various numerical codes which are used to do this. Following this, section 1.3 is a description of various sources of gravitational waves and methods of detecting them. Section 1.4 contains a more in depth description of precessing systems in particular, while section 1.5 describes attempts to model such sources so far. Finally, section 1.6 describes the method commonly used for quantifying the accuracy of such models.

1.1 Gravitational Waves

Gravitational waves are ripples in the curvature of space-time. They are described by the Einstein equation [70, 72]

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{8\pi G}{c^4}T_{\mu\nu}, \quad (1.1)$$

where $R_{\mu\nu}$ is the Ricci tensor, $g_{\mu\nu}$ is the metric tensor, $T_{\mu\nu}$ is the stress-energy tensor, G is the gravitational constant and c is the speed of light in a vacuum. This equation relates the geometry of space-time to matter. Gravitational waves can most simply be seen to emerge from General Relativity in linearised theory [73, 74]. Linearised theory also allows for the derivation of some of the key features of gravitational waves, such as the leading order term in the multipole expansion of gravitational waves and the angular distribution of the radiation. To go beyond linearised theory requires post-Newtonian theory, which is applicable to slow-moving sources in the weak-field region. To solve the Einstein equation in full General Relativity without any such approximations requires numerical relativity, which is discussed in the next section.

1.1.1 Linearised theory

As gravitational waves propagate away from their source, they reach an asymptotic region in which they can be treated as a weak perturbation on top of a flat background metric. In this region we can clearly see the wave-like solutions to the Einstein equation. To do this we expand the Einstein equation about a flat space metric following [116], which gives us the *linearised theory* of gravity.

In linearised theory, the metric $g_{\mu\nu}$ is given by

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.2)$$

where $\eta_{\mu\nu}$ is the flat space Minkowski metric and $h_{\mu\nu}$ is the perturbation; $|h_{\mu\nu}| \ll 1$. Expanding equation 1.1 to linear order in $h_{\mu\nu}$ yields

$$\square \bar{h}_{\mu\nu} + \eta_{\mu\nu} \partial^\rho \partial^\sigma \bar{h}_{\rho\sigma} - \partial^\rho \partial_\nu \bar{h}_{\mu\rho} - \partial^\rho \partial_\mu \bar{h}_{\nu\rho} = -\frac{16\pi G}{c^4} T_{\mu\nu}, \quad (1.3)$$

where $\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h$. We now make use of the gauge freedom in general relativity to choose the *Lorentz gauge*:

$$\partial^\mu \bar{h}_{\mu\nu} = 0, \quad (1.4)$$

which reduces the Einstein equation in linearised gravity to

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}, \quad (1.5)$$

which is in the form of a simple wave equation.

Outside a source of gravitational waves, $T_{\mu\nu} = 0$ so the wave equation becomes $\square \bar{h}_{\mu\nu} = 0$. We can therefore see that gravitational waves travel at the speed of light. We can now make use of the residual gauge freedom to transform in the *transverse-traceless gauge* defined by

$$h^{0\mu} = 0, \quad h_i^i = 0, \quad \partial^i h_{ij} = 0. \quad (1.6)$$

For a plane wave propagating along the z -direction the metric perturbation is now given by

$$h_{\mu\nu}^{\text{TT}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}_{\mu\nu} \cos \left[\omega \left(t - \frac{z}{c} \right) \right], \quad (1.7)$$

where h_+ and h_\times are the plus and cross polarisations of the gravitational wave.

The metric perturbation due to gravitational waves causes a change in the proper distance between two points, which is characterised by the strain h . We define the complex strain to be given by

$$h = h_+ - ih_\times. \quad (1.8)$$

1.1.2 Quadrupole formula

From basic physical arguments we can see that the lowest order term in a multipole expansion of the gravitational wave strain must be the quadrupole term. Mass is conserved to linear order in general relativity, precluding the possibility of a monopole term while general covariance and the conservation of angular momentum to linear

order prevent the presence of a dipole term. This is a general feature and not restricted to the linearised theory. Linearised theory does however provide a useful approximation which can inform us about the key features of this quadrupole term.

Inside the source we no longer have $T_{\mu\nu} = 0$. For astrophysically interesting sources of gravitational waves, the assumptions that went into deriving the linearised theory are no longer valid. Nevertheless we assume that we are in the weak-field regime and that the sources have low velocities so we can still use the wave equation 1.5. This equation can be solved by the method of Green's functions, giving

$$h_{ij}^{\text{TT}}(t, \mathbf{x}) = \frac{4G}{c^4} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \int d^3x' \frac{1}{|\mathbf{x} - \mathbf{x}'|} T_{kl} \left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}, \mathbf{x}' \right), \quad (1.9)$$

where $\Lambda_{ij,kl}$ is the *Lambda tensor* which projects two-tensors onto a plane which is transverse to the direction of propagation $\hat{\mathbf{n}}$. It is given by $\Lambda_{ij,kl} = P_i^k P_j^l - \frac{1}{2} P_{ij} P^{kl}$ where $P_{ij} = \eta_{ij} - n_i n_j$. \mathbf{x} points from the source to the observer at the time at which the gravitational waves are emitted. \mathbf{x}' points from the source to the observer at the time at which the waves are detected.

We now make two assumptions: (i) that the gravitational waves are a long way from the source (weak-field regime) and (ii) that the wavelength of the radiation is much greater than the size of the source (slow-velocity source). The metric perturbation h_{ij} can then be expanded in terms of the momenta of the stress-energy tensor T_{kl} . These can be related to the momenta of the energy density T^{00} :

$$M = \frac{1}{c^2} \int d^3x T^{00}(t, \mathbf{x}), \quad (1.10)$$

$$M^i = \frac{1}{c^2} \int d^3x T^{00} x^i(t, \mathbf{x}), \quad \text{etc.} \quad (1.11)$$

To leading order, the metric perturbation becomes

$$h_{ij}^{\text{TT}}(t, \mathbf{x})_{\text{quad}} = \frac{1}{r} \frac{2G}{c^4} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \ddot{M}^{kl} \left(t - \frac{r}{c} \right) \quad (1.12)$$

$$= \frac{1}{r} \frac{2G}{c^4} \ddot{Q}_{ij}^{\text{TT}} \left(t - \frac{r}{c} \right), \quad (1.13)$$

where $r = |\mathbf{x}|$, $Q_{ij}^{\text{TT}} = \Lambda_{ij,kl} Q^{kl}$ and $Q^{ij} \equiv M^{ij} - \frac{1}{3} \delta^{ij} M_{kk}$, which is traceless by construction. Q^{ij} is the *quadrupole moment*.

We can find the form of the quadrupole formula for an observer far away from the source of the gravitational waves. Considering a Cartesian co-ordinate system in which the waves propagate along the z -direction,

$$P_{ij} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}_{ij}, \quad (1.14)$$

so for an arbitrary matrix A_{kl} ,

$$\Lambda_{ij,kl}A_{kl} = P_i^k P_j^l A_{kl} - \frac{1}{2} P_{ij} P^{kl} A_{km} \quad (1.15)$$

$$= (PAP)_{ij} - \frac{1}{2} P_{ij} \text{Tr}(PA) \quad (1.16)$$

$$= \left(\begin{array}{ccc} \frac{1}{2}(A_{11} - A_{22}) & A_{12} & 0 \\ A_{21} & -\frac{1}{2}(A_{11} - A_{22}) & 0 \\ 0 & 0 & 0 \end{array} \right)_{ij}. \quad (1.17)$$

Substituting this into 1.12 and comparing it to 1.7, the polarisations of a gravitational wave propagating in an inertial frame where the z -direction is aligned with the direction of propagation are given by

$$h_+ = \frac{1}{r} \frac{G}{c^4} (\ddot{M}_{11} - \ddot{M}_{22}), \quad (1.18)$$

$$h_\times = \frac{1}{r} \frac{G}{c^4} \ddot{M}_{12}. \quad (1.19)$$

If instead we consider a system in which a gravitational wave is propagating in a general direction $\hat{\mathbf{n}}$, written with respect to the z -direction using standard spherical polar co-ordinates, then this frame is related to the previous frame by the rotation

$$\mathcal{R} = \begin{pmatrix} \cos \phi & \sin \phi & 0 \\ -\sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & \sin \theta \\ 0 & -\sin \theta & \cos \theta \end{pmatrix}. \quad (1.20)$$

As before, $h_+ = \frac{1}{r} \frac{G}{c^4} (\ddot{M}'_{11} - \ddot{M}'_{22})$ and $h_\times = \frac{1}{r} \frac{G}{c^4} \ddot{M}'_{12}$ where M'_{ij} are calculated in the frame in which the gravitational wave propagates along the z -direction. In terms of the components of the momenta of the energy density in the original frame, the polarisations are given by

$$h_+ = \frac{1}{r} \frac{G}{c^4} \left[\ddot{M}_{11} (\cos^2 \phi - \sin^2 \phi \cos^2 \theta) + \ddot{M}_{22} (\sin^2 \phi - \cos^2 \phi \cos^2 \theta) \right. \\ \left. - \ddot{M}_{33} \sin^2 \theta - \ddot{M}_{12} \sin 2\phi (1 + \cos^2 \theta) \right. \\ \left. + \ddot{M}_{12} \sin \phi \sin 2\theta + \ddot{M}_{23} \cos \phi \sin 2\theta \right], \quad (1.21)$$

$$h_\times = \frac{1}{r} \frac{G}{c^4} \left[(\ddot{M}_{11} - \ddot{M}_{22}) \sin 2\phi \cos \theta + 2\ddot{M}_{12} \cos 2\phi \cos \theta \right. \\ \left. - \ddot{M}_{13} \cos \phi \sin \theta + \ddot{M}_{23} \sin \phi \sin \theta \right]. \quad (1.22)$$

This expression allows the angular distribution of the radiation to be computed if the mass distribution of the source is known.

Following [86], we consider the simple example of two Newtonian bodies of mass M orbiting each other in the $x - y$ plane with separation $2R$ and orbital frequency

ω . These bodies follow the trajectory described by

$$x(t) = R \cos \omega t, \quad (1.23)$$

$$y(t) = R \sin \omega t, \quad (1.24)$$

$$z(t) = 0. \quad (1.25)$$

The non-zero second mass moments are therefore

$$\ddot{M}_{11}(t) = \frac{1}{2} M R^2 (1 + \cos 2\omega t), \quad (1.26)$$

$$\ddot{M}_{22}(t) = \frac{1}{2} M R^2 (1 - \cos 2\omega t), \quad (1.27)$$

$$\ddot{M}_{12}(t) = \frac{1}{2} M R^2 \sin 2\omega t. \quad (1.28)$$

If we consider only the dependence on θ and so take $\phi = 0$ the two polarisations are given by

$$h_+ \propto (1 + \cos^2 \theta), \quad (1.29)$$

$$h_\times \propto \cos \theta. \quad (1.30)$$

Therefore the majority of the power in the gravitational wave is radiated in the direction perpendicular to the orbital plane. This is true for any gravitational wave source formed by two orbiting compact objects, not merely in this simple example.

1.1.3 Post-Newtonian theory

Linearised theory, as described above, assumes that the background space-time curvature and the motion of a gravitational wave source can be treated independently. Treating gravitational waves as a perturbation on top of a flat background metric implies that the motion of the source must be Newtonian. However, for astrophysical sources of gravitational waves this assumption is false. If such sources are slow moving and weakly self-gravitating, their dynamics can be treated using a post-Newtonian formalism instead.

Post-Newtonian theory [36] consists of a series of corrections to linearised theory written as an expansion in powers of $\frac{v}{c}$, where v is the typical internal speed of the gravitational wave source. It is necessary to find general-relativistic corrections to the equations of motion and to compute the gravitational waves emitted by such a source. Furthermore, gravitational waves carry energy away from the source and so influence the equations of motion. Additionally, since a gravitational field is a source of gravitational waves (as general relativity is a non-linear theory) low order terms in the post-Newtonian expansion will produce additional gravitational waves at higher order.

1.1.4 Spherical harmonics

The gravitational wave strain h is often decomposed into a set of modes using the spin -2 weighted spherical harmonics given by

$$Y_{\ell m}^s(\theta, \phi) = (-1)^m \sqrt{\frac{(\ell+m)!(\ell-m)!(2\ell+1)}{4\pi(\ell+s)!(\ell-s)!}} \sin^{2\ell}\left(\frac{\theta}{2}\right) \times \binom{\ell-s}{r} \binom{\ell+s}{r+s-m} (-1)^{\ell-r-s} e^{im\phi} \cot^{2r+s-m}\left(\frac{\theta}{2}\right). \quad (1.31)$$

These modes are given by

$$h_{\ell m} = \int_0^{2\pi} \int_0^\pi h(\theta, \phi) \overline{Y_{\ell m}^{-2}(\theta, \phi)} \sin\theta d\theta d\phi. \quad (1.32)$$

The spin-weighted spherical harmonics form an orthonormal basis. They can be written as a set of traceless symmetric tensors $\mathcal{Y}_{ij}^{\ell m}$ [159]. The five $\ell = 2$ tensors form a basis for the space of traceless symmetric tensors, such as the mass quadrupole Q_{ij} . The majority of the power radiated by gravitational waves is therefore in the $\ell = 2$ modes. From the form of the \mathcal{Y}_{ij}^{2m} we can see that for a binary in the $x - y$ plane with the gravitational wave propagating along the z -direction, the majority of the power will reside in the $(\ell, m) = (2, \pm 2)$ modes.

1.2 Numerical Relativity

The Einstein equation, given in equation 1.1, determines the dynamical evolution of a physical system. However, it can only be solved analytically for very simple systems or under a series of assumptions as described above. Solving the set of 10 differential equations for more complicated systems requires numerical relativity [30].

In order to solve the Einstein equation numerically it is necessary to cast it into a form that is suitable for numerical integration. One such method is known as the “3+1 decomposition” of the Einstein equation [165]. Another is the system of *generalised harmonic co-ordinates* [60, 135].

1.2.1 3+1 decomposition

The basic interpretation of the 3+1 decomposition is that we specify the metric components and their first order derivatives with respect to time everywhere on a three-dimensional space-like hypersurface at a constant time t . Different points on the hypersurface are labelled by their spatial co-ordinates x^i . The metric quantities can now be integrated forward in time and computed, along with their derivatives, on a new spatial hypersurface at a time $t + \delta t$.

These hypersurfaces can be thought of as *foliations* of space-time. If we consider

a 4-dimensional space-time manifold M described by the metric $g_{\mu\nu}$ then we assume that this space-time can be foliated into a family of non-intersecting space-like 3-surfaces Σ , each at constant t . We define the normal to a surface to be $n^\mu \equiv g^{\mu\nu} \nabla_\nu t$. The spatial metric induced on each hypersurface Σ is given by $\gamma_{\mu\nu} = g_{\mu\nu} + n_\mu n_\nu$. The way in which each surface is embedded in the space-time M is described by the *extrinsic curvature*, which calculates how much the direction of n^μ changes across the spatial hypersurface. We define it to be $K_{\mu\nu} \equiv \mathcal{L}_n \gamma_{\mu\nu}$, where \mathcal{L}_n is the Lie derivative along n^μ . The spatial metric $\gamma_{\mu\nu}$ can differ between two different hypersurfaces Σ by an amount proportional to $K_{\mu\nu}$.

The temporal evolution of the spatial co-ordinates are determined by two quantities: the *lapse* and the *shift vector*. The lapse α measures how much proper time elapses between neighbouring surfaces along the normal n^μ to a surface and is defined by $\frac{1}{\alpha^2} \equiv g^{\mu\nu} \nabla_\mu t \nabla_\nu t$. The shift vector β^μ measures the amount by which the spatial co-ordinates are shifted within a surface. Two points with the same spatial co-ordinates on neighbouring hypersurfaces are therefore connected by $t^\mu = \alpha n^\mu + \beta^\mu$.

In the 3+1 decomposition, the Einstein equation is replaced by the *constraint* and *evolution* equations. These are a coupled set of partial differential equations which are first order in time. The constraint equations involve only the spatial metric and the extrinsic curvature as well as their spatial derivatives and describe a three-dimensional hypersurface Σ with gravitational field data $(\gamma_{\mu\nu}, K_{\mu\nu})$ embedded in a 4-dimensional manifold M . The evolution equations determine the evolution of $(\gamma_{\mu\nu}, K_{\mu\nu})$.

In order to be able to stably solve these equations numerically they must be *well-posed*. This means that the solutions to the equations cannot increase more rapidly than exponentially. In order for a system to be well-posed, it can be shown that it must be *strongly hyperbolic* [81]. In order to avoid numerical instabilities, even exponentially growing terms must also be controlled. This further rules out some formulations of the Einstein equation. The most commonly used form of the Einstein equation in numerical relativity are the Baumgarte-Shapiro-Shibata-Nakamura (BSSN) equations [29, 153], which are strongly hyperbolic [143].

1.2.2 Generalised harmonic co-ordinates

The generalised harmonic co-ordinates method instead starts with the original 4-dimensional Einstein equation and introduces a gauge by defining some *gauge source functions*. This yields a non-linear wave equation for the space-time metric $g_{\mu\nu}$ along with a set of constraint equations. This non-linear wave equation can be integrated directly.

This approach differs to the 3+1 decomposition method in several ways [110, 135]. Firstly, the wave equation which must be evolved contains second order time derivatives rather than just the first order terms present in the evolution equations

of the 3+1 decomposition. Further, the choice of an appropriate co-ordinate system is not as clear with this approach since it is determined by the gauge source functions rather than quantities which are directly related to the geometry of the spatial slices.

1.2.3 Waveform extraction

Gravitational wave information is typically extracted from numerical simulations using the Newman-Penrose scalar,

$$\psi_4 = -C_{\alpha\beta\gamma\delta}n^\alpha\bar{m}^\beta n^\gamma\bar{m}^\delta, \quad (1.33)$$

where $C_{\alpha\beta\gamma\delta}$ is the Weyl tensor and n and \bar{m} form part of the null tetrad ℓ , n , m and \bar{m} . ℓ and n are in-going and out-going null vectors respectively while m is constructed from two spatial vectors orthogonal to ℓ and n . The choice of the definition of this tetrad is a convention which varies between numerical codes.

It is useful to be able to relate this quantity to the gravitational wave strain h . In the Bondi gauge, which applies to all the numerical codes mentioned here, the two are related by the simple relation [91]

$$\psi_4 = \ddot{h}. \quad (1.34)$$

However, performing two integrations in the time domain (in order to get h from ψ_4) causes unphysical artefacts to appear in the signal as a result of amplifications of numerical noise. The integration is therefore typically performed in the frequency domain, where one simply divides twice by the angular frequency ω (once for each integration). In order to avoid spectral leakage, *fixed frequency integration* [140] is used where one divides by ω_0 instead of ω for $|\omega| < \omega_0$.

1.2.4 Numerical codes

The two main numerical relativity codes referred to in this thesis are BAM [49] and SpEC [89, 148]. Also referenced are the GT-MAYA [87, 90, 126, 161] and RIT codes [166]. These codes simulate black-hole-binary configurations. BAM, GT-MAYA and the RIT code are all *moving puncture codes* [28, 54]. Moving puncture codes use the BSSN formulation of the 3+1 decomposition of Einstein's equations. SpEC is a pseudo-spectral code which uses a first-order representation of the generalised harmonic co-ordinates. The simulations upon which the majority of work in this thesis are based were performed using BAM, which is described in more detail below.

BAM

As noted above, BAM is a moving puncture code. The initial geometry employed by the puncture method for N black holes is the Brill-Lindquist wormhole topology [47]

with $N + 1$ asymptotically flat ends. The asymptotically flat ends are compactified and form co-ordinate singularities, labelled *punctures*. Since moving puncture codes use the 3+1 decomposition of the Einstein equation, they require initial data for the spatial metric and extrinsic curvature (γ_{ij}, K_{ij}) on a spatial hypersurface Σ with unit normal n^i . These quantities are related to their background space-time analogues by the conformal factor ψ , which diverges at the punctures. In the moving puncture approach, a new variable $\chi = \psi^{-4}$ is defined. χ vanishes at the puncture and is evolved instead of the conformal factor. This gauge choice avoids any finite-differencing of a singularity. BAM uses Bowen-York initial data [41] along with the gauge quantities $\alpha = \psi^{-2}$ and $\beta^i = 0$. The initial data is then evolved using the BSSN system with the gauge chosen to be the ‘‘covariant’’ form of ‘‘1+log’’ slicing [21, 28, 40, 54].

The numerical method employed by BAM utilises the method of lines with finite differencing in space and Runge-Kutta time stepping [133]. Berger-Oliger type adaptive mesh refinement [32] is used. The grid set up of a simulation consists of a hierarchy of nested Cartesian grids. On the lower levels there are two separate grids, each centred on a puncture. Once the grids reach a size at which they would begin to overlap, they are replaced by a single grid which encompasses both punctures and is centred on the centre of motion. The smaller grids track the punctures while the larger ones remain fixed. The location of the punctures are tracked using the shift vector. The grid spacing on level ℓ is given by $h_\ell = \frac{h_0}{2^\ell}$, where h_0 is the grid spacing of the coarsest, outermost level.

The code is expected to be fourth-order accurate. Sixth order finite differencing is used for spatial derivatives of the Einstein equation [92] while sixth order polynomial interpolation is used in the space between refinement levels. The Runge-Kutta time stepping is also fourth order. There are a few sources of second order errors: (i) the initial data solver (though this is run with sufficiently high resolution that it is not expected to impact the evolution), (ii) the implementation of the radiative boundary condition and (iii) the boundary values for intermediate time-levels which come from second order interpolation. Only the third of these sources of error is expected to be non-negligible but it is not expected to dominate at the resolution at which the simulations included in this thesis were performed.

1.3 Gravitational wave sources

There are multiple possible sources of gravitational waves, including stellar sources and processes in the early Universe [65, 144]. Different sources emit gravitational waves at different frequencies. This spectrum of frequencies can be grouped into four distinct frequency bands. The first of these is known as the *high frequency band* ($1\text{Hz} \leq f \leq 10^4\text{Hz}$). This band can be detected with current generation ground-based detectors, such as GEO [18, 69, 115], LIGO [3], Virgo [17] and Kagra [27, 155].

Possible sources that emit in this band include compact stellar-mass binaries (the most promising source), stellar core collapse, rotating neutron stars and stochastic backgrounds. The *low frequency band* ($10^{-5} \text{ Hz} \leq f \leq 1 \text{ Hz}$) may be detectable by future space-based detectors, such as LISA [23] and contains sources such as supermassive black holes and white dwarf binaries. The *very low frequency band* ($10^{-9} \text{ Hz} \leq f \leq 10^{-7} \text{ Hz}$) is emitted by non-stellar sources and may be detectable by pulsar timing arrays [101, 127, 137]. The *ultra low frequency band* ($10^{-18} \text{ Hz} \leq f \leq 10^{-13} \text{ Hz}$) is also emitted by non-stellar sources, such as quantum fluctuations in the early Universe and may be detectable through an imprint left on the cosmic microwave background [100].

In this thesis I will focus on signals from stellar sources– black-hole-binary systems in particular. The signal produced by a black-hole-binary system has three regions: inspiral, merger and ringdown. During the inspiral phase, the binary slowly spirals inwards as a result of the emission of gravitational waves and the amplitude and frequency of the signal increase. In the merger phase the objects plunge towards each other and merge. During ringdown the remnant object settles down to its final state by radiating away any remaining energy. The inspiral can be approximated analytically using post-Newtonian theory while aspects of the ringdown can be calculated using perturbation theory. However, the merger is highly non-linear. It is therefore necessary to use Numerical Relativity to accurately describe the whole signal.

1.3.1 Black hole binaries

A black hole binary consists of two black holes orbiting each other. Labelling the primary black hole 1 and the secondary black hole 2, these black holes have masses m_1 and m_2 where $m_1 \geq m_2$. The binary therefore has mass ratio $q = \frac{m_1}{m_2} \geq 1$. Each black hole can be spinning with spin angular momentum \mathbf{S}_i and the *dimensionless spin* is defined to be $\chi_i = \frac{S_i}{m_i^2}$ which obeys the Kerr limit $|\chi_i| < 1$. The binary therefore has total spin angular momentum $\mathbf{S} = \mathbf{S}_1 + \mathbf{S}_2$. The binary also has orbital angular momentum \mathbf{L} and total angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$. The binary may also have some eccentricity, but for most sources this is expected to have been radiated away a long time before the signal enters the detector’s sensitivity band.

Assuming the system evolves solely through the emission of gravitational waves, it can be described completely by 15 parameters. These can be broken down into 8 *intrinsic parameters* and 7 *extrinsic parameters*. The intrinsic parameters are $\Xi = \{M, \eta, \mathbf{S}_1\}$ where $M = m_1 + m_2$ is the total mass of the binary and $\eta = \frac{q}{(1+q)^2}$ is the symmetric mass ratio. The total mass of the binary M influences only the scale of the signal. There are therefore only 7 fundamental intrinsic parameters. The extrinsic parameters are the luminosity distance to the source, its sky position (polar and azimuthal angles), the polarisation angle, the source inclination and the

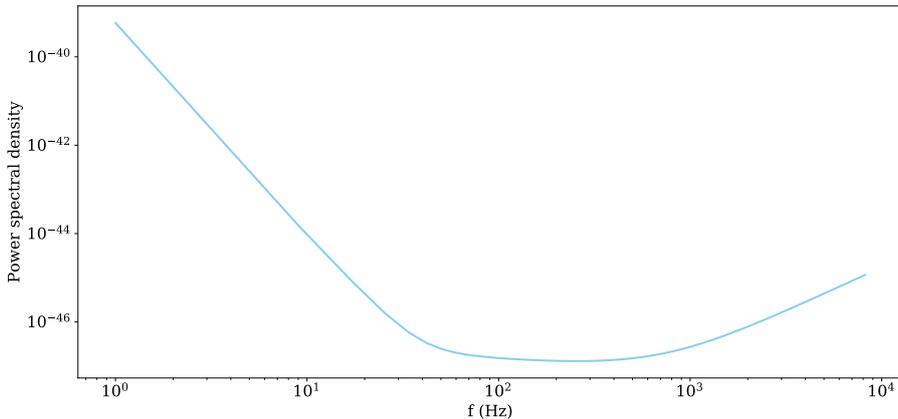


Figure 1.1: Idealised sensitivity curve for aLIGO at design sensitivity.

reference time and phase.

The total mass of the binary affects which part of the signal is “seen” in the detector band [65]. As can be seen from figure 1.1, a signal can be seen in the aLIGO sensitivity band at design sensitivity above ~ 10 Hz. We can estimate that for systems with a total mass of less than $\sim 20M_{\odot}$ most of the signal seen by the detector is expected to be in the inspiral regime. For systems with a mass between $20M_{\odot}$ and $200M_{\odot}$ the merger region dominates, while for systems with a mass above $200M_{\odot}$ but below $1000M_{\odot}$ both the merger and ringdown regions contribute equally. It is therefore necessary to have a model of the entire gravitational wave signal for analysis.

Predictions of the gravitational waveform are necessary to measure the masses and spins of the black holes which make up the binary using “matched-filtering” techniques [22]. It is too computationally expensive to produce the number of gravitational wave templates required for such analyses using Numerical Relativity simulations alone. It is therefore necessary to develop inspiral-merger-ringdown models covering all possible types of black-hole-binary systems.

1.3.2 Ground based interferometers

Compact binary coalescences, such as stellar mass black hole binaries, can be detected with ground based interferometers (as well as other types of detectors). Ground based interferometers are essentially Michelson-Morley interferometers. A laser beam is split by a beam splitter into two beams which are sent down the orthogonal arms of the detector. They are reflected by mirrors at the end of the arms and recombine again at the beam splitter. Any passing gravitational waves distort the relative length of the two arms of the interferometer and so modify the interference pattern of the returning light beams.

In practice the detectors are more complicated than a traditional Michelson-Morley interferometer. In addition to being built on kilometre length scales in order to increase the sensitivity of the detector to small distortions in the lengths of the arms, they also employ a number of sophisticated techniques in order to amplify the signal. These include reflecting the beam back and forth along each of the arms before it recombines in order to further increase the distance it travels as well as techniques for reducing various sources of noise [11] experienced by the detectors [146].

The main way in which the detectors manifest themselves in the work presented in this thesis is through the sensitivity curves of the detector. We use the sensitivity curve for Advanced LIGO at design sensitivity [114], which is shown in figure 1.1. As can be seen from this curve, at design sensitivity the detectors are expected to be sensitive to signals from 20Hz to approximately 8kHz.

1.4 Precession Angles

If the spins on the black holes are not parallel with the orbital angular momentum of the binary, the orientation of the spins as well as that of the orbital angular momentum will vary with time. This causes the orbital plane to *precess* [24, 105]. The precession of the spins and the orbital plane causes the energy in the emitted gravitational wave to be shared between the different waveform modes, resulting in modulations in the gravitational wave amplitude and oscillations in the gravitational wave frequency. This makes it very difficult to produce an analytic mode of the waveforms from precessing binaries.

For a non-precessing system we can define a fixed preferred axis as the direction along which the gravitational radiation is preferentially beamed (the “*radiation axis*”), thus defining a preferred inertial frame in which to view the system. This is not possible for a precessing system. However, it is possible to define a time- or frequency-dependent co-ordinate system which tracks the radiation axis and so tracks the precession of the system. A non-inertial co-ordinate system where the z -axis is aligned with the radiation axis defines a “*co-precessing frame*” [46, 123, 150] (i.e. a frame that precesses along with the binary).

The co-precessing frame as defined above is not unique as there remains an ambiguity in time- or frequency-dependent rotations about the radiation axis. Fixing the rotation about the radiation axis defines the “*minimal-rotation frame*” which is invariant under rotations of the inertial co-ordinates in which the gravitational wave was extracted up to a constant overall rotation [46].

The rotations between an inertial frame and the minimal-rotation frame are described using the $z-y-z$ convention by the Euler angles α, β, γ , two of which are shown in figure 1.2. Methods to calculate these angles are described below. These methods can be applied to the Weyl scalar ψ_4 , the news \dot{h} or the gravitational wave

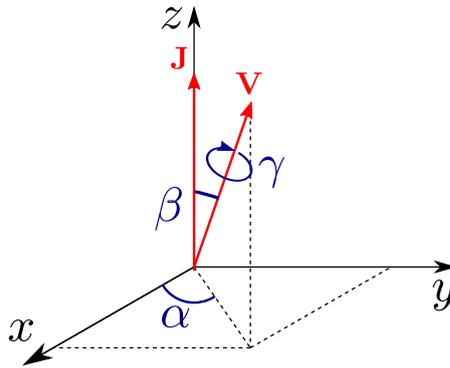


Figure 1.2: The three Euler angles $\{\alpha, \beta, \gamma\}$ which make up the precession angles which describe the transformation from an inertial frame into the co-precessing frame. There are different choices for the definition of \mathbf{V} ; the maximum emission direction, the Newtonian orbital angular momentum and varying orders of the post-Newtonian orbital-angular momentum.

strain h . Following [46] I shall use q to represent these quantities in the inertial frame in which the gravitational wave was extracted and q' to represent them in a co-precessing frame. Similarly these methods can be applied in either the time or frequency domain. For simplicity, I will use the time domain to demonstrate the methods.

1.4.1 Defining a co-precessing frame

Maximisation procedure

A method to track the optimal emission direction of a precessing binary and use this to define a preferred axis was first proposed in [150]. The following follows the treatment presented in this reference. As discussed in section 1.1.4 a gravitational wave signal can be conveniently expressed in terms of spherical harmonics, which are dominated by the $\ell = 2$ modes. For a non-precessing system where the black holes are orbiting in the $x - y$ plane, the dominant modes are the $(\ell, m) = (2, \pm 2)$ modes which are emitted approximately perpendicular to the orbital plane, defining the optimal emission direction. This is as shown in sections 1.1.2 and 1.1.4. Any rotation of the system that is not simply a rotation about this axis will cause mixing between “ m ” modes with a given ℓ value. It is therefore possible to determine the preferred direction of the system from the gravitational wave signal alone by finding the orientation that maximises the amplitude of the $(2, |2|)$ modes. In other words, the aim is to maximise the quantity

$$Q = |q'_{2,2}|^2 + |q'_{2,-2}|^2 = \sum_{m=\pm 2} |q'_{2,m}|^2. \quad (1.35)$$

The direction given by maximising this quantity is shown to be approximately along the direction of the orbital angular momentum of the system.

This quantity can be expressed in terms of inertial frame quantities as well as the two precession angles α ¹ and β as follows. First we expand q in a basis of spin-weighted spherical harmonics

$$q = \sum_{\ell,m} q_{\ell,m} Y_{\ell,m}^{-2}, \quad (1.36)$$

where the spin-weighted spherical harmonics $Y_{\ell,m}^{-2}$ are formed of a part which is dependent on the spin-basis and the regular spherical harmonics $Y_{\ell,m}$. If $\mathbf{R}(\alpha, \beta, \gamma)$ is an arbitrary rotation by the Euler angles α, β, γ then using the $z-y-z$ convention the spherical harmonics obey the transformation law

$$Y_{\ell,m}(\theta', \phi') = \sum_{m'=-\ell}^{\ell} e^{im'\alpha} d_{m'm}^{\ell}(\beta) e^{im\gamma} Y_{\ell,m'}(\theta, \phi), \quad (1.37)$$

where $d_{m'm}^{\ell}(\beta)$ are the Wigner d-matrices given by

$$\begin{aligned} d_{m'm}^{\ell}(\beta) &= \sqrt{(\ell+m)!(\ell-m)!(\ell+m')!(\ell-m')!} \\ &\times \sum_k \frac{(-1)^{k+m'-m}}{k!(\ell+m-k)!(\ell-m'-k)!(m'-m+k)!} \\ &\times \left(\sin \frac{\beta}{2}\right)^{2k+m'-m} \left(\cos \frac{\beta}{2}\right)^{2\ell-2k-m'+m}. \end{aligned} \quad (1.38)$$

An arbitrary rotation \mathbf{R} can also induce a change of spin basis. A quantity η of spin weight s obeys $\eta' = \eta e^{is\chi}$ under a change of spin basis. Therefore under an arbitrary rotation by the Euler angles α, β, γ , the spin-weighted spherical harmonics transform as

$$Y_{\ell,m}^s(\theta', \phi') = e^{is\chi} \sum_{m'=-\ell}^{\ell} e^{im'\alpha} d_{m'm}^{\ell}(\beta) e^{im\gamma} Y_{\ell,m'}(\theta, \phi). \quad (1.39)$$

Inverting equation 1.36 and substituting in the inverse of equation 1.39 gives

$$q_{\ell,m} = \sum_{m'=-\ell}^{\ell} e^{im'\alpha} d_{m'm}^{\ell}(-\beta) e^{im\gamma} q'_{\ell,m}. \quad (1.40)$$

Substituting equation 1.40 into equation 1.35 we find

$$Q(\alpha, \beta) = \sum_{m=\pm 2} \left| \sum_{m'=-\ell}^{\ell} e^{im'\alpha} d_{m'm}^{\ell}(-\beta) q_{\ell,m} \right|^2, \quad (1.41)$$

¹in [150] this angle is referred to as γ .

which depends only on the waveform in the inertial frame and the precession angles α and β .

The method proposed in [150] to find α and β then involves searching over a range of values for α and β to find the global maximum in Q . At the initial time step the direction of the Newtonian orbital angular momentum is used to provide an initial guess of the values of α and β (α_0 and β_0). A search is then performed over the range $(\alpha, \beta) = (\alpha_0 \pm 10^\circ, \beta_0 \pm 10^\circ)$ with angular resolution of 0.1° to find the values of α and β which maximise Q . At subsequent time steps the values of α and β from the previous time step are used as the initial guess and the search is repeated over the range $(\alpha, \beta) = (\alpha_0 \pm 3^\circ, \beta_0 \pm 3^\circ)$.

These angles describe a smoothly time-varying radiation axis which is approximately aligned with the direction of the orbital angular momentum of the system. More will be said on the approximate relationship between this time-varying axis and the orbital angular momentum in Chapter 3. Once the waveform from a precessing binary has been transformed into a frame in which the z -axis is aligned with the radiation axis it resembles a waveforms from a non-precessing system.

Matrix method

An alternative method to find the radiation axis was proposed in [123]. In this method one defines the symmetric tensor

$$\langle L_{(a}L_{b)} \rangle = \sum_{\ell, m, m'} q_{\ell, m'}^* \langle \ell, m' | L_{(a}L_{b)} | \ell, m \rangle q_{\ell, m}, \quad (1.42)$$

composed of the angular momentum operators L_a and their eigenstates $|\ell, m\rangle$. The radiation axis is given by the dominant principle axis of the matrix $L_{(a}L_{b)}$, which is labelled \hat{V} .

As in quantum mechanics we can define the raising and lowering operators

$$L_+ = L_x + iL_y, \quad (1.43)$$

$$L_- = L_x - iL_y, \quad (1.44)$$

which, along with L_z , act on the eigenstates as

$$L_+|\ell, m\rangle = \sqrt{\ell(\ell+1) - m(m+1)}|\ell, m+1\rangle, \quad (1.45)$$

$$L_-|\ell, m\rangle = \sqrt{\ell(\ell+1) - m(m-1)}|\ell, m-1\rangle, \quad (1.46)$$

$$L_z|\ell, m\rangle = m|\ell, m\rangle. \quad (1.47)$$

We can write

$$L^2 = L_x^2 + L_y^2 + L_z^2, \quad (1.48)$$

$$L_+^2 = L_x^2 - L_y^2 + i(L_x L_y + L_y L_x), \quad (1.49)$$

$$L_+ L_z + L_z L_+ = (L_x L_z + L_z L_x) + i(L_y L_z + L_z L_y). \quad (1.50)$$

The tensor $\langle L_{(a} L_{b)} \rangle$ can therefore be written

$$\begin{aligned} L_{(a} L_{b)} &= \begin{pmatrix} L_x^2 & \frac{1}{2}(L_x L_y + L_y L_x) & \frac{1}{2}(L_x L_z + L_z L_x) \\ \frac{1}{2}(L_x L_y + L_y L_x) & L_y^2 & \frac{1}{2}(L_y L_z + L_z L_y) \\ \frac{1}{2}(L_x L_z + L_z L_x) & \frac{1}{2}(L_y L_z + L_z L_y) & L_z^2 \end{pmatrix} \\ &= \begin{pmatrix} L^2 - L_z^2 + \text{Re}(L_+^2) & \frac{1}{2}\text{Im}(L_+^2) & \frac{1}{2}\text{Re}(L_+ L_z + L_z L_+) \\ \frac{1}{2}\text{Im}(L_+^2) & L^2 - L_z^2 - \text{Re}(L_+^2) & \frac{1}{2}\text{Im}(L_+ L_z + L_z L_+) \\ \frac{1}{2}\text{Re}(L_+ L_z + L_z L_+) & \frac{1}{2}\text{Im}(L_+ L_z + L_z L_+) & L_z^2 \end{pmatrix} \\ &= \begin{pmatrix} S_0 + \text{Re}(S_1) & \frac{1}{2}\text{Im}(S_1) & \frac{1}{2}\text{Re}(S_2) \\ \frac{1}{2}\text{Im}(S_1) & S_0 - \text{Re}(S_1) & \frac{1}{2}\text{Im}(S_2) \\ \frac{1}{2}\text{Re}(S_2) & \frac{1}{2}\text{Im}(S_2) & S_3 \end{pmatrix}, \end{aligned} \quad (1.51)$$

where

$$S_0 = \sum_{\ell, m} \left(\ell(\ell + 1) - m^2 \right) |q_{\ell, m}|^2, \quad (1.52)$$

$$S_1 = \sum_{\ell, m} c_{\ell, m} c_{\ell, m+1} q_{\ell, m+2}^* q_{\ell, m}, \quad (1.53)$$

$$S_2 = \sum_{\ell, m} c_{\ell, m} (m + 1) q_{\ell, m+1}^* q_{\ell, m}, \quad (1.54)$$

$$S_3 = \sum_{\ell, m} m^2 |q_{\ell, m}|^2, \quad (1.55)$$

and $c_{\ell, m} = \sqrt{\ell(\ell + 1) - m(m + 1)}$.

The dominant eigenvector of this matrix \hat{V} gives the radiation axis of the system, which can also be described by the precession angles α and β :

$$\hat{V} = \begin{pmatrix} \cos \alpha \sin \beta \\ \sin \alpha \sin \beta \\ \cos \beta \end{pmatrix}. \quad (1.56)$$

The precession angles are therefore given by

$$\alpha = \tan^{-1} \left(\frac{V_y}{V_x} \right) \quad \text{and} \quad \beta = \cos^{-1} \left(\frac{V_z}{|V|} \right). \quad (1.57)$$

If we restrict the tensor $\langle L_{(a} L_{b)} \rangle$ to use only the $\ell = 2$ modes, these angles are

identical to those found using the maximisation procedure.

1.4.2 The minimal-rotation frame

The final Euler angle γ needed to specify the rotation from an inertial frame into a uniquely defined frame which is independent of the initial inertial frame can be defined by imposing the “*minimal-rotation condition*”. This condition was derived in [46] following the method summarised below.

The rotation of a non-inertial frame can be described by the rotation vector $\boldsymbol{\omega}$. Relative to an inertial frame, the time derivative of any vector \mathbf{v} stationary in the rotating frame is given by $\dot{\mathbf{v}} = \boldsymbol{\omega} \times \mathbf{v}$. We want the radiation axis of the system \mathbf{a} to be stationary in the co-precessing frame, so $\dot{\mathbf{a}} = \boldsymbol{\omega} \times \mathbf{a}$. Taking the cross-product with \mathbf{a} and rearranging using the scalar triple product gives $\boldsymbol{\omega} = \mathbf{a} \times \dot{\mathbf{a}} + (\mathbf{a} \cdot \boldsymbol{\omega}) \mathbf{a}$, where the component of $\boldsymbol{\omega}$ is undetermined by this equation. In the non-precessing case $\dot{\mathbf{a}} = 0$ since the radiation axis remains fixed with respect to the inertial frame and $\boldsymbol{\omega} = 0$ so we get the condition that $\mathbf{a} \cdot \boldsymbol{\omega} = 0$. Imposing this condition even in the precessing case gives us the *minimal-rotation condition* since this condition implies $\boldsymbol{\omega}$ has the smallest possible magnitude. This condition is geometrically meaningful since \mathbf{a} is independent of the initial inertial frame.

This condition can be re-expressed in terms of generators of rotations by making use of the relationship between the Lie group of rotations of \mathbb{R}^3 about the origin $SO(3)$ and its Lie algebra $\mathfrak{so}(3)$ which are the generators of rotations. The isomorphism between $\mathfrak{so}(3)$ and standard 3-vectors \mathbb{V}^3 can be represented in the Cartesian frame as

$$\sigma : v^k \mapsto -\epsilon_{jk}^i v^k = \begin{pmatrix} 0 & -v^z & v^y \\ v^z & 0 & -v^x \\ -v^y & v^x & 0 \end{pmatrix}_j^i, \quad (1.58)$$

where $\mathbf{v} = v^k x_k$. The Cartesian basis $(\hat{x}, \hat{y}, \hat{z})$ in \mathbb{V}^3 therefore maps to (X, Y, Z) in $\mathfrak{so}(3)$ where

$$X = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix} \quad Y = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix} \quad Z = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (1.59)$$

which obey the commutation relations

$$[X, Y] = Z, \quad [Y, Z] = X, \quad [Z, X] = Y. \quad (1.60)$$

This mapping also allows us to define the following operations:

- cross product: $\sigma(\mathbf{v} \times \mathbf{w}) = [\sigma(\mathbf{v}), \sigma(\mathbf{w})] = [V, W]$,

- scalar product: $\mathbf{a} \cdot \mathbf{b} = \sigma(\mathbf{a}) \cdot \sigma(\mathbf{b}) = A \cdot B = -\frac{1}{2}A_j^i B_i^j$,
- rotation: $\mathbf{R}\mathbf{v} \mapsto \mathbf{R}\sigma(\mathbf{v})\mathbf{R}^{-1}$.

A vector \mathbf{v}_0 which remains fixed in the rotating frame is given in the inertial frame by $\mathbf{v}(t) = \mathbf{R}(t)\mathbf{v}_0$. This can be written as $V = \mathbf{R}V_0\mathbf{R}^{-1}$. The derivative of this vector is given by

$$\frac{dV}{dt} = [\mathbb{I}, V], \quad (1.61)$$

where

$$\frac{dV}{dt} = \frac{d}{dt} (\mathbf{R}V_0\mathbf{R}^{-1}) = [\dot{\mathbf{R}}\mathbf{R}^{-1}, \mathbf{R}V_0\mathbf{R}^{-1}], \quad (1.62)$$

and

$$[\mathbb{I}, V] = [\mathbb{I}, \mathbf{R}V_0\mathbf{R}^{-1}]. \quad (1.63)$$

Comparing these two expressions we find $\mathbb{I} = \dot{\mathbf{R}}\mathbf{R}^{-1}$ in order for 1.61 to hold for arbitrary \mathbf{v} .

We can define the equivalents of the instantaneous rotation axis $\boldsymbol{\omega}$ and the radiation axis \mathbf{a} as $\mathfrak{so}(3)$ objects: $\mathbb{I} = \sigma(\boldsymbol{\omega})$ and $A = \sigma(\mathbf{a})$. We can therefore rewrite the minimal rotation condition as

$$\mathbb{I} \cdot A = 0. \quad (1.64)$$

In the minimal-rotation frame the radiation axis will be along the z -axis. In the inertial frame we therefore have $\mathbf{a} = \mathbf{R}\hat{z}$ which is written as $A = \mathbf{R}Z\mathbf{R}^{-1}$ in terms of generators. Substituting in the expressions for \mathbb{I} and A into 1.64 we find

$$\mathbf{R}^{-1}\dot{\mathbf{R}} \cdot Z = 0. \quad (1.65)$$

The rotation by the Euler angles α, β, γ which transforms between an inertial frame and the minimal rotation frame is given by

$$\mathbf{R}(\alpha, \beta, \gamma) = \mathbf{R}_\alpha \mathbf{R}_\beta \mathbf{R}_\gamma = e^{\alpha Z} e^{\beta Y} e^{\gamma Z}. \quad (1.66)$$

The expression in generator notation is co-ordinate independent. If we write $\mathbf{R} = \mathbf{R}_{\alpha\beta}\mathbf{R}_\gamma$ then

$$\mathbf{R}^{-1}\dot{\mathbf{R}} = e^{-\gamma Z} \mathbf{R}_{\alpha\beta}^{-1} \dot{\mathbf{R}}_{\alpha\beta} e^{\gamma Z} + \dot{\gamma} Z, \quad (1.67)$$

meaning

$$\dot{\gamma} = - \left(\mathbf{R}_{\alpha\beta}^{-1} \dot{\mathbf{R}}_{\alpha\beta} \right) \cdot Z. \quad (1.68)$$

Expanding $\mathbf{R}_{\alpha\beta}^{-1} \dot{\mathbf{R}}_{\alpha\beta}$ we get

$$\mathbf{R}_{\alpha\beta}^{-1} \dot{\mathbf{R}}_{\alpha\beta} = \mathbf{R}_{\beta}^{-1} \left(\mathbf{R}_{\alpha}^{-1} \dot{\mathbf{R}}_{\alpha} \right) \mathbf{R}_{\beta} + \mathbf{R}_{\beta}^{-1} \dot{\mathbf{R}}_{\beta}. \quad (1.69)$$

If $R(t) \in \mathfrak{so}(3)$ is a curve such that $\mathbf{R}(t) = e^{R(t)}$ then for $R(t) \neq 0$,

$$\mathbf{R}^{-1} \dot{\mathbf{R}} = \dot{R} - \frac{1 - \cos r}{r} [R, \dot{R}] + \frac{r - \sin r}{r^3} \left[R, [R, \dot{R}] \right], \quad (1.70)$$

where r is the magnitude of the non-zero eigenvectors of \mathbf{R} . If $R(t) = 0$, $\mathbf{R}^{-1} \dot{\mathbf{R}} = 0$. For \mathbf{R}_{α} and \mathbf{R}_{β} , $R(t)$ lies along one of the Cartesian basis elements which are constant in time. Therefore $[R, \dot{R}] = 0$ so

$$\mathbf{R}_{\alpha\beta}^{-1} \dot{\mathbf{R}}_{\alpha\beta} = \dot{\alpha} e^{-\beta Y} Z e^{\beta Y} + \dot{\beta} Y. \quad (1.71)$$

For $A, B \in \mathfrak{so}(3)$, $A \neq 0$,

$$e^A B e^{-A} = B + \frac{1 - \cos a}{a^2} [A, [A, B]] + \frac{\sin a}{a} [A, B], \quad (1.72)$$

where a is the magnitude of the non-zero eigenvalues of A . When $A = 0$, $e^A B e^{-A} = B$. So,

$$e^{-\beta Y} Z e^{\beta Y} = Z - (1 - \cos \beta) [Y, [Y, Z]] - \sin \beta [Y, Z] \quad (1.73)$$

$$= (1 + \cos \beta) Z - \sin \beta Y. \quad (1.74)$$

This gives

$$\mathbf{R}_{\alpha\beta}^{-1} \dot{\mathbf{R}}_{\alpha\beta} = \dot{\alpha} \left((1 + \cos \beta) Z - \sin \beta Y \right) + \dot{\beta} Y. \quad (1.75)$$

Substituting this into equation 1.68 we find the minimum rotation condition expressed in terms of the precession angles:

$$\dot{\gamma} = - \dot{\alpha} \cos \beta. \quad (1.76)$$

1.4.3 Physics in the minimal-rotation frame

For a precessing system there are two contributions to the frequency of the binary motion. The first is due to the motion about the orbital-plane axis ω_{orb} while the second is due to the precessional motion, which has frequency $\dot{\alpha} \cos \beta$. The total frequency of the motion is therefore given by $\omega_{\text{inertial}} = \omega_{\text{orb}} + \dot{\alpha} \cos \beta$ [25]. In

the minimal rotation frame, the only contribution to the frequency comes from the motion about the radiation axis ω'_{orb} which coincides with the orbital-plane axis. ω_{orb} is related to ω'_{orb} by $\omega'_{\text{orb}} = \omega_{\text{orb}} + \dot{\gamma} = \omega_{\text{orb}} + \dot{\alpha} \cos \beta$ as the motion about the radiation axis is equal to the motion about the orbital-plane axis plus the original rotation about the orbital-plane axis by an angle γ . Therefore

$$\omega_{\text{inertial}} = \omega'_{\text{orb}}. \quad (1.77)$$

Alternatively this relation can be used to infer what the minimal rotation condition should be. If the binary in the minimum rotation frame is required to have the same angular frequency as in the initial inertial frame then the frequency of the rotation introduced by the final rotation by γ must exactly cancel that introduced by the first two rotations by α and β . Therefore we have that $\gamma = -\dot{\alpha} \cos \beta$.

