



UNIVERSITY OF
BIRMINGHAM

BLACK-HOLE RINGDOWN:

Quasinormal Modes in Numerical-relativity Simulations and
Gravitational-wave Observations

by

ELIOT FINCH

A thesis submitted to the University of Birmingham for the degree of
DOCTOR OF PHILOSOPHY

Institute for Gravitational Wave Astronomy
School of Physics & Astronomy
College of Engineering and Physical Sciences
University of Birmingham

April 2023

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

UNIVERSITY OF BIRMINGHAM

*Abstract*College of Engineering and Physical Sciences
School of Physics & Astronomy

Doctor of Philosophy

**Black-hole Ringdown: Quasinormal Modes in Numerical-relativity Simulations
and Gravitational-wave Observations**

by Eliot Finch

The study of black-hole ringdown, the endpoint of a binary black-hole merger, is a cornerstone of gravitational-wave astronomy. A black hole's characteristic quasinormal-mode spectrum offers the possibility of clean and powerful tests of general relativity, encoding the properties of the remnant black hole (and, in principle, the binary from which it was formed) in the superposition of exponentially-damped sinusoids which constitute the ringdown signal. The observation of gravitational-wave signals is now routine, and the burgeoning field