1.5 Models

As noted in section 1.3, it is necessary to have a complete inspiral-merger-ringdown model of gravitational wave signals for detection and data analysis. There are currently three main families of such waveform models: phenomenological models, effective-one-body models and surrogate models.

1.5.1 Waveform families

The Phenom [20, 78, 85, 94, 102–104, 112, 131, 132, 149, 152] family of models model the key features of the gravitational wave signal using simple functions parameterised by the physical parameters of the binary. The co-efficients in these functions are obtained by fitting the functions to data from numerical simulations and then performing a global fit across the parameter space; we say these co-efficients have been “tuned” to numerical data. These waveforms are constructed in the frequency domain since this is where the data analysis is performed and are therefore fairly computationally inexpensive.

The effective-one-body (EOB) [50, 52, 63, 124, 158] models use a non-perturbative method of resumming the post-Newtonian (PN) equations of motion to obtain accurate expressions for a gravitationally interacting system of two compact objects by mapping it to an effective-one-body problem. The aim is to obtain PN expressions for the equations of motion up to merger and then attach results from perturbation theory of the ringdown. In order to improve the accuracy of the models, calibration co-efficients which have been tuned to numerical relativity are added to the expressions. EOB models are generally computationally expensive since they rely on a series of ordinary differential equations which must be integrated when evaluating the model. In order to resolve this, reduced order models [34, 77] have been made for many EOB approximants.

The surrogate models [35] interpolate between a finite set of numerical relativity waveforms, enabling a waveform to be generated at any point in the parameter space over which the original set of waveforms were generated. This is done by breaking the waveform down into “waveform data pieces” and producing a parametric fit of these pieces using a greedy algorithm. A related method can also be applied to produce reduced order models of existing waveform models in order to increase their efficiency.

1.5.2 PhenomD

PhenomD [94,103] is a model of aligned-spin black-hole-binaries. It was calibrated to 19 waveforms which cover a parameter space of up to mass ratio 18 and a dimensionless spin magnitude of 0.85 (and up to 0.98 for equal mass systems). The waveforms to which the model was calibrated are hybrids of effective-one-body (EOB) waveforms for the inspiral and NR waveforms for the merger and ringdown. Within this calibration region, the model and the NR waveforms typically agree to within 1%. A formal method of quantifying the agreement between two waveforms is defined in section 1.6. Beyond the calibration region the results provided by the model are physically reasonable, but no definitive statements can be made about their accuracy.

The black-hole-binary system being modelled is defined by its total mass $M = m_1 + m_2$, its mass ratio q (or its symmetric mass ratio η) and the spins \mathbf{S}_1 and \mathbf{S}_2 . As noted in section 1.3, the total mass sets the scale and is factored out of the modelling parameters. Since PhenomD is an aligned spin model, it is only concerned with the dimensionless spin components parallel to the orbital angular momentum \mathbf{L} given by $\chi_i = \frac{\mathbf{S}_i \cdot \hat{\mathbf{L}}}{m_i^2}$, where $\chi_i \in [-1, 1]$.

The model provides expressions for the $\ell = 2$, $|m| = 2$ spin-weighted spherical harmonics since these are the dominant modes of aligned spin systems. We can express \tilde{h}_{22} in terms of amplitude A and phase ϕ as

$$\tilde{h}_{22}(f; \Xi) = A(f; \Xi) e^{-i\phi(f; \Xi)}, \quad (1.78)$$

where $\Xi \in (M, \eta, \chi_1, \chi_2)$ are the physical parameters of the binary. PhenomD is comprised of models for A and ϕ .

The model is split into three regions; an inspiral region, a merger-ringdown region and an intermediate region. By construction each of the regions are joined together by C^1 -continuous conditions. The piecewise regions are then joined together using step functions. In the inspiral region the model is based on PN expressions and so parameterised by the masses m_1 and m_2 and dimensionless spins χ_1 and χ_2 of the binary. In the merger-ringdown region the model is parameterised by the symmetric mass ratio η and the effective spin parameter $\hat{\chi}$ where $\hat{\chi}$ is the normalised PN leading

order spin effect on the binary's phase. This is given by

$$\hat{\chi} = \left(1 - \frac{76\eta}{113}\right)^{-1} \chi_{PN}, \quad (1.79)$$

where $\chi_{PN} = \chi_{\text{eff}} - \frac{38\eta}{113}(\chi_1 + \chi_2)$ and $\chi_{\text{eff}} = \frac{m_1\chi_1 + m_2\chi_2}{m_1 + m_2}$. The final black hole is parameterised by the final mass M_f and spin a_f which come from fits to the NR.

Phase

The phase of the model is given by

$$\phi(f; \Xi) = \begin{cases} \phi_{ins} & f < 0.018 \\ \phi_{int} & 0.018 \leq f < 0.5f_{RD}, \\ \phi_{MR} & 0.5f_{RD} < f \end{cases}$$

where f_{RD} is the ringdown frequency and f is given in dimensionless units.

The inspiral phase is modelled by PN approximant TaylorF2 [26, 67, 68] plus the next four higher order PN terms where the co-efficients have been tuned to fit the hybrids used in calibrating the model. These hybrids were constructed using uncalibrated SEOBNRv2 [158]. The inspiral phase is given by:

$$\phi_{ins} = \phi_{TF2} + \frac{1}{\eta} \left(\sigma_0 + \sigma_1 f + \frac{3}{4} \sigma_2 f^{\frac{4}{3}} + \frac{3}{5} \sigma_3 f^{\frac{5}{3}} + \frac{1}{2} \sigma_4 f^2 \right). \quad (1.80)$$

The intermediate and merger-ringdown sections were modelled using the following phenomenological ansätze with the co-efficients calibrated to the NR waveforms:

$$\phi_{int} = \frac{1}{\eta} \left(\beta_0 + \beta_1 f + \beta_2 \ln(f) - \frac{1}{3} \beta_3 f^{-3} \right), \quad (1.81)$$

$$\phi_{MR} = \frac{1}{\eta} \left(\alpha_0 + \alpha_1 f - \alpha_2 f^{-1} + \frac{4}{3} \alpha_3 f^3 + \alpha_4 \arctan \left(\frac{f - \alpha_5 f_{RD}}{f_{damp}} \right) \right), \quad (1.82)$$

where f_{damp} is the ringdown damping frequency. It should be noted that it was in fact the phase derivative that was modelled and this was then integrated to get an expression for the phase.

Amplitude

The amplitude of the model is similarly given by

$$A(f; \Xi) = \begin{cases} A_{ins} & f < 0.018 \\ A_{int} & 0.018 \leq f < 0.5f_{RD} \\ A_{MR} & 0.5f_{RD} < f \end{cases}$$

The inspiral amplitude is modelled using a PN expression for the amplitude plus the next three higher order PN terms with their co-efficients tuned to fit the hybrids:

$$A_{ins} = A_{PN} + A_0 \sum_{i=1}^3 \rho_i f^{\frac{6+i}{3}}. \quad (1.83)$$

The intermediate region uses a phenomenological ansatz with the co-efficients determined using collocation methods to ensure the inspiral and merger-ringdown regions connect smoothly:

$$A_{int} = A_0 \left(\delta_0 + \delta_1 f + \delta_2 f^2 + \delta_3 f^3 + \delta_4 f^4 \right). \quad (1.84)$$

The merger-ringdown region uses a phenomenological ansatz with the co-efficients calibrated to the NR waveforms:

$$A_{MR} = A_0 \gamma_1 \frac{\gamma_3 f_{damp}}{(f - f_{RD})} e^{-\frac{\gamma_2 (f - f_{RD})}{\gamma_3 f_{damp}}}. \quad (1.85)$$

In all of the above expressions the leading order behaviour has been factored out and is given by

$$A_0 = \sqrt{\frac{2\eta}{3\pi^{1/3}}} f^{-\frac{7}{6}}. \quad (1.86)$$

1.5.3 PhenomP

The PhenomP [85, 149, 152] series of models are frequency-domain phenomenological inspiral-merger-ringdown models of precessing black-hole-binaries. These models are based on the principle described in section 1.4 that a precessing waveform can be transformed into an accelerated frame in which the waveform appears to be non-precessing. Conversely, the inverse transformation can be performed on a non-precessing waveform in order to produce a precessing waveform. Therefore the model consists of two distinct parts: a model for the underlying non-precessing aligned-spin binary (hereafter referred to as the co-precessing waveform) and a separate model of the Euler angles $\{\alpha, \beta, \gamma\}$ which capture the precessional motion.

A precessing binary is fully described by seven parameters; the mass ratio q and the six components of the two spins \mathbf{S}_1 and \mathbf{S}_2 (the total mass M sets the overall

scale and so can be factored out). PhenomP captures the basic features of this 7-dimensional parameter space with only three parameters. Two of these are the mass ratio q and the *effective inspiral spin* χ_{eff} which parameterise the underlying aligned-spin model, as discussed for PhenomD above. The third is the *precession spin parameter* χ_p [152] which parameterises the precessional motion of the binary.

The following discussion follows the treatment presented in [152]. The effective inspiral spin, which determines the inspiral rate of the binary, depends only on components of the binary's spins which are parallel to the binary's orbital angular momentum. In order to describe the precession additional information is required from the components of the spins which lie in the plane orthogonal to \mathbf{L} . The leading order PN precession equation [105] is

$$\dot{\mathbf{L}} = \frac{L}{r^3} \left[A_1 \mathbf{S}_1^\perp + A_2 \mathbf{S}_2^\perp \right] \times \hat{\mathbf{L}}, \quad (1.87)$$

where r is the separation of the binary and $A_1 = \left(2 + \frac{2}{3q}\right)$ and $A_2 = \left(2 + \frac{3q}{2}\right)$. From this we can see that the in-plane spin components drive the evolution of the orbital angular momentum and thus the orientation of the binary. The magnitude of these in-plane spin vectors oscillates around a mean value as sometimes they are parallel and add together while at other times they are anti-parallel and cancel out. The average of these contributions is then defined to be

$$S_p := \frac{1}{2} \left[\left(A_1 S_1^\perp + A_2 S_2^\perp \right) + \left| A_1 S_1^\perp - A_2 S_2^\perp \right| \right] \equiv \max \left(A_1 S_1^\perp, A_2 S_2^\perp \right). \quad (1.88)$$

Since the in-plane spin of the smaller black hole is increasingly negligible with increasing mass ratio, all the precessing spin is placed on the larger black hole. The dimensionless precession spin parameter is then defined to be

$$\chi_p = \frac{S_p}{A_1 m_1^2}. \quad (1.89)$$

These approximations do not work in two situations. The first of these is when the precession is dominated by the in-plane spin on the smaller black hole. The second is when there is little or no relative rotation of the in-plane spins (which occurs for $q \approx 1$). In this scenario, the precession spin should be the sum of the two spins so χ_p underestimates the in-plane spin contribution.

This mapping from a seven-dimensional parameter space to a three-dimensional parameter space is not bijective and thus multiple physically different configurations can give the same set of parameters. The parameters $\{q, \chi_{\text{eff}}, \chi_p\}$ therefore define an *approximate equivalence class* of precessing systems. A consequence of this is that it is difficult to distinguish the individual spin vectors in gravitational wave observations; instead it is χ_{eff} and χ_p that can be measured. An advantage of this parameter space reduction is that it reduces the number of numerical relativity

simulations required to produce a calibrated model of the precession angles to an achievable quantity.

For the majority of precessing cases, the direction of the total angular momentum of the system is approximately constant. In such cases, known as *simple precession* [24], the orbital angular momentum and total spin angular momentum precess about this direction. In a small minority of cases where the orbital and spin angular momentum are almost equal in magnitude and opposite in direction we see *transitional precession* [24]. In these cases \mathbf{J} shrinks as the orbital angular momentum is radiated away and at some point crosses the orbital plane, changing sign and thus resulting in a change in direction. The PhenomP models are based on the assumption that the direction of \mathbf{J} is fixed and deal only with simple precession. This is a reasonable assumption since transitional precession is expected to be very rare in the LIGO sensitivity band [24]. The final spin of the remnant black hole is therefore taken to be in the direction of \mathbf{J} .

The final spin of the black hole is given by

$$\mathbf{S}_{\text{fin}} = \mathbf{S}_{\text{fin}}^{\parallel} + \mathbf{S}_{\text{fin}}^{\perp}, \quad (1.90)$$

where $\mathbf{S}_{\text{fin}}^{\parallel}$ and $\mathbf{S}_{\text{fin}}^{\perp}$ are the components of the final spin parallel and perpendicular to the orbital angular momentum [97]. PhenomP assumes the parallel and perpendicular components of the initial spin are conserved but \mathbf{L} evolves as angular momentum is radiated away. The parallel component of the final spin is given by a fit to the final spin of aligned-spin binaries [94] which depends on the symmetric mass ratio η and the effective spin χ_{eff} . The perpendicular component of the spin is unchanged. The magnitude of the final spin used in PhenomP [85] is therefore given by

$$\chi_{fin} = \sqrt{\left(\chi_{fin}^{\parallel}(\eta, \chi_{\text{eff}})\right)^2 + \left(\frac{m_1 \chi_p}{M}\right)^2}. \quad (1.91)$$

The suitability of this formula will be examined in Chapter 4.

The radiated energy and thus the final mass of the system depends only weakly on precession since a precessing waveform can be approximated by a non-precessing one. The radiated energy is given by a fit to the simulations used in the calibration of PhenomD (see equations 3.7 and 3.8 in [94]). The final mass is then given by $M_{fin} = M_{ini} - E_{rad}$.

The underlying co-precessing waveform used in the PhenomP models is based on PhenomD. The inspiral part is given by the equivalent aligned-spin waveform for the precessing case being considered. This is connected to a ringdown part which consists of the aligned spin waveform with the correct value of the final spin as described above.

PhenomPv2: Angle model

PhenomPv2 [38] uses PN expressions for the Euler angles $\{\alpha, \beta, \gamma\}$ derived using the stationary phase approximation for the whole model. These expressions are for single-spin systems and use the mapping defined above to go from a two-spin system to the appropriate single-spin system. Neither PN expressions nor the stationary phase approximation are expected to be accurate through merger and ringdown. However, they have been shown not to significantly degrade the model provided the mass ratio and the spins are not large.

The angular momenta describing the binary are defined in terms of the orthonormal triad $\{\hat{L}_N, \hat{n}, \lambda\}$, where \mathbf{L}_N is the Newtonian orbital angular momentum and \mathbf{n} is the orbital separation vector. The total spin of the system is given by $\mathbf{S} = \mathbf{S}_\parallel + \mathbf{S}_\perp$ where \mathbf{S}_\parallel and \mathbf{S}_\perp are the components of the spin parallel and perpendicular to the orbital angular momentum respectively.

The opening angle β is defined to be the angle between \mathbf{L}_N and \mathbf{J} [24]. It is given by

$$\begin{aligned} \cos \beta &= \frac{J_L}{\sqrt{J_L^2 + J_n^2 + J_\lambda^2}} \\ &= \frac{L + S_\parallel}{\sqrt{(L + S_\parallel)^2 + S_\perp^2}} \\ &= (1 + s^2)^{-\frac{1}{2}} \end{aligned} \tag{1.92}$$

where $s = \frac{S_\perp}{L + S_\parallel}$ and the overall sign is chosen to be positive since for $L + S_\parallel < 0$, \mathbf{J} has changed sign and the approximation that the direction of \mathbf{J} is constant has broken down. In this expression L is given by the standard 2PN non-spinning expression. The spin components are assumed to be fixed and are given by [38]

$$S_\parallel = m_1^2 \chi_1^\parallel + m_2^2 \chi_2^\parallel, \tag{1.93}$$

$$S_\perp = m_1^2 \chi_p. \tag{1.94}$$

The expression for α is derived from the PN expression for $\frac{d\alpha}{dt}$ given in [37]

$$\frac{d\alpha}{dt} = -\frac{\omega_{prec}}{\sin \beta} \frac{J_n}{\sqrt{J_n^2 + J_\lambda^2}} \tag{1.95}$$

where ω_{prec} is defined by $\dot{\hat{L}}_N = -\omega_{prec} \lambda$. The next-to-next-to-leading-order (NNLO) spin-orbit contribution to ω_{prec} is given in [39] and is re-expanded and averaged over one orbit in order to get ω_{prec} as a function of ω , χ_\parallel and χ_\perp only. Radiation reaction is then introduced by allowing the orbital frequency to evolve, using the evolution equation for $\frac{\dot{\omega}}{\omega^2}$ which includes corrections up to 3.5 PN order. Equation 1.95 is

then re-written as

$$\frac{d\alpha}{d\omega} = \frac{1}{\dot{\omega}} \frac{d\alpha}{dt} \quad (1.96)$$

where the RHS can be expanded in terms of ω and the equation can then be integrated up to obtain an expression for $\alpha(\omega)$ as a function of ω , χ_{\parallel} and χ_{\perp} . The expression for α used in PhenomPv2 is then

$$\begin{aligned} \alpha = & -\omega^{-1} \left[\frac{35}{192} + \frac{5\delta}{64m_1} \right] - \omega^{-\frac{2}{3}} \left[\frac{5m_1(3\delta + 7m_1)\chi_{\parallel}}{128M_T^2\eta} \right] \\ & - 5\omega^{-\frac{1}{3}} \left[\frac{(1103 + 824\eta)}{3072} + \frac{\delta(911 + 980\eta)}{7168m_1} + \frac{3\delta^2\eta}{256m_1^2} + \frac{m_1^3(3\delta + 7m_1)\chi_{\perp}^2}{128M_T^4\eta^2} \right] \\ & + 5\omega^{\frac{1}{3}} \left[\frac{8024297}{9289728} + \frac{36\eta}{9289728} (23817 + 85568\eta) \right. \\ & \quad + \frac{\delta}{21676032m_1} (5579117 + 504\eta(-159 + 7630\eta)) \\ & \quad + \frac{\delta^2\eta}{28672m_1^2} (323 + 784\eta) + \frac{3\delta^3\eta^2}{1024m_1^3} \\ & \quad + \frac{m_1^2\chi_{\parallel}^2}{43008M_T^4\eta^2} (504\delta^2\eta + 35m_1^2(19 + 92\eta) + 3\delta m_1(-97 + 812\eta)) \\ & \quad + \frac{m_1(3\delta + 7m_1)}{1536M_T^8\eta^4} \left(-3m_1^6\chi_{\perp}^4 - 96\pi M_T^6\eta^3\chi_{\parallel} \right. \\ & \quad \left. \left. + 2 \left((74\delta + 188m_1) M_T^4\eta^3 + 6m_1^5\chi_{\perp}^2 \right) m_1\chi_{\parallel}^2 \right) \right] \\ & + 5 \ln \omega \left[-\frac{\pi}{48m_1} (3\delta + 7m_1) - \frac{m_1^5\chi_{\perp}^2\chi_{\parallel}}{384M_T^6\eta^3} \right. \\ & \quad \left. + \frac{\chi_{\parallel}}{M_T^2\eta} \left(4032\delta^2\eta + 7m_1^2(599 + 4072\eta) + 3\delta m^2(407 + 7138\eta) \right) \right], \end{aligned} \quad (1.97)$$

where $\chi_{\perp} = \chi_p$ and $\chi_{\parallel} = \frac{1+q}{q}\chi_{\text{eff}}$, m_1 is the mass of the larger black hole (on which this spin is placed), $\delta = m_2 - m_1$, $M_T = m_1 + m_2 = 1$ is the total mass of the system and $\eta = \frac{q}{(1+q)^2}$.

An explicit expression for γ as a function of ω , χ_{\parallel} and χ_{\perp} can similarly be found using the minimal rotation condition $\dot{\gamma} = -\dot{\alpha} \cos \beta$. The expression used for γ in

PhenomPv2 is

$$\begin{aligned}
\gamma = & -\omega^{-1} \left[\frac{35}{192} + \frac{5\delta}{64m_1} \right] - \omega^{-\frac{2}{3}} \left[\frac{5m_1(3\delta + 7m_1)\chi_{\parallel}}{128M_T^2\eta} \right] \\
& - 5\omega^{-\frac{1}{3}} \left[\frac{(1103 + 824\eta)}{3072} + \frac{\delta(911 + 980\eta)}{7168m_1} + \frac{3\delta^2\eta}{256m_1^2} \right] \\
& + 5\omega^{\frac{1}{3}} \left[\frac{8024297}{9289728} + \frac{36\eta}{9289728} (23817 + 85568\eta) \right. \\
& \quad + \frac{\delta}{21676032m_1} (5579117 + 504\eta(-159 + 7630\eta)) \\
& \quad + \frac{\delta^2\eta}{28672m_1^2} (323 + 784\eta) + \frac{3\delta^3\eta^2}{1024m_1^3} \\
& \quad + \frac{m_1^2\chi_{\parallel}^2}{43008M_T^4\eta^2} (504\delta^2\eta + 35m_1^2(19 + 92\eta) + 3\delta m_1(-97 + 812\eta)) \\
& \quad \left. - \frac{\pi m_1\chi_{\parallel}}{16M_T^2\eta} (3\delta + 7m_1) + \frac{m_1^2\chi_{\parallel}^2}{786M_T^4\eta} (3\delta + 7m_1)(75\delta + 188m_1) \right] \\
& + 5 \ln \omega \left[-\frac{\pi}{48m_1} (3\delta + 7m_1) \right. \\
& \quad \left. + \frac{\chi_{\parallel}}{M_T^2\eta} (4032\delta^2\eta + 7m_1^2(599 + 4072\eta) + 3\delta m^2(407 + 7138\eta)) \right].
\end{aligned} \tag{1.98}$$

PhenomPv3: Angle model

PhenomPv3 [102] also uses PN expressions for the Euler angles. These expressions, however, include radiation reaction and spin-orbit and spin-spin effects to leading order in the conservative dynamics and 3.5PN order in the dissipative dynamics (ignoring spin-spin terms). They incorporate two spin effects and so describe generic binary systems. They also do not rely on the stationary phase approximation. The expressions were derived in the development of the first closed-form analytic inspiral waveform model for generically precessing binaries presented in [56, 58]. Radiation-reaction effects were introduced by performing a perturbative expansion to the analytic solution to the conservative precession equations presented in [79] in the ratio of the precession and radiation-reaction timescales [57, 106], known as multiple scale analysis [31]. The frequency domain expressions were derived using shifted uniform asymptotics [107].

The two-spin expressions for the precession angles are given in [58]. In the following discussion, all appendices referenced are in [58]. α , denoted ϕ_z in [58], is given by

$$\alpha = \alpha_{-1} + \alpha_0, \tag{1.99}$$

where

$$\alpha_{-1} = \sum_{n=0}^5 \langle \Omega_\alpha \rangle^{(n)} \alpha^{(n)} + \alpha_{-1}^0, \quad (1.100)$$

$$\alpha_0 = \frac{C_\phi}{\dot{\psi}} \frac{\sqrt{n_c}}{n_c - 1} \arctan \left[\frac{(1 - \sqrt{n_c}) \tan \psi}{1 + \sqrt{n_c} \tan^2 \psi} \right] + \frac{D_\phi}{\dot{\psi}} \frac{\sqrt{n_d}}{n_d - 1} \arctan \left[\frac{(1 - \sqrt{n_d}) \tan \psi}{1 + \sqrt{n_d} \tan^2 \psi} \right], \quad (1.101)$$

where $\langle \Omega_\alpha \rangle^{(n)}$ and $\alpha^{(n)}$ are given in Appendix D and α_{-1}^0 is an integration constant. C_ϕ , D_ϕ , n_c and n_d are given in Appendix B,

$$\dot{\psi} = \frac{A}{2} \sqrt{S_+^2 + S_3^2}, \quad (1.102)$$

$$\psi = \psi_0 - \frac{3g_0}{4} \delta m v^{-3} \left(1 + \psi_1 v + \psi_2 v^2 \right), \quad (1.103)$$

where S_+ , S_- and S_3 are the three roots of the evolution equation for the total spin magnitude S , v is the PN expansion parameter, δm is the mass difference, g_0 is given in Appendix C, A is given in Appendix B, ψ_1 and ψ_2 are given in Appendix C. β , denoted θ_L in [58], is given by

$$\cos \beta = \hat{\mathbf{J}} \cdot \hat{\mathbf{L}} = \frac{J^2 + L^2 - S^2}{2JL}, \quad (1.104)$$

where J , L and S are the magnitudes of the total angular momentum, the orbital angular momentum and the spin angular momentum respectively. Finally, γ , denoted ζ in [58], is given by

$$\gamma = \gamma_{-1} + \gamma_0, \quad (1.105)$$

where

$$\gamma_{-1} = \eta v^{-3} \sum_{n=0}^5 \langle \Omega_\gamma \rangle^{(n)} v^n + \gamma_{-1}^0, \quad (1.106)$$

$$\gamma_0 = \frac{A_\beta}{\dot{\psi}} (C_\phi + D_\phi) + 2d_0 \frac{B_\beta}{\dot{\psi}} \left(\frac{C_\phi}{s_d - d_2} - \frac{D_\phi}{s_d + d_2} \right), \quad (1.107)$$

where $\langle \Omega_\gamma \rangle^{(n)}$ and $\gamma^{(n)}$ are given in Appendix F and γ_{-1}^0 is an integration constant, η is the symmetric mass ratio,

$$A_\beta = \frac{J^2 + L^2 - S_+^2}{2JL}, \quad (1.108)$$

$$B_\beta = \frac{S_+^2 - S_-^2}{2JL}, \quad (1.109)$$

and d_0 , d_2 and s_d are given in Appendix B.

The expressions for α and γ involve a series expansion in terms of the gravitational wave frequency. The order at which these terms are truncated affects the accuracy of the model; if the expansion order is too low then the inspiral expression will not be sufficiently accurate while if it is too high then the expressions will be inaccurate at high frequencies. For the implementation in PhenomPv3, these expressions are truncated at fifth order, which is the second highest-order available. For β the 3PN expression which includes spin-orbit terms was used for the magnitude of the orbital angular momentum.

1.6 Matches

We quantify the disagreement between two waveforms h_1 and h_2 using the standard inner product weighted by the power spectral density of the detector $S_n(f)$ [64]:

$$\langle h_1 | h_2 \rangle = 4\text{Re} \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(f)} \mathrm{d}f. \quad (1.110)$$

The match is given by the inner product between two normalised waveforms optimised over various sets of parameters Θ . For non-precessing waveforms containing only the (2,2) mode, we optimise over time shifts and the phase of the waveform since these have no physical effect on the signal; a time shift corresponds only to a change in the merger time of the binary while a change in the phase corresponds to a rotation of the black holes in the orbital plane. For more generic waveforms, we choose different sets of parameters to optimise over depending on what we are studying. The match is given by

$$M(h_1, h_2) = \max_{\Theta} \langle h_1 | h_2 \rangle \quad (1.111)$$

and the mismatch by

$$\mathcal{M}(h_1, h_2) = 1 - M(h_1, h_2). \quad (1.112)$$

1.6.1 Non-precessing matches

For non-precessing matches neglecting higher order modes it is only necessary to optimise the inner product over time and phase shifts. To optimise over time shifts we use the fact that the Fourier transform of any time-shifted function g is given by $\mathcal{F}[g(t - t_0)] = e^{2\pi i f t_0} \mathcal{F}[g(t)]$. For an arbitrary time shift of h_2 with respect to h_1 , the inner product is therefore given by

$$\langle h_1 | h_2 \rangle = 4\text{Re} \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(f)} e^{-2\pi i f t_0} \mathrm{d}f. \quad (1.113)$$

To optimise over time shifts we can calculate the inverse Fourier transform

$$\mathcal{I}(t_0) = \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(f)} e^{-2\pi i f t_0} df = \mathcal{F}^{-1} \left[\frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(f)} \right] (t_0). \quad (1.114)$$

This can be evaluated for any t_0 to give a value of the inner product for every possible time shift. Finding the maximum value of the inverse Fourier transform gives the inner product optimised over time shifts (as well as the corresponding time shift).

Introducing a phase shift φ_0 the inner product becomes

$$\langle h_1 | h_2 \rangle = 4\text{Re} \left[e^{i\varphi_0} \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(f)} e^{-2\pi i f t_0} df \right] = 4\text{Re} \left[e^{i\varphi_0} \mathcal{I} \right], \quad (1.115)$$

where once the optimisation over time shifts has occurred \mathcal{I} is given by a single complex number. Optimising over the phase shift therefore involves finding the value of φ_0 for which the real part of the $e^{i\varphi_0} \mathcal{I}$ is maximised. The non-precessing match is then finally $M(h_1, h_2) = 4|\mathcal{I}|$.

1.6.2 Precessing matches

The match calculation for precessing waveforms is more involved—two waveforms can no longer differ from each other by a simple phase shift φ_0 since they contain additional modes which have differing phases. In order to identify the correct quantity to optimise over for precessing waveforms, we note that for non-precessing waveforms, optimising over the phase of the binary is equivalent to optimising over the polarisation of the binary. We can write a non-precessing waveform as $h = A e^{i\varphi}$ where A and φ are the phase of the (2,2) mode, since this is assumed to be the dominant contribution to the waveform. A change in the polarisation of the system therefore introduces a phase shift of φ_0 . For precessing systems there is not the same degeneracy between optimisation over the waveform polarisation and over phase shifts, but phase optimisations are related to optimisations with respect to the polarisation angle. For precessing waveforms we therefore optimise the match over polarisation and time shifts. Further to this we also vary the inclination of the template waveform while keeping the inclination of the signal waveform fixed. The following treatment of matches for precessing-binary signals closely follows the presentation in Appendix B of [152].

Optimisation over waveform polarisation and time shifts

It can be shown (e.g. in [145]) that the detector response to a gravitational wave signal $h = h_+ - i h_\times$ of polarisation ψ is given by

$$h_R(t) = \cos 2\psi h_+(t) + \sin 2\psi h_\times(t) = \text{Re} \left[h(t) e^{2i\psi} \right]. \quad (1.116)$$

The Fourier transform of this is

$$\tilde{h}_R(f) = \frac{1}{2} \left[\tilde{h}(f) e^{2i\psi} + \tilde{h}^*(-f) e^{-2i\psi} \right], \quad (1.117)$$

using the property of Fourier transforms that $\mathcal{F}[g^*(t)] = \tilde{g}^*(-f)$. In order to calculate the match between two waveforms— the first detected with polarisation ψ , h^1 , and the second with polarisation σ , h^2 — we compute the inner product weighted by the detector response defined by

$$\langle h^1 | h^2 \rangle = 2 \int_{-\infty}^{\infty} \frac{\tilde{h}^1(f) \tilde{h}^{2*}(f)}{S_n(|f|)} df, \quad (1.118)$$

which can be rearranged to give an expression in the form of equation 1.110 using the fact that $\tilde{g}(f) = \tilde{g}^*(-f)$. Substituting equation 1.117 into equation 1.118 we find the inner product between the waveforms h^1 and h^2 is given by

$$\begin{aligned} \langle h_R^1 | h_R^2 \rangle &= \frac{1}{2} \int_{-\infty}^{\infty} \frac{df}{S_n(|f|)} \times \\ &\quad \left[\tilde{h}^1(f) e^{2i\psi} + \tilde{h}^{1*}(-f) e^{-2i\psi} \right] \left[\tilde{h}^2(f) e^{2i\sigma} + \tilde{h}^{2*}(-f) e^{-2i\sigma} \right]^* \end{aligned} \quad (1.119)$$

$$\begin{aligned} &= \text{Re} \left[\int_{-\infty}^{\infty} \frac{\tilde{h}^1(f) \tilde{h}^{2*}(f)}{S_n(|f|)} df e^{2i(\psi-\sigma)} \right] \\ &\quad + \text{Re} \left[\int_{-\infty}^{\infty} \frac{\tilde{h}^1(f) \tilde{h}^2(-f)}{S_n(|f|)} df e^{2i(\psi+\sigma)} \right], \end{aligned} \quad (1.120)$$

using the property that $\int_{-\infty}^{\infty} x(f) df = \int_{-\infty}^{\infty} x(-f) df$ for any function x . Similarly, the norm of a waveform is given by

$$\|h_R\|^2 = \langle h_R | h_R \rangle \quad (1.121)$$

$$= \frac{1}{2} \int_{-\infty}^{\infty} \frac{1}{S_n(|f|)} \left[\tilde{h}(f) e^{2i\psi} + \tilde{h}^*(-f) e^{-2i\psi} \right] \left[\tilde{h}(f) e^{2i\psi} + \tilde{h}^*(-f) e^{-2i\psi} \right]^* df \quad (1.122)$$

$$\begin{aligned} &= \int_{-\infty}^{\infty} \frac{|\tilde{h}(f)|^2}{S_n(|f|)} df \\ &\quad + \frac{1}{2} \left[\int_{-\infty}^{\infty} \frac{\tilde{h}(f) \tilde{h}(-f)}{S_n(|f|)} df e^{4i\psi} + \int_{-\infty}^{\infty} \frac{\tilde{h}^*(f) \tilde{h}^*(-f)}{S_n(|f|)} df e^{-4i\psi} \right]. \end{aligned} \quad (1.123)$$

We wish to optimise this calculation over the polarisation of the waveform. In order to simplify the optimisation over polarisation we can write the components of equations 1.119 and 1.123 as complex numbers with a real amplitude and phase.

Substituting

$$Ae^{i\sigma_A} = \int_{-\infty}^{\infty} \frac{\tilde{h}^{2*}(f)}{S_n(|f|)} \left[\tilde{h}^1(f) e^{2i\psi} + \tilde{h}^{1*}(-f) e^{-2i\psi} \right] df, \quad (1.124)$$

$$B_1 = \int_{-\infty}^{\infty} \frac{|\tilde{h}(f)|^2}{S_n(|f|)} df, \quad B_2 e^{i\sigma_B} = \int_{-\infty}^{\infty} \frac{\tilde{h}(f) \tilde{h}(-f)}{S_n(|f|)} df, \quad (1.125)$$

into equations 1.119 and 1.123 we find

$$\langle h_R^1 | h_R^2 \rangle = A \cos(2\sigma - \sigma_A), \quad (1.126)$$

$$\|h_R^2\|^2 = B_1 + B_2 \cos(4\sigma + \sigma_B), \quad (1.127)$$

which gives

$$\begin{aligned} \left\langle \frac{h_R^1}{\|h_R^2\|} \middle| \frac{h_R^2}{\|h_R^2\|} \right\rangle &= \frac{A \cos(2\sigma - \sigma_A)}{\|h_R^2\| \sqrt{B_1 + B_2 \cos(4\sigma + \sigma_B)}} \\ &= \frac{A (\cos \sigma_A + \sin \sigma_A t_\sigma)}{\|h_R^2\| \sqrt{B_1 (1 + t_\sigma^2) + B_2 (\cos \sigma_B (1 - t_\sigma^2) - 2 \sin \sigma_B t_\sigma)}}, \end{aligned} \quad (1.128)$$

where $t_\sigma = \tan 2\sigma$ and $\|h_R^2\|$ is independent of σ . The value of σ which maximises this expression is given by

$$\tan 2\sigma = \frac{B_1 \sin \sigma_B + B_2 \sin(\sigma_A + \sigma_B)}{B_1 \cos \sigma_B - B_2 \cos(\sigma_A + \sigma_B)}, \quad (1.130)$$

which gives the inner product maximised over polarisation as

$$\max_{\sigma} \left\langle \frac{h_R^1}{\|h_R^2\|} \middle| \frac{h_R^2}{\|h_R^2\|} \right\rangle = \frac{A}{\|h_R^2\|} \sqrt{\frac{B_1 - B_2 \cos(\sigma_B + 2\sigma_A)}{B_1^2 - B_2^2}}. \quad (1.131)$$

The optimisation over time shifts follows the same principle as for the non-precessing case. In order to calculate the precessing match we take the inverse Fourier transform of the integrand in each of the quantities A , B_1 and B_2 in order to find the optimisation over time and then use equation 1.131 to find the optimisation over the polarisation angle.

Variation over inclination

Finally, for a precessing waveform we no longer consider only the (2,2) mode. Consequently, changing the orientation of the system with respect to the detector results in a different combination of the modes.

As described in section 1.1.4, the waveform can be written as

$$h(t) = \sum_{\ell,m} Y_{\ell,m}^{-2}(\theta, \phi) h^{\ell,m}(t). \quad (1.132)$$

Substituting this in equation 1.120, we find that

$$\begin{aligned} \langle h_R^1 | h_R^2 \rangle &= \text{Re} \left[\sum_{\ell,m} -2Y_{\ell,m}^* \int_{-\infty}^{\infty} \frac{\tilde{h}^1(f) (\tilde{h}_{\ell,m}^2)^*(f)}{S_n(|f|)} df e^{2i(\psi-\sigma)} \right] \\ &+ \text{Re} \left[\sum_{\ell,m} -2Y_{\ell,m}^* \int_{-\infty}^{\infty} \frac{\tilde{h}^1(f) (\tilde{h}_{\ell,m}^2)^*(-f)}{S_n(|f|)} df e^{2i(\psi+\sigma)} \right], \end{aligned} \quad (1.133)$$

and from equation 1.123

$$\|h_R^2\|^2 = \int_{-\infty}^{\infty} \frac{|\tilde{h}(f)|^2}{S_n(|f|)} df + \text{Re} \left[\int_{-\infty}^{\infty} \frac{\tilde{h}(f) \tilde{h}(-f)}{S_n(|f|)} df e^{4i\sigma} \right] \quad (1.134)$$

$$\begin{aligned} &= \sum_{\ell,m,m'} -2Y_{\ell,m} -2Y_{\ell,m'}^* \int_{-\infty}^{\infty} \frac{\tilde{h}_{\ell,m}^2(f) (\tilde{h}_{\ell,m}^2)^*(f)}{S_n(|f|)} df \\ &+ \text{Re} \left[\sum_{\ell,m,m'} -2Y_{\ell,m} -2Y_{\ell,m'}^* \int_{-\infty}^{\infty} \frac{\tilde{h}_{\ell,m}^2(f) \tilde{h}_{\ell,m}^2(-f)}{S_n(|f|)} df e^{4i\sigma} \right]. \end{aligned} \quad (1.135)$$

Defining the vectors

$$A_{\ell,m}^{(1)} = \int_{-\infty}^{\infty} \frac{\tilde{h}^1(f) (\tilde{h}_{\ell,m}^2)^*(f)}{S_n(|f|)} df, \quad A_{\ell,m}^{(2)} = \int_{-\infty}^{\infty} \frac{\tilde{h}^1(-f) (\tilde{h}_{\ell,m}^2)^*(f)}{S_n(|f|)} df, \quad (1.136)$$

and the following matrices

$$B_{\ell,m,m'}^{(1)} = \int_{-\infty}^{\infty} \frac{\tilde{h}_{\ell,m}^2(f) (\tilde{h}_{\ell,m'}^2)^*(f)}{S_n(|f|)} df, \quad B_{\ell,m,m'}^{(2)} = \int_{-\infty}^{\infty} \frac{\tilde{h}_{\ell,m}^2(f) \tilde{h}_{\ell,m'}^2(-f)}{S_n(|f|)} df, \quad (1.137)$$

and combining the spherical harmonics into a vector \mathbf{Y} we can write

$$\langle h_R^1 | h_R^2 \rangle = \sum_{\ell} \text{Re} \left[A'_{\ell} e^{-2i\sigma} \right] = \sum_{\ell} (A' \cos(2\sigma - \sigma_{A'}))_{\ell}, \quad (1.138)$$

$$\|h_R^2\|^2 = \sum_{\ell} \left[(B'_1)_{\ell} + \text{Re} \left[(B'_2 e^{i\sigma_{B'_2}})_{\ell} e^{4i\sigma} \right] \right] = \sum_{\ell} \left[(B'_1)_{\ell} + B'_2 \cos(4\sigma + \sigma_{B'})_{\ell} \right], \quad (1.139)$$

where

$$\left(A'e^{i\sigma_{A'}}\right)_\ell = \left(e^{2i\psi}\mathbf{A}^{(1)}_\ell + e^{-2i\psi}\mathbf{A}^{(2)}_\ell\right) \cdot \mathbf{Y}_\ell^*, \quad (1.140)$$

$$\left(B'_1\right)_\ell = \mathbf{Y}_\ell^T \mathbf{B}^{(1)}_\ell \mathbf{Y}_\ell^*, \quad (1.141)$$

$$\left(B'_2 e^{i\sigma_{B'_2}}\right)_\ell = \mathbf{Y}_\ell^T \mathbf{B}^{(2)}_\ell \mathbf{Y}_\ell. \quad (1.142)$$

The precessing match optimised over polarisation and time shifts can now be calculated as described above.

1.6.3 Power weighted matches

When performing matches between NR waveforms we are limited by the length of the signal. The signal is often just the last few orbits prior to merger while for detections of low mass binaries ground-based interferometers will detect many inspiral cycles prior to merger. Therefore when assessing the accuracy of these NR waveforms for a range of binary masses we perform *power weighted matches* [120]. In this case we consider the match between two waveforms over the frequency range $[f_{min}, f_{max}]$, where the portion of the waveform in range $[f_{min}, f_{mid}]$ is represented by a post-Newtonian model of the signal and the portion in the range $[f_{mid}, f_{max}]$ is given by an NR waveform. From equation 1.111 we can see that for two waveforms h_1 and h_2 the match is given by

$$M(h_1, h_2) = \frac{1}{\|h_1\| \|h_2\|} \max_{\Theta} \left[4\text{Re} \int_{f_{min}}^{f_{max}} \frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(f)} df \right], \quad (1.143)$$

where $\|h\|$ is the norm of the waveform. If we assume (i) that the main contribution to the match comes from the dephasing between the two waveforms while their norms can be considered to be identical and (ii) any error in the model for the inspiral portion of the waveform is negligible, then we have that

$$\begin{aligned} M(h_1, h_2) &= \frac{1}{\|h\|^2} \max_{\Theta} \left[4\text{Re} \int_{f_{min}}^{f_{mid}} \frac{|\tilde{h}|^2}{S_n(f)} df + 4\text{Re} \int_{f_{mid}}^{f_{max}} \frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(f)} df \right] \\ &= \frac{1}{\|h\|^2} \left[\|h\|_{(f_{min}, f_{mid})}^2 + \|h\|_{(f_{mid}, f_{max})}^2 M(h_1, h_2)_{(f_{mid}, f_{max})} \right], \end{aligned} \quad (1.144)$$

where $\|h\|^2 = \|h\|_{(f_{min}, f_{mid})}^2 + \|h\|_{(f_{mid}, f_{max})}^2$. The mismatch over the complete waveform is therefore given by

$$\begin{aligned} \mathcal{M}(h_1, h_2) &= \frac{\|h\|_{(f_{mid}, f_{max})}^2}{\|h\|^2} \times \mathcal{M}(h_1, h_2)_{(f_{mid}, f_{max})} \\ &\equiv \frac{\|h_{\text{NR}}\|^2}{\|h_{\text{model}}\|^2 + \|h_{\text{NR}}\|^2} \times \mathcal{M}_{\text{NR}}, \end{aligned} \quad (1.145)$$

where the subscripts indicate where the match and overlap have been calculated using the post-Newtonian model or the NR waveform. This gives an upper bound on the value for the match since it assumes that the mismatch between the inspiral portion of the waveform is negligible.

Chapter 2

BAM catalogue

2.1 Introduction

Following the detections of gravitational waves to date by the LIGO and Virgo collaborations [6, 7, 9, 10, 13–16], we expect an increasing number of detections as the detectors get increasingly sensitive. Among these detections we expect some that will show clear evidence of precession. Precession is most easily detectable for systems with large mass ratios and for those with high spins. We therefore want a waveform model of precessing systems which is accurate, even in these more extreme regions of parameter space. The catalogue detailed in this chapter was produced in order to assist in the creation of such a model.

In producing the first precessing model that has been tuned to Numerical Relativity (NR) waveforms, we wish to capture the dominant precession effects first. We therefore consider single spin systems which obey simple precession. Consequently, the catalogue contains single spin configurations where the spin is placed on the larger black hole and neglects two spin configurations and those with a varying azimuthal spin component, since the effects from these kinds of systems are expected to be subdominant. Such single spin configurations can be parameterised by three quantities; their mass ratio, the magnitude of the spin on the larger black hole and the spin angle between the spin and the orbital angular momentum of the system.

Since precession effects are most noticeable at high mass ratios, the catalogue must extend up to high mass ratios. The earlier tuned non-precessing model PhenomD [94, 103] was based on a catalogue containing systems up to mass ratio 18. Producing the number of precessing simulations with mass ratio 18 required to calibrate the precessing model that will be developed in Chapter 4 is computationally prohibitive. The catalogue extends up to mass ratio 8. In order to confidently capture the dependence of precession effects on mass ratio, we produced simulations at four different mass ratios up to mass ratio 8 that are approximately equally spaced in symmetric mass ratio. We already have aligned and anti-aligned waveforms at these mass ratios and for non-aligned spin configurations we chose five equally spaced

values for the spin angle excluding 0° and 180° . Finally, we expect the dependence of the precession effects on the spin magnitude to be approximately linear— for example consider the leading order spin dependence of the PN expressions for α and γ given in equations 1.97 and 1.98. We therefore chose only two values of the spin magnitude.

In this chapter we present a catalogue of 40 NR waveforms for use in a tuned precessing waveform model. We assess the accuracy of the waveforms that comprise this catalogue in order to place an upper bound on the accuracy of the model. We first quantify the numerical errors in the waveform quantities calculated from the simulation using Richardson extrapolation. We then estimate the mismatch between NR waveforms and the theoretical “true” analytical waveform.

2.2 Catalogue

The catalogue consists of 40 waveforms produced using the BAM code, which is described in section 1.2.4 of the Introduction. Each of these are single spin precessing waveforms where the spin is placed on the larger black hole. The black holes are labelled by their masses m_1 and m_2 , with total mass $M = m_1 + m_2$ and mass ratio $q = \frac{m_1}{m_2} > 1$. The spin on the larger black hole has magnitude χ and spin angle θ , which is taken to be the angle between the spin and the orbital angular momentum of the system. The waveforms span the parameter space with $q \in [1, 2, 4, 8]$, $\chi \in [0.4, 0.8]$ and $\theta (^\circ) \in [30, 60, 90, 120, 150]$. More details are given in tables 2.1–2.2. The values for χ , θ , the effective spin χ_{eff} and the precession spin parameter χ_p given in the tables are calculated from the initial parameters at a specified reference frequency used when setting up the simulation and at the relaxed time, which is defined below, in brackets. The eccentricity e is measured over the region $200 - 1000M$ using the method described in [93]. The orbital frequency ω_{orb} is calculated from the dynamics at the relaxed time. The number of orbits N_{orb} is calculated from the relaxed time until the peak in the (2,2) mode of the ψ_4 waveform.

The relaxed time is the point in the simulation after which an initial burst of “junk radiation” has died away. We estimate the timescale on which the junk radiation affects the waveform data by calculating the damping period t_D of the quasi-normal modes of a single perturbed black hole with mass m_1 and spin 0.8. We chose a fixed value for the spin since it has little effect on the damping frequency. The value of 0.8 was chosen since this is the highest spin included in the catalogue (which corresponds to the longest damping time). The relaxed time is then given by $t_{\text{rel}} = t_{\text{max}} + t_D$, where t_{max} is the time at which the maximum in the junk radiation occurs.

The values of the parameters used in the production of the initial data for the simulations (such as the masses of the black holes, their spins and separation as well as the starting frequency) are specified at a large separation. The PN equations of

Name	q	χ	θ	χ_{eff}	χ_p	D/M	e ($\times 10^{-3}$)	$M\omega_{\text{orb}}$	N_{orb}
bam19-1	1	0.400 (0.400)	30.004 (29.924)	0.173 (0.174)	0.200 (0.200)	11.573 (11.338)	1.547	0.023	10.001
bam19-2	1	0.400 (0.400)	60.007 (59.851)	0.100 (0.101)	0.346 (0.346)	11.569 (11.340)	2.318	0.023	9.706
bam19-3	1	0.400 (0.400)	90.041 (89.832)	-0.000 (0.001)	0.400 (0.400)	11.510 (11.285)	4.034	0.023	9.091
bam19-4	1	0.400 (0.400)	120.039 (119.837)	-0.100 (-0.099)	0.346 (0.347)	11.572 (11.330)	2.971	0.023	8.714
bam19-5	1	0.400 (0.400)	150.013 (149.888)	-0.173 (-0.173)	0.200 (0.201)	11.619 (11.366)	2.356	0.023	8.451
bam19-6	1	0.800 (0.801)	30.029 (29.845)	0.346 (0.347)	0.400 (0.399)	11.560 (11.311)	2.145	0.023	10.947
bam19-7	1	0.800 (0.801)	60.041 (59.736)	0.200 (0.202)	0.693 (0.692)	11.617 (11.371)	2.041	0.023	10.281
bam19-8	1	0.800 (0.801)	90.071 (89.690)	-0.000 (0.002)	0.800 (0.801)	11.571 (11.308)	2.070	0.023	9.098
bam19-9	1	0.800 (0.801)	120.059 (119.668)	-0.200 (-0.198)	0.692 (0.696)	11.585 (11.309)	1.734	0.023	8.285
bam19-10	1	0.800 (0.801)	150.049 (149.801)	-0.347 (-0.346)	0.399 (0.403)	11.566 (11.269)	1.148	0.023	7.596
bam19-11	2	0.400 (0.400)	30.040 (29.921)	0.231 (0.231)	0.200 (0.200)	11.551 (11.381)	1.396	0.023	11.136
bam19-12	2	0.400 (0.400)	60.089 (59.833)	0.133 (0.134)	0.347 (0.346)	11.632 (11.469)	1.832	0.023	10.783
bam19-13	2	0.400 (0.400)	90.117 (89.809)	-0.001 (0.001)	0.400 (0.400)	11.630 (11.436)	1.298	0.023	9.922
bam19-14	2	0.400 (0.400)	120.086 (119.815)	-0.134 (-0.133)	0.346 (0.347)	11.579 (11.395)	1.930	0.023	8.947
bam19-15	2	0.400 (0.400)	150.063 (149.903)	-0.231 (-0.231)	0.200 (0.201)	11.634 (11.431)	1.501	0.023	8.501
bam19-16	2	0.800 (0.802)	30.096 (29.880)	0.461 (0.463)	0.401 (0.399)	11.453 (11.223)	2.761	0.023	12.454
bam19-17	2	0.800 (0.802)	60.149 (59.726)	0.265 (0.269)	0.694 (0.692)	11.472 (11.261)	1.989	0.023	11.433
bam19-18	2	0.800 (0.802)	90.235 (89.712)	-0.002 (0.003)	0.800 (0.802)	11.571 (11.369)	2.798	0.023	10.093
bam19-19	2	0.800 (0.802)	120.188 (119.676)	-0.268 (-0.264)	0.691 (0.696)	11.544 (11.288)	1.680	0.023	8.312
bam19-20	2	0.800 (0.802)	150.135 (149.836)	-0.462 (-0.462)	0.398 (0.403)	11.626 (11.339)	1.369	0.023	7.227

Table 2.1: Properties of twenty of the waveforms included in the catalogue with mass ratios 1 and 2.

Name	q	χ	θ	χ_{eff}	χ_p	D/M	e ($\times 10^{-3}$)	$M\omega_{\text{orb}}$	N_{orb}
bam19-21	4	0.400 (0.400)	30.102 (30.069)	0.277 (0.277)	0.201 (0.201)	10.582 (10.553)	1.404	0.026	11.963
bam19-22	4	0.400 (0.400)	60.179 (60.115)	0.159 (0.160)	0.347 (0.347)	10.742 (10.712)	1.239	0.025	11.204
bam19-23	4	0.400 (0.400)	90.196 (90.057)	-0.001 (-0.000)	0.400 (0.400)	11.552 (11.545)	1.443	0.023	12.353
bam19-24	4	0.400 (0.400)	120.176 (120.063)	-0.161 (-0.160)	0.346 (0.346)	11.570 (11.519)	1.439	0.023	10.958
bam19-25	4	0.400 (0.400)	150.104 (150.049)	-0.277 (-0.277)	0.199 (0.200)	11.544 (11.463)	2.033	0.023	9.810
bam19-26	4	0.800 (0.801)	30.192 (30.144)	0.553 (0.554)	0.402 (0.402)	10.104 (10.027)	1.350	0.027	13.245
bam19-27	4	0.800 (0.802)	60.354 (60.256)	0.317 (0.318)	0.695 (0.696)	10.482 (10.376)	0.749	0.026	12.121
bam19-28	4	0.800 (0.802)	90.397 (90.195)	-0.004 (-0.002)	0.800 (0.802)	11.487 (11.365)	1.755	0.023	12.329
bam19-29	4	0.800 (0.802)	120.363 (120.185)	-0.324 (-0.323)	0.690 (0.693)	11.579 (11.461)	2.013	0.023	9.987
bam19-30	4	0.800 (0.801)	150.213 (150.120)	-0.555 (-0.556)	0.397 (0.399)	11.640 (11.478)	1.162	0.023	8.153
bam19-31	8	0.400 (0.400)	30.153 (29.866)	0.307 (0.309)	0.201 (0.199)	9.542 (9.670)	1.252	0.030	13.206
bam19-32	8	0.400 (0.400)	60.266 (59.911)	0.176 (0.178)	0.347 (0.346)	9.689 (9.840)	2.109	0.029	12.265
bam19-33	8	0.400 (0.400)	90.305 (90.026)	-0.002 (-0.000)	0.400 (0.400)	10.107 (10.210)	1.079	0.028	11.785
bam19-34	8	0.400 (0.400)	120.261 (120.093)	-0.179 (-0.178)	0.345 (0.346)	10.476 (10.567)	1.342	0.026	11.143
bam19-35	8	0.400 (0.400)	150.148 (150.074)	-0.308 (-0.308)	0.199 (0.200)	10.766 (10.847)	1.492	0.026	10.659
bam19-36	8	0.800 (0.802)	30.305 (29.666)	0.614 (0.619)	0.404 (0.397)	8.964 (8.970)	1.010	0.032	15.870
bam19-37	8	0.800 (0.801)	60.538 (59.418)	0.350 (0.362)	0.697 (0.690)	9.372 (9.362)	1.364	0.030	14.223
bam19-38	8	0.800 (0.801)	90.620 (89.777)	-0.008 (0.003)	0.800 (0.801)	10.021 (10.005)	2.143	0.028	11.924
bam19-39	8	0.800 (0.802)	120.513 (120.255)	-0.361 (-0.359)	0.689 (0.693)	10.879 (10.907)	2.753	0.025	11.113
bam19-40	8	0.800 (0.802)	150.286 (150.199)	-0.618 (-0.618)	0.397 (0.398)	11.453 (11.426)	0.915	0.024	10.181

Table 2.2: Properties of twenty of the waveforms included in the catalogue with mass ratios 4 and 8.

motion are then evolved to get the values at the reference frequency ω_{orb} . We use a combination of the Hamilton equations of motion in the standard Taylor expanded form (as in [84, 93]) and the EOB form [50, 66], as has been done previously—such as in [94, 136]. More details on these equations are given in Appendix A of [136]. This method of obtaining these parameters does not enable one to specify the exact values at the reference frequency. In order to do this an iterative algorithm was developed [98] where the parameters are specified, the PN equations evolved and the parameters at the reference time are checked to see if they satisfy the requirements of the simulation. If the spin is not in the desired orientation at the reference time then the discrepancy between the actual spin direction and the desired direction is noted and the initial spin is rotated by that amount. The PN equations are then evolved again and the resulting spin orientation is checked. This process is repeated until the correct spin orientation occurs at the starting frequency of the simulation, to within a tolerance of 1% of the starting frequency.

The phase error in the waveform accumulates as the simulation is evolved. The most effective way to reduce this effect is to increase the resolution at which the simulation is performed. However this also decreases the speed at which the simulation runs and so increases the amount of time it takes for the simulation to be completed. In addition, increasing the resolution increases the memory requirement of the simulation (which scales as N^3 where N is the number of grid points in each direction). With the computational resources available, we require that the resolution of the simulations that make up the catalogue run on no more than 256 cores and take on the order of one month of continuous run time. In order to ensure the phase error is reasonably small while maintaining a realistic resolution we require the simulations to have merged by $\sim 2000M$. Additionally, since our goal is to produce a model in the frequency domain, we would ideally like all our waveforms to cover the same frequency range and so have the same starting frequency. These requirements determine the reference frequency of the simulations. An initial choice of $M\omega_{\text{orb}} = 0.023$ was made. For some high mass ratio cases with large aligned spins this leads to simulations with lengths in excess of 3000M, so for these cases the starting frequency was increased.

We want to be able to assess the accuracy of these waveforms since this will place a lower bound the accuracy of anything which relies on this data—such as the precessing waveform model presented in Chapter 4. The simulations which comprise the catalogue were performed with a grid spacing $d = 0.0104$ and the waveforms were extracted at a distance $R_{\text{ext}} = 90M$ away from the source.

2.3 Accuracy assessment waveforms

In order to assess the accuracy of the data which comprise this catalogue we studied a subset of four of the configurations described in tables 2.1 and 2.2. These four config-

urations are $(q, \chi, \theta) = \{(4, 0.4, 60), (4, 0.8, 120), (8, 0.4, 30), (8, 0.8, 150)\}$. The set of simulations used in the accuracy analysis of the $(4, 0.4, 60)$ case were performed with a lower starting frequency of $M\omega_{\text{orb}} = 0.023$ to provide an assessment of the error present in a longer simulation.

The two main sources of error in our waveforms are the finite resolution of the simulation and the finite radius at which the data are extracted. In order to assess the effect of the finite resolution, we performed a set of three simulations with low, medium and high resolution for each of the four configurations listed above. We also performed an additional simulation with very high resolution for the $(8, 0.8, 150)$ configuration. The resolution of each of these simulations is given in table 2.3. We extracted the waveform data at $R_{\text{ext}} = \{50, 60, 70, 80, 90\}M$ which were all on the same refinement level.

Resolution	d (grid spacing)	N (number of grid points)
Low	0.0125	80
Medium	0.0104	96
High	0.00833	120
Very high	0.00694	144

Table 2.3: Properties of the various resolutions at which the simulations were performed.

In quantifying the error in the waveforms due to these two sources we focus on estimating the mismatch between the medium resolution waveforms extracted at a distance of $90M$ from the source and an estimate of the “true” waveform at infinitely good resolution and infinitely far from the source. This estimate is obtained by extrapolating through the mismatches between different resolutions and extraction radii. We calculated the convergence order of the BAM code with respect to the resolution and extraction radius then used this to extrapolate the mismatch. We also used Richardson extrapolation to estimate the truncation error due to resolution and extraction radius.

2.4 Quantifying Waveform Accuracy

2.4.1 Richardson Extrapolation

A quantity q calculated at finite resolution or extraction radius can be given by

$$q^* = q(\Delta) + e_i \Delta^i, \quad (2.1)$$

where Δ is the expansion parameter ($\frac{1}{N}$ for resolution or $\frac{1}{R_{\text{ext}}}$ for extraction radius), e is the finite order error and i is the order at which the error contributes. In this chapter the quantity q we are considering is the waveform extracted from the

numerical simulation. We therefore have that

$$q^* = q(\Delta) + e_n \Delta^n + \mathcal{O}(\Delta^{n'}), \quad (2.2)$$

where $e_n \Delta^n$ is the leading order error contribution, n is the convergence order of the simulation and $n' > n$.

Considering two waveforms computed using different expansion parameters Δ_1 and Δ_2 , we can solve the two simultaneous equations that arise from equation 2.2 to give

$$q^* = R(\Delta_1, \Delta_2) + \mathcal{O}\left(\left(\frac{\Delta_1}{\Delta_2}\right)^{n'}\right), \quad (2.3)$$

where

$$R(\Delta_1, \Delta_2) = \frac{\left(\frac{\Delta_1}{\Delta_2}\right)^n q(\Delta_2) - q(\Delta_1)}{\left(\frac{\Delta_1}{\Delta_2}\right)^n - 1}, \quad (2.4)$$

is the *Richardson extrapolation* [59] of $q(\Delta_1)$. $R(\Delta_1, \Delta_2)$ has a higher order error due to the truncation of the expansion in Δ than $q(\Delta_1)$. $R(\Delta_1, \Delta_2) - q(\Delta_1)$ gives the truncation error of the quantity q .

2.4.2 Convergence

Considering now three waveforms computed with expansion parameters $\Delta_1 > \Delta_2 > \Delta_3$ we can eliminate q^* in equation 2.3. Neglecting higher order error terms, the ratio of the difference between two sets of numerical waveforms with expansion parameter $\Delta_1 > \Delta_2 > \Delta_3$ is then given by

$$\mathcal{C} = \frac{q(\Delta_1) - q(\Delta_2)}{q(\Delta_2) - q(\Delta_3)} = \frac{\Delta_1^n - \Delta_2^n}{\Delta_2^n - \Delta_3^n}. \quad (2.5)$$

This relation holds for features of a waveform, such as its amplitude and phase, but not for more involved quantities such as the match.

To find how the ratio of two matches between waveforms of differing expansion parameter depends on the expansion parameter, we must examine how the match depends on the amplitude and phase of the waveform. From equation 1.111 we can see that the match goes as

$$M \sim \text{Re} \left[\frac{1}{\mathcal{N}_1 \mathcal{N}_2} \int h_1(f) h_2^*(f) df \right], \quad (2.6)$$

where \mathcal{N}_i are the normalisation of each of the waveforms respectively. We take h_1 to be the waveform containing either the amplitude or phase error and h_2 to be the “true” waveform.

If we assume the true waveform to be normalised then a waveform containing some amplitude error ΔA is given by

$$h(f) = (1 + \Delta A(f)) e^{i\phi}. \quad (2.7)$$

Substituting equation 2.7 into 2.6 we find

$$\begin{aligned} M &\sim \left[\int (1 + \Delta A)^2 \right]^{-\frac{1}{2}} \operatorname{Re} \int (1 + \Delta A) \, df \\ &= (1 + 2B + C)^{-\frac{1}{2}} (1 + B) \\ &\simeq 1 - B^2 - \frac{1}{2}C(1 + B), \end{aligned} \quad (2.8)$$

where $B = \int \Delta A \, df$, $C = \int (\Delta A)^2 \, df$ and we have assumed that ΔA is small in order to make the approximation in the final step. The mismatch, as defined in equation 1.112, therefore goes as $B^2 + \frac{1}{2}C$ and so is dominated by the square of the amplitude error.

Similarly for a normalised waveform which contains some phase error $\Delta\phi$,

$$h(f) = e^{i(\phi^*(f) + \Delta\phi(f))}, \quad (2.9)$$

where ϕ^* is the “true” phase. Substituting this expression into equation 2.6 we find

$$\begin{aligned} M &\sim \operatorname{Re} \left[\int e^{i\Delta\phi} \, df \right] \\ &\simeq \operatorname{Re} \int \left(1 + i\Delta\phi - \frac{1}{2}(\Delta\phi)^2 \right) \, df \\ &= \int \left(1 - \frac{1}{2}(\Delta\phi)^2 \right) \, df, \end{aligned} \quad (2.10)$$

where again we have assumed that the error in the phase is small in order to perform the expansion in the middle step. The mismatch is therefore dominated by the square of the phase error.

Since the mismatch is dominated by the square of the error in the amplitude and phase and, as shown in equation 2.2, the numerical errors in the amplitude and phase vary as Δ^n , we have that

$$\mathcal{M}(h_1, h_2) \approx \kappa (\Delta_1^n - \Delta_2^n)^2. \quad (2.11)$$

From this we can see that the ratio of the mismatch \mathcal{M} between two pairs of waveforms is given by

$$\frac{\mathcal{M}(A : B)}{\mathcal{M}(B : C)} = \frac{(\Delta_A^n - \Delta_B^n)^2}{(\Delta_B^n - \Delta_C^n)^2}. \quad (2.12)$$

2.4.3 Combining mismatches

From equation 2.12 we can see that if we have three waveforms with expansion parameters Δ_A , Δ_B and Δ_C respectively, then we can solve a set of simultaneous equations to find the relation between the mismatch for each of the sets of waveforms. This implies that for three waveforms with $\Delta_A > \Delta_B > \Delta_C$, the correct way to combine the mismatches $\mathcal{M}(A : B)$ and $\mathcal{M}(B : C)$ in order to get the mismatch $\mathcal{M}(A : C)$ is given by

$$\mathcal{M}(A : C) = \left(\sqrt{\mathcal{M}(A : B)} + \sqrt{\mathcal{M}(B : C)} \right)^2. \quad (2.13)$$

This is useful for, among other things, combining the mismatch due to different sources of error in our numerical waveforms.

2.5 Waveform errors

In order to estimate the numerical error in the waveform quantities due to the finite resolution of the simulation and the finite radius at which the data were extracted, we performed Richardson extrapolation. This requires an estimate of the convergence order of the code with respect to these quantities. We first processed the data, removing the junk radiation from the waveform in the inertial frame in which the simulation was performed. We aligned the waveforms at merger, where merger is defined to be the time at which the quantity

$$A^2 = \sum_{m=-2}^2 |A_{2m}(t)|^2, \quad (2.14)$$

is maximised and resampled using a constant time step of 0.1M. We then rotated the waveform into the co-processing frame and aligned the co-processing phases at merger. The quantities for which we are interested in quantifying the numerical error are the amplitude and phase of the (2,2) mode in the co-processing frame as well as two of the Euler angles α and β required to rotate the waveform from the inertial frame into the co-processing frame. The Euler angles were calculated using the matrix method outlined in section 1.4.1 of the Introduction.

The standard way to perform a convergence test with respect to the resolution is to perform a set of three simulations where the resolution improves by a factor of two between each of the simulations. This is computationally prohibitive—the high resolution run would use 2^6 times as memory as the low resolution run. Similarly we cannot extract a waveform at three different radii on the same level a reasonable distance from the source if each of the extraction radii is twice as far away from the source as the previous one.

In previous work, clean convergence has rarely been demonstrated for both finite-

difference (such as BAM [49, 94]) and pseudo-spectral (such as SpEC [147]) simulations. Since we do not expect to see clean convergence, we estimated the convergence order of the code to be the value of n for which the quantity

$$\delta = (\phi(\Delta_1) - \phi(\Delta_2)) - \mathcal{C}(\phi(\Delta_2) - \phi(\Delta_3)), \quad (2.15)$$

is minimised, where $\mathcal{C} = \frac{\Delta_1^n - \Delta_2^n}{\Delta_2^n - \Delta_3^n}$ as in equation 2.5. This was done for both waveforms of varying resolution and extraction radius for the $q = 8$, $\chi = 0.8$, $\theta = 150^\circ$ configuration. We used the medium, high and very high resolution waveforms and the waveforms extracted at $R_{\text{ext}} = \{90, 70, 60\}$ M. The results of this analysis are shown in figure 2.1. From section 1.2.4 in the Introduction, we see that the code can be either second, fourth or sixth order convergent, depending on the dominant source of error. If our set of waveforms do not all clearly lie within the convergence regime we will not see clean convergence. From figure 2.1 we can see that averaged across the waveform, our results would imply fifth order convergence. We find the sixth order spatial differencing dominates the error at early times and fourth order time stepping dominates at late times as it accumulates faster. We therefore make the conservative assumption that the code is consistent with fourth order finite-differencing. As expected, we see the errors to fall off as $\frac{1}{R_{\text{ext}}}$.

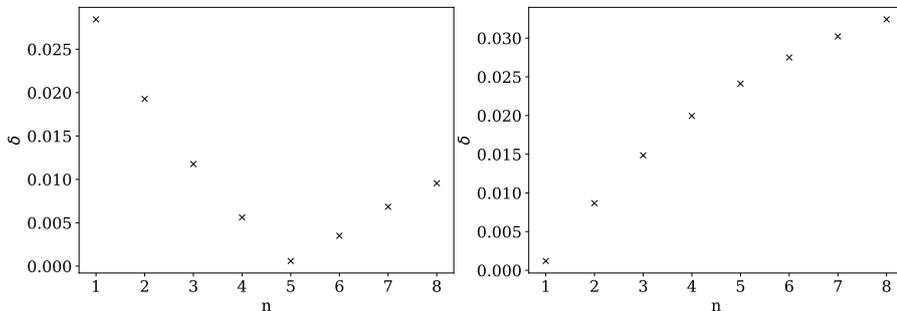


Figure 2.1: The value δ as given by equation 2.15 as a function of convergence order. The left panel shows the results for waveforms of differing resolution. The right panel shows the results for waveforms at differing extraction radii. The configuration $q = 8$, $\chi = 0.8$, $\theta = 150^\circ$ was used in this analysis.

We then calculated the Richardson extrapolated values of the amplitude, phase, α and β using the formula given in equation 2.4. These were used to estimate the error in each of these quantities. The error in the waveform quantities does not monotonically increase with decreasing resolution since not all of the resolutions lie in the convergence regime. The error in the low and medium resolution waveforms is therefore estimated from the combination of the error between these waveforms and the very high resolution waveform and between the very high resolution waveform and the “true” waveform. An example of the growth of the error as a function of resolution is shown in figure 2.2 and as a function of the extraction radius in figure

2.3. In these examples we show the error in the phase. As can be seen from figure 2.2, the dephasing of the waveform due to the finite resolution is ~ 0.1 radians for the medium resolution simulation. Similarly, from figure 2.3, the dephasing due to the finite extraction radius is ~ 0.4 radians for the waveform extracted at 90M. The total phase error in the waveform is therefore estimated to be about 0.4 radians by combining the errors in quadrature.

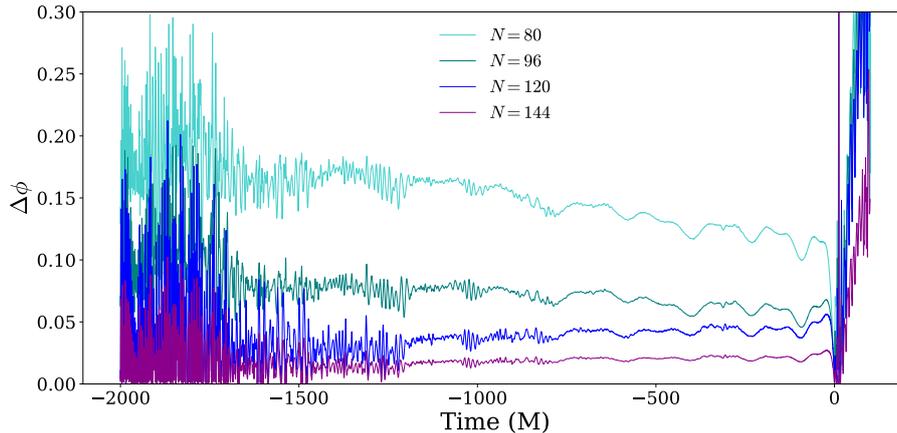


Figure 2.2: Variation of the absolute error in the time domain co-processing phase with simulation resolution. The phases have been aligned at merger and are measured in radians.

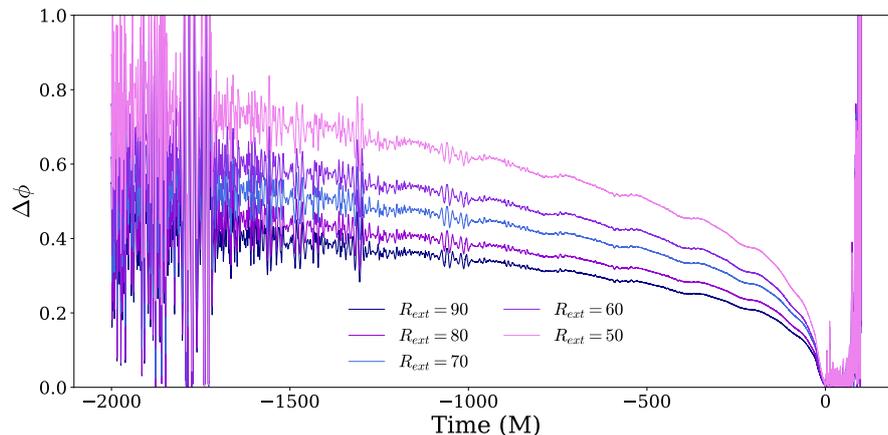


Figure 2.3: Variation of the absolute error in the time domain co-processing phase with extraction radius. The phases have been aligned at merger and are measured in radians.

For each of the cases considered in this chapter, the relative error in the amplitude and phase of the co-processing waveform are of the order of a few percent, while the relative error in the precession angles is around half a percent. The exact values of these errors for the $q = 8$, $\chi = 0.8$, $\theta = 150^\circ$ configuration are given in table

	% Error (1s.f.)		
	Resolution	Extraction Radius	Total
ϕ	0.2	1	1
A	0.5	8	8
α	0.3	0.03	0.3
β	0.5	0.1	0.5

Table 2.4: Relative error in the waveform quantities compared with the Richardson extrapolated quantities for the $q = 8$, $\chi = 0.8$, $\theta = 150^\circ$ configuration.

2.4. These are representative of the results for all cases. This is relevant for the tuned precessing model presented in Chapter 4 since it implies that the model for the precession angles cannot be accurate to more than 0.5%.

The errors in the amplitude and the precession angles are affected by the dephasing in the waveform. Therefore, although these results are a good diagnostic for the reliability of the code and a good way to compare accuracy between different simulations performed with the same code, they are difficult to translate into meaningful measures of the accuracy from waveform modelling or other gravitational wave applications. In order to get a more meaningful estimate of the accuracy of the waveform we performed the mismatch analysis presented in the following section.

2.6 Matches

The waveform quantities examined in the previous section are the traditional quantities used when estimating the convergence order and accuracy of a NR code. When assessing the accuracy of a waveform it is more useful to look at the mismatch. In the following section we estimate the mismatch between waveforms at finite resolution and extraction radius with the “true” waveform which is infinitely well resolved and infinitely far from the source.

Since these are precessing configurations, we calculated precessing matches as described in section 1.6.2. The matches used the aLIGO sensitivity curve shown in figure 1.1. They were performed over the frequency range from 10Hz to $0.11Mf$, which corresponds to the end of the NR waveform. In order to see how the match varied over a range of total masses that might be observed by current ground based detectors, we further calculated the power-weighted match as is described in 1.6.3, using PhenomPv3 as the model for the low frequency part of the waveform. We then calculated the mismatch as given by 1.112. These matches use all the $\ell = 2$ modes.

We performed matches between the high resolution waveforms and all other resolutions available for a given configuration. Each of the waveforms involved in these matches were extracted at $R_{\text{ext}} = 90\text{M}$. These results are shown in figure 2.4. We also performed matches between waveforms extracted at $R_{\text{ext}} = 90\text{M}$ and each

of the other extraction radii on the same level for the medium resolution simulations for each configuration. These results are shown in figure 2.5. From figure 2.5 we can see that the matches generally follow a trend implying the waveform is being extracted sufficiently far from the source so that we are in the convergence regime. The matches at between waveforms extracted at $R_{\text{ext}} = 90M$ and $R_{\text{ext}} = 50M$ do not follow this trend. This is not true for the mismatches with respect to resolution shown in figure 2.4 where the matches between (i) the low and high resolutions and (ii) the medium and high resolutions clearly do not follow a trend for most of the configurations. This implies that the low and medium resolutions do not lie in the convergence regime.

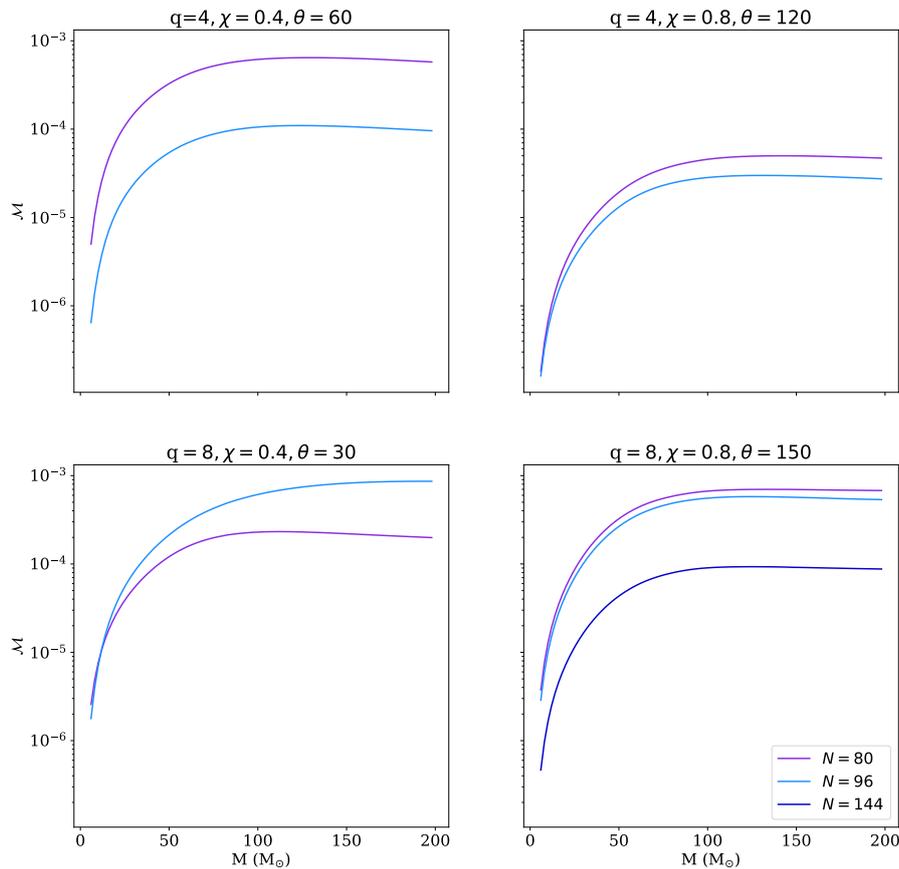


Figure 2.4: Mismatch between waveforms at varying resolution against the high resolution waveform as a function of the total mass of the system.

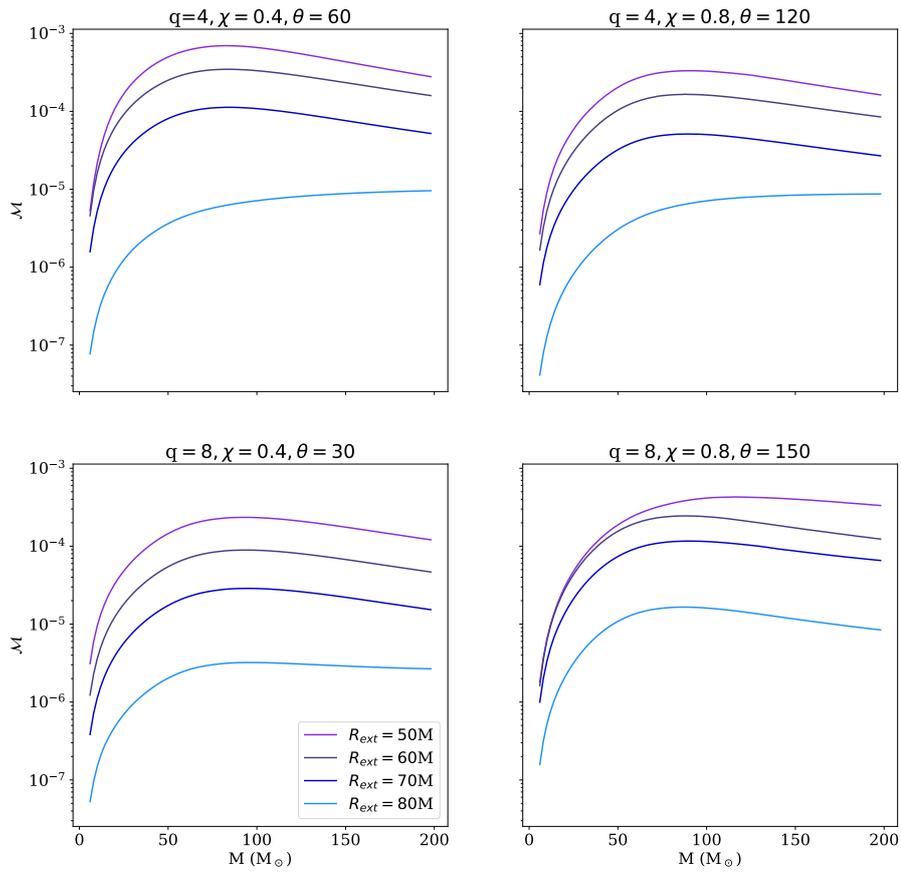


Figure 2.5: Mismatch between waveforms at varying extraction radii and the waveform extracted at $R_{\text{ext}} = 90M$ as a function of the total mass of the system.

2.6.1 Convergence order

The low and medium resolutions do not lie in the convergence regime. While the low resolution lies far outside the convergence regime, the medium resolution approaches the convergence regime. We therefore do not expect analyses involving the medium resolution waveform to show exact convergence but it can be used to provide an estimate of the convergence order. This is demonstrated clearly in figure 2.6 where we show the mismatch between the medium and high and the high and very high resolutions, as well as the prediction of the medium to high resolution mismatch based on the high to very high resolution mismatch using equation 2.12 for varying convergence order. From this analysis it is clear that the mismatch is closest to being fourth order convergent. This analysis could only be done for the case ($q = 8$, $\chi = 0.8$, $\theta = 150^\circ$) since this is the only case for which we have the very high resolution run.

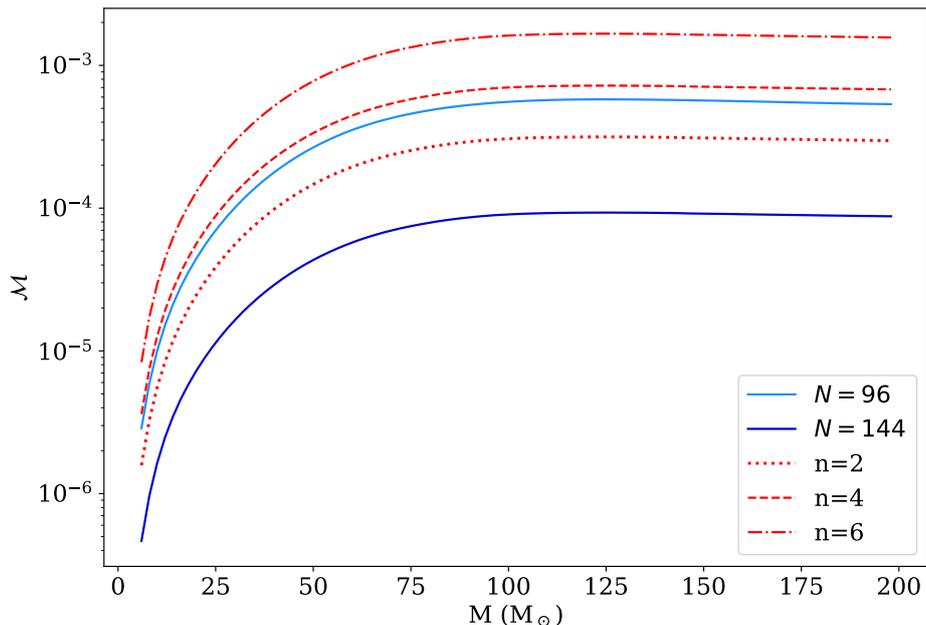


Figure 2.6: Mismatches demonstrating fourth order convergence of the BAM code with respect to resolution. The solid lines show the calculated mismatch while the dotted red lines show the predicted mismatch for varying convergence order. This is for the $q = 8$, $\chi = 0.8$, $\theta = 150^\circ$ configuration.

Since the waveforms extracted at varying extraction radii mostly lie within the convergence regime, we calculated the ratio of the mismatch between each of the pairs of waveforms from different extraction radii using equation 2.12 for varying convergence order. For each of the four configurations we investigated, it was found that the results were most consistent with first order convergence. This is demon-

strated in figure 2.7 where the solid lines show the calculated mismatch between two waveforms of different extraction radii and the dotted red line shows the expected value of the match for first order convergence.

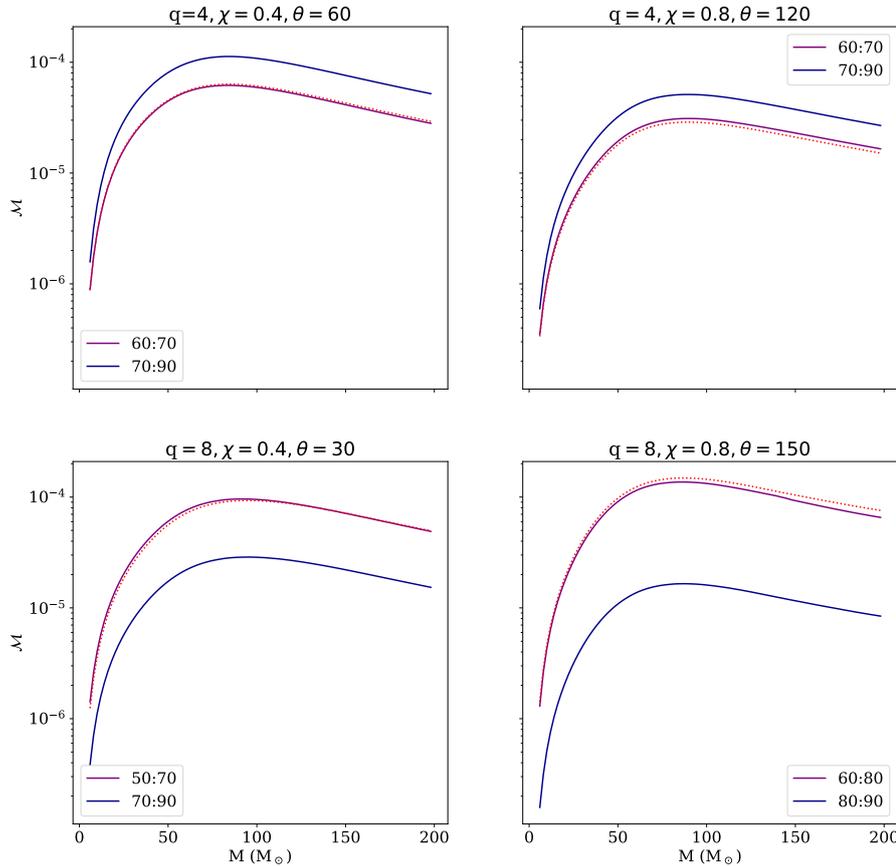


Figure 2.7: Mismatches demonstrating first order convergence of the BAM code with respect to extraction radius. The solid lines show the calculated mismatch, while the dotted red line shows the predicted mismatch for the pair of waveforms shown in purple based on the mismatch in blue, assuming first order convergence.

Not all the waveforms from the different extraction radii show perfect convergence for every configuration. The mismatch between $R_{\text{ext}} = 80$ and $R_{\text{ext}} = 90$ often does not follow the trend— we expect this is because the mismatch between these waveforms is so small ($\mathcal{O}(10^{-6})$) that it is very sensitive to any data processing performed in the course of calculating the match. The mismatch between $R_{\text{ext}} = 50$ and $R_{\text{ext}} = 90$ also often does not follow the trend and we do not expect it to hold for small extraction radii.

The convergence order calculated using this method agrees with the estimate

calculated in the previous section; the code is approximately fourth order convergent with respect to resolution and first order convergent with respect to the extraction radius.

2.6.2 Extrapolation

Having estimated the convergence order of the code n , we can calculate the convergence relation of the mismatches shown in figures 2.4 and 2.5. The convergence relation takes the form

$$\mathcal{M}(\Delta : \Delta_{\text{ref}}) = \kappa (\Delta^n - \Delta_{\text{ref}}^n)^2, \quad (2.16)$$

where Δ_{ref} is the value of the expansion parameter at the reference resolution ($N = 120$) or extraction radius ($R_{\text{ext}} = 90M$) and κ is a co-efficient to be calculated. The mismatch between the reference waveform and the “true” waveform is therefore given by $\kappa \Delta_{\text{ref}}^{2n}$.

We first look at the mismatch due to the finite resolution of the simulation. We assume that the two highest resolution simulations we performed ($N = 120$ and $N = 144$) lie in the convergence regime but we know the two lower resolution simulations do not. Assuming fourth order convergence, we therefore calculate κ_{res} using

$$\kappa_{\text{res}} = \frac{\mathcal{M}(\Delta_{144} : \Delta_{120})}{(144^4 - 120^4)^2}. \quad (2.17)$$

From figure 2.8, which shows the convergence relation for the mismatches calculated for a system with total mass $100M_{\odot}$, we can see that this appears to be a reasonable assumption. From κ_{res} we can estimate the mismatch between the high or very high resolution waveforms with an infinitely well resolved waveform. However, we want to know the mismatch for the medium resolution runs since this is the resolution that was used to perform the simulations for the catalogue of waveforms presented in tables 2.1 and 2.2. Since this resolution does not lie in the convergence regime (and the phase error does not improve monotonically from the medium resolution to the high and very high resolution waveforms) we cannot simply use the calculated convergence relation in order to estimate the mismatch between a waveform at this resolution and the “true” waveform. Instead we use the formula given in equation 2.13 to add the mismatch between the medium resolution and the very high resolution waveforms to the mismatch between the very high resolution waveform and

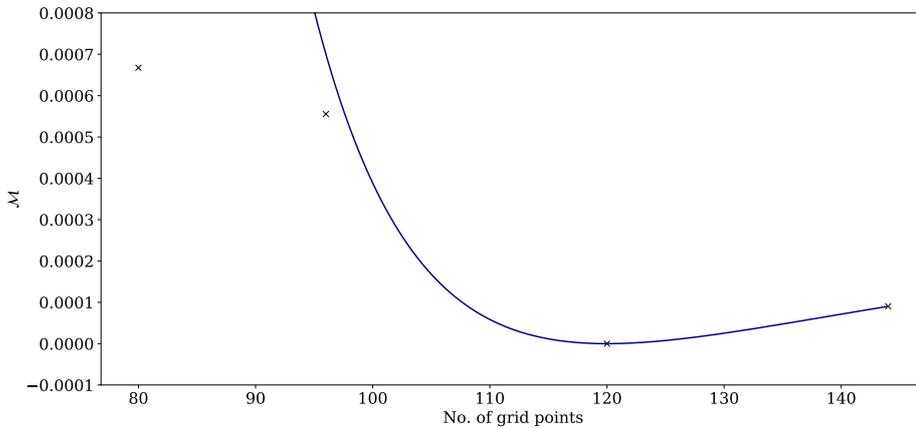


Figure 2.8: Variation of the mismatch with resolution for a $q = 8$, $\chi = 0.8$, $\theta = 150^\circ$ system with a total mass of $100M$. The mismatches shown are with respect to the high resolution run and so the mismatch is 0 at $N = 120$. The line shows the relation in equation 2.16 with κ_{res} assuming fourth order convergence calculated using equation 2.17 and is consistent with all resolutions except the lowest one at $N = 80$.

the “true” waveform:

$$\begin{aligned} \mathcal{M}(\Delta_{96} : \Delta_\infty) &= \left(\sqrt{\mathcal{M}(\Delta_{96} : \Delta_{144})} + \sqrt{\mathcal{M}(\Delta_{144} : \Delta_\infty)} \right)^2 \\ &= \left(\sqrt{\mathcal{M}(\Delta_{96} : \Delta_{144})} + \sqrt{\frac{\kappa_{\text{res}}}{144^4}} \right)^2. \end{aligned} \quad (2.18)$$

The result of this extrapolation procedure is shown in figure 2.9. We could only perform this calculation for the case ($q = 8$, $\chi = 0.8$, $\theta = 150^\circ$) since this is the only case for which we have the very high resolution run. However, from figure 2.4 we can see that the mismatch between the medium and high resolution runs is the worst for this case, so this estimate should give an upper bound for the mismatch between the medium resolution run and the “true” waveform. The maximum mismatch between a medium resolution waveform and an infinitely well resolved one is 6.0×10^{-4} .

We next examine the mismatch due to the finite distance from the source at which the waveform is extracted. To calculate the first order convergence relation with respect to the extraction radius, we performed a fit through each of the mismatches which were found to follow the convergence relation. This is demonstrated for mismatches between waveforms of different extraction radii and the waveform at $R_{\text{ext}} = 90M$ in figure 2.10, for a system with total mass $100M_\odot$. This fit gives the value of κ_{ext} for every value of the total mass of the system. From this we can calculate the mismatch between the waveform at $R_{\text{ext}} = 90M$ and the “true” waveform from $\mathcal{M}(\Delta_{90} : \Delta_\infty) = \frac{\kappa_{\text{ext}}}{90^2}$.

The mismatch between the waveform extracted at $R_{\text{ext}} = 90M$ and the “true”

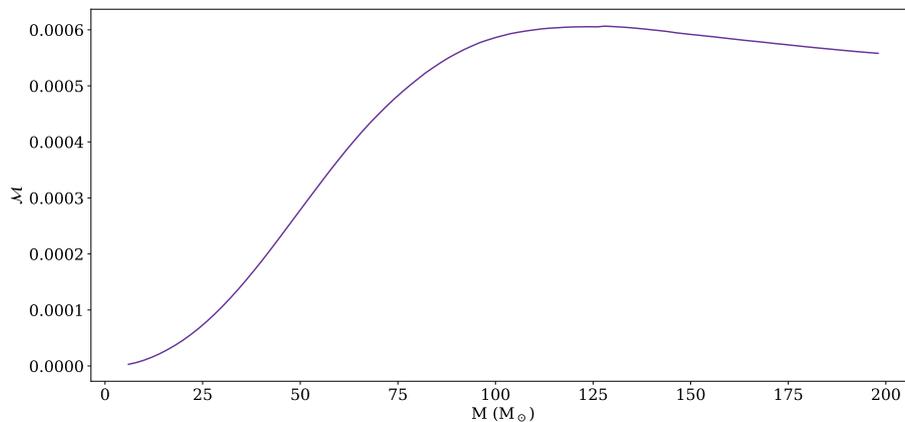


Figure 2.9: Projected mismatch between a waveform extracted at a resolution of $N = 96$ and one that is infinitely well resolved.

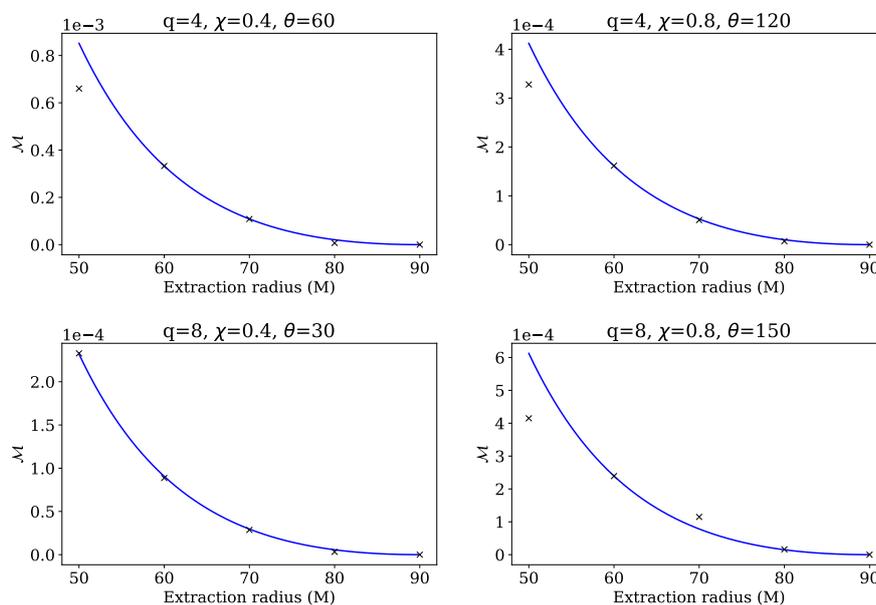


Figure 2.10: Variation of the mismatch with extraction radius for a system with a total mass of $100M$. The mismatches shown are with respect to the waveform extracted at $R_{\text{ext}} = 90M$. The line shows the fit based on equation 2.15 assuming first order convergence. It is consistent with all extraction radii except at $R_{\text{ext}} = 50M$.

waveform is shown in figure 2.11. The configuration which gives the greatest mismatch is $q = 4$, $\chi = 0.4$, $\theta = 60$ since, as noted above, this simulation was much longer than the others and so has greater opportunity to accumulate phase error between the two waveforms. The maximum mismatch between a waveform at $R_{\text{ext}} = 90$ and at $R_{\text{ext}} \rightarrow \infty$ is taken to be 1.4×10^{-3} .

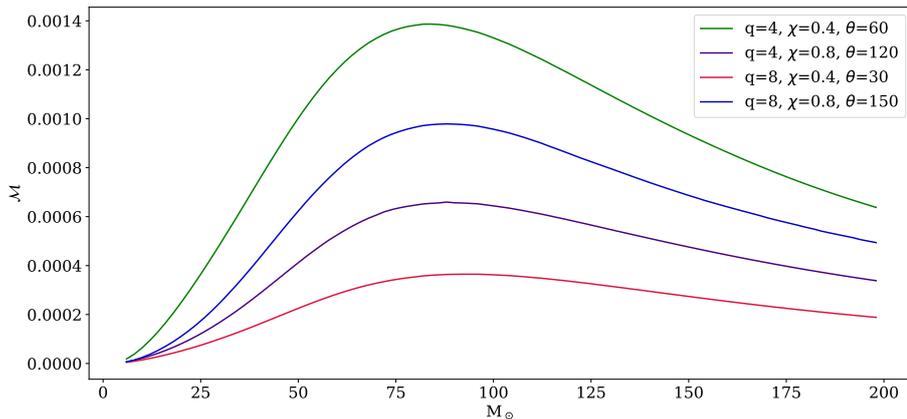


Figure 2.11: Projected mismatch between a waveform extracted at a radius of $R_{\text{ext}} = 90M$ and one extracted infinitely far away.

We estimate the mismatch between our medium resolution waveform extracted at $R_{\text{ext}} = 90M$ and the true waveform using

$$\mathcal{M} = \left(\sqrt{\mathcal{M}_{\text{resolution}}} + \sqrt{\mathcal{M}_{\text{extraction radius}}} \right)^2. \quad (2.19)$$

A conservative estimate of the mismatch between a waveform extracted at a finite distance of $90M$ from the source for a simulation performed with a grid spacing $d = 0.0104$ and the theoretical ‘analytical’ solution is therefore 3.83×10^{-3} . This provides a limit on the mismatch error of the waveforms used to tune the precessing model of 0.4% .

2.7 Conclusion

We have produced a catalogue of 40 single spin precessing NR waveforms for use in the construction of a tuned precessing model. The catalogue covers more extreme parts of the parameter space (those with high mass ratio and spin magnitude) where a tuned precessing model will be most useful. The catalogue does not include two-spin cases, cases where the azimuthal spin component varies or cases that exhibit transitional precession since these cases are expected to contribute only to subdominant precession effects, or be extremely rare observations.

We have assessed the convergence order of the BAM code for these simulations

and found that it is fourth order convergent with respect to resolution and first order convergent with respect to extraction radius. From this we have been able to estimate the error in the waveform quantities computed from the simulation using Richardson extrapolation. Most importantly, we can see that the error in the precession angles α and β (which are used in the production of the tuned precessing model) is $\mathcal{O}(10^{-1})\%$. We also estimated the mismatch between the waveforms produced by simulations performed at the finite resolution ($d = 0.0104$) and extracted at the finite distance ($R_{\text{ext}} = 90\text{M}$) used to produce the waveforms in the catalogue and the “true” infinitely well resolved waveform extracted infinitely far away from the source. From this we can see that the NR waveforms agree with the “true” waveform to within 0.4%. We therefore find that the limiting accuracy of the input waveforms for the tuned precessing model is 0.4%.

Chapter 3

Orbital dynamics

3.1 Introduction

In recent years the LIGO and Virgo detectors [3, 5, 17] have made the first observations of binary black hole (BBH) systems, through measurements of their gravitational wave emission [6, 7, 9, 10, 13–16]. The properties of the black holes can be measured by comparing the signal against theoretical gravitational wave models [4, 8], which are informed in part by NR solutions of Einstein’s equations for the last orbits and merger of two black holes (see, e.g., the review [83]). NR waveforms have also been used to assess the systematic errors of the gravitational wave measurements [12]. To use NR waveforms as proxy signals one must specify the binary’s orientation and orbital phase at a particular time or signal frequency. There is an inherent ambiguity in doing this, because the binary’s dynamics cannot be directly related to the waveform. The purpose of this chapter is to define an effective binary orientation and phase, which can be calculated directly from the waveform, and compare it against the co-ordinate dynamics in NR simulations.

The general theory of relativity predicts gravitational waves that travel at the speed of light, c .¹ In principle, we can relate the dynamics of two orbiting black holes to a gravitational wave signal a distance d away, through a retarded time, $t_{\text{GW}} = d/c$. This is possible in a PN calculation [36], where the signal can be calculated explicitly from point-particle dynamics. An equivalent identification has not been rigorously defined for solutions of the full nonlinear Einstein equations, which are calculated numerically. The proper distance from the source to the observer is not a well-defined concept. We lack unique definitions of mass, angular momentum and centre-of-mass in general relativity [156]; in a numerical simulation the binary dynamics depend non-trivially on the gauge (co-ordinate) conditions used in the evolution of Einstein’s equations; proper distances depend on the dynamical curvature across the intervening spacetime; and gravitational waves are only rigorously defined at null infinity. In practice, these formal ambiguities lead to negligible uncertainties

¹Throughout this chapter we will adopt geometric units, $G = c = 1$.

in gravitational wave signal modelling and source measurements; see, for example, section IV.B of [142] in the case of waveform modelling, and [45] for a discussion of retarded times in NR simulations.

The situation is different when we wish to use NR waveforms as proxy signals. A binary configuration is specified by the black-hole masses and spin magnitudes, but also by the binary orientation, orbital phase, and spin directions at a particular time or frequency during the binary’s inspiral. Now we must relate the dynamics to the signal. Given the above, we are forced to make approximations. One way to do this is to define an approximate retarded time. Another is to note that during the inspiral the frequency of the dominant signal harmonic is, to a good approximation, twice the orbital frequency, and to map the dynamics at each orbital frequency to the corresponding signal frequency. A similar mapping can be made using the orbital and signal phases, although the two approaches will not give identical results, as we discuss in section 3.4.1.

In this chapter, we take a different approach. We define a binary orientation and phase with respect to the gravitational wave signal only. The starting point is the earlier work in [150], which proposed studying the direction of maximum gravitational wave emission (or dominant emission direction), which was called the direction of “quadrupole alignment” (QA). The results in [150] suggested that the QA direction may track the direction of $\hat{\mathbf{L}}$. If $\hat{\mathbf{L}}$ is calculated using a PN approximation, then the leading-order (Newtonian) contribution is the normal to the orbital plane, which exhibits nutation, but when all known PN terms are included, the full $\hat{\mathbf{L}}$ precesses smoothly. In the NR example studied in [150], the QA direction precessed smoothly without nutation and agreed well with the (appropriately time-shifted) direction of $\hat{\mathbf{L}}$. This led the authors to suggest that the QA direction may track the orbital angular momentum, rather than the orbital-plane direction. More recent work has shown that this direction varies between different radiation frames, and also depends on whether the direction is calculated using the gravitational wave strain h , the Bondi news (the first time derivative of h), or the Weyl scalar ψ_4 (the second time derivative) [44, 113, 119]. Nonetheless, in general these differences are small, and any given choice of the QA direction provides us with an ideal means to define a proxy to the binary orientation with respect to the gravitational wave signal alone. Since the gravitational wave signal is the only invariant observable we have access to, this orientation provides a robust measure to identify and compare simulations.

As described in more detail in section 1.4 in the Introduction, the first QA definitions [123, 150] specified only the two Euler angles needed to transform into a frame that tracks the precession. A third Euler angle is also needed to uniquely specify the phase (up to an overall constant). A method to calculate the third angle is given in [46], completing the definition of a co-precessing frame. In this work we use that procedure to define a proxy orbital phase, Φ , from the gravitational wave signal,

which in turn allows us to define a proxy orbital separation unit vector $\hat{\mathbf{n}}$, which we compare with those quantities calculated directly from the orbital dynamics. Once again, we show that this does not provide an exact mapping to the phase calculated directly from the dynamics, even if time shifts and gauge effects could be removed; but Φ does serve as a phase that is in principle gauge invariant and uniquely defined. An alternative prescription to find $\hat{\mathbf{n}}$ from the gravitational wave signal is suggested in [42] although to our knowledge its efficacy for NR waveforms has not been explored.

To connect our work to the practical problem of constructing proxy gravitational wave signals from NR waveforms, we describe our work and results using the notation and conventions of the NR Injection Infrastructure [151], which provides a consistent way to go from waveforms produced using a variety of NR gravitational wave codes to waveforms that are suitable for injections as a “discrete” waveform approximant for use with the LIGO Algorithm Library (LAL). The LAL framework requires injected waveforms to be in a frame that describes the wave propagation from the source to gravitational wave detectors on Earth. The NR Injection Infrastructure rotates the waveforms into this format. These rotations require the unit orbital angular momentum of the binary, $\hat{\mathbf{L}}$, and the unit separation vector of the two black holes, $\hat{\mathbf{n}}$. The unit separation vector $\hat{\mathbf{n}}$ can be constructed from the normal to the orbital plane and the orbital phase; our approach will be to define $\hat{\mathbf{n}}$ from $\hat{\mathbf{L}}$ and Φ . These quantities are currently calculated using the dynamics information provided by a simulation. To relate these dynamical data to the gravitational wave signal, one either uses an estimate of the retarded time t_{GW} (provided along with the NR waveform, and corresponding to Format 1 in [151]), or maps the orbital frequency $\Omega(t)$ to the signal frequency (Formats 2 and 3). The method we propose is equivalent to mapping the orbital *phase* to that of the signal, and without any of the gauge ambiguities of the black-hole co-ordinate dynamics.

The chapter is organised as follows. In section 3.2 we describe the rotations performed by the NR Injection Infrastructure. In section 3.3 we summarise the procedure to find the unit orbital angular momentum, which is described in more detail in Refs. [46, 123, 150], and describe how to also find the co-precessing phase and the unit separation vector from the waveform. Section 3.4 describes the various co-ordinate ambiguities associated with these calculations. In section 3.5 we compare $\hat{\mathbf{L}}$, Φ , and $\hat{\mathbf{n}}$, which have been calculated from the waveform, with those found from the dynamics. We also discuss how the different choices of time shift affect this comparison and show why it is important to ensure a consistent choice is used.

3.2 Frame conventions

In this section we summarise three co-ordinate systems used to specify gravitational wave signals. We follow the conventions and notations used in the Numerical Rela-

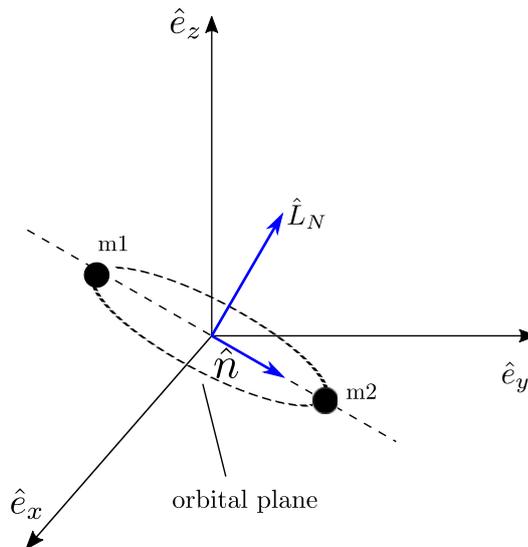


Figure 3.1: The binary properties in the NR simulation frame (black) at a time t_{ref} . The binary is then rotated to the LAL source frame (blue) where $\hat{\mathbf{z}}$ is parallel to the (Newtonian) orbital angular momentum $\hat{\mathbf{L}}_N$ at time t_{ref} and $\hat{\mathbf{x}}$ is aligned along $\hat{\mathbf{n}}$.

tivity Injection Infrastructure [151]. Gravitational wave signals are represented by the gravitational-wave strain, which corresponds to the metric perturbation h_{ij}^{TT} . Numerical simulations calculate the Weyl scalar ψ_4 , from which h_{ij}^{TT} can be found by integrating twice with respect to time [140]. In numerical simulations this perturbation is extracted far from the orbiting black holes, where the spacetime is approximately flat. This region of spacetime is known as the *wave zone* [159]. The waves are extracted at a retarded time t_{GW} . In the wave zone, a Cartesian coordinate system $(\hat{\mathbf{e}}_x, \hat{\mathbf{e}}_y, \hat{\mathbf{e}}_z)$ is used. This co-ordinate system can be related to polar co-ordinates $(\hat{\mathbf{e}}_r, \hat{\mathbf{e}}_\theta, \hat{\mathbf{e}}_\phi)$. The strain can then be decomposed into modes in a basis of spin-weighted spherical harmonics, ${}^{-2}Y_{\ell m}$, and is written as

$$h^{\text{NR}}(t_{\text{GW}}; \theta, \phi) = h_+^{\text{NR}} - ih_\times^{\text{NR}} \quad (3.1)$$

$$= \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} H_{\ell m}(t_{\text{GW}}) {}^{-2}Y_{\ell m}(\theta, \phi), \quad (3.2)$$

where the extracted gravitational wave modes can be expressed as

$$H_{\ell m}(t_{\text{GW}}) = A_{\ell m}(t_{\text{GW}}) e^{-i\Phi_{\ell m}(t_{\text{GW}})}. \quad (3.3)$$

We adopt the convention that for a binary orbiting counter-clockwise in the plane defined by $\hat{\mathbf{e}}_x \times \hat{\mathbf{e}}_y$, $\Phi_{22}(t_{\text{GW}})$ is a monotonically increasing function. This is the opposite convention to that used in [151].

Once the gravitational wave has been extracted and decomposed as described above it needs to be prepared for injection. This involves transforming the waveform from the frame in which it has been generated (the *NR simulation frame*) into the frame in which the binary is viewed from Earth. This is done in two stages. First, the waveform is rotated into a frame defined by certain properties of the binary at a given reference time. The choice of this frame is arbitrary but must be consistent between injections. A set of conventions in defining this frame, known as the *LAL source frame* [1,151], are therefore chosen. These conventions are described below. In this frame, waveforms generated by a particular binary should be the same regardless of the code used to generate them or the choice of co-ordinate system in the original simulation. From this intermediate frame, the waveform is then rotated into the final frame, the *wave frame*, defined by the relationship between the binary and the observer. The NR simulation frame and LAL source frames are shown in figure 3.1.

In the NR simulation frame one can define the separation vector of the two black holes as the direction from body 2 to body 1 (where body 1 is the heavier object) given by

$$\mathbf{n} = \mathbf{r}_1 - \mathbf{r}_2, \quad (3.4)$$

where \mathbf{r}_i is the position of the centre of the i th body. The Newtonian orbital angular momentum of the binary can be defined as

$$\mathbf{L}_N = \mathbf{L}_1 + \mathbf{L}_2 = \sum_{i=1}^2 m_i (\mathbf{r}_i \times \mathbf{v}_i), \quad (3.5)$$

where m_i is the mass and \mathbf{v}_i the velocity of the i th object. In moving-puncture codes, \mathbf{r}_i will be the puncture positions and m_i will be given by the apparent horizon masses, and in excision codes \mathbf{r}_i will be the co-ordinate centres of the apparent horizons and m_i will be the Christodoulou masses [61].

The LAL source frame is defined as the frame where the co-ordinate axes satisfy the following equalities

$$\hat{\mathbf{x}}^{\text{ref}} \equiv \hat{\mathbf{n}} \quad (3.6)$$

$$\hat{\mathbf{y}}^{\text{ref}} \equiv \hat{\mathbf{L}}_N \times \hat{\mathbf{n}} \quad (3.7)$$

$$\hat{\mathbf{z}}^{\text{ref}} \equiv \hat{\mathbf{L}}_N \quad (3.8)$$

at a reference epoch defined either by a reference time t_{ref} or a reference orbital frequency $\Omega_{\text{ref}}^{\text{orb}}$ where $\Omega^{\text{orb}}(t_{\text{ref}}) = \Omega_{\text{ref}}^{\text{orb}}$. Choosing a different reference epoch will in general produce a different source frame.

Finally, the waveform is rotated into the *wave frame*. In this frame the $\hat{\mathbf{Z}}$ axis points towards the observer along the line of sight while the $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$ vectors are

orthogonal to the line of sight. The intersection of the orbital plane with the XY axis is referred to as the line of ascending node. The transformation to the wave frame is given in [151]; for the remainder of this chapter we will work in either the inertial NR simulation frame or the co-precessing (quadrupole-aligned) frame, as described in section 3.3.

Ambiguities in this procedure arise from the NR simulation data. The NR simulation frame is the co-ordinate system in which the numerical simulation was performed. The physical interpretation of the co-ordinates in the NR simulation frame depends on the co-ordinates of the initial data, and on the gauge conditions used during the numerical evolution. If simulations with two different codes, using different initial-data constructions and different gauge conditions, are used to simulate the same physical system, then in principle we expect the asymptotic gravitational-wave signals to be the same, but the black-hole dynamics in the respective NR simulation frames may not be. We aim to circumvent these ambiguities in the method that we propose in the following sections.

3.3 Determining $\hat{\mathbf{L}}_{\mathbf{N}}$ and $\hat{\mathbf{n}}$

Currently the NR Injection Infrastructure requires $\hat{\mathbf{n}}$ and $\hat{\mathbf{L}}_{\mathbf{N}}$ to be provided as part of the metadata associated with each waveform. It is recommended that these quantities are calculated using equations 3.4 and 3.5 respectively (and then normalising). The positions and velocities of the black holes required to calculate these quantities come from the dynamics of the binary. This information, along with the spins of the black holes, forms part of the metadata provided with each NR waveform. There are several sources of ambiguity in the resulting choice of LAL source frame (as defined via equations 3.6–3.8). One is the gauge dependence of the co-ordinate dynamics and spin measurements. (Broadly speaking, codes that use variants of the generalized-harmonic evolution system [134, 135], like SpEC [147, 157], use harmonic-like co-ordinates [110], while moving-puncture codes [28, 54] use ADMTT-like co-ordinates [96]. For an example of one comparison between these co-ordinates, see Appendix D of [48]). The black-hole dynamics information can be mapped to the waveform using either a retarded time, or relating the gravitational wave frequency with the orbital frequency. Data provided in Format 1 of the NR Injection Infrastructure specifies the dynamics at the reference time only. The choice of time shift that relates the reference time in the waveform to the reference time in the dynamics quantities is made by the group supplying the waveform. Data provided in Formats 2 and 3 provide the dynamics information for the entire time evolution of the waveform. The dynamics quantities are then mapped to the waveform by finding the value of $\hat{\mathbf{n}}$ and $\hat{\mathbf{L}}_{\mathbf{N}}$ at the time at which the gravitational wave frequency is equal to the orbital frequency. If the retarded time is used, then a further ambiguity arises from the definition of retarded time t_{GW} used by a particular NR group.

These ambiguities could be resolved by finding $\hat{\mathbf{n}}$ and $\hat{\mathbf{L}}$ from the waveform. The NR waveform includes some error due to extraction at a finite co-ordinate radius, or due to approximate extrapolation to infinity (as discussed in the previous chapter), but in general exhibits far less gauge variation than the dynamics. We now describe a procedure to do this.

3.3.1 Determining $\hat{\mathbf{L}}$ from the waveform

It has already been shown that the approximate direction of the orbital angular momentum of a binary can be found from the waveform of the emitted gravitational waves using a variety of methods [123, 150]. [123] describes how this quantity can be found from the dominant principal axis of the quadrupolar part of the radiation axis. We use the Quadrupole Alignment procedure [150] described in section 1.4.1 of the Introduction; the two methods can be shown to be equivalent [46] when the method of [123] is restricted to the $\ell \leq 2$ modes. This procedure finds the frame in which $|\psi_{4,22}|^2 + |\psi_{4,2-2}|^2$ is maximised. In this frame $\hat{\mathbf{z}} \parallel \hat{\mathbf{L}}$. This transformation requires two angles (α and β), which define the rotation into a co-precessing frame, i.e. a frame that precesses along with the binary. In order to uniquely define this frame (up to an overall constant rotation, corresponding to a constant phase shift in the waveform in the co-precessing frame) we apply the minimum rotation condition [46], as defined in equation 1.76 in the Introduction. This angle is determined up to an integration constant, which corresponds to a constant rotation. A time-dependent rotation $\mathbf{R}(\gamma\beta\alpha)$ can then be performed between the inertial frame in which the simulation was performed and the co-precessing frame using the three Euler angles (α, β, γ). Using the $z - y - z$ convention, the $\psi_{4,\ell m}$ modes obey the transformation law

$$\psi_{4,\ell m}^{QA} = \sum_{m'=-\ell}^{\ell} e^{im'\gamma} d_{m'm}^{\ell}(-\beta) e^{im\alpha} \psi_{4,\ell m'} \quad (3.9)$$

where $\psi_{4,\ell m'}$ are the modes in the NR simulation frame and $\psi_{4,\ell m}^{QA}$ are the modes in the co-precessing (quadrupole-aligned) frame. $d_{m'm}^{\ell}$ are the Wigner d -matrices [80, 164].

The co-precessing frame rotates with the orbital angular momentum in order to ensure $\hat{\mathbf{L}}$ remains parallel to the z -axis at all times. Since $\hat{\mathbf{L}}$ is approximately perpendicular to the orbital plane, the orbital plane remains approximately in the xy -plane in the co-precessing frame.

3.3.2 Determining orbital phase and $\hat{\mathbf{n}}$ from the waveform

During the early inspiral of a non-precessing binary, the orbital phase of the binary can be found from the phase of the waveform, using

$$\omega_0^{\ell m} = m \frac{d}{dt} \phi_{\text{orb}}(t_0), \quad (3.10)$$

where $\omega_0^{\ell m}$ is the angular frequency of $\psi_{4,\ell m}$ and ϕ_{orb} is the orbital phase of the binary in the orbital plane. PN corrections to this relation are small [36], the differences between the phases of h and ψ_4 are also small [43], and this approximation holds to high accuracy even up until a few orbits before merger (see, e.g., equation 7 of [51]). Consequently, the orbital phase of the binary is half that of the phase of the $\psi_{4,22}$ mode. The phase of a $\psi_{4,\ell m}$ mode, $\Phi_{\ell m}^{\psi_4}$, is the unwrapped argument of the complex time series $\psi_{4,\ell m}$ given by

$$\psi_{4,\ell m} = A_{\ell m}^{\psi_4} e^{-i\Phi_{\ell m}^{\psi_4}}. \quad (3.11)$$

As stated above, the phase of the (2, 2) mode is a monotonically increasing function. Therefore, once the orbital phase has been calculated the unit separation vector is given by

$$\hat{\mathbf{n}} = \begin{pmatrix} \cos \phi_{\text{orb}} \\ -\sin \phi_{\text{orb}} \\ 0 \end{pmatrix} \approx \begin{pmatrix} \cos \left[\frac{1}{2} (\Phi_{22}^{\psi_4} + \Phi_0) \right] \\ -\sin \left[\frac{1}{2} (\Phi_{22}^{\psi_4} + \Phi_0) \right] \\ 0 \end{pmatrix}, \quad (3.12)$$

where Φ_0 is the orbital phase offset, which depends on the conventions used in the NR code used to produce the simulation. It is 0 if the phase of ψ_4 is 0 mod 2π when the black holes are on the x -axis, and $\pm\pi$ if this occurs when they are on the y -axis.

For a precessing binary, a similar procedure can be performed by rotating the waveform into the co-precessing frame described in section 3.3.1. The orbital phase can now be estimated from the co-precessing waveform phase $\Phi_{22}^{\psi_4,QA}$ in the same manner as above. This is similar to the way the orbital phase is computed in equation 6 in [35] except here we use just the phase of the (2,2) mode while [35] uses the average of the phase of the (2,2) and (2,-2) modes. For the remainder of this chapter, we will define the orbital phase, as estimated from the waveform, as $\Phi = (\Phi_{22}^{\psi_4,QA} + \Phi_0)/2$.

The unit separation vector can then be found as for a non-precessing waveform. It then needs to be rotated back into the NR simulation frame using the angles α , β and γ found above. Since these angles were defined using the $z - y - z$ convention, the rotations required to rotate a vector from the Quadrupole Aligned frame to the

NR simulation frame are

- rotate by γ about the z-axis
- then rotate by β about the y-axis
- then rotate by α about the z-axis.

This is given by

$$\begin{aligned}
 & \begin{pmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos \beta & 0 & -\sin \beta \\ 0 & 1 & 0 \\ \sin \beta & 0 & \cos \beta \end{pmatrix} \begin{pmatrix} \cos \gamma & -\sin \gamma & 0 \\ \sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \hat{n}_x^{\text{QA}} \\ \hat{n}_y^{\text{QA}} \\ \hat{n}_z^{\text{QA}} \end{pmatrix} = \\
 & \begin{pmatrix} \cos \alpha \left(\cos \beta \left(\cos \gamma \hat{n}_x^{\text{QA}} - \sin \gamma \hat{n}_y^{\text{QA}} \right) - \sin \beta \hat{n}_z^{\text{QA}} \right) - \sin \alpha \left(\sin \gamma \hat{n}_x^{\text{QA}} + \cos \gamma \hat{n}_y^{\text{QA}} \right) \\ \sin \alpha \left(\cos \beta \left(\cos \gamma \hat{n}_x^{\text{QA}} - \sin \gamma \hat{n}_y^{\text{QA}} \right) - \sin \beta \hat{n}_z^{\text{QA}} \right) + \cos \alpha \left(\sin \gamma \hat{n}_x^{\text{QA}} + \cos \gamma \hat{n}_y^{\text{QA}} \right) \\ \sin \beta \left(\cos \gamma \hat{n}_x^{\text{QA}} - \sin \gamma \hat{n}_y^{\text{QA}} \right) + \cos \beta \hat{n}_z^{\text{QA}} \end{pmatrix}.
 \end{aligned} \tag{3.13}$$

Since the waveform is rotated into the co-precessing frame by the three Euler angles, the orientation of $\hat{\mathbf{n}}^{\text{QA}}$ is determined up to a constant phase based on the choice of integration constant when calculating γ . However, when rotating $\hat{\mathbf{n}}^{\text{QA}}$ into the NR simulation frame, the rotation by γ removes this ambiguity, meaning $\hat{\mathbf{n}}$ is uniquely determined in the NR simulation frame regardless of the choice of integration constant.

3.3.3 Determining the co-precessing orbital phase

Alternatively, one can rotate the unit separation vector $\hat{\mathbf{n}}_{\mathbf{d}}$ (calculated from the positions of the black holes) into the co-precessing frame. This involves performing the above rotations in the reverse order using the Euler angles calculated from the Newtonian orbital angular momentum. The co-precessing orbital phase can then easily be calculated.

Since the Euler angle γ is found using integration a constant is introduced into the co-precessing phases. We determined this constant using the fact that $\arccos(\hat{\mathbf{n}}_{\mathbf{w}} \cdot \hat{\mathbf{n}}_{\mathbf{d}}) = \Phi - \phi_{\text{orb}}$.

3.3.4 Code Conventions

Several convention choices enter into the calculation of ψ_4 . These determine the relationship between the phase of ψ_4 and the orbital phase of the binary, i.e. they determine the orbital phase offset Φ_0 given in equation (3.12). The three relevant choices here are the sign convention in the definition of the Riemann and Weyl tensors, the definition of ψ_4 itself and the choice of origin of the azimuthal angle φ of

the spherical co-ordinates. The first two of these differences introduce an ambiguity in the definition of ψ_4 of $\psi_4 \rightarrow e^{i\psi_0}\psi_4$. The third introduces the ambiguity $\psi_{4,\ell m} \rightarrow e^{im\varphi_0}\psi_{4,\ell m}$ [53]

An example of the effect of different choices in these conventions is the difference in the phase of ψ_4 calculated by identical simulations produced using the BAM [49, 92] and SpEC [147, 157] codes. These have been explained in [53]. The two codes use the opposite sign convention in the definition of the Riemann and Weyl tensors. Additionally, a different choice of null tetrad is made when defining ψ_4 ; in the BAM code, ψ_4 is defined via $\psi_4 = -C_{\alpha\mu\beta\nu}n^\mu n^\nu \bar{m}^\alpha \bar{m}^\beta$ [53], while in the SpEC code ψ_4 is defined by $\psi_4 = -C_{\alpha\mu\beta\nu}\ell^\mu \ell^\nu \bar{m}^\alpha \bar{m}^\beta$ (see [128] and figure 4.3.1 of [62]). $(\ell^\mu, m^\mu, \bar{m}^\mu, n^\mu)$ is an appropriate null tetrad where ℓ and n are ingoing and outgoing null vectors respectively and $-\ell \cdot n = 1 = m \cdot \bar{m}$. $C_{\alpha\mu\beta\nu}$ is the Weyl tensor. These choices produce a phase offset of π ($\psi_0 = -1$) between ψ_4 calculated by BAM and by SpEC at equivalent points in the waveform for an identical simulation. The choice of the origin of φ can differ between simulations. However it seems that on the whole the choices made by BAM and SpEC do not introduce any additional phase offset.

These different choices of conventions mean that for BAM the phase of ψ_4 is (0 mod 2π) when the two black holes are on the x -axis (of the co-precessing frame) whereas for SpEC this happens when the two black holes are on the y -axis. Consequently, in order to calculate a value of $\hat{\mathbf{n}}$ which agrees with the dynamics information provided along with a simulation,

$$\hat{\mathbf{n}}_{\text{BAM}}^{\text{QA}} = \begin{pmatrix} \cos \Phi_{22}^{\psi_4, \text{QA}} \\ -\sin \Phi_{22}^{\psi_4, \text{QA}} \\ 0 \end{pmatrix}, \quad (3.14)$$

while

$$\hat{\mathbf{n}}_{\text{SXS}}^{\text{QA}} = \begin{pmatrix} -\sin \Phi_{22}^{\psi_4, \text{QA}} \\ -\cos \Phi_{22}^{\psi_4, \text{QA}} \\ 0 \end{pmatrix}. \quad (3.15)$$

The GT-MAYA [87, 90, 161] and RIT [166] codes appear to use the same conventions as the SpEC code. These conventions are also used when producing the PN waveforms outlined in [25, 36].

In general, a consistent convention for Φ_0 must be chosen. A choice of $\Phi_0 = \frac{\pi}{2}$ agrees with the PN convention. This will give a consistent definition of $\hat{\mathbf{n}}$ from the waveform, regardless of the convention choice of the NR code which determines the dynamics of the simulation. The individual code conventions need to be taken into account only when we wish to compare back to the co-ordinate dynamics of the original NR simulation.

3.4 Co-ordinate ambiguities

In this section we illustrate two co-ordinate ambiguities that we referred to earlier. The first is in the definition of the retarded time t_{GW} ; we consider typical choices of retarded time that have been used in numerical-relativity studies, and also the retarded times implied by aligning either the gravitational wave phase or frequency with the corresponding quantity calculated from the dynamics. The second is in the estimate of the orbital-plane orientation. We use a PN example to illustrate these ambiguities, in particular differences in the QA direction calculated using ψ_4 and h . The differences in these directions are nonetheless small, as we illustrate with both PN and NR examples.

3.4.1 Retarded time

As mentioned above, there is an ambiguity when relating information about the binary dynamics calculated at the source of the simulation to waveform information extracted at some finite co-ordinate distance from the source. Different groups use different conventions to define the relationship between the time at the source t and the retarded time t_{GW} . The two methods most commonly used are (i) to treat the spacetime as if it were flat and (ii) to assume the propagation time is given by the tortoise co-ordinate, as in, for example, [45] (although here an additional correction to the co-ordinate time is included). The two choices of retarded time can be summarised as

$$(i) \ t_{\text{GW}} = t + R_{ex}, \quad (3.16)$$

$$(ii) \ t_{\text{GW}} = t + R_{ex} + 2M \ln \left| \frac{R_{ex}}{2M} - 1 \right|, \quad (3.17)$$

where M is the initial total mass of the system and R_{ex} is the co-ordinate radius at which the gravitational wave signal was extracted from the NR simulation.

A further choice of the retarded time can be defined as the value of t_{GW} where the phase of the waveform is twice the orbital phase of the binary at a time t . This gives a time-dependent time shift between retarded time t_{GW} and source time t . An equivalent time shift can be used to align the angular frequency of the waveform with the orbital angular frequency.

These different conventions mean the metadata provided with the waveforms used in the NR Injection Infrastructure are not defined in a consistent manner. The method described in section 3.3 to find $\hat{\mathbf{L}}$ and $\hat{\mathbf{n}}$ removes this ambiguity and provides a consistent way of defining $\hat{\mathbf{L}}$ and $\hat{\mathbf{n}}$ for all waveforms. This method is equivalent to using the time shift that aligns the phase of the waveform with the orbital phase at each time step, in the co-precessing frame. This time-dependent time shift can then be used to also report the spins in a consistent manner by relating the values calculated by the NR simulation to the equivalent point in the waveform.

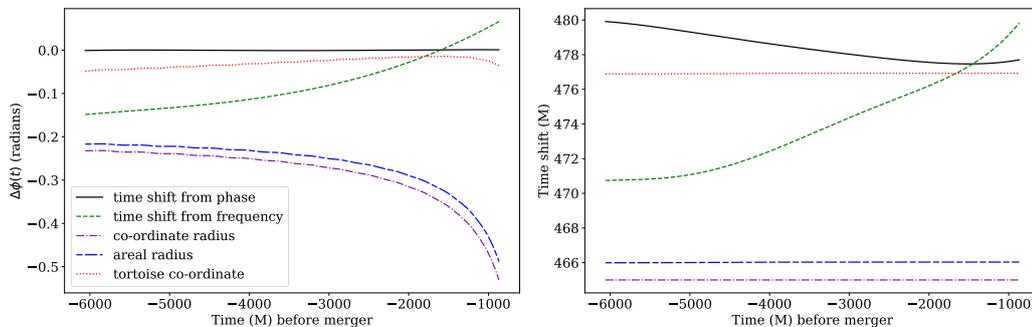


Figure 3.2: Simulation SXSBBH0152 ($q = 1$, $M\omega_{22}^{\text{start}} = 0.0297$, $\chi_1 = (0, 0, 0.6) = \chi_2$). The left hand plot shows the difference between the orbital phase estimate from the gravitational wave signal, Φ , and the orbital phase ϕ_{orb} for each of the time shifts shown in the right hand plot.

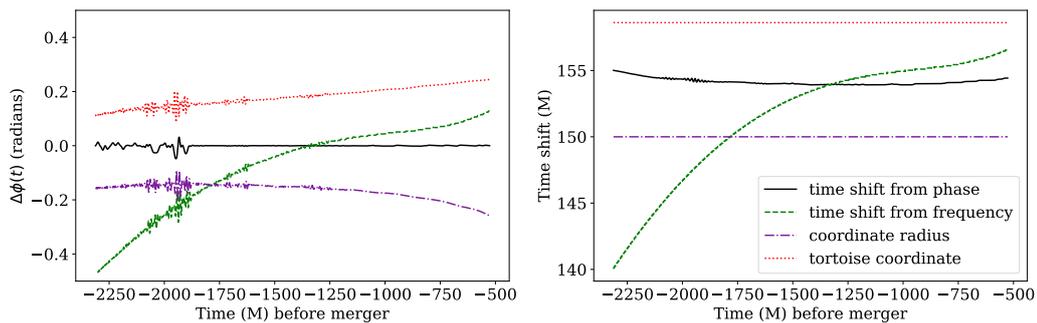


Figure 3.3: The same quantities as in figure 3.2, but for the BAM q8 simulation ($q = 8$, $M\omega_{22}^{\text{start}} = 0.0625$, $\chi_1 = (0, 0, 0.756)$, $\chi_2 = (0, 0, 0.0945)$).

We will illustrate the difference between these choices with two waveforms from non-precessing binaries. One is an equal-mass-binary waveform selected from the SXS catalog of SpEC waveforms [2], and the other is a mass-ratio 1:8 binary simulated with the BAM code. (These are the SXS 152 and BAM q8 configurations listed in table 3.1 in section 3.5). We denote the orbital phase of the two black holes by $\phi_{orb}(t)$, and the corresponding phase of the gravitational-wave signal by Φ . As described in section 3.3.2, the orbital phase is given by the change in position of the separation vector in the orbital plane and the gravitational-wave phase is calculated in the co-precessing frame. For each choice of retarded time t_{GW} , we calculate the phase difference $\Delta\phi(t) = \Phi(t_{\text{GW}}) - \phi_{orb}(t)$. Figures 3.2 and 3.3 show the results for several choices of retarded time. For the SXS waveform, we consider three choices of retarded time: as defined by the co-ordinate extraction radius, by the areal radius of the extraction sphere, $R_{\text{areal}} = \sqrt{A/4\pi}$, where A is the proper area of the extraction sphere [2, 45], and by the tortoise co-ordinate calculated from the areal radius. (The tortoise co-ordinate choice was used to produce the Format 1 metadata for SXS waveforms). For the co-ordinate and areal-radius choices, we see that the phase difference can be as large as 0.5 rad $1000M$ before merger. The phase difference when using the tortoise co-ordinate is much smaller, but still non-zero. By construction the phase difference is zero for the time shift obtained by aligning the orbital and gravitational wave phases. The signal propagation times implied by each choice are shown in the right panel. We see that the propagation time varies with the waveform-based choices, but that is not surprising, given the gauge-dependent nature of the co-ordinate dynamics. The areal radius has not been calculated for the BAM waveform, so figure 3.3 shows results only for the co-ordinate extraction radius, and the tortoise co-ordinate calculated using this value. We again see that the phase difference is smallest when using the tortoise co-ordinate. The variation in the time shift required to align the phases is comparable between the SXS and BAM waveforms.

Based on the results in figures 3.2 and 3.3, we see that the tortoise co-ordinate provides the best phase alignment between the dynamics and gravitational wave signal for both codes. We also find that the results based on our procedure give similar agreement. This procedure has the additional advantages that it can be applied agnostically to all NR waveforms, and is based directly on the gauge-invariant gravitational wave signal.

3.4.2 Orbital plane nutation

In precessing configurations the orbital plane exhibits nutation that is not present in the direction of the full PN orbital angular momentum. [150] showed that the QA direction calculated from ψ_4 also precesses smoothly, suggesting that this method may be a better approximation to the direction of the orbital angular momentum,

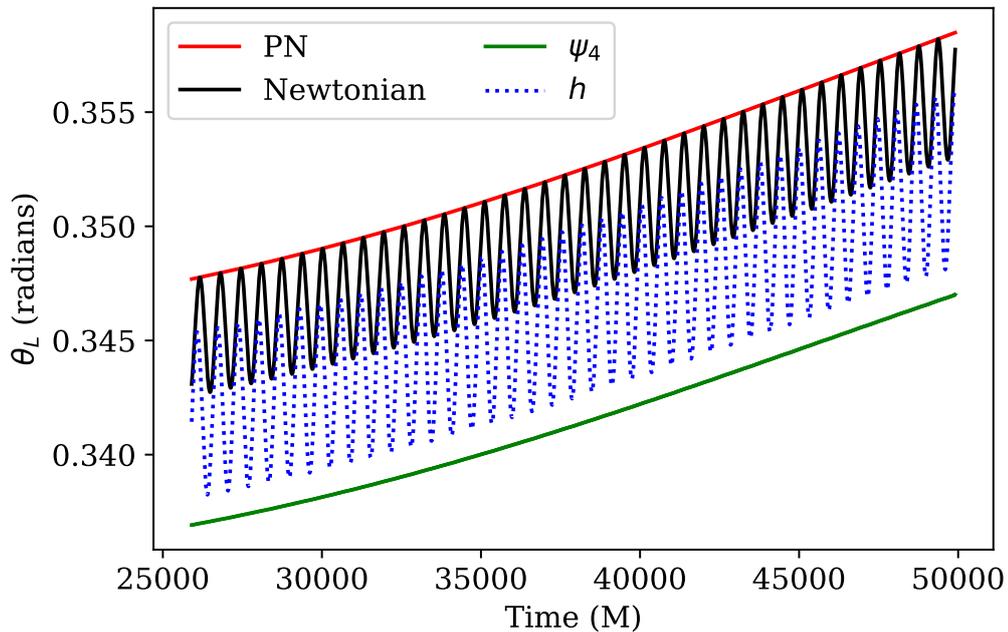


Figure 3.4: θ_L calculated from the Newtonian orbital angular momentum (i.e. the normal to the orbital plane), the post-Newtonian orbital angular momentum, and from ψ_4 and the gravitational wave strain. (See text for discussion).

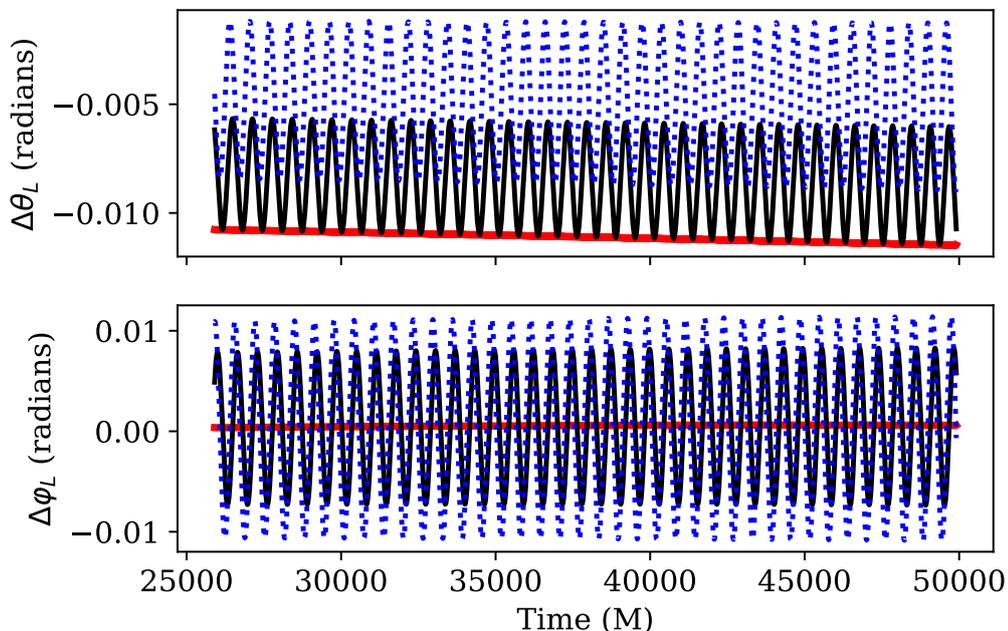


Figure 3.5: Difference between the calculation of θ_L and φ_L from ψ_4 , and that calculated from $\hat{\mathbf{L}}_{\mathbf{N}}$ (black line), $\hat{\mathbf{L}}_{\mathbf{PN}}$ (red line) and $\hat{\mathbf{L}}_{\mathbf{h}}$ (dashed blue) for a post-Newtonian waveform with $q = 3$, $\chi = 0.75$ on the larger black hole, on average in the orbital plane.

than to the orbital plane. We consider a PN example, and illustrate that although this identification does not hold, the QA direction is nonetheless a good approximation to the binary orientation. For an NR configuration we also quantify the differences between the ψ_4 estimates of the binary orientation and phase, and those calculated from the orbital dynamics, and show that they are small.

We first consider the example of a PN waveform; the waveform was constructed using the expressions for the $\ell = 2$ modes given in [25]. These expressions require as inputs the positions, momenta and spins of the black holes, which were found by by evolving the full PN equations of motion including non-spinning terms up to 3PN order and up to next-to-leading order contributions to the spin terms. More details of the method to construct this waveform are summarised in [149]. The PN waveform has the advantage that there is no time shift required between the waveform and the dynamics, removing the retarded-time ambiguity. Our example is a mass-ratio 1:3 system, where the larger black hole has a spin of $\chi = S/m^2 = 0.75$, and the spin lies on average in the orbital plane. We consider a 25 000 M -long segment of a PN waveform for this system; the orbital angular frequency range is 0.00491 – 0.00525. As well as calculating the direction of maximum emission for this waveform, using both h and ψ_4 , we calculate the Newtonian orbital angular momentum (as in equation 3.5) and the post-Newtonian orbital angular momentum (given by $\hat{\mathbf{L}}_{\text{PN}} = \mathbf{n} \times \mathbf{p}$, where \mathbf{p} is the momentum of the system [149]).

In a simple-precession configuration, the orbital angular momentum precesses around the total angular momentum, \mathbf{J} , and the precession can be described by the opening angle θ_L between the orbital and total angular momenta, and the cumulative precession angle φ_L . In figure 3.4 we show four calculations of θ_L : the Newtonian orbital angular momentum direction, $\hat{\mathbf{L}}_{\text{N}}$ (solid black line), which exhibits nutation; the post-Newtonian angular momentum direction, $\hat{\mathbf{L}}_{\text{PN}}$ (solid red line), which precesses smoothly; and the QA estimates calculated from ψ_4 (solid green line) and h (dashed blue line). From this figure we make several observations. (1) θ_L calculated from ψ_4 precesses smoothly, but does *not* agree with the direction of $\hat{\mathbf{L}}_{\text{PN}}$. (2) θ_L calculated from h exhibits nutation, but does not agree with the direction of $\hat{\mathbf{L}}_{\text{N}}$. We note that if we calculate the PN amplitude using only leading-order contributions (our full PN waveform used the amplitudes from [25]), then the QA θ_L calculated from h agrees perfectly with that of $\hat{\mathbf{L}}_{\text{N}}$ (which we expect by construction), but θ_L calculated from ψ_4 still precesses smoothly. This suggests that the apparent agreement between the QA and $\hat{\mathbf{L}}_{\text{PN}}$ directions in [150] was due only to the use of ψ_4 in the QA procedure, with differences masked by gauge ambiguities, and in general these directions should not be expected to be identical, as is also seen in [44]. Note also that [44] shows that the nutation in the h -based calculation is reduced if one includes PN signal amplitude terms that account for the mode asymmetries that lead to out-of-plane recoil, but some nutation does remain.

In figure 3.5 we show the difference between the maximum gravitational wave

emission direction $\hat{\mathbf{L}}_{\psi_4}$ as calculated from ψ_4 , and the Newtonian orbital angular momentum direction, $\hat{\mathbf{L}}_{\mathbf{N}}$, the post-Newtonian angular momentum direction, $\hat{\mathbf{L}}_{\mathbf{PN}}$ and the maximum gravitational wave emission direction $\hat{\mathbf{L}}_{\mathbf{h}}$. We see that although there are differences between different estimates $\Delta\theta_{L_N}$ and $\Delta\varphi_{L_N}$ are oscillatory while $\Delta\theta_{L_{PN}}$ and $\Delta\varphi_{L_{PN}}$ are smoothly varying. This is because $\hat{\mathbf{L}}_{\mathbf{N}}$ shows nutation while $\hat{\mathbf{L}}_{\psi_4}$ and $\hat{\mathbf{L}}_{\mathbf{PN}}$ do not [150]. Additionally, $\hat{\mathbf{L}}_{\mathbf{h}}$ shows nutation. Note that similar behaviour is seen for NR simulations in [113], which considers strain, ψ_4 , and also the Bondi news, $N = \dot{h}$. We used ψ_4 to calculate both $\hat{\mathbf{L}}$ and $\hat{\mathbf{n}}$ in all subsequent examples. Although $\hat{\mathbf{L}}_{\psi_4}$ and $\hat{\mathbf{L}}_{\mathbf{PN}}$ agree well, they are not equal. This may be due to differing PN orders in the description of the dynamics and of the waveform; whether the quantities converge with higher order PN treatments remains to be studied. Note that the nutation in the dynamics can be removed by using an orbit-averaged PN treatment, in which case it is the gravitational wave-based precession that exhibits nutation [119], but this is not consistent with the fully general-relativistic results of NR simulations, as the later examples will illustrate.

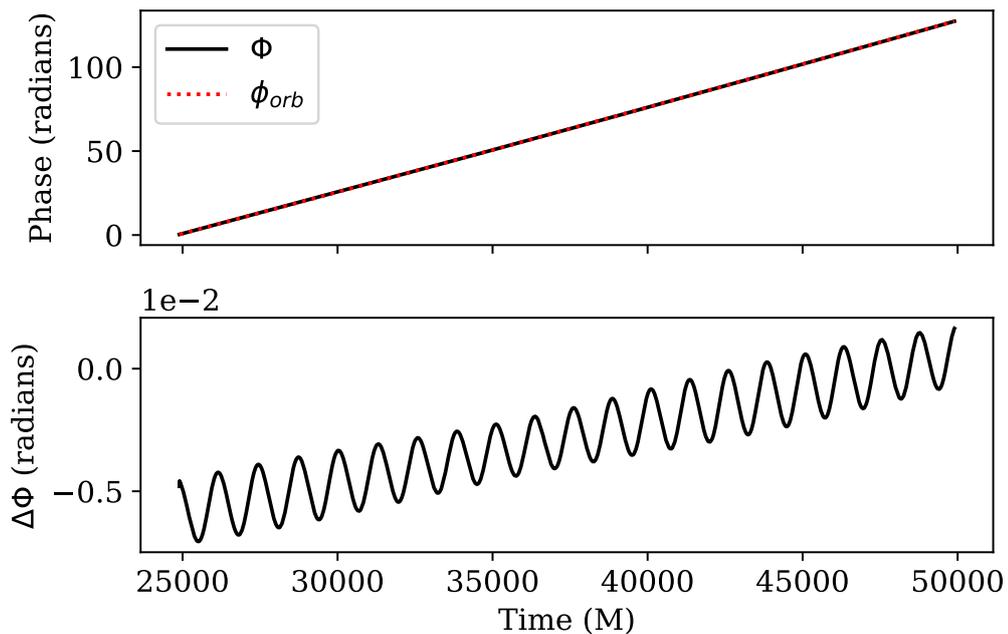


Figure 3.6: PN waveform ($q = 3$, $\chi = 0.75$ on the larger black hole, on average in the orbital plane). A comparison of half the co-precessing gravitational wave phase Φ and the orbital phase ϕ_{orb} .

We also compared the orbital phase of the waveform with the co-precessing gravitational wave phase. The orbital phase was found by integrating the orbital frequency from the PN equations and setting the integration constant using the method described in section 3.3.3. The result of this comparison is shown in figure 3.6. As can be seen, they agree very well over the whole 25000M of inspiral.

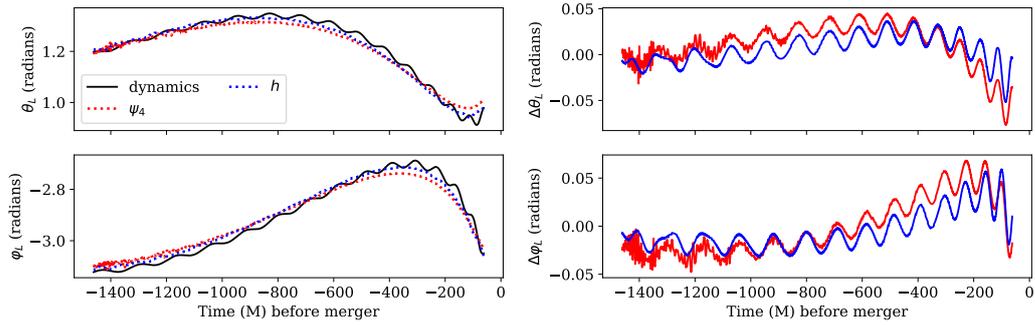


Figure 3.7: BAM q1 ($q = 1$, $M\omega_{22}^{\text{start}} = 0.0354$, $\chi_1 = (0, -0.8, 0)$, $\chi_2 = 0$). On the left is shown the evolution of θ_L and φ_L as calculated for the Newtonian orbital angular momentum direction from the dynamics and for the direction of maximum emission from the waveform. The right hand side shows the difference between the quantities calculated from the dynamics and those calculated using ψ_4 (red) and those calculated using strain (blue). The time shift used for this comparison is that obtained by aligning the orbital and signal phases.

Figure 3.7 shows a similar comparison for an NR simulation, the BAM q1 configuration in table 3.1. Here the co-processing orbital phase is found as described in section 3.3.3. The quantities calculated using the dynamics are time-shifted assuming a flat space time (the time shift described by equation 3.16). We again see that the Newtonian dynamics exhibit nutation that is not present in the maximum emission direction calculated from gravitational wave signal. In this case, the magnitude of the nutation present in the direction calculated from the strain is much smaller than seen for the PN example in figure 3.4. The reason for this is not clear but one possible explanation is that h is calculated from ψ_4 for NR simulations using fixed frequency integration and this process smoothes out the oscillations.

3.5 Numerical Comparisons

In this section we compare our gravitational wave and dynamics based calculations of the co-processing phases for NR waveforms produced using a representative set of current codes. This is complicated by the ambiguities that we discussed in the previous section, but a direct comparison provides us with a general sense of how well these different estimates agree, and whether our method gives physically reasonable results. The NR waveforms that we used are summarised in table 3.1; these are either private BAM simulations, or simulations available through the SXS, Georgia Tech, RIT and LVC-NR catalogues [88, 95, 118].

As for the non-precessing case, we first compared the co-ordinate phases found from the waveform and the orbital motion. We chose to align the phases using the static time shift provided with the waveform metadata, i.e. the time shift suggested by the group that produced the NR simulations. For the BAM, Georgia Tech and

Simulation	q	$M\omega_{22}^{\text{start}}$	χ_1	χ_2	Precession cycles before merger
Non-precessing waveforms					
SXS 152	1	0.0297	(0, 0, 0.6)	(0, 0, 0.6)	-
BAM q8	8	0.0625	(0, 0, 0.756)	(0, 0, 0.0945)	-
Precessing waveforms					
SXS 58	5	0.0316	(0.5, 0, 0)	(0, 0, 0)	1.00
BAM q1	1	0.0354	(0.000896, -0.8, -0.000412)	(0, 0, 0)	0.609
GT0718	2.5	0.0628	(-0.187, 0.221, 0.526)	(0.0996, 0.541, 0.239)	0.673
RIT 0168	2	0.0389	(-0.438, 0.716, -0.104)	(0.173, -0.373, -0.301)	0.778

Table 3.1: List of simulations used in numerical comparisons. $M\omega_{22}^{\text{start}}$ gives the gravitational wave frequency at the start of the waveform. The number of precession cycles is estimated from the number of orbits \mathbf{LN} completes around $\mathbf{J} = \mathbf{LN} + \mathbf{S}$.

RIT waveforms, this is the value of the co-ordinate extraction radius, R_{ex} . For the SpEC waveforms it is the tortoise co-ordinate calculated from the areal radius; see section 3.4.1. The phases agree well, as can be seen from figure 3.8. The discrepancy between the two values arises predominantly from the static time shift used to compare them.

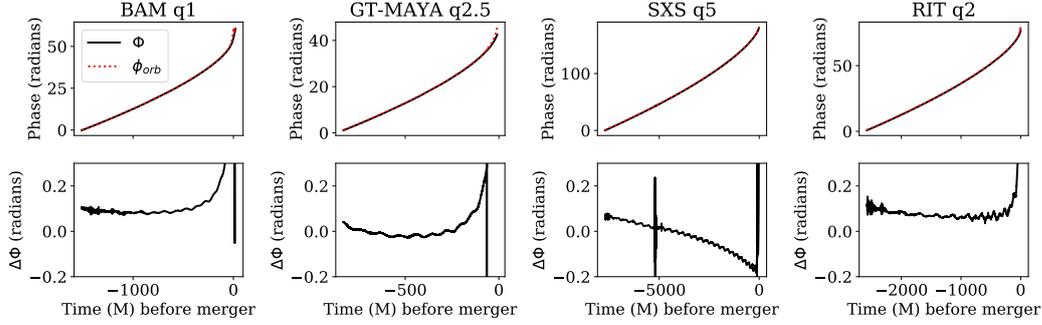


Figure 3.8: A comparison of the orbital phase Φ estimated from the gravitational wave signal, and the orbital phase ϕ_{orb} calculated from the dynamics, for the four precessing waveforms listed in table 3.1. The comparison was made using the time shifts provided by the group that produced each NR simulation.

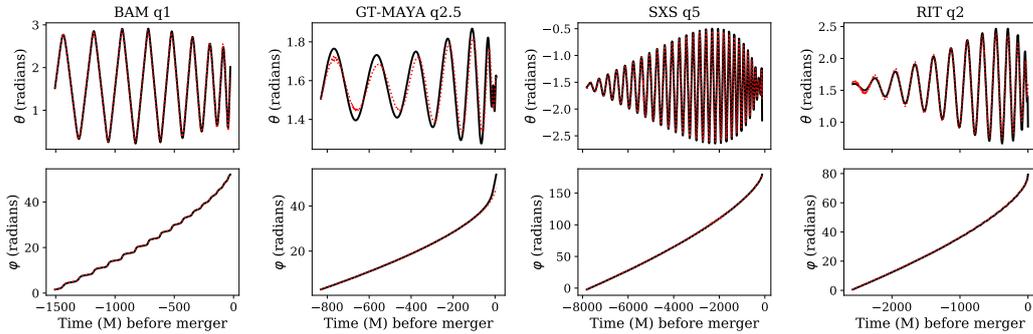


Figure 3.9: A comparison of $\hat{\mathbf{n}}$ calculated from the dynamics (solid black) and using the waveform phase (dashed red) for the four precessing waveforms listed in table 3.1. This comparison was made using the time shifts provided by the group that produced each NR simulation.

We also compared the gravitational wave and dynamics based estimates of $\hat{\mathbf{n}}$ (i.e. $\hat{\mathbf{n}}_w$ and $\hat{\mathbf{n}}_d$) in order to observe what impact differences in phase estimates had on the quantities that are directly used by the NR Injection Infrastructure. We calculated θ , the angle between $\hat{\mathbf{n}}$ and the z -axis, and ϕ , the cumulative angle between the projection of $\hat{\mathbf{n}}$ in the xy plane and the x -axis. We then found the difference in the quantities calculated from the dynamics (d) and those calculated from the waveform (w), given by $\Delta\theta = \theta_w - \theta_d$ and $\Delta\phi = \phi_w - \phi_d$. These are shown in figures 3.9 and 3.10. For the BAM q1 waveform, \mathbf{J} is almost along the x -direction, leading to large oscillations in the orientation of the orbital plane with respect to the z -axis, and

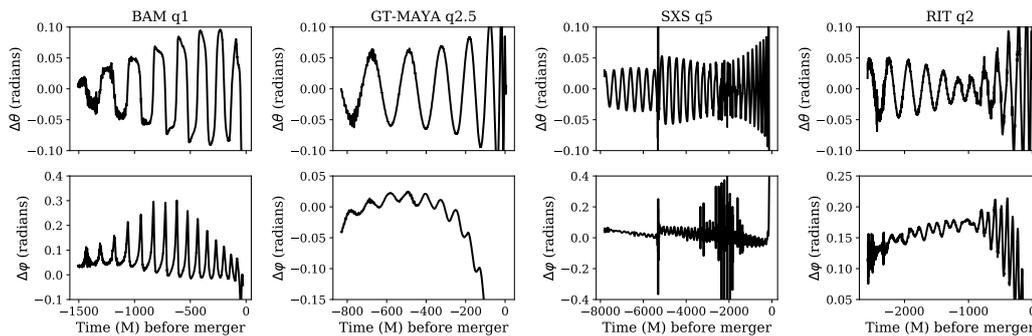


Figure 3.10: A comparison of the difference between θ and φ (shown in figure 3.9) calculated from the dynamics and from the waveform for the four precessing waveforms listed in table 3.1. This comparison was made using the time shifts provided by the group that produced each NR simulation.

consequently also in θ . In the other simulations \mathbf{J} is approximately aligned in the z -direction, leading to smaller oscillations. We can see from the general agreement between the gravitational wave and dynamics based quantities, that our method to find $\hat{\mathbf{n}}$ is reliable regardless of the simulation’s initial configuration. The SXS, RIT and GATech simulations all have $\hat{\mathbf{L}}$ approximately aligned in the z -direction at the beginning of the simulation. For the SXS and RIT waveform this accounts for the growth in the amplitude of the θ oscillations with time. The oscillations may not change much in amplitude for the GATech waveform because it is shorter than the SXS and RIT waveforms and so experiences less of a precession cycle.

We see that our method of estimating the orbital phase Φ and unit separation $\hat{\mathbf{n}}$ using the waveform reproduces the value calculated from the dynamics to reasonable accuracy. The levels of disagreement are consistent with the retarded-time and coordinate ambiguities, and the approximations inherent in the QA procedure.

3.6 Conclusion

We have extended previous work, which calculates a variant of the orbital angular momentum $\hat{\mathbf{L}}$ based entirely on the gravitational wave signal [46, 123, 150], to also calculate an effective orbital phase, Φ . These can be used to prescribe the binary orientation and orbital phase when using NR waveforms as proxy gravitational wave signals. The most immediate application is through the NR Injection Infrastructure used by the LIGO-Virgo collaboration [151], and we follow the same notation and conventions. Our method makes it possible to orient the source without reference to the gauge-dependent binary dynamics, or a retarded time, which lacks a unique definition. The results of this method are in principle gauge invariant (up to finite-extraction-radius errors in the NR waveforms), and can be used agnostically on all current BBH NR waveforms.

As part of the validation of our method, we have compared the results to those found from the co-ordinate dynamics. The differences between the two approaches are consistent with ambiguities in the definition of the retarded time, and the smoother precession of the gravitational wave-based calculation of precession as compared to that from the orbital dynamics.

We note that the current NR Injection Infrastructure does not specify a choice of several conventions in the NR wave extraction (see section 3.3.4). In calculating the orbital phase it is necessary to take into account the choice of conventions used in extracting the NR waveforms.

The remaining dynamical quantities that are not considered in our method are the individual black-hole spin vectors, and the separation between the two black holes. The separation is not used as an observable in gravitational wave astronomy applications. Potential extensions of our method to include the time-evolution of the spin vectors is left to future work.

Given that our method provides a unique, gauge-invariant measure of $(\hat{\mathbf{L}}(t), \hat{\mathbf{n}}(t))$ to prescribe binary configurations, we recommend it as the standard measure of these quantities in the NR Injection Infrastructure. However, we note that at the level of precision required for current analyses, results between different methods for calculating these dynamics quantities agree well. It also provides a useful check of the metadata attached to a precessing waveform calculated using more traditional methods, such as for the waveforms presented in Chapter 2 and used for the model developed in the following chapter.

Chapter 4

Tuned precessing model

4.1 Introduction

The LIGO and Virgo collaborations are detecting an ever increasing number of gravitational wave events [6, 7, 9, 10, 13–16]. None of these detections have so far shown conclusive evidence of precession. This is in part due to the difficulty of measuring precession [75, 76], especially for BBH systems where the masses of the two black holes are close to equal. Recently some events have been measured with a clear non-zero spin [16]. The properties of the binary are determined by the process and environment in which it forms. Systems formed through dynamical processes tend to have isotropically oriented spins [117, 129, 154] while those formed from isolated massive binaries (such as through common envelope evolution [33, 130]) generally have aligned spins [99, 160]. If the binaries have an isotropic spin distribution then a vanishingly small fraction will have aligned spins, meaning the majority of systems will have some degree of precession. Since we do not know which formation channel is the most prevalent it is necessary to model generic spin binaries.

As detectors get more sensitive we expect to detect an increasing number of gravitational wave events. As a consequence of this, we expect to detect an increasing number of more unusual events, such as those with particularly loud signal-to-noise ratio (SNR), higher mass ratios or larger spin magnitudes. In these cases, the detection of precession is more likely. It is therefore imperative to have accurate inspiral-merger-ringdown (IMR) models of precessing signals across a wide parameter space.

To date there have been several different families of precessing waveforms. These can be divided into two main approaches: modelling the underlying non-precessing waveform and the precession effects separately [85, 102, 124] and producing a surrogate model of the complete precessing waveform [162]. The first of these approaches has been used by two different waveform families: Phenom and EOB. More detail is given on these different families in section 1.5 of the Introduction. All of these models have been used in the analysis of events detected by LIGO and Virgo, al-

though all have their limitations. The Phenom and EOB models do not have the precessing component of the model tuned to NR and rely on expressions for the precession angles that are known to be valid only during the inspiral part of the waveform while the surrogate models can only be used for events with a high total mass.

In this chapter we present a phenomenological model for the precession angles used to “twist up” a non-precessing waveform to produce a precessing waveform as described in section 1.4 of the Introduction. This model is tuned to NR waveforms of single spin systems with mass ratios from 1 to 8. A spin of magnitude 0.4 or 0.8 is placed on the larger black hole inclined at an angle from 30° to 150° to the orbital angular momentum of the binary. The model is a full IMR model and employs PN expressions for the angles below the frequencies for which we had NR data available.

4.2 Model Outline

4.2.1 PhenomP: Underlying Principles

The principles and approximations that underpin current models of precessing BBH signals are described in section 1.4 of the Introduction. A description of current Phenom models is given in section 1.5. Both of these are summarised here for ease of reference and to highlight key features relevant to the work that is presented in the following chapter.

The PhenomP set of precessing waveform models uses the fact that a precessing waveform can be modelled by a co-precessing waveform, which does not exhibit amplitude and phase modulations due to precession effects, “twisted up” [85]. This allows the precessing waveform to be broken down into two separate components: a model for the underlying co-precessing waveform and a model for the Euler angles required to perform the time- or frequency-dependent rotation in order to “twist up” the waveform into the inertial frame, using equation 1.40 [149]. For a BBH system, the majority of the power radiated through gravitational waves is emitted in a direction approximately perpendicular to the orbital plane [119, 150]. By tracking the direction of maximum emission we can therefore track the precession of the orbital plane. In the co-precessing frame, the system does not precess so the rotation from the inertial frame to the co-precessing frame can be defined using the direction of maximum emission.

The rotation between an inertial frame in which a waveform is precessing and the non-inertial co-precessing frame is described by three Euler angles $\{\alpha, \beta, \gamma\}$. These angles describe the precessional motion of the direction of maximum emission \mathbf{V} about a given axis. For a consistent definition of the angles, we define the *precession angles* where we take this axis to be along the direction of the total angular momentum in the inertial frame. These angles can be found from the waveform as described

in section 1.4 [123, 150]. Two of these precession angles, α and β , are shown in figure 1.2 in the Introduction. β is the angle between \mathbf{V} and \mathbf{J} while α is the angle between the projection of \mathbf{V} into a plane perpendicular to \mathbf{J} and some constant axis. The third precession angle γ is fixed by the ‘minimal-rotation condition’ [46] in order to ensure the orbital velocity of the waveform is invariant under the transformation. The minimal-rotation condition means γ is given by equation 1.76 once the other two angles have been found.

Current precessing models make a number of approximations when modelling precessing systems. The first set of approximations are made in modelling the underlying co-precessing waveform. During inspiral, the in-plane spin components have only a small effect on the phase of the waveform in the co-precessing frame. This leads to the first of these assumptions; the co-precessing waveform can be approximated by the equivalent aligned spin waveform [126, 149, 150]. This provides a simple mapping to a family of non-precessing waveforms. The final spin of the remnant black hole will not be the same as for the equivalent inspiral aligned-spin waveform. The ringdown part of the co-precessing waveform is therefore modelled by an aligned-spin waveform where the final spin has been adjusted to an approximation of the final spin of the precessing system. These two parts are connected by an intermediate region to form the model for the co-precessing waveform. PhenomP versions 1-3 [85, 102] use PhenomD [94, 103] as the underlying aligned spin model while PhenomPv3HM [104] uses PhenomHM [112] which includes higher modes. Similarly, PhenomXPHM [131] uses PhenomXHM [78], a higher mode model which extends PhenomXAS [132]. This approximation is also made by the SEOB family of models— for example, SEOBNRv4PHM [124] employs SEOBNRv4HM [63] as the underlying aligned spin model.

The second set of approximations are made in modelling the Euler angles used to twist up the co-precessing waveform. To date, precessing models have used expressions for these angles which describe the precessional dynamics of the system—the motion of the orbital angular momentum \mathbf{L} about the total angular momentum \mathbf{J} . \mathbf{L} can either be the Newtonian value or a PN description of varying order. The direction of \mathbf{L} is approximately, although not exactly, the same as the direction of maximum emission [42, 82, 150]. Further to this, current versions of PhenomP use a model for these angles derived from PN theory and apply this model through inspiral, merger and ringdown. This is despite the fact that we do not expect them to be accurate beyond late inspiral. More details are given on the angles used by PhenomPv3 in the Introduction in section 1.5.3. We will see how these angles compare with those from the NR simulations in later sections of this chapter.

The final approximation made by the PhenomP family is that the direction of \mathbf{J} is fixed and taken to be in the z -direction. The model is therefore only valid for cases of simple precession where the direction of \mathbf{J} is approximately constant and neglects transitional precession.

In this chapter we present an improved model for the precession angles and leave improvements to the underlying co-precessing model to future work. Here we present a model for the angles which describes the precession of the optimum emission direction itself rather than the precession of \mathbf{L} . Since these angles can be calculated from the waveform itself (using the method described in section 1.4.1 of the introduction) they are the appropriate angles to put the waveform into the co-precessing frame. Details on how we get PN expressions for these angles are given in section 4.4.1. The model for the angles presented here is a full IMR model, where the inspiral angles still use the PN expressions but the merger-ringdown angles are now given by a phenomenological ansatz where the co-efficients have been tuned to NR. We still make the assumption that the direction of \mathbf{J} remains fixed.

4.2.2 Tuned angle model: A Summary

Here we present a precessing waveform model where the precession angles are tuned to NR in the merger-ringdown region, while the inspiral angles are calculated from PN theory [58], as in PhenomPv3 [102]. We tuned these angles to the catalogue of 40 NR simulations which was presented in Chapter 2. This catalogue covers the parameter space of single spin configurations with $q \in [1, 2, 4, 8]$, $\chi \in [0.4, 0.8]$ and $\theta(^{\circ}) \in [30, 60, 90, 120, 150]$ where q is the mass ratio, χ is the dimensionless spin magnitude and θ is the angle between the orbital angular momentum and the spin angular momentum. We present the results of using our IMR description of the precession angles to “twist up” the non-precessing model PhenomD, as was done for previous versions of PhenomP [85, 102]. However, these angles can be used to “twist up” any aligned-spin model, such as PhenomHM [112] or PhenomXHM [78]. Additionally, some of the techniques presented here are applicable to *any* “twisted” model, such as the improvements in the PN expression for β described in section 4.4.1. The goal of this work is to produce an accurate description of the precession angles and we leave improvements to the underlying aligned spin co-precessing waveforms to future work.

The model of the precession angles presented here is a frequency-domain model. The precession angles are the same for the strain h and for ψ_4 in the frequency-domain, but not in the time-domain. As stated in equation 1.34 in the Introduction, $\psi_4 = \ddot{h}$. If the modes of the gravitational wave strain are given by

$$h_{\ell m}(t) = A_{\ell m}(t) e^{-im\phi(t)}, \quad (4.1)$$

then the ψ_4 modes are

$$\psi_{4,\ell m}(t) = A'_{\ell m}(t) e^{-im\phi'(t)}, \quad (4.2)$$

where

$$A'_{\ell m} = \sqrt{\left(\ddot{A} - m^2 \dot{\phi}^2 A\right)^2 + m^2 \left(2\dot{\phi}\dot{A} + \ddot{\phi}A\right)^2}, \quad (4.3)$$

$$\phi' = \phi + \frac{1}{m} \arctan \left(\frac{m \left(2\dot{\phi}\dot{A} + \ddot{\phi}A\right)}{\ddot{A} - m^2 \dot{\phi}^2 A} \right). \quad (4.4)$$

The distribution of power between the modes is therefore different for h and for ψ_4 in the time domain. This means the precession angles required to transform the waveform into the co-precessing frame in which the majority of the power resides in the (2,2) mode will be different. By contrast, in the frequency domain we have

$$\tilde{\psi}_4 = \text{F.T.} [\psi_4] = \text{F.T.} [\ddot{h}] = -\omega^2 \tilde{h}, \quad (4.5)$$

where $\omega = 2\pi f$ and f is the gravitational wave frequency. Since ω is an overall factor in front of all the modes at a given frequency the direction which maximises $|\tilde{h}|^2$ will be the same as that which maximises $|\omega^2 \tilde{h}|^2$. The precession angles will therefore be the same for h and for ψ_4 . The frequency domain is therefore the natural regime in which to work.

The merger-ringdown angles presented in this model are tuned to NR. During the modelling procedure we first transformed the NR waveforms into a frame in which \mathbf{J} is aligned along the z -direction at all times. The final model therefore makes the approximation that the direction of \mathbf{J} is fixed, as is done in other PhenomP models. The full signal includes asymmetries between the $+m$ and $-m$ modes, which are not present in non-precessing systems [44, 48]. By using a non-precessing model for the co-precessing waveform we are, by construction, ignoring mode asymmetries in our model. Therefore, for consistency, we symmetrise the NR waveforms and associated precession angles. The waveforms were symmetrised using the method described in [44].

The calibration co-efficients of the phenomenological expressions used to describe the merger-ringdown angles can be parameterised by the symmetric mass ratio η , the dimensionless spin magnitude χ and the cosine of the angle between the orbital angular momentum and the spin angular momentum $\cos \theta$. The model was tuned to simulations with a single spin placed on the larger black hole where χ and $\cos \theta$ are easily defined. As with previous models, a mapping from a two-spin system to a single spin system is therefore required.

We define the *effective spin parameter* χ_{eff} [19, 142] which parameterises the rate of inspiral and the *precession spin parameter* χ_p [152] which parameterises the precessional motion of the binary. We expect χ_{eff} to be approximately constant

throughout the inspiral [138]. χ_{eff} and χ_p for a generic two-spin system are given by

$$\chi_{\text{eff}} = \frac{m_1 \chi_1^{\parallel} + m_2 \chi_2^{\parallel}}{m_1 + m_2}, \quad (4.6)$$

$$\chi_p = \frac{S_p}{A_1 m_1^2}, \quad (4.7)$$

where $S_p = \max(A_1 S_1^{\perp}, A_2 S_2^{\perp})$ and S_i^{\parallel} and S_i^{\perp} are the individual spin components parallel and perpendicular to the orbital angular momentum. χ_{eff} parameterises the spin parallel to the orbital angular momentum while χ_p parameterises the spin perpendicular to the orbital angular momentum in the plane of the binary.

The single spin system which corresponds to a two spin system defined by \mathbf{S}_1 and \mathbf{S}_2 has a spin $(\chi_{\perp}, 0, \chi_{\parallel})$ placed on the larger black hole where

$$\chi_{\parallel} = \frac{M \chi_{\text{eff}}}{m_1}, \quad (4.8)$$

$$\chi_{\perp} = \chi_p. \quad (4.9)$$

χ_{\parallel} is parallel to the orbital angular momentum of the system while χ_{\perp} lies in the orbital plane. The total spin magnitude χ and the angle between the orbital and spin angular momenta are then given by

$$\chi = \sqrt{\chi_{\parallel}^2 + \chi_{\perp}^2}, \quad (4.10)$$

$$\cos \theta = \frac{\chi_{\parallel}}{\chi}. \quad (4.11)$$

These reduce to the correct values for the cases to which we tuned the model and also correctly re-weight two-spin cases and cases where the spin is predominantly on the smaller black hole.

4.3 Calculating NR precession angles

Using the gravitational wave output of 40 NR simulations, we perform low level data processing which prepares each waveform to be used in the calculation of the optimal emission direction and the frequency-domain precession angles $\{\alpha(f), \beta(f), \gamma(f)\}$. We calculate the optimal emission direction using the rotationally invariant eigenvalue method described in section 1.4.1 in the Introduction with only the $\ell = 2$ modes. Since each simulation is performed in a different inertial frame and exhibits different numerical artefacts, the data must undergo pre-processing before the precession angles can be calculated. The result of this pre-processing is a set of precession angles that describe the simple precession of each of the BBH configurations with no imprint of mode asymmetries and no numerical artefacts at early times.

4.3.1 Data processing

For each NR simulation, spin weight -2 spherical harmonic multipole moment data, as defined in equation 1.32 in the Introduction, are stored for the radiative Weyl scalar. Each $\psi_{4,\ell m}$ time series contains multipole moment data for inspiral, merger and ringdown.

Before the calculation of the frequency domain optimal emission direction, the data are processed in two ways. First, spurious inspiral (“junk”) radiation, due to imperfect initial data, is windowed away, using a window function that increases from zero to one over the course of three gravitational wavelengths. It is found that when windowing over more than two wavelengths, the choice of (smooth) window function has no significant effect on results. For simplicity, a standard Hann window is used [121]. The window starts at the first peak in the real part of $\psi_{4,22}$ such that the following peak is less than or equal to the largest duration between peaks in the time series. This most often results in less than $200 M$ of contaminated inspiral data being tapered away. The window is applied equally to the real and complex parts of ψ_4 for all multipoles. Similarly, post-ringdown data are windowed such that the Hann window turns off to the right between the point where numerical noise overtakes exponential decay, and the end of the data file. The inspiral and post-ringdown windowing reduces unphysical broadband frequency domain power.

Second, the time domain data are zero-padded to the right such that the frequency domain step size, in geometric units, is less than 5×10^{-4} . This ensures the frequency-domain features will be sufficiently well resolved.

4.3.2 Conventions for feature alignment of NR angles

We impose processing choices and conventions towards the alignment of features between different NR cases, and the ultimate connection with angles from PhenomPv3. These choices and conventions are summarized below.

Simple Precession

For consistency with PhenomPv3 and its related PN model, we place all NR waveforms in a frame where the system angular momentum direction is fixed throughout coalescence. Therefore, we begin our pre-modelling feature-alignment in a frame where the radiative angular momentum direction, $\mathbf{J}(t)$, is at all times along the z -direction. This convention amounts to a minor modification of the NR data, whose radiative $\mathbf{J}(t)$ varies by at most $\sim 6^\circ$ from its initial direction. The use of this $\mathbf{J}(t)$ -frame is accomplished by first computing the angular momentum emission rate from the multipole moments according to [141], where the source’s angular momentum

rate, $d\mathbf{J}/dt$, is given by

$$\begin{aligned} \frac{dJ_x}{dt} &= \lim_{r \rightarrow \infty} \frac{r^2}{32\pi} \operatorname{Im} \left\{ \sum_{\ell, m} \int_{-\infty}^t \int_{-\infty}^{t'} \psi_{\ell m} dt'' dt' \right. \\ &\quad \left. \times \int_{-\infty}^t \left(f_{\ell, m} \psi_{\ell, m+1}^* + f_{\ell, -m} \psi_{\ell, m-1}^* \right) dt' \right\}, \end{aligned} \quad (4.12)$$

$$\begin{aligned} \frac{dJ_y}{dt} &= \lim_{r \rightarrow \infty} \frac{r^2}{32\pi} \operatorname{Re} \left\{ \sum_{\ell, m} \int_{-\infty}^t \int_{-\infty}^{t'} \psi_{\ell m} dt'' dt' \right. \\ &\quad \left. \times \int_{-\infty}^t \left(f_{\ell, m} \psi_{\ell, m+1}^* - f_{\ell, -m} \psi_{\ell, m-1}^* \right) dt' \right\}, \end{aligned} \quad (4.13)$$

$$\begin{aligned} \frac{dJ_z}{dt} &= \lim_{r \rightarrow \infty} \frac{r^2}{16\pi} \operatorname{Im} \left\{ \sum_{\ell, m} m \int_{-\infty}^t \int_{-\infty}^{t'} \psi_{\ell m} dt'' dt' \right. \\ &\quad \left. \times \int_{-\infty}^t \psi_{\ell, m}^* dt' \right\}, \end{aligned} \quad (4.14)$$

where $f_{\ell, m} = \sqrt{\ell(\ell+1) - m(m+1)}$ and $*$ denotes complex conjugation. We use the convention that $\operatorname{Im}(a + ib) = b$, for a and b real. To minimize the impact of low frequency numerical noise, integrals in equations 4.12–4.14 are performed using the fixed frequency method [140].

Numerical integration of equations 4.12–4.14 enables the calculation of $\mathbf{J}(t)$. We chose the integration constant so that $\mathbf{J}(0) = \mathbf{J}_{\text{ADM}}$, as defined by the simulation's initial data. \mathbf{J}_{ADM} is the ADM value of the total angular momentum [41]. While this choice of integration constant results in $\mathbf{J}(t)$ that is sensitive to the removal of early time spurious radiation, we find that the resulting final spin of the system agrees closely with that estimated from the remnant's apparent horizon. Disagreement is at most 5%. Thus, we conclude that the removal of system junk radiation is a small source of error in $\mathbf{J}(t)$. This is consistent with the expectation that the junk radiation carries minimal angular momentum, since individual boosted or spinning black holes are axisymmetric and therefore do not radiate angular momentum.

We use the $\mathbf{J}(t)$ to define Euler angles consistent with [46]'s minimal rotation condition, hence ensuring that our $\mathbf{J}(t)$ -frame transformation is invertible and independent of the simulation's initial frame.

Symmetrisation

For consistency with PhenomPv3, we enforce that the co-precessing multipole moments of our model obey the same symmetry properties of their non-precessing counterparts. This means that we neglect to model $\pm m$ mode asymmetries. Although the asymmetric contributions are weak, there is some evidence that they are necessary for non-biassed measurements of precessing systems [98], and we plan to

model these contributions in future work.

We use the following symmetrisation procedure. The $\mathbf{J}(t)$ -frame waveform is rotated into the co-precessing frame using the time-domain precession angles calculated using the matrix method described in section 1.4.1 of the Introduction. In this frame the $(2, \pm 2)$ modes dominate and for consistency with previous non-precessing/co-precessing models we consider only these modes. We define the time-domain symmetrised $(2, 2)$ mode to be

$$\psi_{4,22}^{\text{sym}} = \frac{1}{2} \left(\psi_{4,22} + \psi_{4,2-2}^* \right), \quad (4.15)$$

which is consistent with the definition given in [44]. The $(2, -2)$ mode is then given by

$$\psi_{4,2-2}^{\text{sym}} = \left(\psi_{4,22}^{\text{sym}} \right)^*. \quad (4.16)$$

At this stage there remain asymmetries between the positive and negative frequency domain behaviour of the optimal emission direction angles used to rotate the co-precessing waveform back into the $\mathbf{J}(t)$ -frame. We Fourier transform the symmetrised waveform in the $\mathbf{J}(t)$ -frame and calculate the precession angles using the matrix method. We enforce symmetry in negative and positive frequency angles by averaging positive and negative frequency angles via

$$\bar{\alpha}(f) = \frac{1}{2} (\alpha_{-}(|f|) + \alpha_{+}(|f|)) \quad (4.17)$$

$$\bar{\beta}(f) = \frac{1}{2} (\beta_{-}(|f|) + \beta_{+}(|f|)) \quad (4.18)$$

$$\bar{\gamma}(f) = \frac{1}{2} (\gamma_{-}(|f|) + \gamma_{+}(|f|)). \quad (4.19)$$

In equations 4.17–4.19, subscripts $-$ and $+$ denote angles at positive and negative frequencies, respectively.

In equations 4.15–4.19 we effectively average over subdominant effects, leaving their inclusion in modelling to future investigation.

4.4 Inspiral angles

The inspiral angles used in the model are based on the PhenomPv3 angles described in section 1.5 of the Introduction. The inspiral expression used for α is exactly the same as in PhenomPv3. We have modified the expression for β as described below. Figure 4.1 shows the PhenomPv3 angles for the case $(q, \chi, \theta) = (8, 0.8, 60)$. In this figure, as elsewhere in this chapter unless otherwise stated, the value for the angles is given in radians. As can be seen, α and γ agree well with the NR data at low frequencies but do not capture all the features present in the data as the system approaches merger. While the expression for β is closer to the NR data during

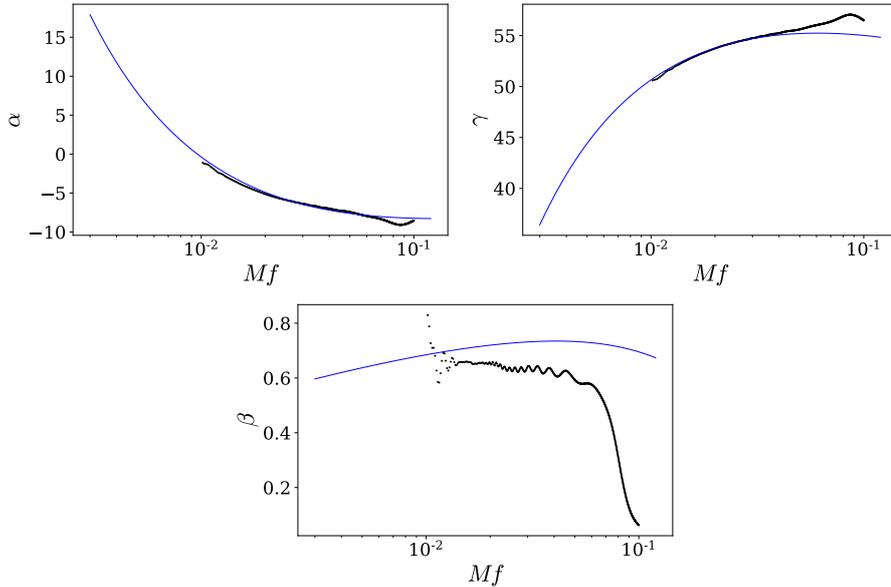


Figure 4.1: Comparison of the post-Newtonian expressions for each of the precession angles (blue line) with the NR data (black dots) for the case with $(q, \chi, \theta) = (8, 0.8, 60)$.

the inspiral section of the NR waveform and diverges towards merger, there always appears to be an offset between the PN prediction and the actual value. The cause and our solution to this discrepancy are described in the following sections.

4.4.1 Higher-order PN corrections to β

In the quadrupole approximation the maximum gravitational wave signal power is emitted perpendicular to the orbital plane, and therefore the angles that describe the precession dynamics of the orbital plane are the same as those associated with the co-precessing frame of the gravitational wave signal [46, 123, 150]. For the full signal, this identification is only approximate [44, 82, 119, 150], and we expect the approximation to be less accurate at higher frequencies. Our modelling approach is based on applying a frequency-dependent rotation to a model of the waveform in the co-precessing frame, and as such the rotation angles should be those associated with the *signal*. However, all current models [85, 102, 125, 158] use the angles associated with the *dynamics*.

Figure 4.1 shows that the PhenomPv3 PN *dynamics* angles α and γ from the inspiral match well with the NR *signal* angles for the merger and ringdown. This suggests that it is possible to identify the dynamics and signal values for these two angles for the purposes of our model. However, this is not the case for the inclination angle of the orbital plane. In this section we will refer to the inclination angle of

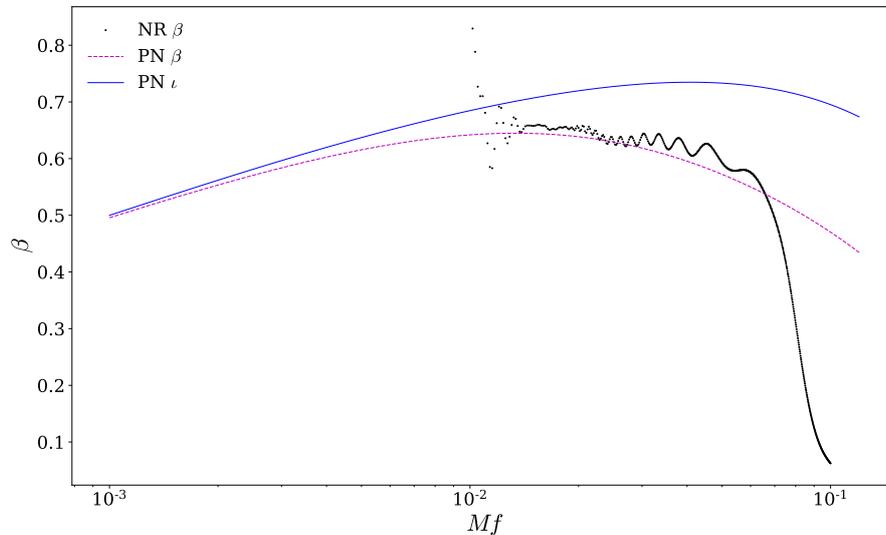


Figure 4.2: Inclination angles for the $(q = 8, \chi = 0.8, \theta = 60^\circ)$ configuration. Shown is the NR inclination angle of the QA frame, β (black dots); the PN inclination angle of the orbital plane, ι (solid blue) and the approximate QA angle β as a function of ι (dashed magenta). See text for details.

the orbital plane with respect to \mathbf{J} as ι , and continue to denote the inclination angle of the co-precessing frame by β . Figure 4.2 shows ι and β for the $(q = 8, \chi = 0.8, \theta = 60^\circ)$ configuration, where ι is calculated from PhenomPv3, and β is from the corresponding NR simulation. Although we expect the PN and NR results to agree at sufficiently low frequencies (where the signal is weaker and the quadrupole approximation is more applicable), they clearly do not agree well over the frequency range of our current NR results.

Fortunately, we have access to PN signal amplitudes beyond the quadrupole approximation, and can use these to calculate a more accurate estimate of the signal β . One way to do this would be to calculate a full PN waveform, e.g., from the model in [58], and apply the quadrupole-alignment procedure to calculate β . However, this will be much more computationally expensive than the current PhenomPv3 approximant, and it is possible to obtain a sufficiently accurate result with a simpler approach.

To illustrate our approach, consider the rotation from a co-precessing signal that contains only the $(\ell = 2, |m| = 2)$ modes, $h_{2,\pm 2}^{\text{NP}}$, to produce a precessing-binary signal in the inertial frame. We will focus on only the $(2, 2)$ and $(2, 1)$ modes, and only the

angles ι, α (since the additional phase rotation γ will not affect our argument),

$$h_{2,2} = e^{-2i\alpha} \left(\cos^4 \left(\frac{\iota}{2} \right) h_{2,2}^{\text{NP}} + \sin^4 \left(\frac{\iota}{2} \right) h_{2,-2}^{\text{NP}} \right), \quad (4.20)$$

$$h_{2,1} = -2e^{-i\alpha} \left(\cos^3 \left(\frac{\iota}{2} \right) \sin \left(\frac{\iota}{2} \right) h_{2,2}^{\text{NP}} - \cos \left(\frac{\iota}{2} \right) \sin^3 \left(\frac{\iota}{2} \right) h_{2,-2}^{\text{NP}} \right). \quad (4.21)$$

The non-precessing modes can be written as,

$$h_{2,\pm 2}^{\text{NP}} = A e^{\mp 2i\Phi}, \quad (4.22)$$

where A and Φ are the time/frequency-dependent amplitude and orbital phase. When ι is small, $h_{2,2}^{\text{NP}}$ makes the strongest contribution to the precessing-waveform modes, and we see that ι determines the relative amplitude of $h_{2,2}$ and $h_{2,1}$. We can isolate the $e^{-2i\Phi}$ term as follows,

$$\bar{h}_{2,2} = \frac{1}{2\pi} \int_0^{2\pi} h_{2,2} e^{2i\Phi} d\Phi \quad (4.23)$$

$$= A e^{-2i\alpha} \cos^4 \left(\frac{\iota}{2} \right), \quad (4.24)$$

$$\bar{h}_{2,1} = -2A e^{-i\alpha} \cos^3 \left(\frac{\iota}{2} \right) \sin \left(\frac{\iota}{2} \right). \quad (4.25)$$

From these we can readily calculate that the inclination ι is

$$\iota = 2 \tan^{-1} \left(\frac{|\bar{h}_{2,2}|}{2|\bar{h}_{2,1}|} \right). \quad (4.26)$$

At leading (quadrupole) order, ι is the precession angle β .

If we now use higher-order PN amplitude expressions [25], then the dynamics inclination angle ι will no longer be the same as the signal inclination angle β , but the expression above will give us an estimate of the orbit-averaged β . Note that the angles in [58] as used by PhenomPv3 are also orbit-averaged (i.e., nutation effects are absent), so this is a consistent treatment.

The mode expressions in [25] are given in terms of the orbital phase Φ , the precession angles α and ι , and the spin components. For the spin components, we make an approximate reduction to our single-spin systems as follows. The inclination of the spin from the z -axis is the spin's inclination from the orbital angular momentum vector, θ , minus the inclination of the orbital angular momentum from the z -axis, ι . The projection of the spin onto the x - y plane rotates at a frequency $\dot{\alpha}$, and so the azimuthal angle is given as $(\alpha + \alpha_0)$; the choice of the initial precession phase

α_0 is unimportant, because the orbital averaging will remove the α dependence in equations 4.24 and 4.25, so we are free to choose $\alpha_0 = 0$. The final result, for a given configuration, depends only on the dynamics inclination ι as a function of frequency; we use the PhenomPv3 expression for $\iota(f)$.

In [25] the amplitudes are expanded in powers of $v = (\pi f)^{1/3}$. We define $\delta = (m_2 - m_1)/(m_1 + m_2)$ (where $m_2 < m_1$), $\eta = m_1 m_2 / (m_1 + m_2)^2$, $\chi_s = (\chi_1 + \chi_2)/2$, $\chi_a = (\chi_1 - \chi_2)/2$, and so,

$$\begin{aligned}\chi_{s/a,x} &= \pm \chi \sin(\theta - \iota) \cos(\alpha), \\ \chi_{s/a,y} &= \pm \chi \sin(\theta - \iota) \sin(\alpha), \\ \chi_{s/a,z} &= \pm \cos(\theta - \iota).\end{aligned}\tag{4.27}$$

If we substitute these into the PN mode expressions for $h_{2,2}$ and $h_{2,1}$, and then apply equation 4.26, we obtain the relatively simple expression,

$$\beta = 2 \tan^{-1} \left(\frac{\sec(\iota/2) (a_1 v^2 + a_2 v^3)}{b_1 + b_2 v^2 + b_3 v^3} \right),\tag{4.28}$$

where

$$\begin{aligned}a_1 &= (55\eta - 65) \sin \iota, \\ a_2 &= 7\chi(6\delta - 5\eta - 6)(2 \cos \iota - 1) \sin \theta \\ &\quad + (66\chi(\delta + \eta - 1) \cos \theta + 84\pi) \sin \iota, \\ b_1 &= 42 \cos \left(\frac{\iota}{2} \right), \\ b_2 &= (55\eta - 107) \cos \left(\frac{\iota}{2} \right), \\ b_3 &= 14\chi(-6\delta + 5\eta + 6) \sin \theta \sin \left(\frac{\iota}{2} \right) \\ &\quad + 56 \cos \left(\frac{\iota}{2} \right) ((\delta + \eta - 1)\chi \cos \theta + 3\pi/2).\end{aligned}\tag{4.29}$$

Figure 4.2 also shows the modified $\beta(\iota)$ for the ($q = 8$, $\chi = 0.8$, $\theta = 60^\circ$) configuration. We see the PN inspiral $\beta(\iota)$ now matches well with the NR merger-ringdown results. We find similar results across the parameter space that we have considered, and therefore to calculate β in our model, we use equation 4.28 in conjunction with ι from PhenomPv3 as calculated in Refs. [58, 102].

4.4.2 Two spin β

The PhenomPv3 expression for ι evaluated for a two spin system show oscillations which become unphysically large through late inspiral and towards merger and ringdown which are not seen in the precession angles calculated for two spin NR systems.

These oscillations also complicate connecting the inspiral expression to the tuned merger-ringdown ansatz. We therefore taper these oscillations so we recover the value and gradient of beta for an equivalent single spin system at the point at which we wish to connect the inspiral and merger-ringdown parts of the model.

For a system described by two spins \mathbf{S}_1 and \mathbf{S}_2 we define a mapping to the appropriate single spin system described by $\mathbf{S}'_1 = \chi_p (\sin \theta, 0, \cos \theta)$ and $\mathbf{S}'_2 = (0, 0, 0)$. We evaluate the PhenomPv3 expression for ι for both of these configurations and so find the oscillations introduced by the two-spin effects by taking the difference between the two:

$$\iota_{osc} = \iota(\mathbf{S}_1, \mathbf{S}_2) - \iota(\mathbf{S}'_1, \mathbf{S}'_2). \quad (4.30)$$

We then apply a taper to these oscillations that ensures ι will tend to the single spin value and gradient at a given frequency f_c and add them back to the single spin value. Our final two spin expression for ι is given by

$$\iota = \iota(\mathbf{S}'_1, \mathbf{S}'_2) + \cos^2\left(\frac{2\pi f}{4f_c}\right) \times \iota_{osc}. \quad (4.31)$$

f_c is the frequency at which the inspiral expression for β is connected to the merger-ringdown expression. It is defined below in equation 4.48.

Furthermore, for a two spin system the magnitude of the in-plane spin component oscillates. In order to perform the correct rescaling to obtain β as described in the previous section we need to obtain an estimate for the frequency-dependent value of this component. To do this we assume that the component of the spins parallel to the orbital angular momentum $S_{||}$ remains fixed and we use the 3PN expression for \mathbf{L} that is used by the PhenomPv3 expressions to calculate ι in the first place. The in-plane spin component S_{\perp} is therefore given by

$$S_{\perp} = (L + S_{||}) \tan \iota, \quad (4.32)$$

where $S_{||} = S_1^{||} + S_2^{||}$. Substituting this expression for S_p in 4.7 we get a value for χ_p . We then calculate χ and $\cos \theta$ as described in equations 4.6 to 4.11 and these values are used to rescale ι .

The effect of this treatment can be seen in figure 4.3, which shows β for SXS1397 (the intrinsic properties of which are given for Case 24 in table 4.1). The PN expression for the angle captures the oscillations seen at low frequency very well. However these oscillations do not continue to high frequency and are greatly over-estimated by the full two-spin PN expression. Tapering the oscillations to the single spin value at the connection frequency resolves this issue well. For $f > f_c$ the PN expression is replaced by the merger-ringdown expression described in the following section so the behaviour of the PN angles here are not an issue.

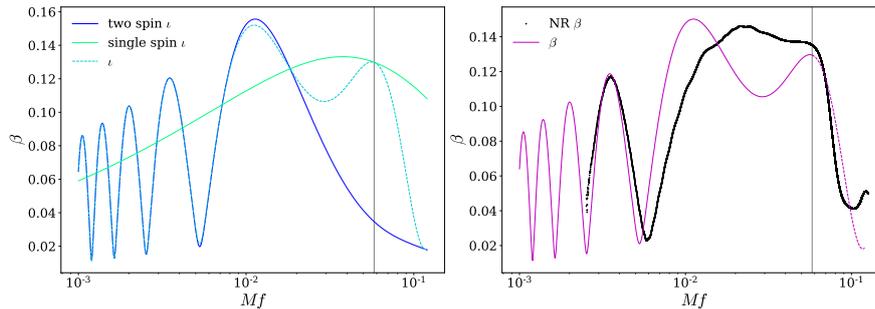


Figure 4.3: Various options for the PN expression for the opening angle. The left-hand panel shows the PN value of ι for a two spin system (blue) and for the equivalent single spin system (green) calculated using the expressions used in PhenomPv3. In light blue is shown the effect of tapering the two-spin oscillations to the single-spin value at the connection frequency f_c , shown as a grey vertical line. In the right-hand panel the value for β used in the model (pink) is compared with the NR value of β found for this case. The configuration shown is Case 24 in table 4.1.

4.5 Merger-ringdown angle model

The PN expressions for the precession angles cannot be reliably extended through merger and ringdown and when compared with the NR angles clearly do not capture the features present at high frequency. We therefore present a phenomenological description of the precession angles α and β in the merger-ringdown regime. We describe the functional form of the angles and produce a global fit for each of the co-efficients of the relevant ansatz. This provides a frequency domain description of the precession angles across the parameter space.

4.5.1 Alpha

The morphology of the merger-ringdown part of α is qualitatively very similar to that of the phase derivative, seen in [103]. α shows a $\frac{1}{f}$ fall off with a Lorentzian dip centred around what is approximately to ringdown frequency of the BBH system. This prompted the ansatz

$$\alpha(f) - \langle \alpha(f) \rangle = \frac{A_1}{f} + \frac{A_2 \sqrt{A_3}}{A_3 + (f - A_4)^2}, \quad (4.33)$$

where A_1 , A_2 , A_3 and A_4 are free co-efficients.

The fitting region is based around the Lorentzian dip; it is defined to be the range $Mf_{dip} - 0.0225 \leq Mf \leq Mf_{dip} + 0.0075$ where M is the total mass of the binary and f_{dip} is the frequency at which α reaches its minimum. The global fit for α within this fitting region has a mean percentage error of 8.68% across the 40 calibration waveforms. Some example comparisons of the result of these fits with the NR value for α are shown in figure 4.4.

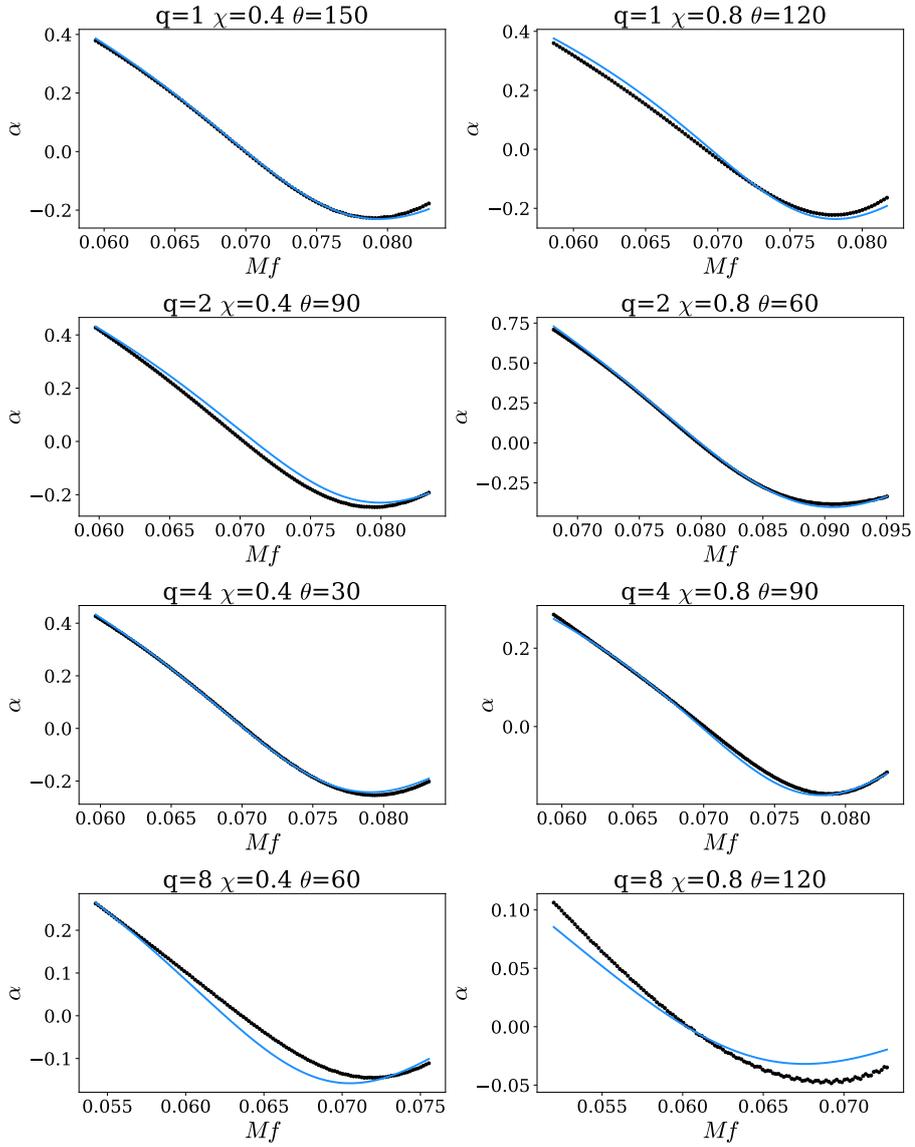


Figure 4.4: Comparison of the phenomenological ansatz presented in equation 4.33 with the NR data over the frequency range to which the co-efficients in the ansatz were tuned for a selection of the cases which comprise the NR catalogue.

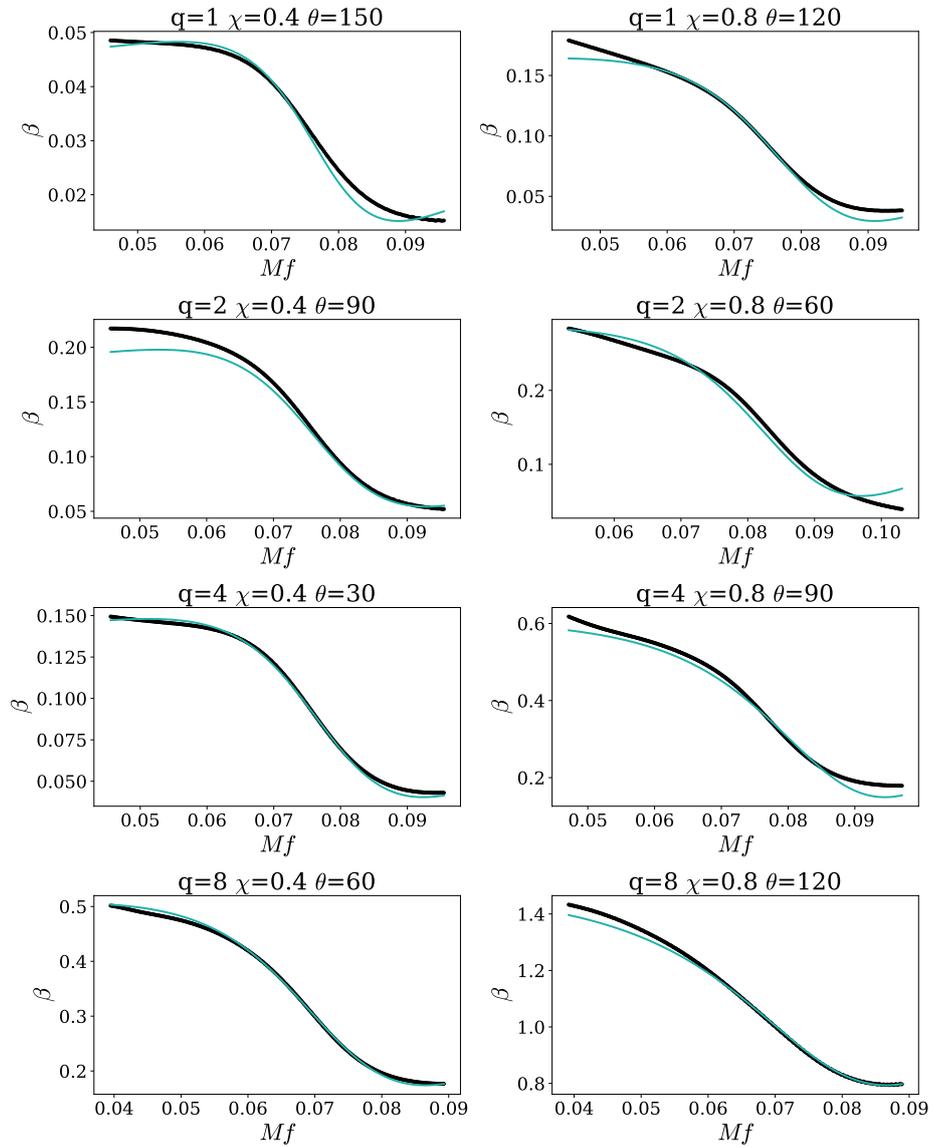


Figure 4.5: Comparison of the phenomenological ansatz presented in equation 4.34 with the NR data over the frequency range to which the co-efficients in the ansatz were tuned for a selection of the cases which comprise the NR catalogue.

4.5.2 Beta

During merger and ringdown, β drops rapidly as the dominant emission direction relaxes to its final direction. The ansatz used to describe β is therefore chosen to grow at low frequencies (as seen in the PN expressions), turnover at the correct frequency, capture the drop and finally tend asymptotically towards the constant value to which the dominant emission direction relaxes. The ansatz we chose to describe this behaviour is

$$\beta(f) - \langle \beta(f) \rangle = \frac{B_1 + B_2 f + B_3 f^2}{1 + B_4 (f + B_5)^2}, \quad (4.34)$$

where B_1, B_2, B_3, B_4 and B_5 are free co-efficients. We also present a global fit for $B_0 = \langle \beta(f) \rangle$.

The fitting region for β is centred around the inflection point in the turnover f_{inf} ; $Mf \in Mf_{inf} \pm 0.03$. Within this fitting region, the global fit for β shows a mean percentage error of 8.08% across the 40 calibration waveforms. Some example comparisons of the result of these fits with the NR value for β are shown in figure 4.5.

4.6 The phenomenological co-efficients

The two ansätze given above which describe the merger-ringdown behaviour of α (equation 4.33) and β (equation 4.34) have 10 free co-efficients between them. Each of these co-efficients was fitted across the three-dimensional parameter space described by the symmetric mass ratio η , the dimensionless spin magnitude χ and the cosine of the angle between the orbital angular momentum and the spin angular momentum $\cos\theta$.

4.6.1 The fitting procedure

The optimum value of each of the co-efficients for each waveform in the calibration set was found by fitting the relevant ansatz to the NR data using the non-linear least-squares fitting function `curve_fit` from the python package `scipy` [163]. This function uses the Levenberg-Marquardt algorithm to perform the least-squares fitting. We then performed a three-dimensional fit of each of the co-efficients using the fitting algorithm `mvpfit` [111]. This gives each of the co-efficients as a polynomial expansion in η , χ and $\cos\theta$. Since we have 40 calibration waveforms, the maximum possible number of terms which can appear in these expressions is 39 in order to avoid over fitting. The fits are restricted so that the highest order term in each dimension is one less than the total number of data points in that dimension. Since the value of each of the co-efficients in the ansatz is to some extent dependent on the value of each of the other co-efficients, we found a global fit for each co-efficient

in turn, re-fitting the ansatz to the data while keeping the co-efficients which had already been fitted fixed.

The general expression for each co-efficient is

$$\Lambda^i = \sum_{p=0}^3 \sum_{q=0}^1 \sum_{r=0}^4 \lambda_{pqr}^i \eta^p \chi^q \cos^r \theta, \quad (4.35)$$

where $\Lambda \in [A, B]$ are the co-efficients in the ansatz describing α and β respectively and $i \in [1, 2, 3, 4]$ and $[0, 1, 2, 3, 4, 5]$ respectively. The λ_{lmn}^i give the co-efficients of the polynomial expansion of the multi-dimensional fits of Λ_i . This expression has a maximum of 39 terms. The values of the polynomial co-efficients λ_{lmn}^i are given in Appendix A.

The variation of the co-efficients for α and β across the parameter space can be seen in figures 4.6 and 4.7 respectively. These figures show that the co-efficients vary smoothly across the parameter space. As can be seen from the residual plots above the fit surfaces, the global fits agree closely with the values of the co-efficients found from fitting the ansatz to each individual simulation.

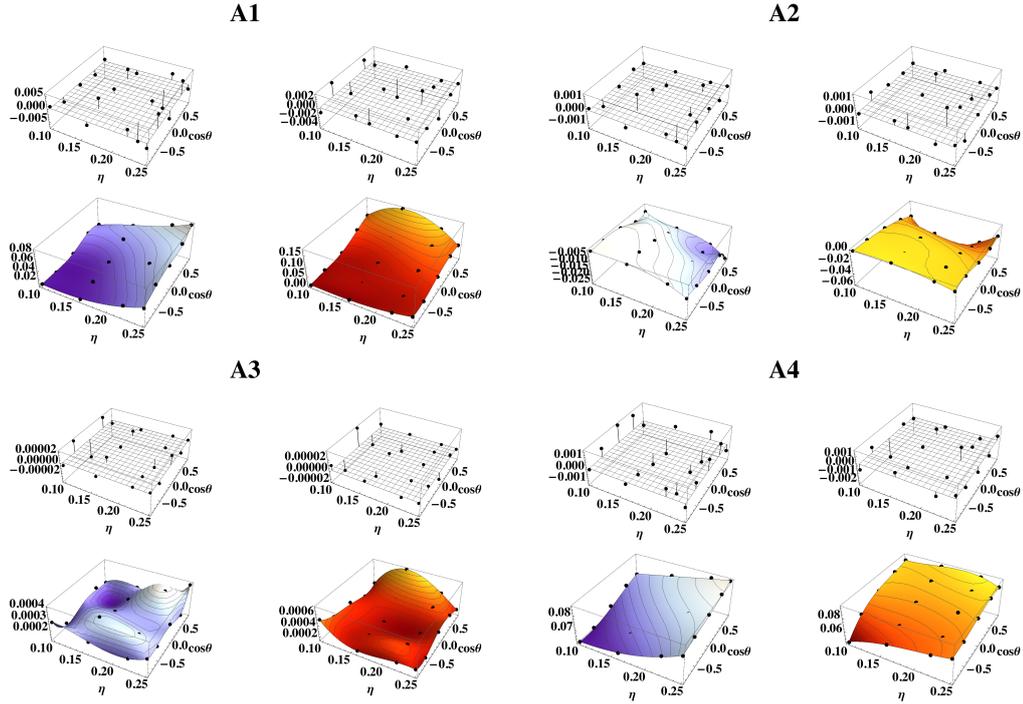


Figure 4.6: Comparison of the fits for each of the co-efficients for the ansatz for α given in equation 4.33 with the co-efficients found from the data as described in section 4.6.1. The fits are shown as two-dimensional surfaces covering the parameter space described by η and $\cos \theta$. On the left in blue are the fits for the simulations with $\chi = 0.4$ and on the right in red are the fits for $\chi = 0.8$. Above each of these surfaces are shown the residuals.

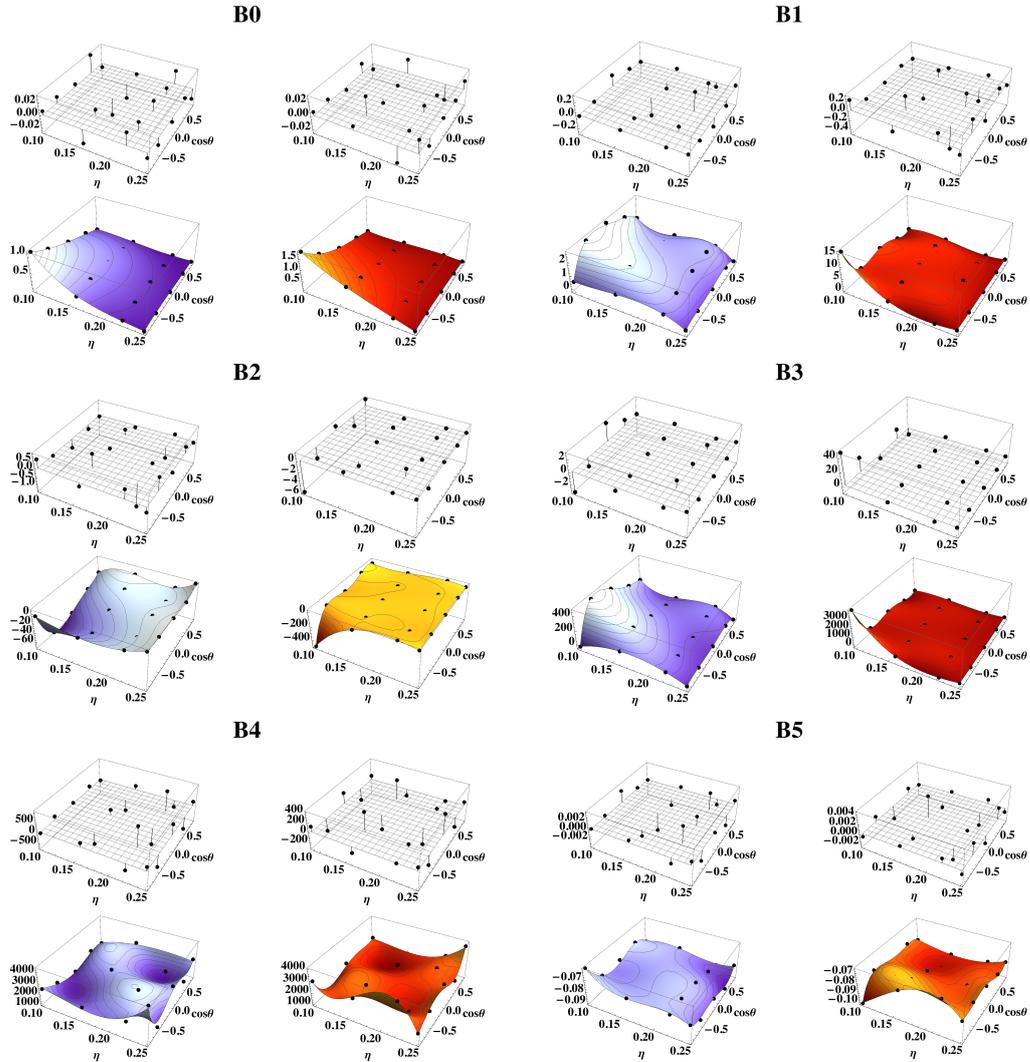


Figure 4.7: Comparison of the fits for each of the co-efficients for the ansatz for β given in equation 4.34 with the co-efficients found from the data as described in section 4.6.1. The fits are shown as two-dimensional surfaces covering the parameter space described by η and $\cos\theta$. On the left in blue are the fits for the simulations with $\chi = 0.4$ and on the right in red are the fits for $\chi = 0.8$. Above each of these surfaces are shown the residuals.

4.6.2 Key features of the merger-ringdown model

The key features of the morphology of α are the location, width and height of the Lorentzian dip. These are described by the parameters A_4 , $2\sqrt{A_3}$ and $\frac{A_2}{\pi\sqrt{A_3}}$ respectively. Figures 4.8 to 4.10 show how these vary over the three-dimensional parameter space of η , χ and $\cos\theta$ over which the modelling was performed.

As can be seen from figure 4.8, the dip occurs at higher frequencies for smaller mass ratios (larger η), and for smaller values of the opening angle θ . This trend is most pronounced for larger spin magnitudes. Since the location of this feature approximately corresponds to the ringdown frequency of the system this is as expected. Close to aligned spin (small θ) we see the hang up effect [55] by which the system merges at a higher frequency since more angular momentum must be radiated away, and the opposite is seen for anti-aligned spin. This effect is most pronounced for larger spin magnitudes.

Figure 4.9 shows that the dip is deepest for small mass ratios and shallowest for large mass ratios. This disparity is most marked for cases with large spin magnitude. This is also seen with the phase derivative of non-precessing waveforms— for example the same trend is seen with the equivalent co-efficient α_4 in the model for the phase derivative used in PhenomD [103].

From figure 4.10 we can see that the width of the dip has a more complicated variation across the parameter space. In general it is widest for systems with aligned spins and that the difference between aligned and anti-aligned systems is most pronounced for systems with high spins. In PhenomD, the width of the dip in the phase derivative was found to be represented by the ringdown damping frequency to a good approximation [103]. This same identification cannot be made for the width of the dip in α seen here using the same final mass and spin fits that are used in PhenomD to calculate the ringdown damping frequency.

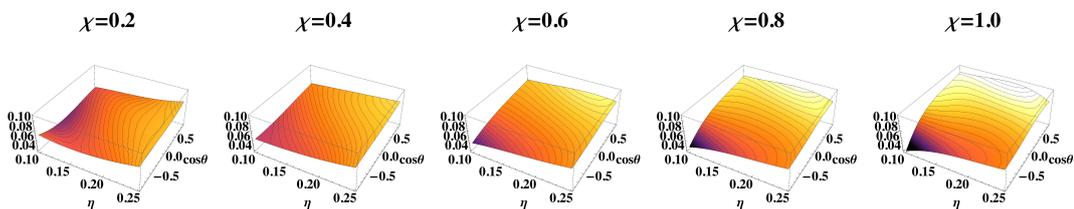


Figure 4.8: Variation in the location of the centre of the Lorentzian dip seen in α across the three-dimensional parameter space over which the modelling was performed.

The key features of β are the mean value B_0 , the location of the inflection point f_{inf} in the turnover, the width of the turnover and the value of the asymptote at high frequencies β_{RD} .

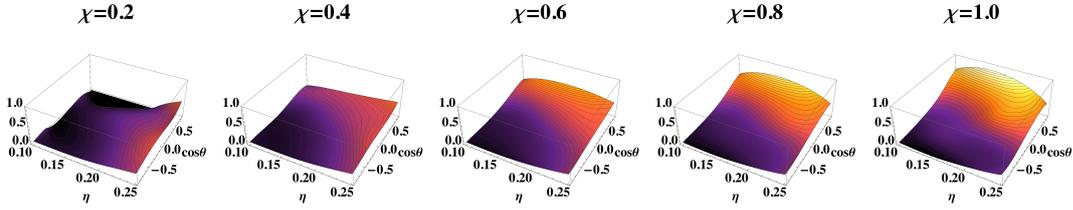


Figure 4.9: Variation in the height of the Lorentzian dip seen in α across the three-dimensional parameter space over which the modelling was performed.

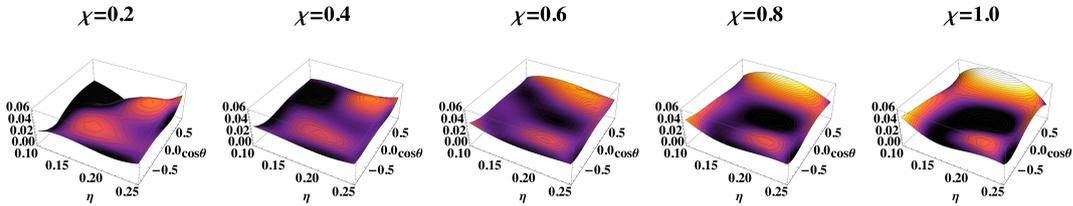


Figure 4.10: Variation in the width of the Lorentzian dip seen in α across the three-dimensional parameter space over which the modelling was performed.

Figure 4.11 shows that B_0 , the mean value of β close to merger, grows as the mass ratio and spin magnitude increase. This is expected as β is approximately equal to the angle between $\mathbf{J} = \mathbf{L} + \mathbf{S}$ and \mathbf{L} . For a fixed mass ratio, increasing χ increases the value of $|\mathbf{S}|$ and similarly for fixed χ , increasing the mass ratio also increases the value of $|\mathbf{S}| = m_1^2 \chi = q^2 \chi / (1 + q)^2$ when normalising to $M = 1$. For constant \mathbf{L} and orientation of \mathbf{S} , increasing the magnitude of \mathbf{S} increases the angle between \mathbf{L} and \mathbf{J} , thus increasing β . At small spin magnitudes and mass ratios, β increases as the angle θ between \mathbf{L} and \mathbf{S} increases towards 90° and then decreases again. In these cases, $|\mathbf{L}| > |\mathbf{S}|$ so this corresponds to the angle between \mathbf{L} and \mathbf{J} increasing until \mathbf{L} and \mathbf{S} are perpendicular and then decreasing until \mathbf{L} and \mathbf{S} are anti-aligned and \mathbf{J} is aligned along \mathbf{L} . For higher spin magnitudes and mass ratios, close to merger $|\mathbf{L}| < |\mathbf{S}|$, so the angle between \mathbf{L} and \mathbf{J} increases as θ increases until \mathbf{L} and \mathbf{S} are anti-aligned and \mathbf{J} is now also anti-aligned with \mathbf{L} . In this case β simply increases as θ increases.

The inflection point of the ansatz modelling the turnover in β gives a measure of the location of the drop. It is also used to determine which of the three morphologies shown in figure 4.13 the ansatz in β is taking and to select the correct part of the ansatz to ensure β always displays a drop at merger. The inflection points of an expression occur at the roots of the second derivative of the expression. The second

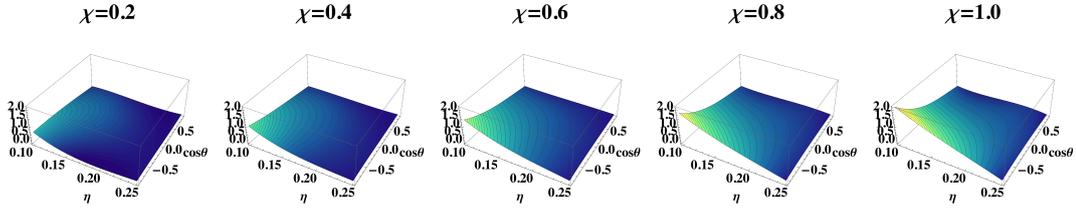


Figure 4.11: Variation in the amplitude of the merger-ringdown part of β across the three-dimensional parameter space over which the modelling was performed.

derivative of 4.34 takes the form

$$\beta''(f) = \frac{af^3 + bf^2 + cf + d}{\left(1 + B_4(B_5 + f)^2\right)^3}, \quad (4.36)$$

where a , b , c and d are functions of the fitting co-efficients B_1 , B_2 , B_3 , B_4 and B_5 . In order to find the roots of this cubic we re-write it in the form of a depressed cubic

$$x'^3 + px' + q = 0, \quad (4.37)$$

where

$$x' = x + \frac{b}{3a}, \quad (4.38)$$

$$p = \frac{3ac - b^2}{3a^2}, \quad (4.39)$$

$$q = \frac{2b^3 - 9abc + 27a^2d}{27a^3}. \quad (4.40)$$

In the case where this expression has three real roots, these are given by

$$x' = 2\sqrt{-\frac{p}{3}} \cos \left[\frac{1}{3} \arccos \left(\frac{3q}{2p} \sqrt{-\frac{3}{p}} \right) - \frac{2n\pi}{3} \right], \quad (4.41)$$

where $n = 0, 1, 2$.

We wish to be able to define a single, smoothly varying inflection point which tracks the location of the turnover in β during merger across the parameter space. As the co-efficients of the cubic vary, the morphology of equation 4.34 changes, as shown in figure 4.13. For $a < 0$ we have the morphology shown in the central panel of the figure. We therefore select the central root which is the only one with a negative gradient. For $a > 0$, we have the morphology shown in the outer panels. For this morphology we need to differentiate between the two outer roots, which both have a negative gradient. This is determined by the ‘shift’ of the roots, $\frac{b}{3a}$. In

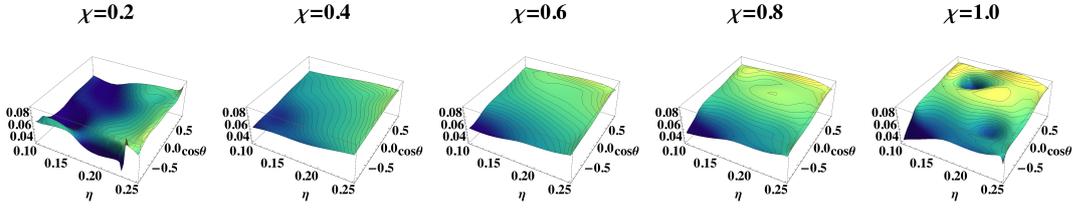


Figure 4.12: Variation in the location of the inflection point in the merger-ringdown part of β across the three-dimensional parameter space over which the modelling was performed.

cases where $\frac{b}{3a} > h$ we choose the first root (as seen in the left-hand panel) while in cases where $\frac{b}{3a} < h$ we choose the final root (as seen in the right-hand panel), where $h = m \cos \theta + c$ is the value of the oblique asymptote in the expression for $\frac{b}{3a}$.

In the case where we have complex roots, two of the roots will be in the complex plane while one will be on the real axis. In this case we select the only real root.

We also consider the case where $a = 0$ and the second derivative is a quadratic. In this case we have only one root with a negative gradient, which is the desired root. Finally, we consider the case where both $a = 0$ and $b = 0$. Here we have only one root which gives us the desired inflection point.

From figure 4.12 we can see that enforcing these conditions gives us a smoothly varying value of the inflection point across the parameter space. This value is highest for systems where the spin is close to being aligned with the orbital angular momentum and has a larger magnitude. This is expected as the inflection point is a measure of the location of the turnover in β which we expect to correspond to the frequency at which the black holes merge. This occurs at higher frequencies for systems with higher spin values and smaller angles between the orbital angular

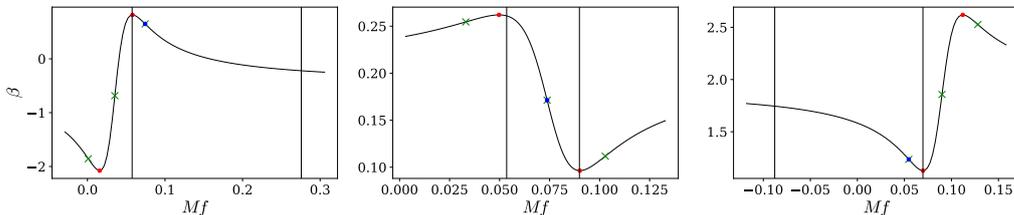


Figure 4.13: Possible morphologies of the ansatz given by equation 4.34 depending on the values taken by the co-efficients in different regions of the parameter space. From left to right the panels show systems with $(q, \chi, \theta) = (8, 0.2, 155)$, $(2.5, 0.4, 90)$ and $(5, 0.8, 160)$. The red dots mark the extrema, the green crosses show the inflection points and the blue dot indicates the inflection point chosen as described in section 4.6.2. The points of maximum curvature around this inflection point are shown by the black lines, which give a measure of the width of the turnover.

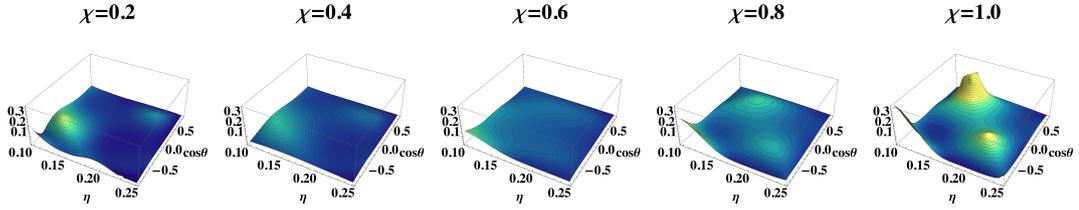


Figure 4.14: Variation in the width of the turnover of β across the three-dimensional parameter space over which the modelling was performed.

momentum and the spin. This is due to the “hang-up” effect of the spin-orbit contribution to the PN phasing [36], which becomes much more pronounced closer to merger, e.g., [55]. This is qualitatively similar to the variation of the location of the Lorentzian dip in α (shown in figure 4.8) as both are related to the merger frequency of the binary.

The width of the turnover is defined to be the distance between the extrema of the curvature given by

$$C(f) = \frac{\beta''(f)}{\left(1 + \beta'^2(f)\right)^{3/2}}, \quad (4.42)$$

adjacent to the inflection point shown in figure 4.12. The variation of the width is shown in figure 4.14.

We define β_{RD} using two different methods depending on the morphology of the ansatz for β . In cases where β tends towards an asymptote immediately following the inflection point selected following the method described above, we choose β_{RD} to be equal to the value of this asymptote. In cases where β reaches a local minimum immediately following the inflection point we choose β_{RD} to be equal to the value β reaches at this point. Following this definition we see from figure 4.15 that β_{RD} varies smoothly across the parameter space, tending close to zero for systems with small mass ratios and small spin angles θ , regardless of spin magnitude. It is only for cases with high spin magnitude and large values of θ that β_{RD} has a significant value. This value increases as the dimensionless spin magnitude increases. This quantity corresponds to the angle between the optimal emission direction of the final black hole $\mathbf{V}_{\mathbf{f}}$ and the total angular momentum $\mathbf{J}_{\mathbf{f}}$. The optimal emission direction need not be related to the direction of the black holes’ spin, which will determine the direction of the final black hole’s total angular momentum. For one possible explanation of this effect, consider perturbing a stationary spinning BH, such that gravitational waves were emitted predominantly at an angle ζ to \mathbf{S} . Given that the BH is spinning, the perturbation is dragged around the BH, meaning the optimal

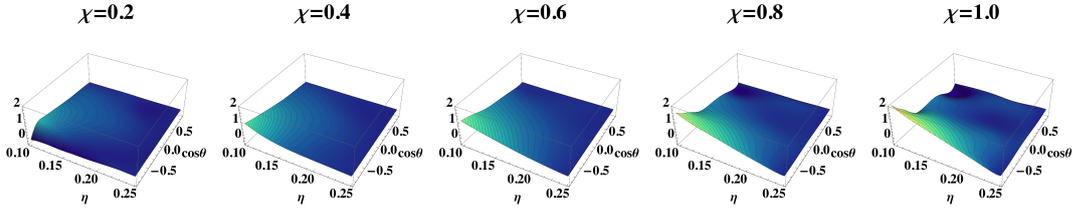


Figure 4.15: Variation in the final value towards which β tends during ringdown across the three-dimensional parameter space over which the modelling was performed.

emission direction would precess around \mathbf{S} , but the opening angle ζ would stay the same. The discussion in sections III.C. and III.D. of [122] explains this effect in terms of the superposition of quasi-normal modes emitted by the final black hole during ringdown. They show that if the (2,0) mode is negligible then β will decay to zero, otherwise a variety of behaviours are possible. From our NR data we find β decays to a constant non-zero value. As noted above, this value is largest for systems with high anti-aligned spin. These are systems where, prior to merger, β is largest since $|\mathbf{S}| > |\mathbf{L}|$ and so \mathbf{L} (and therefore the optimum emission direction) is anti-aligned with \mathbf{J} .

4.7 Connecting inspiral angles to merger-ringdown expressions

The expressions for the precession angles for the two distinct regions, inspiral and merger-ringdown, are connected so that connection is smooth and the full IMR expression for the angles agrees with the NR data over the entirety of the region for which it is available. The method used to connect the two regions was different for each angle.

4.7.1 Connection method

For α the regions are connected using an interpolating function of the form

$$\alpha_{\text{interp}}(f) = a_0 f^2 + a_1 f + a_2 + \frac{a_3}{f}, \quad (4.43)$$

defined over the frequency range $[f_1, f_2]$. The co-efficients of this expression are chosen so that

1. $\alpha_{\text{interp}}(f_1) = \alpha_{\text{PN}}(f_1)$ and $\alpha_{\text{interp}}(f_2) = \alpha_{\text{PN}}(f_2)$ since there is freedom in an overall constant offset in α ,

2. $\alpha'_{\text{interp}}(f_1) = \alpha'_{\text{PN}}(f_1)$ and $\alpha'_{\text{interp}}(f_2) = \alpha'_{\text{MR}}(f_2)$ in order to ensure the two parts are connected continuously.

α_{PN} is the PN expression used for α in the inspiral regime. α_{MR} is the merger-ringdown ansatz given in equation 4.34. The co-efficients are given by

$$\begin{aligned}
 a_0 &= \frac{1}{D} \left[2(f_1\alpha_1 - f_2\alpha_2) - (f_1 - f_2) \left((f_1\alpha'_1 + f_2\alpha'_2) + (\alpha_1 - \alpha_2) \right) \right], \\
 a_1 &= \frac{1}{D} \left[3(f_1 + f_2)(f_1\alpha_2 - f_2\alpha_1) \right. \\
 &\quad \left. + (f_1 - f_2) \left((f_1 + 2f_2)(f_1\alpha'_1 + \alpha_1) + (2f_1 + f_2)(f_2\alpha'_2 + \alpha_2) \right) \right], \\
 a_2 &= \frac{1}{D} \left[6f_1f_2(f_1\alpha_1 - f_2\alpha_2) \right. \\
 &\quad \left. + (f_1 - f_2) \left(f_2(2f_1 + f_2)(f_1\alpha'_1 + \alpha_1) + f_1(f_1 + 2f_2)(f_2\alpha'_2 + \alpha_2) \right) \right], \\
 a_3 &= \frac{1}{D} \left[f_1f_2^2(f_2 - 3f_1)\alpha_1 - f_1^2f_2(f_1 - 3f_2)\alpha_2 \right. \\
 &\quad \left. + f_1f_2(f_1 - f_2) \left(f_2(f_1\alpha'_1 + \alpha_1) + f_1(f_2\alpha'_2 + \alpha_2) \right) \right], \tag{4.44}
 \end{aligned}$$

where α_i and α'_i are the value of α and its derivative at the limits of the frequency range and $D = (f_2 - f_1)^3$.

The frequency region over which the interpolation was performed was chosen to be as small as possible. The lower frequency limit was chosen to be the highest frequency for which the inspiral expressions agreed with the NR data while the upper frequency limit was chosen to be the lower limit for which the fitted merger-ringdown expressions still agreed well with the NR data. Since the PN expressions for the angles agree well with the NR data over most of the waveform for most of the parameter space, there is a wide range of frequency values over which the interpolation could be performed. We choose the frequency range to be defined in terms of the location l of the dip: $f_1 = \frac{2l}{7} = \frac{2A_4}{7}$ and $f_2 = \frac{l}{3} = \frac{A_4}{3}$.

For β the agreement between the PN expression and the NR data is not as good. Even including the higher order amplitude corrections described in section 4.4.1, the starting frequency of the NR simulations is not low enough to cover the region in which the PN expression closely matches the data for all cases. We therefore cannot use the same interpolation method as was employed for connecting the inspiral and merger-ringdown expressions for α . Instead, we employ a rescaling function which leaves the PN expression invariant at low frequencies but ensures it smoothly connects with the merger-ringdown value of β at the connection frequency f_c . This rescaling function is given by

$$k(f) = 1 + b_1f + b_2f^2, \tag{4.45}$$

which tends to one at low frequencies thus leaving the PN expression unchanged. In order to ensure the value of β and its derivative match at the connection frequency, the co-efficients b_1 and b_2 are given by

$$b_1 = -\frac{1}{\beta_1^2 f_c} \left[-2\beta_1 (\beta_2 - \beta_1) + (\beta_1 \beta_2' - \beta_2 \beta_1') f_c \right], \quad (4.46)$$

$$b_2 = -\frac{1}{(\beta_1 f_c)^2} \left[\beta_1 (\beta_2 - \beta_1) - (\beta_1 \beta_2' - \beta_2 \beta_1') f_c \right], \quad (4.47)$$

where β_1 and β_1' are the value of β and its derivative at the connection frequency given by the original PN expressions while β_2 and β_2' are the values from the merger-ringdown expression at the connection frequency.

The definition of the connection frequency depends on the morphology of β for a particular case. As has been seen previously, in some parts of the parameter space β rises steadily until just before merger then turns over and drops rapidly. However, in other parts of the parameter space this turnover is much more gradual and begins at much lower frequencies. Our ansatz for β captures both of these morphologies well. In cases where the turnover occurs within the fitting region, we define the connection frequency f_c as the frequency at which the merger-ringdown part has a particular gradient $d\beta_c$. This is found by expanding the gradient of the curve about the maximum as a Taylor series. We find the connection frequency is given by

$$f_c = f_{max} + \frac{1}{\beta'''} \left[-\beta'' + \sqrt{\beta''^2 + 4\beta'''' d\beta_c} \right], \quad (4.48)$$

where f_{max} is the frequency at which the maximum occurs, $d\beta_c = \beta'(f_c - f_{max})$ is the gradient of the curve at the connection frequency, and β'' and β''' are the second and third derivatives of β evaluated at the maximum. The gradient at which we wish to perform the connection varies across the parameter space. We therefore define the connection frequency by

$$d\beta_c = 2.5 \times 10^{-4} \times d\beta_{inf}^2, \quad (4.49)$$

where $d\beta_{inf}$ is the gradient at the inflection point. In cases where the turnover is not present within the fitting region we instead define the connection frequency to be the lower frequency limit of the fitting region, thus ensuring β is still falling at this frequency. In this case, $f_c = f_{inf} - 0.03$, where f_{inf} is the inflection point.

4.7.2 Full IMR expressions

The expressions describing the precession angles in each of the different regions are connected using piece-wise C^1 -continuous functions.

The full IMR expression for α is

$$\alpha_{\text{IMR}}(f) = \begin{cases} \alpha_{\text{PN}} & 0 \leq f < f_1 \\ \alpha_{\text{interp}} & f_1 \leq f < f_2 \\ \alpha_{\text{MR}} & f_2 \leq f \end{cases}$$

where α_{PN} , α_{interp} and α_{MR} are the PN expression used to describe α during inspiral, the interpolating function used to describe the late inspiral angles in the region f_1 to f_2 and the phenomenological ansatz which has been tuned to NR to describe the merger-ringdown angles respectively.

In cases where β has a minimum immediately following the inflection point chosen as described in section 4.6.2, the full IMR expression for β is

$$\beta_{\text{IMR}}(f) = \begin{cases} k\beta_{\text{PN}} & 0 \leq f < f_c \\ \beta_{\text{MR}} & f_c \leq f < f_f \\ \beta_{\text{RD}} & f_f \leq f \end{cases},$$

while in cases where β tends towards an asymptote immediately following the inflection point, the full IMR expression for β is

$$\beta_{\text{IMR}}(f) = \begin{cases} k\beta_{\text{PN}} & 0 \leq f < f_c \\ \beta_{\text{MR}} & f_c \leq f \end{cases}.$$

β_{PN} is the PN expression for β including the higher-order amplitude corrections discussed in 4.4.1, k is the rescaling function applied to these expressions as outlined above, β_{MR} is the phenomenological ansatz which has been tuned to NR in the merger-ringdown regime and β_{RD} is the constant value of β to which the system settles down as discussed in section 4.6.2. In this model it is given by the value of β at the minimum in the merger-ringdown expression. f_f is correspondingly given by the frequency at which the minimum occurs.

γ is then calculated over the entirety of the frequency range for which the waveform is produced by enforcing the minimal-rotation condition given in equation 1.76. The decision to do this rather than produce a separate model for γ was made as it was found that γ must be very accurate in order to consistently transform between an inertial frame and the co-precensing frame. The very small discrepancy between the expression for γ presented in [58] and the numerically calculated value required by the minimal-rotation condition is sufficient to seriously degrade the model. This discrepancy is exacerbated here since we are no longer using the dynamical expression for β presented in [58].

The full model of these angles is shown for two examples in very different parts of the parameter space in figures 4.16 and 4.17.

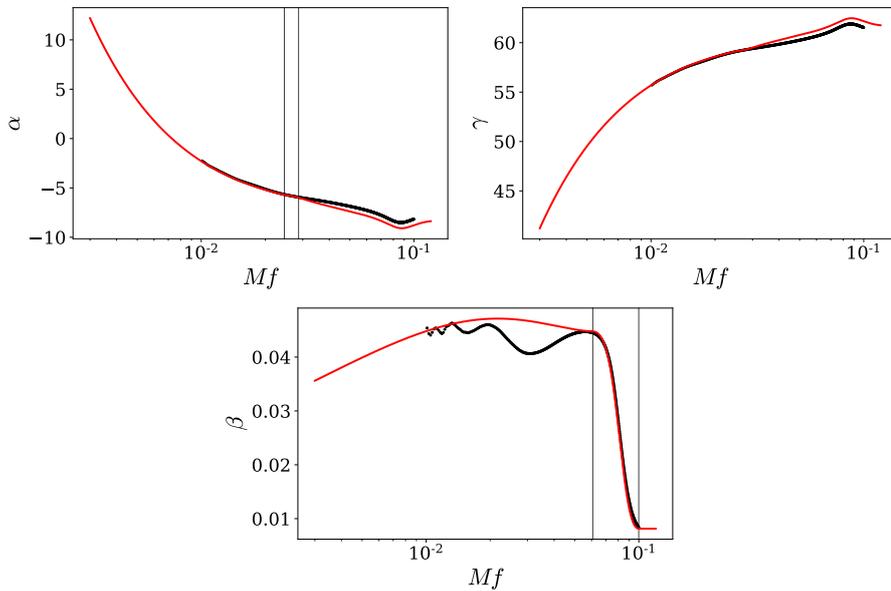


Figure 4.16: Comparison of the complete model for each of the precession angles (red line) with the NR data (black dots) for the case with $(q, \chi, \theta) = (1, 0.4, 30)$. The vertical black lines show the connection frequencies for α and β .

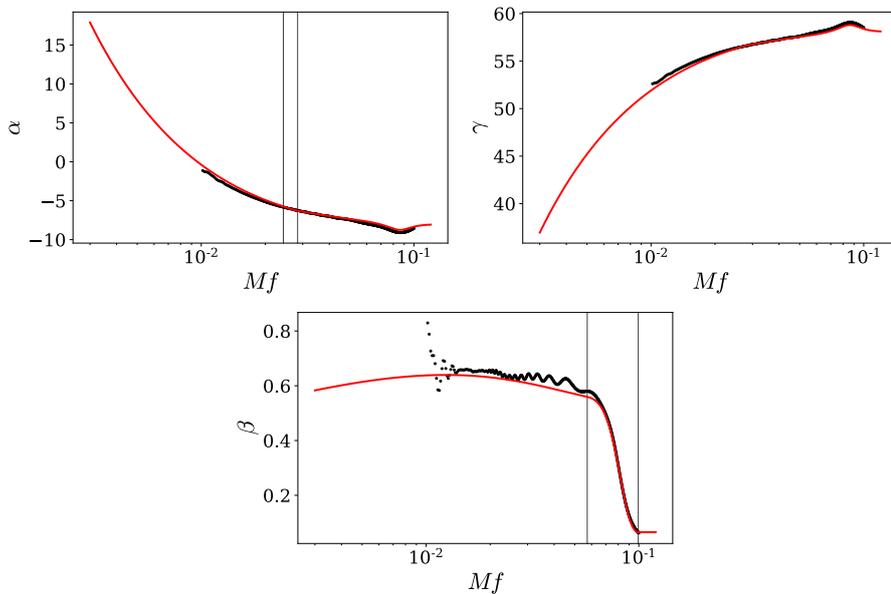


Figure 4.17: Comparison of the complete model for each of the precession angles (red line) with the NR data (black dots) for the case with $(q, \chi, \theta) = (8, 0.8, 60)$. The vertical black lines show the connection frequencies for α and β .

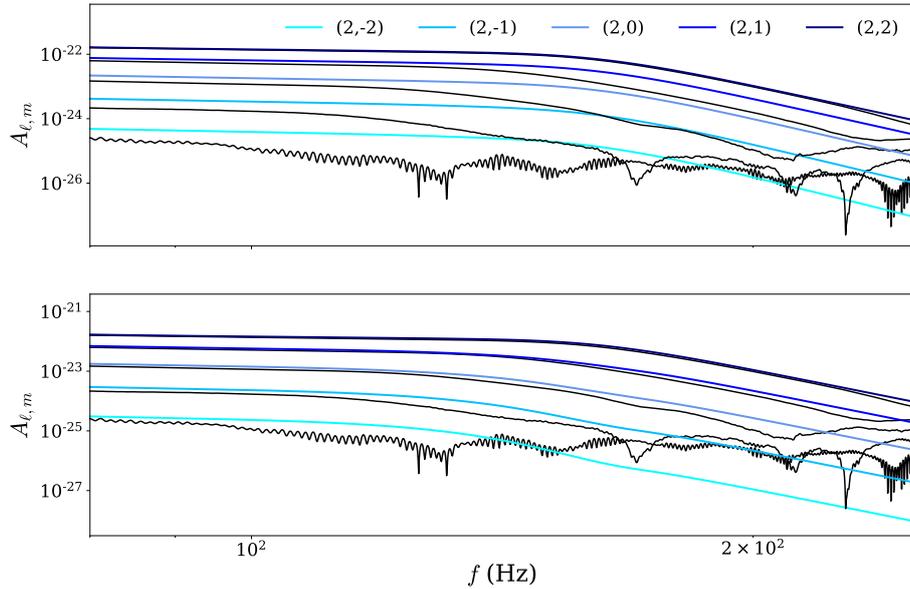


Figure 4.18: A comparison of the $\ell = 2$ modes between the NR data (shown in black on both panels) with PhenomPv3 (top panel) and the new model (bottom panel). PhenomPv3 does not show the hierarchy in the turnover frequency if each of the $\ell = 2$ modes, which is now seen in our new model. The system plotted here has $q = 4$, $\chi = 0.4$, $\theta = 90^\circ$ with a total mass of 100 solar masses at a distance of 100Mpc.

4.8 Model validation: qualitative behaviour

The improved model for the precession angles presented here enables us to capture key features of precessing waveforms which were not seen previously. These features include the hierarchy in the turnover frequency of the $\ell = 2$ modes and the time-domain behaviour of the precession angles.

4.8.1 Hierarchy in the turnover frequency of the $\ell = 2$ modes

From equation 4.26 we can see that β is approximately given by the ratio of the amplitude of the (2,2) and (2,1) modes. The rapid drop in β therefore implies that the amplitude of the (2,1) mode must have decreased relative to the (2,2) mode and so the (2,1) mode will begin to experience ringdown decay before the (2,2) mode. Once both modes are decaying exponentially (at roughly the same rate) β levels off. This trend continues for all of the $\ell = 2$ modes. The mode hierarchy that we see is consistent in the NR waveforms with this.

By capturing the turnover in β in our model we reproduce the hierarchy in the turnover frequency of the $\ell = 2$ modes, as seen in the lower panel of figure 4.18. This feature has not been modelled in previous precessing Phenom models, which is

shown in the upper panel of figure 4.18 where each of the $\ell = 2$ modes turnover at the same frequency.

4.8.2 Time domain comparisons

Accurately modelling the merger-ringdown features of the angles in the frequency domain has also enabled us to reproduce key features of the angles in the time domain after merger. The time domain angles were calculated by “twisting up” the frequency-domain co-precessing NR waveform using the frequency-domain model for the angles. We then performed an inverse Fourier transform on the resulting frequency-domain precessing waveform and calculated the precession angles in the time domain using the matrix method. Since this procedure involves an inverse Fourier transform of the NR data, the resulting angles contain some noise artefacts. These could be reduced with additional tuning but this is not required for the basic check of the phenomenology performed here.

The results of this calculation are shown in figures 4.19 and 4.20. Most notably, the time domain α now continues rising after merger rather than levelling off. If α takes a constant value it implies the precession of the optimum emission direction has stopped. As has been noted previously [122], this is clearly not seen in the NR data. This is a further feature of the precessional motion captured by the model presented here. Additionally, the time domain β now drops rapidly instead of continuing to rise. This shows we have managed to capture the closing up of the opening angle as the angular momentum is radiated away through gravitational wave emission.

4.8.3 Behaviour beyond calibration region

As with any tuned model, beyond the calibration region there is no guarantee of the accuracy of the model for the angles. However, we want to ensure that they do not display pathological or physically incorrect behaviour.

For α we see pathological behaviour for $A_3 < 0$ and physically incorrect behaviour for $A_1 < 0$ (α would decrease as a function of frequency) or $A_2 > 0$ (the dip in α would have the wrong sign). As it is only a small region of parameter space in which this might happen, we enforce the conditions that $A_1, A_3 > 0$ and $A_2 < 0$ by taking the absolute value of the co-efficients with the appropriate sign. For A_2 we replace any positive values with zero.

For β we see unphysical behaviour when the final value drops below zero. For cases with $\beta_{\text{RD}} < 0$ we therefore revert to the value for β calculated in section 4.4.1 rescaled to the value the merger-ringdown ansatz would take at the connection frequency. We see pathological behaviour for $A_4 \lesssim 0$. Physically incorrect behaviour starts to emerge when A_4 drops below $\mathcal{O}(10^2)$. In order to avoid such behaviour we require $A_4 \geq 150$ and replace the fitted value of A_4 by 150 where it falls below this value. Since $A_4 \sim 10^3$ across the majority of the parameter space this concern

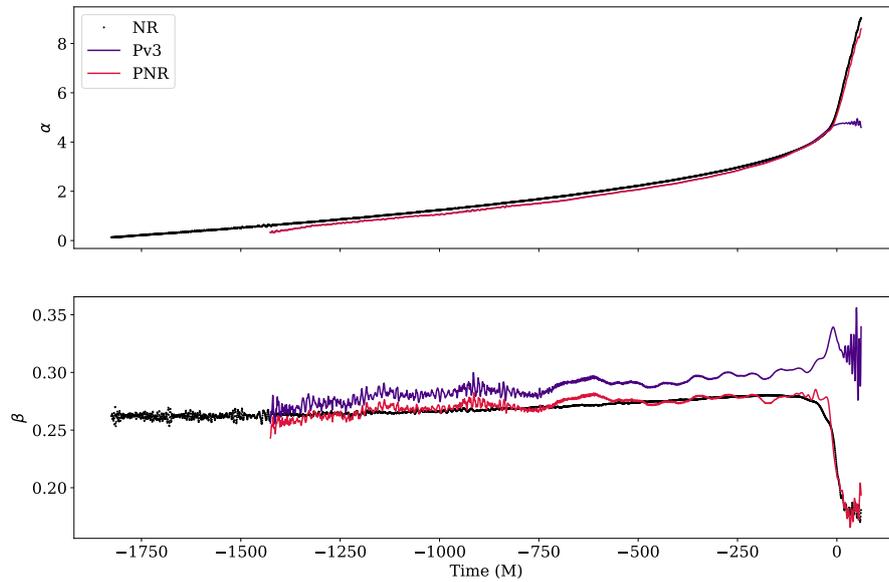


Figure 4.19: Comparison of the time domain precession angles from both the PhenomPNR and PhenomPv3 models with the NR data. These angles are for the case with $(q, \chi, \theta) = (4, 0.4, 60)$.

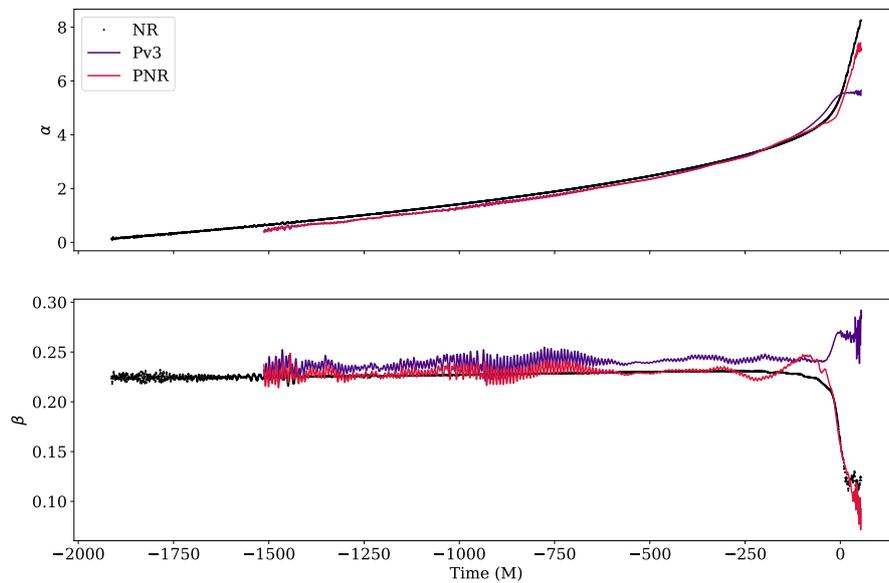


Figure 4.20: Comparison of the time domain precession angles from both the PhenomPNR and PhenomPv3 models with the NR data. These angles are for the case with $(q, \chi, \theta) = (8, 0.4, 30)$.

only arises for very extreme configurations ($\chi \approx 1$) where the accuracy of the model cannot be guaranteed.

4.9 Model validation: Matches

We now wish to test the accuracy of our angle model in the context of gravitational wave signal analysis. To do this we calculate the match (using the method detailed in section 1.6 of the Introduction) between the NR waveform and a model for the precessing waveform for a given configuration as a test of the accuracy of the model. In order to assess the accuracy of the angle model itself we model the precessing waveform by twisting up the co-precessing NR waveform with the model angles. We also assess the accuracy of a complete model produced by twisting up modified PhenomD waveforms using the angle model. The mismatch calculated between the NR waveform and the complete model will contain errors introduced by inaccuracies in the underlying co-precessing model in addition to inaccuracies in the angle model. Since we do not aim to model mode asymmetries in this work, our model does not capture them. We therefore perform matches testing the angle model using the symmetrised NR waveform (in both the \mathbf{J} -aligned and co-precessing frame). The matches assessing the accuracy of the full model are performed against both symmetrised NR waveforms and the full non-symmetrised waveforms in order to assess the effect of not including mode asymmetries.

4.9.1 Verification waveforms

We performed matches against the calibration waveforms as well as an additional set of verification waveforms. The additional set of waveforms used to test the accuracy of the model for configurations for which it was not tuned is given in table 4.1. This set includes two spin configurations which belong to the SXS catalogue.

4.9.2 Matches: Accuracy of the angle model

In order to test the accuracy of the angle model we constructed a set of precessing waveforms by calculating the symmetrised frequency-domain co-precessing NR waveform containing only the $\ell = 2$ modes and “twisting” this waveform up with the model for the angles. We constructed two sets of precessing waveforms in this fashion; the first using the model for the angles presented in this chapter and the second using the PhenomPv3 angles in order to quantify the effect of modelling the merger-ringdown behaviour of the angles. We then calculated the precessing match between these waveforms and the symmetrised NR waveforms comprising only the $\ell = 2$ modes. These matches were calculated at a fixed total mass $M = 100M_{\odot}$ and distance 100 Mpc. For ease of presentation, the full precessing match is calculated

Case No.	Simulation ID	M_{ref}	q	χ	θ ($^\circ$)	χ_1	χ_2
01	BAM	1.17×10^{-3}	1.000	0.600	90	(0.212, 0.561, 0.003)	(0, 0, 0)
02	BAM	1.16×10^{-3}	1.000	0.600	150	(0.150, 0.261, -0.519)	(0, 0, 0)
03	SXS0077	7.31×10^{-4}	1.500	0.500	90	(0.500, 0, 0)	(0, 0, 0)
04	SXS0097	7.20×10^{-4}	1.500	0.500	90	(0.500, 0, 0)	(0, 0, 0)
05	SXS0080	7.31×10^{-4}	1.500	0.500	90	(0.500, 0, 0)	(0, 0, 0)
06	SXS0018	6.18×10^{-4}	1.500	0.500	90	(0.500, 0, 0)	(0, 0, 0)
07	SXS0092	7.34×10^{-4}	1.500	0.500	150	(0.250, 0, -0.433)	(0, 0, 0)
08	BAM	1.15×10^{-3}	2.000	0.600	30	(0.192, 0.230, 0.520)	(0, 0, 0)
09	BAM	1.17×10^{-3}	2.000	0.600	60	(0.265, 0.446, 0.302)	(0, 0, 0)
10	BAM	1.14×10^{-3}	2.000	0.600	90	(0.303, 0.518, 0.003)	(0, 0, 0)
11	BAM	1.17×10^{-3}	2.000	0.600	120	(0.303, 0.424, -0.298)	(0, 0, 0)
12	BAM	1.13×10^{-3}	2.000	0.600	150	(0.169, 0.250, -0.519)	(0, 0, 0)
13	SXS0033	8.79×10^{-4}	3.000	0.500	30	(0.250, 0, 0.433)	(0, 0, 0)
14	SXS0035	6.70×10^{-4}	3.000	0.500	90	(0.500, 0, 0)	(0, 0, 0)
15	BAM	1.27×10^{-3}	4.000	0.600	60	(0.425, 0.299, 0.301)	(0, 0, 0)
16	SXS1109	8.00×10^{-4}	5.000	0.500	90	(0.500, 0, 0)	(0, 0, 0)
17	SXS0058	8.00×10^{-4}	5.000	0.500	90	(0.500, 0, 0)	(0, 0, 0)
18	BAM	1.55×10^{-3}	8.000	0.600	30	(0.296, 0.059, 0.519)	(0, 0, 0)
19	SXS0161	7.24×10^{-4}	1.0	0.792	139	(0.52, -0.003, -0.299)	(0.520, -0.003, -0.299)
20	SXS0115	8.02×10^{-4}	1.07	0.244	74	(0.022, 0.006, -0.204)	(0.212, -0.139, 0.29)
21	SXS0116	8.01×10^{-4}	1.08	0.212	53	(-0.028, 0.097, 0.033)	(0.153, -0.105, 0.102)
22	SXS0124	8.00×10^{-4}	1.26	0.393	41	(0.227, -0.121, 0.07)	(0.043, -0.115, 0.286)
23	SXS0102	7.30×10^{-4}	1.5	0.5	90	(0.500, 0, 0)	(0.500, 0, 0)
24	SXS1397	4.02×10^{-4}	1.561	0.299	111	(0.224, 0.092, -0.176)	(-0.437, -0.159, 0.109)
25	SXS0135	8.04×10^{-4}	1.64	0.186	128	(-0.093, 0.064, 0.024)	(0.109, -0.232, -0.228)
26	SXS0143	8.02×10^{-4}	1.92	0.441	28	(-0.107, -0.042, 0.436)	(-0.200, -0.386, -0.090)
27	SXS0144	8.03×10^{-4}	1.94	0.214	146	(0.074, 0.095, -0.288)	(-0.054, -0.011, 0.215)
28	SXS0049	8.82×10^{-4}	3.0	0.527	72	(0.500, 0, 0)	(0, 0, 0.500)
29	SXS1160	8.74×10^{-4}	3.0	0.658	63	(0.517, 0.281, 0.379)	(-0.357, -0.412, -0.251)
30	SXS0165	1.42×10^{-3}	6.0	0.93	125	(0.739, 0.185, -0.496)	(-0.190, 0.049, -0.228)

Table 4.1: Additional configurations used to verify the accuracy of the model in cases where it had not been tuned.

for a range of inclinations as described in section 1.6.2 of the Introduction, but we plot only the mean of these values for the majority of cases.

From figure 4.22 we can see that the matches between the model waveforms and the NR for each of the calibration waveforms are above 0.99 across the majority of the parameter space. The only cases for which this is not true are in the most extreme corner of the parameter space we modelled; cases with $q = 8$, $\chi = 0.8$ and $\theta \geq 90^\circ$. In these cases we find the PN expressions used for α during inspiral deviate from the NR waveform at reasonably low frequencies— in the case of $q = 8$, $\chi = 0.8$, $\theta = 120$ before the start of the NR waveform (as shown in figure 4.21). Improving the model for these cases would require a model for the intermediate region between where the PN expression ceases to be accurate and where the current model begins, which in itself would require longer NR waveforms to be produced. However, despite the matches falling below the level we might wish, they still show significant improvement on the previous model.

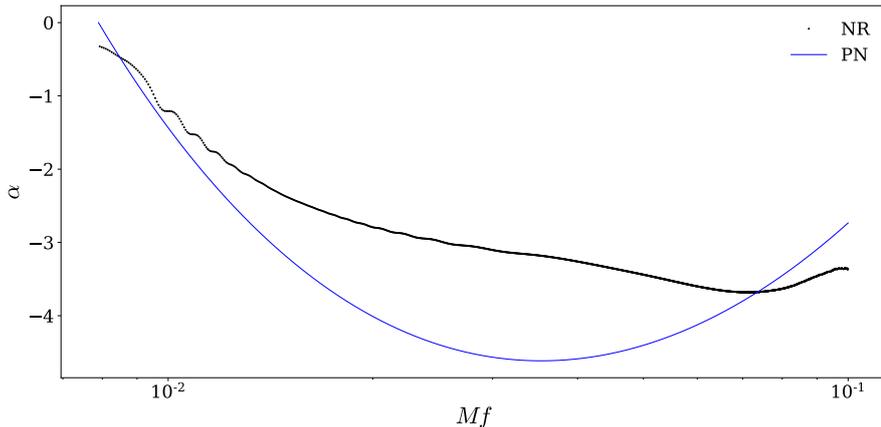


Figure 4.21: Comparison of the value of α from the PN expression with the value calculated from the NR waveform. In order to see a region over which the two values agree well we would need a longer NR waveform. In order to have a complete IMR model for α in this case we would also require a model of the ‘intermediate’ region which was tuned to the NR data.

The best matches are seen in the least extreme part of parameter space; namely for low mass ratio systems. However this is the region of parameter space where existing models for the angles already perform reasonably well. The biggest improvement in the matches as a result of the improved model for the angles is seen at higher mass ratios, particularly for larger θ . A selection of these cases are demonstrated in figure 4.23. As expected, the matches against symmetrised NR waveforms are symmetric about an inclination of 90° .

Figure 4.24 shows the matches between the model waveforms and a set of single-spin cases which the model was not tuned to, while figure 4.25 shows the matches for a set of two-spin cases. Since none of the single spin cases are towards the more

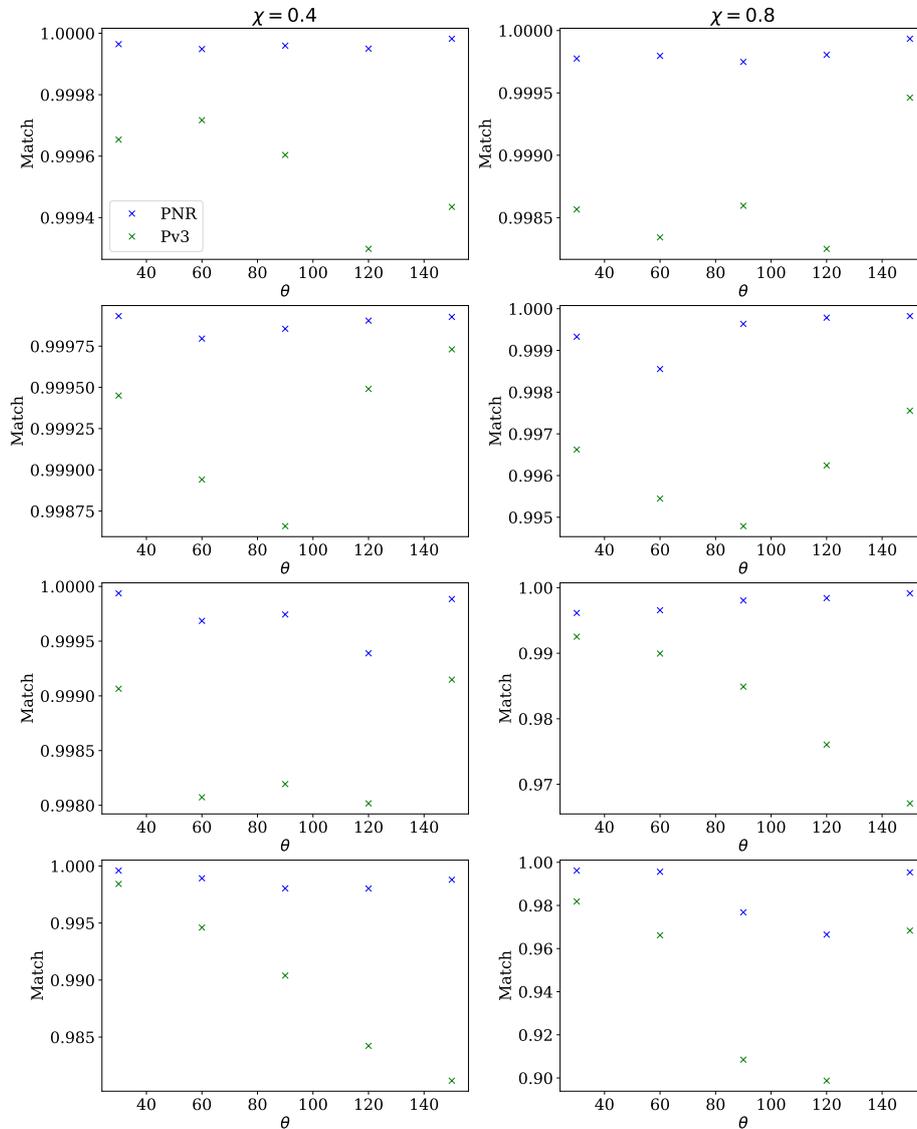


Figure 4.22: Precessing matches for each of the 40 NR waveforms used to calibrate the model, averaged over inclination. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (blue) and twisted up with the angle model used by PhenomPv3 (green). The plots show cases with $\chi = 0.4$ on the left and $\chi = 0.8$ on the right. Starting from the top, each of the rows show $q = 1, 2, 4$ and 8 respectively.

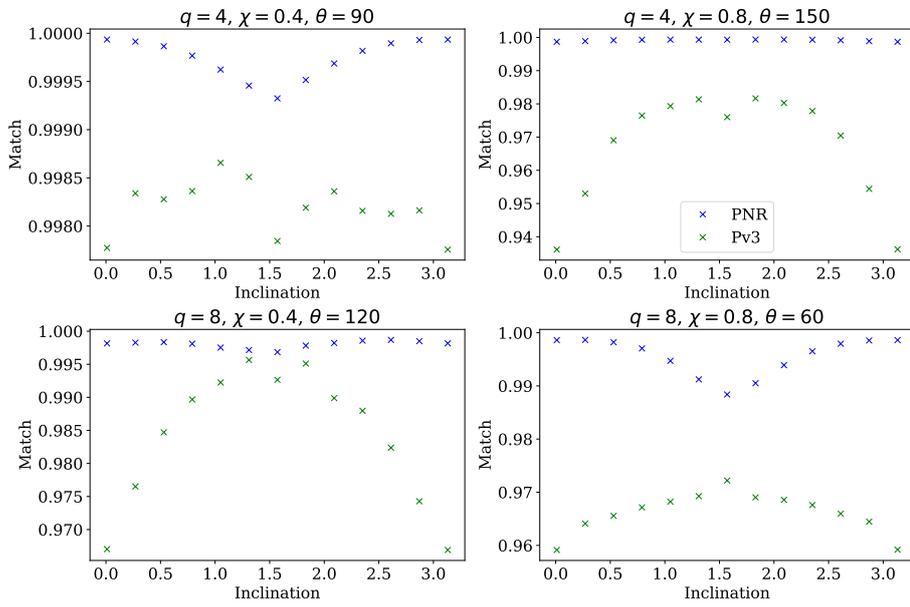


Figure 4.23: Precessing matches for four of the 40 NR waveforms used to calibrate the model. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (blue) and twisted up with the angle model used by PhenomPv3 (green).

extreme part of the parameter space, all the matches are above 0.99. As a general trend we can see that the improvement in the match obtained by modelling the merger-ringdown region is greatest towards for more extreme cases; particularly for higher mass ratios. The majority of the single spin cases are also for fairly low mass ratios and spin magnitudes. We see good matches for these cases with both the new model for the angles and the PhenomPv3 angles. However, for two cases (29 and 30) which have higher mass ratios and spin magnitudes we see a measurable improvement in the match. The good matches for these two-spin cases also provides an important verification of the mapping from the two-spin case to the appropriate single-spin configuration for the merger-ringdown part of the model described in section 4.2.2 as well as the tapering of the oscillations apparent in the PN two-spin expression for β , which is described in section 4.4.2.

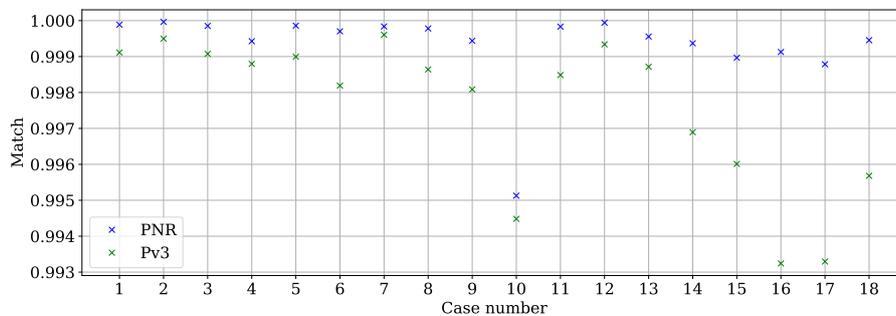


Figure 4.24: Precessing matches for the single spin verification NR waveforms, averaged over inclination. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (blue) and twisted up with the angle model used by PhenomPv3 (green).

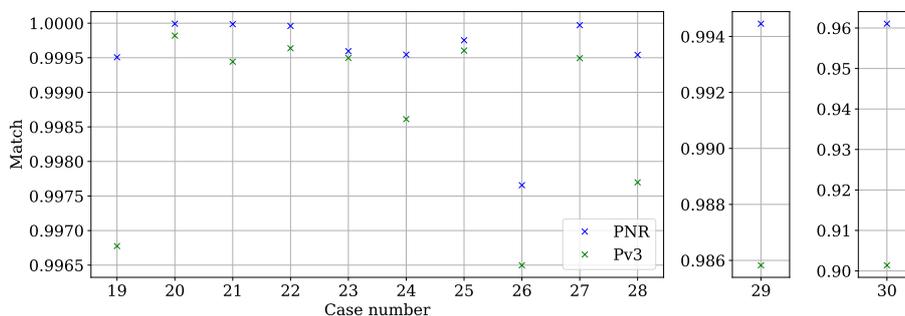


Figure 4.25: Precessing matches for the two spin verification NR waveforms, averaged over inclination. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and the co-precessing NR waveform twisted up with the angle model presented here (blue) and twisted up with the angle model used by PhenomPv3 (green).

4.9.3 Matches: Accuracy of the full model

The complete precessing model is made by “twisting up” a model for the underlying co-precessing waveform using a model for the precession angles. Here we use the modified PhenomD model as the co-precessing waveform, as has been used in previous versions of PhenomP. The modified model has adjusted the ringdown part of the waveform to have the correct final spin for the remnant of the precessing system as described in section 1.5.3 of the Introduction. PhenomPNR is formed by “twisting up” this waveform using the model for the angles presented earlier in the chapter. We show the match between PhenomPNR and the symmetrised $\ell = 2$ modes of the NR waveform in figure 4.26 and the match with the unsymmetrised $\ell = 2$ modes of the NR waveform in figure 4.27. The match between PhenomPv3 and the same NR waveforms is shown for reference.

As we can see from figure 4.26, the match between the complete model and the NR waveforms is significantly worse than the matches testing the angle model shown in figure 4.22. In most cases, this is dominated by inaccuracies in the co-precessing model, as can be seen in figure 4.28. In cases with high mass ratio, high spin and a large opening angle, the co-precessing model performs well. Inaccuracies in the approximation that the co-precessing waveform can be modelled by an aligned-spin model in the strong field regime were first discussed in [139]. The main inaccuracy we observed in the model is in the estimation of the location of the dip in the co-precessing phase derivative. For cases where the co-precessing model performs well, the dip in the co-precessing phase derivative occurs at lower frequencies meaning less of the inspiral part of the signal contributes. This reduces the de-phasing between the model and the co-precessing NR waveform. It can also be seen from figure 4.26 that PhenomPv3 performs better in certain regions of the parameter space than PhenomPNR, which is not seen in the matches testing the angle model alone. We think this is because the model for the angles used in PhenomPv3 is much smoother than the model used in PhenomPNR and does not include any of the merger-ringdown features. This means the change in the phase of the waveform is less marked between inspiral and merger-ringdown so the overall match of the two waveforms is better, whereas with the current version of PhenomPNR we can obtain a good match during either inspiral or merger-ringdown but not over the whole of the waveform. We expect this will be improved by an improved model of the co-precessing phase (and amplitude) but this is left for future work.

Comparing figures 4.26 and 4.27, we can see that the exclusion of mode asymmetries in our co-precessing model and in the model for the angles further degrades the match. The largest degradation occurs for large mass ratios where precession effects are most likely to be strongest and we are therefore most likely to be able to identify precession in the signal. In order to have a fully accurate precessing model we therefore need to include the effects of mode asymmetries. This is left to future

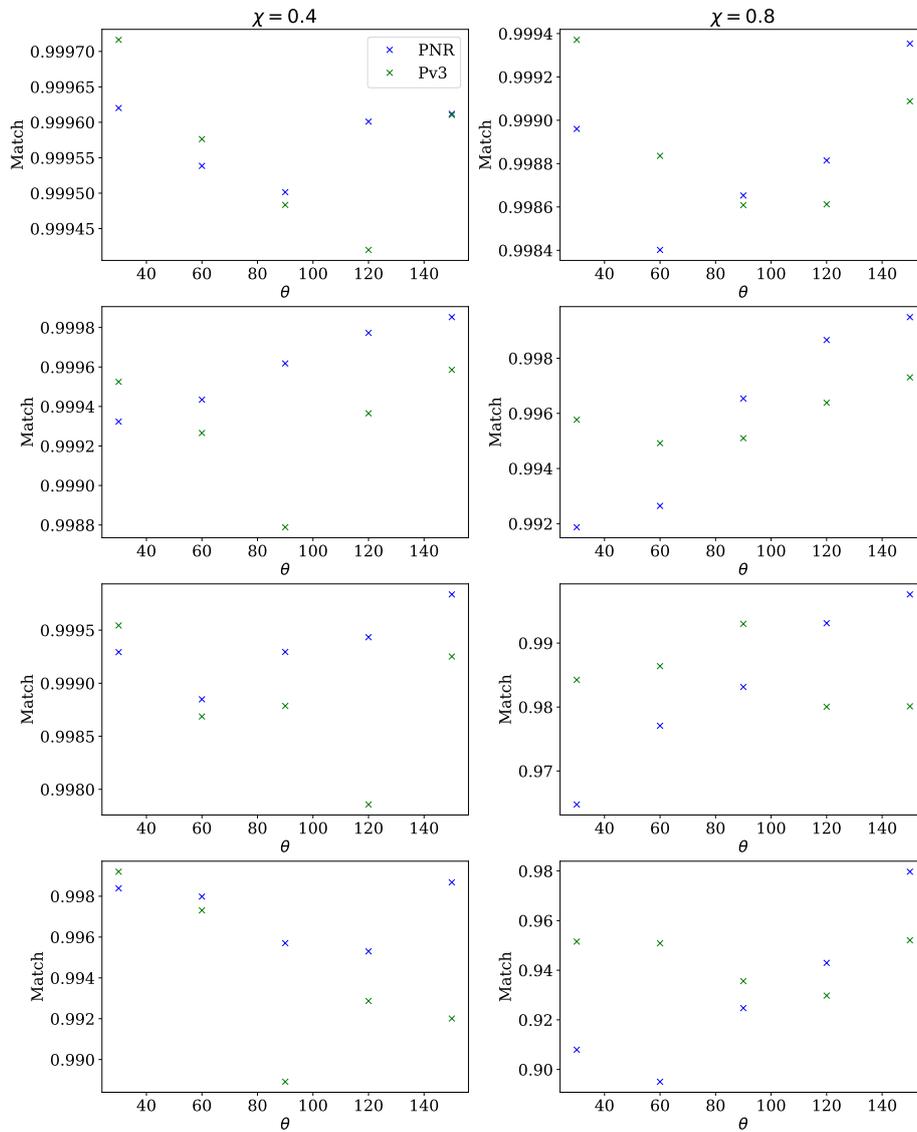


Figure 4.26: Precessing matches for each of the 40 NR waveforms used to calibrate the model, averaged over inclination. These matches are between the symmetrised NR waveforms in the \mathbf{J} -aligned frame and complete precessing model— PhenomPNR is shown in blue and PhenomPv3 is shown in green.

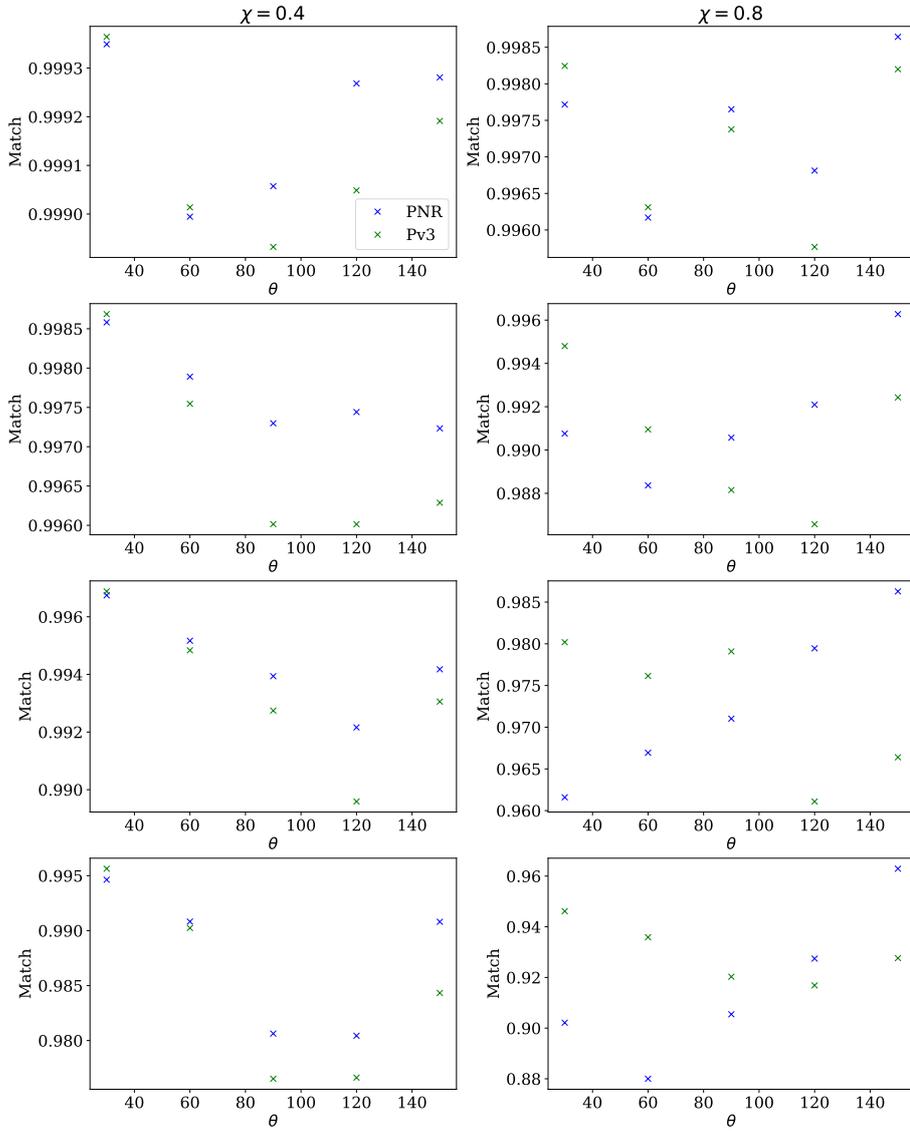


Figure 4.27: Precessing matches for each of the 40 NR waveforms used to calibrate the model, averaged over inclination. These matches are between the unsymmetrised NR waveforms in the \mathbf{J} -aligned frame and complete precessing model—PhenomPNR is shown in blue and PhenomPv3 is shown in green.

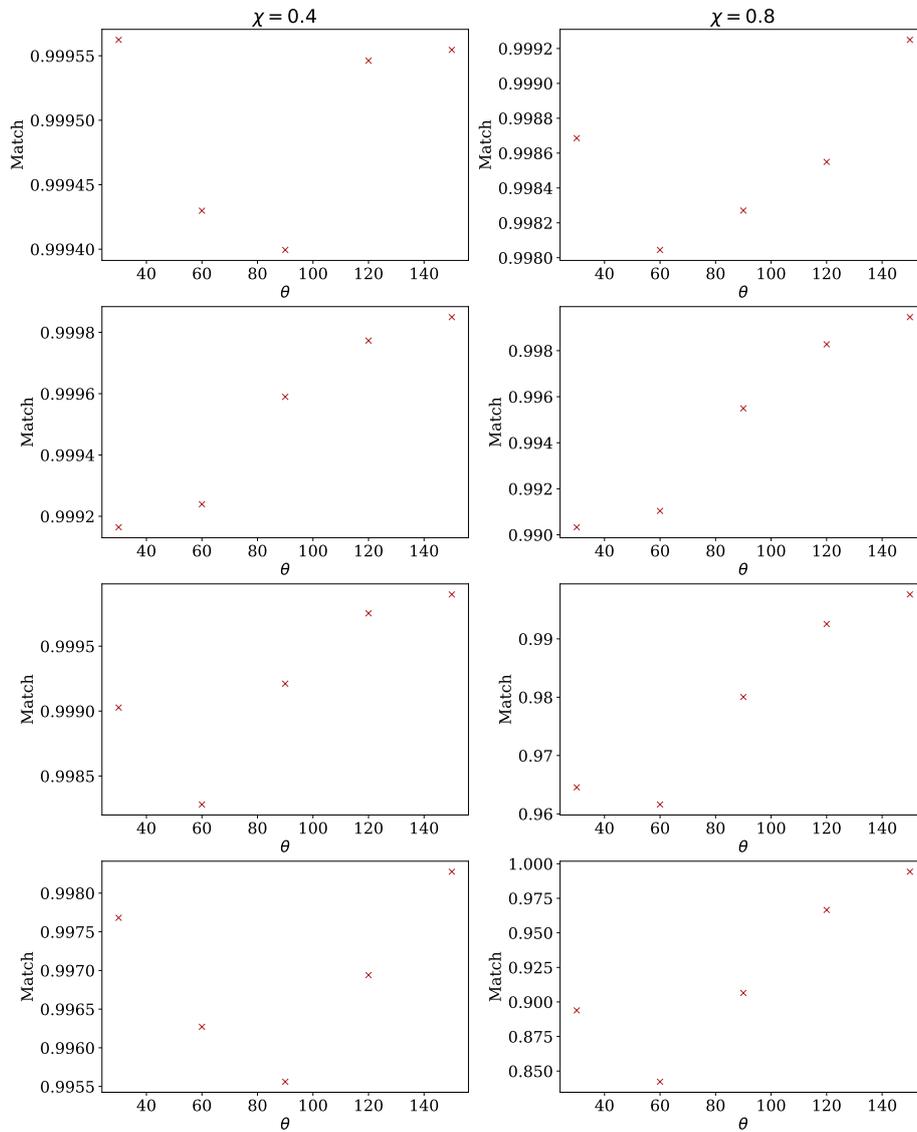


Figure 4.28: Non-precessing matches for each of the 40 NR waveforms used to calibrate the model. These matches are between the symmetrised co-precessing NR waveforms in the \mathbf{J} -aligned frame and modified PhenomD which provides the co-precessing model employed by PhenomP versions 1-3.

work.

4.10 Conclusion

We have produced a complete IMR model for the precession angles $\{\alpha, \beta, \gamma\}$. This model uses the most accurate PN expressions for the angles available, including a new technique to calculate β that describes the precessional motion of the direction of optimal emission rather than of the orbital angular momentum. This technique is also presented here. These expressions are connected to a phenomenological expression for the merger-ringdown angles, enforcing C^1 -continuity. The merger-ringdown expressions have been tuned to a catalogue of 40 NR simulations of single spin systems covering the parameter space of mass ratios from one to eight, spin magnitudes of 0.4 and 0.8 and an inclination angle between the spin and the orbital angular momentum between 30° and 150° , where the spin has been placed on the larger black hole. The expressions are therefore single spin expressions. By making use of a non-bijective mapping between the six spin components required to describe a two spin system and the two required to describe a single spin system we have produced an IMR model for the precession angles of a two spin system. We leave to future work the production of a model for the angles that employs merger-ringdown expressions tuned to a catalogue of NR simulations containing two spin systems.

This model for the angles can be used to “twist up” any co-precessing waveform model to produce a model for precessing systems. The matches presented in section 4.9.2 show the result of using angle model to twist up the co-precessing NR waveform. From this, we can see that across the majority of the parameter space the model for the angles reproduces the precessing waveform with a very high degree of accuracy. However, section 4.9.3 it is shown that when using modified PhenomD as the underlying co-precessing model the matches are much poorer. This is a result of the poor agreement between the modified PhenomD model and the co-precessing NR waveform, especially in the more extreme regions of the parameter space. A complete, accurate model of precessing systems will also need to include higher modes and mode asymmetries. These improvements are left to a future work.

The improved model of the angles allows us to capture features of the precessing waveform which have not previously been modelled. The key features that have been captured by this model are the hierarchy in $\ell = 2$ modes seen in section 4.8.1 and the merger-ringdown features seen in the time-domain angles in section 4.8.2.

Chapter 5

Conclusion

As we settle into the era of gravitational wave astronomy with an ever increasing number of gravitational wave events [6, 7, 9, 10, 13–16] detected by the LIGO [3] and Virgo [17] detectors, we need to be prepared to analyse the wide range of systems which may be detected in the future. The detectors themselves are increasing in sensitivity. In the fourth observing run of the current generation of ground-based gravitational wave detectors (O4), which is scheduled to start in late 2021 or early 2022, the LIGO detectors will be nearing design sensitivity and Virgo will have completed phase 1 of their final upgrade. This increases both the signal-to-noise ratio at which events will be detected and the volume of the Universe observable by the detectors— between the third and fourth observing runs the search volume for BBH systems is expected to go from 3.4×10^8 Mpc yr to 1.5 Gpc yr. The number of detectors is also increasing with the addition of KAGRA [27, 155] for O4 and IndIGO [108] in approximately 2025. This will hopefully lead to an increased fraction of time for which two or more detectors are operating and so an increased number of confident detections. The full four-detector network expected to be operational by O4 is predicted to make 79_{-44}^{+89} confident BBH detections as opposed to the 17_{-11}^{+22} predicted during the third observing run and the 10 BBH signals detected in the first and second observing runs [109].

We expect that in detecting more events we will consequently encounter more events in extreme parts of the parameter space— signals from systems with high mass ratios and high spin magnitudes. These are systems for which precession effects will be greater and so more easily detectable. The greater signal-to-noise ratio at which some of these events will be detected will also enable us to make more detailed inferences about the properties of the system emitting the gravitational waves. This means we will be able to detect more subtle properties, such as the spins and— crucially— the amount of precession. We therefore need to be able to confidently identify, describe and model the gravitational wave signals emitted by precessing black-hole-binaries.

The main focus of this thesis has been to provide a precision description of pre-

cessing BBH systems. We have made crucial steps towards achieving this. Firstly, we have produced a catalogue of 40 NR waveforms from single spin precessing systems which systematically covers previously unexplored parts of the precessing parameter space. Secondly, we have developed a method to reliably and accurately relate the dynamical properties of a BBH system to the waveform quantities for data from NR simulations. Finally, we developed the first model of signals from precessing BBH waveforms where the precession effects have been tuned to NR.

In Chapter 2 we presented the NR catalogue of precessing systems which forms the basis of the tuned precessing model presented in Chapter 4. We assessed the accuracy of these waveforms and found that they agree with the extrapolated “true” signal to within 0.4%. This provides an estimate of the limit on the mismatch error of the waveforms used to tune the precessing model.

In Chapter 3 we provided a method for prescribing the configuration of the BBH system at a given point in the waveform. It has been previously shown how to calculate the direction of optimal emission from a waveform [149], which provides the direction of the orbital angular momentum to a good approximation. In this chapter we presented a method for finding the unit separation vector of the black holes from the waveform. We also showed how this method can be used to calculate the time shift between the dynamical and waveform quantities extracted from a simulation. This enables us to confidently relate the black holes’ spins to a given point in a waveform although the value of the spins will still contain ambiguities due to gauge choices. Being able to accurately and reliably specify the configuration of a BBH system at a given point in a waveform has two main uses. Firstly it is useful when using NR waveforms as templates to compare against detected signals. Secondly it is useful when developing, testing and using a waveform model that has been tuned to NR waveforms as it enables us to know the exact system we are studying.

In Chapter 4 we presented a complete IMR model of the Euler angles $\{\alpha, \beta, \gamma\}$ that describe the precession of the direction of maximum emission of a BBH system—the precession angles. The model for α and β uses expressions derived from PN theory for the inspiral part of the waveform and phenomenological expressions for the merger-ringdown part where the co-efficients in the expressions have been tuned to NR. The inspiral expression for β employs a technique to get the inclination of the direction of maximum emission of the gravitational waves as opposed to the inclination of the orbital angular momentum of the system. This technique can be applied to any PN expression for β derived using the orbital angular momentum and provides a more appropriate value of β since it is the precession of the waveform, not the dynamical quantities that we wish to track. The phenomenological expression used to describe the merger-ringdown angles agree with the NR data to within 10% on average. In the model for the precession angles presented here we calculate γ numerically using the minimal rotation condition since the very small inaccuracies

in the closed form PN expressions for gamma are sufficient to have a serious impact on the accuracy of the final precessing model.

Using this model for the precession angles we can produce a waveform model for precessing systems by “twisting up” a non-precessing waveform with the angles using the method first described in [149]. We found that if we assume our model for the underlying co-precessing waveform is completely accurate (which can be achieved by using the NR co-precessing waveform) then the resulting precessing waveform is accurate to within 1% across the majority of the precessing parameter space (except for systems with mass ratio 8, spin magnitude 0.8 and spin angle greater than or equal to 90°) and within 3% across the whole of the single spin space to which the model was tuned. This is a good test of the accuracy of the angle model itself. We see significant improvement in the accuracy of the angle model in the most extreme parts of parameter space (mass ratios greater than 4 and spin magnitudes around 0.8) as well as a reasonable improvement across the whole parameter space. This is an important achievement since, although there may not be many gravitational wave events detected with such extreme properties, the more extreme regions of parameter space still need to be reliably sampled while analysing the properties of less extreme events. The LVC has already detected a system with a mass ratio of around 3 which showed strong evidence of non-zero spin [16]. It is therefore already necessary to have a reliable precessing model for systems up to mass ratios of 8 at least.

The results presented in Chapter 4 also show that the current co-precessing model is not sufficient to translate these improvements in the model for the precession angles into improvements in the overall waveform model. An improved model for the co-precessing waveform which has been tuned to NR simulations is therefore necessary to obtain an accurate precessing model at high mass ratios and spin magnitudes. Such a model is left to future work.

Although modelling precessing systems has been greatly simplified by the realisation that they can be modelled by twisting up non-precessing waveforms using a model for the precession angles, it remains an extremely complicated task with many outstanding challenges. In order to have a complete waveform model for precessing systems we must also incorporate higher order modes (with $\ell \geq 2$), mode asymmetries and two-spin effects. These features will all need to be calibrated to NR through merger and ringdown.

To produce a complete precessing model will therefore require the production of a tuned co-precessing model for each of the higher modes, which incorporates mode asymmetries, as well as a tuned model for the angles which includes two-spin systems. As a first attempt to incorporate higher modes into this tuned precessing model we could produce a tuned model of the (2,2) mode in the co-precessing frame and then extend this to a higher-mode co-precessing waveform model using the same techniques employed in PhenomHM [112]. Although this is not as accurate as

a fully tuned higher-mode co-precessing model it would enable us to take advantage of the improvements to the model for the precession angles presented here. These improvements increase the accuracy of the precessing waveform model, particularly at higher mass ratios where higher modes are most likely to be important.

The current tuned model for the angles could be improved by calibrating it to an increased number of NR waveforms which cover the single spin parameter space more densely. A particular improvement might come from simulations with an increased number of spin magnitudes in order to extend the current fits beyond their linear dependence on χ . Furthermore, if we are extending the model to higher modes then we should use the modes which compose the co-precessing model to calculate the precession angles for consistency. We do not expect the inclusion of these higher order modes to greatly impact the morphology of the precession angles so the model presented here may remain sufficient. However, the effect should be investigated. Finally, as was seen in Chapter 4, the NR waveforms used to tune the model are not long enough to match up well with the PN angles at low frequencies. This is true for β for the majority of cases and for α in the most extreme part of the parameter space. To be confident we have produced the most accurate model possible will therefore require longer NR waveforms and a tuned model for the intermediate region connecting the PN expressions to the current merger-ringdown model or more accurate PN expressions.

Incorporating two-spin effects when modelling the angles also poses a significant challenge. We do not see the oscillations in the precession angles for simulations of the length of those presented in the catalogue in Chapter 2. They are only visible in the SXS waveform shown in figure 4.3, which has a starting frequency around a third of that chosen for our BAM simulations. Producing a catalogue of simulations of an equivalent length with the BAM code would be computationally prohibitive as a result of the high resolution required to limit the dephasing of the waveform. The tapering of the oscillations is not seen in the PN expression for the precession angles, rather the amplitude of the oscillations grows monotonically towards merger. A tuned model of the precession angles for two-spin systems will therefore require either PN expressions which are more accurate to higher frequencies and show the tapering of the oscillations or a tuned model of the envelope used to taper the oscillations. This would require the production of a set of two-spin NR waveforms over the specific frequency range over which the oscillations taper off. We will also need to produce a higher dimensional parameter space fit of the merger-ringdown model presented in Chapter 4 that can be tuned to two-spin NR simulations. This would require a much larger catalogue of simulations than that presented in Chapter 2 as we would need to systematically explore this higher dimensional space.

Before embarking on the production of such a comprehensive model we should first assess the accuracy of the existing model to see whether all the effects listed above need to be incorporated. The hierarchy as to which effects have the greatest

impact can then be established. From the analysis presented so far, it appears that a tuned model for the (2,2) mode of the co-precessing waveform will produce the largest improvement across the majority of the parameter space. We will need to perform further studies to assess whether including higher modes or mode asymmetries produce the next greatest improvement and what the impact of having a tuned higher mode co-precessing model would be. A tuned model for the intermediate region for the precession angles is clearly required to improve the model in the more extreme regions of parameter space. Such studies will be useful to guide where the most effort should be applied to the development of models in the near future.

It may also be worthwhile extending the parameter space coverage of the current model. This would require the production of further NR simulations at higher mass ratios or spin magnitudes. Increasing the parameter space coverage up to mass ratio 18 would require the production of a further ten simulations. Assuming these simulations would be of a similar length to the simulations which make up the current catalogue, this would require around 350000 CPU hours per simulation which, while costly, is not impossibly prohibitively so. However, the inclusion of additional spin magnitude simulations might begin to make it unfeasible. Since we are more likely to see systems with lower mass ratios and a range of spin magnitudes it is probably more valuable to focus on filling in the existing parameter space first.

Alternative modelling techniques have been developed which do not have problems incorporating the wide range of features enumerated above range of features—such as the surrogate model NRSur7dq4 [162]. However, these techniques face other challenges. For example, NRSur7dq4 can only be used for the analysis of high mass systems where we detect only the last few orbits before merger since this is the length of the NR simulations used to build the model. Overcoming this will require the production of a large catalogue of hybrid waveforms for two spin precessing systems, which must include higher modes and mode asymmetries. This entails its own difficulties. Thus, it is necessary to continue to develop a complete precessing model in order to meet the challenges posed by ever more sensitive gravitational wave detectors.

Appendices

Appendix A

Ansatz co-efficients

In this appendix we provide tables containing the values of the co-efficients λ_{pqr}^i of the global fits described by equation 4.35. We provide fits for each of the co-efficients in the merger-ringdown expressions for α (equation 4.33) and β (equation 4.34). These fits take the form of a polynomial expansion in each of the three model parameters. They were tuned to a catalogue of 40 NR simulations and so we limit them to a total of 39 terms with the highest order term in each dimension restricted to one less than the number of simulations in that dimension. As can be seen from the tables, none of the fits require the full 39 terms.

	χ_{lmn}^i					
	100	010	001	110	101	011
A_1	$9.93648310 \times 10^{-2}$	$2.38962937 \times 10^{-1}$	$-2.87751861 \times 10^{-1}$	-4.32424432	5.91396924	1.11424866
A_2	$1.23482795 \times 10^{-1}$	$-3.02796689 \times 10^{-2}$	$-1.44303539 \times 10^{-1}$	$1.70717689 \times 10^{-1}$	2.76946661	-
A_3	$-3.83556507 \times 10^{-3}$	$-1.43919236e \times 10^{-4}$	$1.63656220 \times 10^{-2}$	$2.01559774 \times 10^{-2}$	$-3.07053277 \times 10^{-1}$	$-2.37862800 \times 10^{-2}$
A_4	$6.17416517 \times 10^{-1}$	$1.22210404 \times 10^{-1}$	$-6.52661264 \times 10^{-2}$	-1.13171539	1.31553469	$1.71589699 \times 10^{-1}$
B_0	1.23668139×10	1.58433091	-3.50288841	-2.82778560×10	5.94143472×10	2.80604806
B_1	2.67212028×10	2.34455952×10	3.09790853×10	-3.62120408×10^2	-7.26258395×10^2	-6.33417778×10
B_2	-1.21934795×10^3	-4.94271843×10^2	-7.01770441×10^2	8.32961622×10^3	1.88061237×10^4	1.48731179×10^3
B_3	1.02593191×10^4	2.52076450×10^3	3.73557463×10^3	-4.73357823×10^4	-1.19883104×10^5	-8.15733657×10^3
B_4	-3.29820780×10^4	-6.99841886×10^3	-1.98160698×10^5	2.29474908×10^5	3.97103544×10^6	3.37837240×10^5
B_5	-1.88421360	$-1.12733891 \times 10^{-1}$	$-4.87705135 \times 10^{-1}$	2.71645116	1.04911098×10	1.25553043

Table A.1: Table containing co-efficients of the three-dimensional polynomial descriptions of the ansätze for α and β . Terms in this table are up to linear order in up to two parameters.

	χ_{lmm}^i				
	200	002	210	201	102
A_1	-1.63986502	2.16240581 $\times 10^{-1}$	2.75334818 $\times 10$	-3.66840816 $\times 10$	-3.15838708
A_2	-1.35989152	-3.75968094 $\times 10^{-1}$	4.98904724 $\times 10^{-1}$	-1.73408329 $\times 10$	6.68046879
A_3	6.45587933 $\times 10^{-2}$	7.94852752 $\times 10^{-3}$	-2.15122495 $\times 10^{-1}$	1.80924724	-1.65162853 $\times 10^{-1}$
A_4	-1.21465632	1.02792395 $\times 10^{-1}$	2.99504477	-7.71328765	-8.75326916 $\times 10^{-1}$
E_0	-1.00945399 $\times 10^2$	1.49199121	1.81694072 $\times 10^2$	-3.11849123 $\times 10^2$	-3.10016540 $\times 10$
B_1	-3.35586115 $\times 10^2$	9.80203695 $\times 10$	2.32320527 $\times 10^3$	4.60875997 $\times 10^3$	-1.36704446 $\times 10^3$
B_2	1.33153740 $\times 10^4$	-2.37866492 $\times 10^3$	-5.56186978 $\times 10^4$	-1.27563018 $\times 10^5$	3.15275462 ⁴
B_3	-1.07722901 $\times 10^5$	1.52372005 $\times 10^4$	3.29748581 $\times 10^5$	8.69020418 $\times 10^5$	-1.92966679 $\times 10^5$
B_4	5.88395888 $\times 10^5$	-1.18024617 $\times 10^4$	-1.92995181 $\times 10^6$	-2.47067181 $\times 10^7$	4.95347903 $\times 10^5$
B_5	1.48326489 $\times 10$	5.49422204 $\times 10^{-3}$	-2.24028760 $\times 10$	-6.55564486 $\times 10$	1.33895041

Table A.2: Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are up to second order in one of the parameters.

	χ_{lmn}^i					
	012	111	211	112	202	212
A_1	$-7.84827961 \times 10^{-1}$	-2.13618400×10	1.33628279×10^2	1.47783939×10	9.20125054	-6.76291766×10
A_2	$2.90881928 \times 10^{-1}$	$-8.59278587 \times 10^{-2}$	$5.37382649 \times 10^{-1}$	-4.07337829	-3.66975850	1.19335319×10
A_3	$2.63097499 \times 10^{-3}$	$4.44438465 \times 10^{-1}$	-2.62191122	$-7.64030325 \times 10^{-2}$	1.03465414	$6.86525190 \times 10^{-1}$
A_4	$-2.08966322 \times 10^{-1}$	-3.08676437	1.89803776×10	1.73765869	-	-
B_0	$8.62019757 \times 10^{-1}$	-6.75619620×10	4.03362767×10^2	4.83506219	1.48893790×10^2	-3.59215343×10
B_1	-2.39177733×10^2	1.48280710×10^3	-9.55873351×10^3	3.22546006×10^3	5.98218735×10^3	-1.33456721×10^4
B_2	5.77871418×10^3	-3.87126740×10^4	2.64619180×10^5	-7.33825641×10^4	-1.27514860×10^5	2.73758436×10^5
B_3	-3.64233761×10^4	2.46287167×10^5	-1.78873149×10^6	4.38286877×10^5	7.20710347×10^5	-1.46542108×10^6
B_4	4.67301183×10^4	-6.68827767×10^6	4.12063148×10^7	-1.29034950×10^6	-4.15256858×10^6	9.35966011×10^6
B_5	$7.63599392 \times 10^{-2}$	-2.58222581×10	$1.57286446e \times 10^2$	-3.07637547	-1.68667698×10	3.18892783×10

Table A.3: Table containing co-efficients of the three-dimensional polynomial descriptions of the ansätze for α and β . Terms in this table are up to second order in up to two parameters.

	χ_{lmm}^i				
	300	003	310	301	311
A_1	8.50016617	$5.28800867 \times 10^{-1}$	-5.58515402×10	6.75553207×10	-2.56408417×10^2
A_2	2.59287090	$1.01068701 \times 10^{-1}$	-3.14251436	3.56937882	-3.89849982
A_3	$-1.85917435 \times 10^{-1}$	$-1.88347274 \times 10^{-2}$	$5.59776161 \times 10^{-1}$	-3.39546523	4.94682122
A_4	-	$-3.95760269 \times 10^{-2}$	-1.30995276	1.36170318×10	-3.61072461×10
B_0	2.06283255×10^2	$1.86905353 \times 10^{-1}$	-3.66109299×10^2	5.17984009×10^2	-7.12420699×10^2
B_1	9.12415221×10^2	5.48552893×10	-4.94762928×10^3	-8.72529266×10^3	1.84557604×10^4
B_2	-3.37510784×10^4	-2.07225678×10^3	1.20115453×10^5	2.52531546×10^5	-5.31556200×10^5
B_3	2.67009593×10^5	1.77227221×10^4	-7.22222293×10^5	-1.78671167×10^6	3.72063365×10^6
B_4	-1.67215611×10^6	2.87874772×10^5	4.50452602×10^6	4.82669714×10^7	-7.99202851×10^7
B_5	-3.48537052×10	$2.58717680 \times 10^{-1}$	5.41079392×10	1.25431282×10^2	-2.96991961×10^2

Table A.4: Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are up to third order in one of the parameters.

	103	013	113	303	313
A_1	-1.05391807×10	-2.02762153	3.84438396×10	-8.08530688×10	3.38629060×10^2
A_2	-2.20772502	$3.40332041 \times 10^{-1}$	-5.78526769	-4.95219780×10	-1.21679151×10
A_3	$3.62088194 \times 10^{-1}$	$2.20593210 \times 10^{-2}$	$-4.35824891 \times 10^{-1}$	4.31073186	-5.630333300
A_4	$1.75903119 \times 10^{-1}$	$8.76230904 \times 10^{-2}$	$-4.47686868 \times 10^{-1}$	-	-
B_0	$-3.99481147 \times 10^{-1}$	$-9.44845879 \times 10^{-1}$	3.59712669	-	-
B_1	-4.84002368×10^2	-1.31840721×10^2	1.22200530×10^3	-	-
B_2	1.92089364×10^4	4.76345870×10^3	-4.55173492×10^4	-	-
B_3	-1.67761173×10^5	-3.97356491×10^4	3.84937394×10^5	-	-
B_4	-5.86843630×10^6	-5.09807946×10^5	1.02493128×10^7	-7.41737401×10^7	1.27336012×10^8
B_5	-8.05027134	-1.15356426	2.74518930×10	-1.23217145×10^2	3.67558344×10^2

Table A.5: Table containing co-efficients of the three-dimensional polynomial descriptions of the ansätze for α and β . Terms in this table are first or third order in each of the parameters.

	χ_{lmm}^i			
	302	312	203	213
A_1	-1.14765308	8.344542229×10	5.55215756×10	-2.08811138×10^2
A_2	6.47343086×10	-1.15428928	1.88775744×10	2.19710667×10
A_3	-2.00432953	-1.72512581	-2.21330072	2.77777363
A_4	7.87646332	-1.52679543×10	$1.14102095 \times 10^{-1}$	-
B_0	-1.92355950×10^2	-	-	-
B_1	-8.24870177×10^3	1.69394440×10^4	1.03039971×10^3	-2.73983879×10^3
B_2	1.55856499×10^5	-2.87822063×10^5	-4.30362417×10^4	1.04953516×10^5
B_3	-7.57276455×10^5	1.16423268×10^6	3.83773256×10^5	-8.99013339×10^5
B_4	9.08484160×10^6	-1.94937957×10^7	3.72507341×10^7	-6.44042482×10^7
B_5	4.55618396×10	-8.53804187×10	5.88901393×10	-1.83043318×10^2

Table A.6: Table containing co-efficients of the three-dimensional polynomial descriptions of the ansätze for α and β . Terms in this table are second or third order in two of the parameters.

	χ_{lmn}^i			
	004	104	014	114
A_1	$-7.37754924 \times 10^{-2}$	$4.80391818 \times 10^{-1}$	$-2.52063832 \times 10^{-2}$	-
A_2	$4.35571781 \times 10^{-1}$	-8.32909027	$-3.12146628 \times 10^{-2}$	-
A_3	$-1.30032985 \times 10^{-2}$	$2.85756568 \times 10^{-1}$	-	-
A_4	-	-	-	-
B_0	-	-	-	-
B_1	-1.43262966×10^2	1.51969817×10^3	3.49409199×10^2	-3.77177292×10^3
B_2	4.26635635×10^3	-4.49334022×10^4	-1.03654996×10^4	1.10710343×10^5
B_3	-3.16544257×10^4	3.31818656×10^5	7.61639390×10^4	-8.06983699×10^5
B_4	-1.56100479×10^4	1.03626462×10^5	2.29027760×10^4	-1.28688699×10^5
B_5	$2.61015904e \times 10^{-2}$	$-4.80330864 \times 10^{-3}$	$-1.48337653 \times 10^{-1}$	$6.35734496 \times 10^{-1}$

Table A.7: Table containing co-efficients of the three-dimensional polynomial descriptions of the ansätze for α and β . Terms in this table are fourth order in spin angle and linear in symmetric mass ratio.

	χ_{lmm}^i			
	204	214	304	314
A_1	-	-	-	-
A_2	5.25014699×10	-	-1.06684764×10^2	3.73661688
A_3	-1.92324332	-	3.97777025	-
A_4	-	-	-	-
B_0	-	-	-	-
B_1	-3.84079816×10^3	9.58470993×10^3	-	-
B_2	1.12833261×10^5	-2.79098443×10^5	-	-
B_3	-8.29794558×10^5	2.02272637×10^6	-	-
B_4	-	-	-	-
B_5	-	-	-	-

Table A.8: Table containing co-efficients of the three-dimensional polynomial descriptions of the co-efficients of the ansätze for α and β . Terms in this table are fourth order in spin angle and up to third order in symmetric mass ratio.

Bibliography

- [1] LALSimulation: Inspiral simulation packages. http://software.ligo.org/docs/lalsuite/lalsimulation/group__lalsimulation__inspiral.html.
- [2] SXS gravitational waveform database: Help and documentation. <https://data.black-holes.org/waveforms/documentation.html>.
- [3] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015.
- [4] B. P. Abbott et al. Binary Black Hole Mergers in the first Advanced LIGO Observing Run. *Phys. Rev.*, X6(4):041015, 2016.
- [5] B. P. Abbott et al. GW150914: The Advanced LIGO Detectors in the Era of First Discoveries. *Phys. Rev. Lett.*, 116(13):131103, 2016.
- [6] B. P. Abbott et al. GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence. *Phys. Rev. Lett.*, 116(24):241103, 2016.
- [7] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016.
- [8] B. P. Abbott et al. Properties of the Binary Black Hole Merger GW150914. *Phys. Rev. Lett.*, 116(24):241102, 2016.
- [9] B. P. Abbott et al. GW170608: Observation of a 19-solar-mass Binary Black Hole Coalescence. *Astrophys. J.*, 851(2):L35, 2017.
- [10] B. P. Abbott et al. GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence. *Phys. Rev. Lett.*, 119(14):141101, 2017.
- [11] Benjamin P. Abbott et al. Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93(11):112004, 2016. [Addendum: *Phys.Rev.D* 97, 059901 (2018)].
- [12] Benjamin P. Abbott et al. Effects of waveform model systematics on the interpretation of GW150914. *Class. Quant. Grav.*, 34(10):104002, 2017.

-
- [13] Benjamin P. Abbott et al. GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2. *Phys. Rev. Lett.*, 118(22):221101, 2017.
- [14] B.P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev. X*, 9(3):031040, 2019.
- [15] B.P. Abbott et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4M_{\odot}$. *Astrophys. J. Lett.*, 892:L3, 2020.
- [16] R. Abbott et al. GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses. 4 2020.
- [17] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015.
- [18] C. Affeldt et al. Advanced techniques in GEO 600. *Class. Quant. Grav.*, 31(22):224002, 2014.
- [19] P. Ajith et al. Inspiral-merger-ringdown waveforms for black-hole binaries with non-precessing spins. *Phys. Rev. Lett.*, 106:241101, 2011.
- [20] Parameswaran Ajith et al. Phenomenological template family for black-hole coalescence waveforms. *Class. Quant. Grav.*, 24:S689–S700, 2007.
- [21] Miguel Alcubierre, Bernd Bruegmann, Peter Diener, Michael Koppitz, Denis Pollney, Edward Seidel, and Ryoji Takahashi. Gauge conditions for long term numerical black hole evolutions without excision. *Phys. Rev. D*, 67:084023, 2003.
- [22] Bruce Allen, Warren G Anderson, Patrick R Brady, Duncan A Brown, and Jolien DE Creighton. Findchirp: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Physical Review D*, 85(12):122006, 2012.
- [23] Pau Amaro-Seoane et al. Laser Interferometer Space Antenna. 2 2017.
- [24] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries. *Phys. Rev. D*, 49:6274–6297, 1994.
- [25] K.G. Arun, Alessandra Buonanno, Guillaume Faye, and Evan Ochsner. Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms. *Phys. Rev. D*, 79:104023, 2009. [Erratum: *Phys.Rev.D* 84, 049901 (2011)].

- [26] KG Arun, Bala R Iyer, Bangalore S Sathyaprakash, and Pranesh A Sundararajan. Parameter estimation of inspiralling compact binaries using 3.5 post-newtonian gravitational wave phasing: The nonspinning case. *Physical Review D*, 71(8):084008, 2005.
- [27] Yoichi Aso, Yuta Michimura, Kentaro Somiya, Masaki Ando, Osamu Miyakawa, Takanori Sekiguchi, Daisuke Tatsumi, Hiroaki Yamamoto, KAGRA Collaboration, et al. Interferometer design of the kagra gravitational wave detector. *Physical Review D*, 88(4):043007, 2013.
- [28] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter. Gravitational wave extraction from an inspiraling configuration of merging black holes. *Phys. Rev. Lett.*, 96:111102, 2006.
- [29] Thomas W Baumgarte and Stuart L Shapiro. Numerical integration of einsteins field equations. *Physical Review D*, 59(2):024007, 1998.
- [30] Thomas W Baumgarte and Stuart L Shapiro. *Numerical relativity: solving Einstein's equations on the computer*. Cambridge University Press, 2010.
- [31] Carl M Bender and Steven A Orszag. *Advanced mathematical methods for scientists and engineers I: Asymptotic methods and perturbation theory*. Springer Science & Business Media, 2013.
- [32] Marsha J. Berger and Joseph Oliger. Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations. *J. Comput. Phys.*, 53:484, 1984.
- [33] Hans A. Bethe and G.E. Brown. Evolution of binary compact objects which merge. *Astrophys. J.*, 506:780–789, 1998.
- [34] Jonathan Blackman, Scott E. Field, Chad R. Galley, Bla Szilgyi, Mark A. Scheel, Manuel Tiglio, and Daniel A. Hemberger. Fast and Accurate Prediction of Numerical Relativity Waveforms from Binary Black Hole Coalescences Using Surrogate Models. *Phys. Rev. Lett.*, 115(12):121102, 2015.
- [35] Jonathan Blackman, Scott E Field, Mark A Scheel, Chad R Galley, Christian D Ott, Michael Boyle, Lawrence E Kidder, Harald P Pfeiffer, and Béla Szilágyi. Numerical relativity waveform surrogate model for generically precessing binary black hole mergers. *Physical Review D*, 96(2):024058, 2017.
- [36] Luc Blanchet. Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Rev. Rel.*, 17:2, 2014.
- [37] Luc Blanchet, Alessandra Buonanno, and Guillaume Faye. Tail-induced spin-orbit effect in the gravitational radiation of compact binaries. *Phys. Rev. D*, 84:064041, 2011.

-
- [38] Alejandro Bohe, Mark Hannam, Sascha Husa, Frank Ohme, Michael Pürrer, and Patricia Schmidt. LIGO Document T1500603-v4.
- [39] Alejandro Bohe, Sylvain Marsat, Guillaume Faye, and Luc Blanchet. Next-to-next-to-leading order spin-orbit effects in the near-zone metric and precession equations of compact binaries. *Class. Quant. Grav.*, 30:075017, 2013.
- [40] C. Bona, J. Masso, E. Seidel, and J. Stela. First order hyperbolic formalism for numerical relativity. *Phys. Rev. D*, 56:3405–3415, 1997.
- [41] Jeffrey M. Bowen and Jr. York, James W. Time asymmetric initial data for black holes and black hole collisions. *Phys. Rev. D*, 21:2047–2056, 1980.
- [42] Michael Boyle. Angular velocity of gravitational radiation from precessing binaries and the corotating frame. *Physical Review D*, 87(10):104006, 2013.
- [43] Michael Boyle, Duncan A. Brown, Lawrence E. Kidder, Abdul H. Mroue, Harald P. Pfeiffer, Mark A. Scheel, Gregory B. Cook, and Saul A. Teukolsky. High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions. *Phys. Rev.*, D76:124038, 2007.
- [44] Michael Boyle, Lawrence E. Kidder, Serguei Ossokine, and Harald P. Pfeiffer. Gravitational-wave modes from precessing black-hole binaries. 2014.
- [45] Michael Boyle and Abdul H Mroué. Extrapolating gravitational-wave data from numerical simulations. *Physical Review D*, 80(12):124045, 2009.
- [46] Michael Boyle, Robert Owen, and Harald P. Pfeiffer. A geometric approach to the precession of compact binaries. *Phys. Rev.*, D84:124011, 2011.
- [47] Dieter R. Brill and Richard W. Lindquist. Interaction energy in geometrostatics. *Phys. Rev.*, 131:471–476, 1963.
- [48] Bernd Bruegmann, Jose A. Gonzalez, Mark Hannam, Sascha Husa, and Ulrich Sperhake. Exploring black hole superkicks. *Phys. Rev.*, D77:124047, 2008.
- [49] Bernd Bruegmann, Jose A. Gonzalez, Mark Hannam, Sascha Husa, Ulrich Sperhake, and Wolfgang Tichy. Calibration of Moving Puncture Simulations. *Phys. Rev. D*, 77:024027, 2008.
- [50] A. Buonanno and T. Damour. Effective one-body approach to general relativistic two-body dynamics. *Phys. Rev. D*, 59:084006, 1999.
- [51] Alessandra Buonanno, Gregory B. Cook, and Frans Pretorius. Inspiral, merger and ring-down of equal-mass black-hole binaries. *Phys. Rev.*, D75:124018, 2007.

- [52] Alessandra Buonanno and Thibault Damour. Transition from inspiral to plunge in binary black hole coalescences. *Phys. Rev. D*, 62:064015, 2000.
- [53] Juan Calderón Bustillo, Alejandro Bohé, Sascha Husa, Alicia M Sintés, Mark Hannam, and Michael Pürrer. Comparison of subdominant gravitational wave harmonics between post-newtonian and numerical relativity calculations and construction of multi-mode hybrids. *arXiv preprint arXiv:1501.00918*, 2015.
- [54] Manuela Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Phys. Rev. Lett.*, 96:111101, 2006.
- [55] Manuela Campanelli, C.O. Lousto, and Y. Zlochower. Spinning-black-hole binaries: The orbital hang up. *Phys. Rev. D*, 74:041501, 2006.
- [56] Katerina Chatziioannou, Antoine Klein, Neil Cornish, and Nicolas Yunes. Analytic Gravitational Waveforms for Generic Precessing Binary Inspirals. *Phys. Rev. Lett.*, 118(5):051101, 2017.
- [57] Katerina Chatziioannou, Antoine Klein, Nicolás Yunes, and Neil Cornish. Gravitational waveforms for precessing, quasicircular compact binaries with multiple scale analysis: Small spin expansion. *Physical Review D*, 88(6):063011, 2013.
- [58] Katerina Chatziioannou, Antoine Klein, Nicols Yunes, and Neil Cornish. Constructing Gravitational Waves from Generic Spin-Precessing Compact Binary Inspirals. *Phys. Rev. D*, 95(10):104004, 2017.
- [59] E Ward Cheney and David R Kincaid. *Numerical mathematics and computing*. Cengage Learning, 2012.
- [60] Yvonne Choquet-Bruhat. The cauchy problem. In L Witten, editor, *Gravitation: An Introduction to Current Research*, chapter 13, page 168. John Wiley and Sons, New York, 1962.
- [61] D. Christodoulou. Reversible and irreversible transformations in black hole physics. *Phys. Rev. Lett.*, 25:1596–1597, 1970.
- [62] Tony Chu. *Numerical simulations of black-hole spacetimes*. California Institute of Technology, 2012.
- [63] Roberto Cotesta, Alessandra Buonanno, Alejandro Boh, Andrea Taracchini, Ian Hinder, and Serguei Ossokine. Enriching the Symphony of Gravitational Waves from Binary Black Holes by Tuning Higher Harmonics. *Phys. Rev. D*, 98(8):084028, 2018.

-
- [64] Curt Cutler and Eanna E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral wave form? *Phys. Rev.*, D49:2658–2697, 1994.
- [65] Curt Cutler and Kip S. Thorne. An Overview of gravitational wave sources. In *Proceedings, 16th International Conference on General Relativity and Gravitation: Durban, South Africa, July 15-21, 2001*, pages 72–111, 2013.
- [66] Thibault Damour. Coalescence of two spinning black holes: an effective one-body approach. *Phys. Rev. D*, 64:124013, 2001.
- [67] Thibault Damour, Bala R. Iyer, and B.S. Sathyaprakash. A Comparison of search templates for gravitational waves from binary inspiral. *Phys. Rev. D*, 63:044023, 2001. [Erratum: *Phys.Rev.D* 72, 029902 (2005)].
- [68] Thibault Damour, Bala R. Iyer, and B.S. Sathyaprakash. A Comparison of search templates for gravitational waves from binary inspiral - 3.5PN update. *Phys. Rev. D*, 66:027502, 2002.
- [69] KL Dooley, JR Leong, Thomas Adams, C Affeldt, A Bisht, C Bogan, J Degallaix, C Gräf, S Hild, J Hough, et al. Geo 600 and the geo-hf upgrade program: successes and challenges. *Classical and Quantum Gravity*, 33(7):075009, 2016.
- [70] Albert Einstein. The field equations of gravitation. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1915:844–847, 1915.
- [71] Albert Einstein. Fundamental ideas of the general theory of relativity and the application of this theory in astronomy. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 315, 1915.
- [72] Albert Einstein. On the general theory of relativity. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1915:799–801, 1915.
- [73] Albert Einstein. Approximative Integration of the Field Equations of Gravitation. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1916:688–696, 1916.
- [74] Albert Einstein. On gravitational waves. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1:154, 1918.
- [75] Stephen Fairhurst, Rhys Green, Mark Hannam, and Charlie Hoy. When will we observe binary black holes precessing? 8 2019.
- [76] Stephen Fairhurst, Rhys Green, Charlie Hoy, Mark Hannam, and Alistair Muir. The two-harmonic approximation for gravitational waveforms from precessing binaries. 8 2019.

- [77] Scott E. Field, Chad R. Galley, Jan S. Hesthaven, Jason Kaye, and Manuel Tiglio. Fast prediction and evaluation of gravitational waveforms using surrogate models. *Phys. Rev. X*, 4(3):031006, 2014.
- [78] Cecilio Garca-Quirs, Marta Colleoni, Sascha Husa, Hector Estells, Geraint Pratten, Antoni Ramos-Buades, Maite Mateu-Lucena, and Rafel Jaume. IMRPhenomXHM: A multi-mode frequency-domain model for the gravitational wave signal from non-precessing black-hole binaries. 1 2020.
- [79] Davide Gerosa, Michael Kesden, Ulrich Sperhake, Emanuele Berti, and Richard O’Shaughnessy. Multi-timescale analysis of phase transitions in precessing black-hole binaries. *Phys. Rev. D*, 92:064016, 2015.
- [80] Joshua N Goldberg, Alan J MacFarlane, Ezra T Newman, Fritz Rohrlich, and EC George Sudarshan. Spin-s spherical harmonics and δ . *Journal of Mathematical Physics*, 8(11):2155–2161, 1967.
- [81] Carsten Gundlach and Jose M. Martin-Garcia. Hyperbolicity of second-order in space systems of evolution equations. *Class. Quant. Grav.*, 23:S387–S404, 2006.
- [82] Eleanor Hamilton and Mark Hannam. Inferring black-hole orbital dynamics from numerical-relativity gravitational waveforms. *Phys. Rev.*, D98(8):084018, 2018.
- [83] Mark Hannam. Modelling gravitational waves from precessing black-hole binaries: Progress, challenges and prospects. *Gen. Rel. Grav.*, 46:1767, 2014.
- [84] Mark Hannam, Sascha Husa, Frank Ohme, Doreen Muller, and Bernd Bruegmann. Simulations of black-hole binaries with unequal masses or nonprecessing spins: Accuracy, physical properties, and comparison with post-Newtonian results. *Phys. Rev. D*, 82:124008, 2010.
- [85] Mark Hannam, Patricia Schmidt, Alejandro Boh, Lela Haegel, Sascha Husa, Frank Ohme, Geraint Pratten, and Michael Pürrer. Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms. *Phys. Rev. Lett.*, 113(15):151101, 2014.
- [86] James B Hartle. *Gravity: An introduction to Einsteins general relativity*. American Association of Physics Teachers, 2003.
- [87] James Healy, Janna Levin, and Deirdre Shoemaker. Zoom-whirl orbits in black hole binaries. *Physical review letters*, 103(13):131101, 2009.
- [88] James Healy, Carlos O Lousto, Yosef Zlochower, and Manuela Campanelli. The rit binary black hole simulations catalog. *Classical and Quantum Gravity*, 34(22):224001, 2017.

-
- [89] Daniel A. Hemberger, Mark A. Scheel, Lawrence E. Kidder, Bla Szilgyi, Geoffrey Lovelace, Nicholas W. Taylor, and Saul A. Teukolsky. Dynamical Excision Boundaries in Spectral Evolutions of Binary Black Hole Spacetimes. *Class. Quant. Grav.*, 30:115001, 2013.
- [90] Frank Herrmann, Ian Hinder, Deirdre Shoemaker, and Pablo Laguna. Unequal mass binary black hole plunges and gravitational recoil. *Classical and Quantum Gravity*, 24(12):S33, 2007.
- [91] Ian Hinder et al. Error-analysis and comparison to analytical models of numerical waveforms produced by the NRAR Collaboration. *Class. Quant. Grav.*, 31:025012, 2014.
- [92] Sascha Husa, Jose A. Gonzalez, Mark Hannam, Bernd Bruegmann, and Ulrich Sperhake. Reducing phase error in long numerical binary black hole evolutions with sixth order finite differencing. *Class. Quant. Grav.*, 25:105006, 2008.
- [93] Sascha Husa, Mark Hannam, Jose A. Gonzalez, Ulrich Sperhake, and Bernd Bruegmann. Reducing eccentricity in black-hole binary evolutions with initial parameters from post-Newtonian inspiral. *Phys. Rev. D*, 77:044037, 2008.
- [94] Sascha Husa, Sebastian Khan, Mark Hannam, Michael Pürrer, Frank Ohme, Xisco Jimnez Forteza, and Alejandro Boh. Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal. *Phys. Rev. D*, 93(4):044006, 2016.
- [95] Karan Jani, James Healy, James A Clark, Lionel London, Pablo Laguna, and Deirdre Shoemaker. Georgia tech catalog of gravitational waveforms. *Classical and Quantum Gravity*, 33(20):204001, 2016.
- [96] Piotr Jaranowski and Gerhard Schafer. Third postNewtonian higher order ADM Hamilton dynamics for two-body point mass systems. *Phys. Rev.*, D57:7274–7291, 1998. [Erratum: *Phys. Rev.*D63,029902(2001)].
- [97] Xisco Jimnez-Forteza, David Keitel, Sascha Husa, Mark Hannam, Sebastian Khan, and Michael Pürrer. Hierarchical data-driven approach to fitting numerical relativity data for nonprecessing binary black holes with an application to final spin and radiated energy. *Phys. Rev. D*, 95(6):064024, 2017.
- [98] Chinmay Kalaghatgi. *Modelling and Studying Gravitational Waves From Binary-Black-Hole Mergers*. PhD thesis, School of Physics and Astronomy, Cardiff, United Kingdom, 2019.
- [99] Vassiliki Kalogera. Spin orbit misalignment in close binaries with two compact objects. *Astrophys. J.*, 541:319–328, 2000.

- [100] Marc Kamionkowski and Andrew H. Jaffe. Detection of gravitational waves from inflation. *ICTP Lect. Notes Ser.*, 3:415–429, 2001.
- [101] M. Kerr. From Binaries to Backgrounds: Pulsar Timing and General Relativity. In *2nd LeCosPA Symposium: Everything about Gravity, Celebrating the Centenary of Einstein’s General Relativity*, pages 19–28, 2017.
- [102] Sebastian Khan, Katerina Chatziioannou, Mark Hannam, and Frank Ohme. Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects. *Phys. Rev. D*, 100(2):024059, 2019.
- [103] Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürrer, Xisco Jimnez Forteza, and Alejandro Boh. Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. *Phys. Rev. D*, 93(4):044007, 2016.
- [104] Sebastian Khan, Frank Ohme, Katerina Chatziioannou, and Mark Hannam. Including higher order multipoles in gravitational-wave models for precessing binary black holes. *Phys. Rev. D*, 101(2):024056, 2020.
- [105] Lawrence E. Kidder. Coalescing binary systems of compact objects to post-Newtonian 5/2 order. 5. Spin effects. *Phys. Rev. D*, 52:821–847, 1995.
- [106] Antoine Klein, Neil Cornish, and Nicolás Yunes. Gravitational waveforms for precessing, quasicircular binaries via multiple scale analysis and uniform asymptotics: The near spin alignment case. *Physical Review D*, 88(12):124015, 2013.
- [107] Antoine Klein, Neil Cornish, and Nicols Yunes. Fast Frequency-domain Waveforms for Spin-Precessing Binary Inspirals. *Phys. Rev. D*, 90:124029, 2014.
- [108] LIGO-India. LIGO-India, Proposal of the Consortium for Indian Initiative in Gravitational-wave Observations (IndIGO). <https://dcc.ligo.org/LIGO-M1100296/public>.
- [109] KAGRA Collaboration LIGO Scientific Collaboration, Virgo Collaboration. Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO, Advanced Virgo and KAGRA. <https://dcc.ligo.org/LIGO-P1200087/public>.
- [110] Lee Lindblom, Mark A. Scheel, Lawrence E. Kidder, Robert Owen, and Oliver Rinne. A New generalized harmonic evolution system. *Class. Quant. Grav.*, 23:S447–S462, 2006.
- [111] L. London and E. Fauchon-Jones. On modeling for Kerr black holes: Basis learning, QNM frequencies, and spherical-spheroidal mixing coefficients. *Class. Quant. Grav.*, 36(23):235015, 2019.

-
- [112] Lionel London, Sebastian Khan, Edward Fauchon-Jones, Cecilio Garca, Mark Hannam, Sascha Husa, Xisco Jimnez-Forteza, Chinmay Kalaghatgi, Frank Ohme, and Francesco Pannarale. First higher-multipole model of gravitational waves from spinning and coalescing black-hole binaries. *Phys. Rev. Lett.*, 120(16):161102, 2018.
- [113] Carlos O. Lousto and Yosef Zlochower. Black hole binary remnant mass and spin: A new phenomenological formula. *Phys. Rev.*, D89(10):104052, 2014.
- [114] LSC. LIGO Document T1800044-v5. <https://dcc.ligo.org/LIGO-T1800044/public>.
- [115] Harald Luck et al. The upgrade of GEO600. *J. Phys. Conf. Ser.*, 228:012012, 2010.
- [116] Michele Maggiore. *Gravitational waves: Volume 1: Theory and experiments*, volume 1. Oxford university press, 2008.
- [117] Ilya Mandel, Duncan A. Brown, Jonathan R. Gair, and M.Coleman Miller. Rates and Characteristics of Intermediate-Mass-Ratio Inspirals Detectable by Advanced LIGO. *Astrophys. J.*, 681:1431–1447, 2008.
- [118] Abdul H Mroué, Mark A Scheel, Béla Szilágyi, Harald P Pfeiffer, Michael Boyle, Daniel A Hemberger, Lawrence E Kidder, Geoffrey Lovelace, Serguei Ossokine, Nicholas W Taylor, et al. Catalog of 174 binary black hole simulations for gravitational wave astronomy. *Physical Review Letters*, 111(24):241104, 2013.
- [119] Evan Ochsner and Richard O’Shaughnessy. Asymptotic frame selection for binary black hole spacetimes II: Post-Newtonian limit. *Phys. Rev.*, D86:104037, 2012.
- [120] Frank Ohme, Mark Hannam, and Sascha Husa. Reliability of complete gravitational waveform models for compact binary coalescences. *Phys. Rev. D*, 84:064029, 2011.
- [121] Alan V. Oppenheim, Ronald W. Schafer, and John R. Buck. *Discrete-Time Signal Processing (2nd Ed.)*. Prentice-Hall, Inc., USA, 1999.
- [122] R. O’Shaughnessy, L. London, J. Healy, and D. Shoemaker. Precession during merger: Strong polarization changes are observationally accessible features of strong-field gravity during binary black hole merger. *Phys. Rev. D*, 87(4):044038, 2013.
- [123] R. O’Shaughnessy, B. Vaishnav, J. Healy, Z. Meeks, and D. Shoemaker. Efficient asymptotic frame selection for binary black hole spacetimes using asymptotic radiation. *Phys. Rev.*, D84:124002, 2011.

- [124] Serguei Ossokine et al. Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation. 4 2020.
- [125] Yi Pan, Alessandra Buonanno, Andrea Taracchini, Lawrence E. Kidder, Abdul H. Mrou, Harald P. Pfeiffer, Mark A. Scheel, and Bla Szilgyi. Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism. *Phys. Rev.*, D89(8):084006, 2014.
- [126] L. Pekowsky, R. O’Shaughnessy, J. Healy, and D. Shoemaker. Comparing gravitational waves from nonprecessing and precessing black hole binaries in the corotating frame. *Phys. Rev. D*, 88(2):024040, 2013.
- [127] Delphine Perrodin et al. Pulsar observations with European telescopes for testing gravity and detecting gravitational waves. In *14th Marcel Grossmann Meeting on Recent Developments in Theoretical and Experimental General Relativity, Astrophysics, and Relativistic Field Theories*, volume 2, pages 1843–1849, 2017.
- [128] 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. *Classical and Quantum Gravity*, 24(12):S59, 2007.
- [129] Simon F. Portegies Zwart and Steve L.W. McMillan. The Runaway growth of intermediate-mass black holes in dense star clusters. *Astrophys. J.*, 576:899–907, 2002.
- [130] Simon F. Portegies Zwart and Lev R. Yungelson. Formation and evolution of binary neutron stars. *Astron. Astrophys.*, 332:173–188, 1998.
- [131] Geraint Pratten et al. Let’s twist again: computationally efficient models for the dominant and sub-dominant harmonic modes of precessing binary black holes. 4 2020.
- [132] Geraint Pratten, Sascha Husa, Cecilio Garcia-Quiros, Marta Colleoni, Antoni Ramos-Buades, Hector Estelles, and Rafel Jaume. Setting the cornerstone for the IMRPhenomX family of models for gravitational waves from compact binaries: The dominant harmonic for non-precessing quasi-circular black holes. 1 2020.
- [133] William H Press, Saul A Teukolsky, William T Vetterling, and Brian P Flannery. Numerical recipes in c, 1988.
- [134] Frans Pretorius. Evolution of binary black hole spacetimes. *Phys. Rev. Lett.*, 95:121101, 2005.
- [135] Frans Pretorius. Numerical relativity using a generalized harmonic decomposition. *Class. Quant. Grav.*, 22:425–452, 2005.

-
- [136] Michael Pürrer, Sascha Husa, and Mark Hannam. An Efficient iterative method to reduce eccentricity in numerical-relativity simulations of compact binary inspiral. *Phys. Rev. D*, 85:124051, 2012.
- [137] Wenzer Qin, Kimberly K. Boddy, Marc Kamionkowski, and Liang Dai. Pulsar-timing arrays, astrometry, and gravitational waves. *Phys. Rev. D*, 99(6):063002, 2019.
- [138] Etienne Racine. Analysis of spin precession in binary black hole systems including quadrupole-monopole interaction. *Phys. Rev. D*, 78:044021, 2008.
- [139] Antoni Ramos-Buades, Patricia Schmidt, Geraint Pratten, and Sascha Husa. Validity of common modeling approximations for precessing binary black holes with higher-order modes. *Phys. Rev. D*, 101(10):103014, 2020.
- [140] Christian Reisswig and Denis Pollney. Notes on the integration of numerical relativity waveforms. *Class. Quant. Grav.*, 28:195015, 2011.
- [141] Milton Ruiz, Miguel Alcubierre, Darío Núñez, and Ryoji Takahashi. Multiple expansions for energy and momenta carried by gravitational waves. *General Relativity and Gravitation*, 40:1705–1729, 2008.
- [142] L. Santamaria et al. Matching post-Newtonian and numerical relativity waveforms: systematic errors and a new phenomenological model for non-precessing black hole binaries. *Phys. Rev.*, D82:064016, 2010.
- [143] Olivier Sarbach, Gioel Calabrese, Jorge Pullin, and Manuel Tiglio. Hyperbolicity of the BSSN system of Einstein evolution equations. *Phys. Rev. D*, 66:064002, 2002.
- [144] Bangalore Suryanarayana Sathyaprakash and Bernard F Schutz. Physics, astrophysics and cosmology with gravitational waves. *Living reviews in relativity*, 12(1):2, 2009.
- [145] B.S. Sathyaprakash and B.F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Rev. Rel.*, 12:2, 2009.
- [146] Peter R. Saulson. Gravitational wave detection: Principles and practice. *Comptes Rendus Physique*, 14:288–305, 2013.
- [147] Mark A. Scheel, Michael Boyle, Tony Chu, Lawrence E. Kidder, Keith D. Matthews, and Harald P. Pfeiffer. High-accuracy waveforms for binary black hole inspiral, merger, and ringdown. *Phys. Rev.*, D79:024003, 2009.
- [148] 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. *Physical Review D*, 74(10):104006, 2006.
-

- [149] Patricia Schmidt, Mark Hannam, and Sascha Husa. Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and non-precessing inspiral signals. *Phys. Rev.*, D86:104063, 2012.
- [150] Patricia Schmidt, Mark Hannam, Sascha Husa, and P. Ajith. Tracking the precession of compact binaries from their gravitational-wave signal. *Phys. Rev.*, D84:024046, 2011.
- [151] Patricia Schmidt, Ian W Harry, and Harald P Pfeiffer. Numerical relativity injection infrastructure. *arXiv preprint arXiv:1703.01076*, 2017.
- [152] Patricia Schmidt, Frank Ohme, and Mark Hannam. Towards models of gravitational waveforms from generic binaries: Ii. modelling precession effects with a single effective precession parameter. *Physical Review D*, 91(2):024043, 2015.
- [153] Masaru Shibata and Takashi Nakamura. Evolution of three-dimensional gravitational waves: Harmonic slicing case. *Physical Review D*, 52(10):5428, 1995.
- [154] Steinn Sigurdsson and Lars Hernquist. Primordial black holes in globular clusters. *Nature*, 364:423–425, 1993.
- [155] Kentaro Somiya. Detector configuration of KAGRA: The Japanese cryogenic gravitational-wave detector. *Class. Quant. Grav.*, 29:124007, 2012.
- [156] László B. Szabados. Quasi-Local Energy-Momentum and Angular Momentum in General Relativity. *Living Rev. Rel.*, 12:4, 2009.
- [157] Bla Szilgyi. Key Elements of Robustness in Binary Black Hole Evolutions using Spectral Methods. *Int. J. Mod. Phys.*, D23(7):1430014, 2014.
- [158] Andrea Taracchini et al. Effective-one-body model for black-hole binaries with generic mass ratios and spins. *Phys. Rev. D*, 89(6):061502, 2014.
- [159] Kip S Thorne. Multipole expansions of gravitational radiation. *Reviews of Modern Physics*, 52(2):299, 1980.
- [160] AV Tutukov and LR Yungelson. The merger rate of neutron star and black hole binaries. *Monthly Notices of the Royal Astronomical Society*, 260(3):675–678, 1993.
- [161] Birjoo Vaishnav, Ian Hinder, Frank Herrmann, and Deirdre Shoemaker. Matched filtering of numerical relativity templates of spinning binary black holes. *Physical Review D*, 76(8):084020, 2007.
- [162] Vijay Varma, Scott E. Field, Mark A. Scheel, Jonathan Blackman, Davide Gerosa, Leo C. Stein, Lawrence E. Kidder, and Harald P. Pfeiffer. Surrogate

- models for precessing binary black hole simulations with unequal masses. *Phys. Rev. Research.*, 1:033015, 2019.
- [163] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, CJ Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [164] Eugene Wigner. *Group theory: and its application to the quantum mechanics of atomic spectra*, volume 5. Academic Press, Inc., 1959.
- [165] James W York Jr. Kinematics and dynamics of general relativity. *Sources of gravitational radiation*, pages 83–126, 1979.
- [166] Yosef Zlochower, John G Baker, Manuela Campanelli, and Carlos O Lousto. Accurate black hole evolutions by fourth-order numerical relativity. *Physical Review D*, 72(2):024021, 2005.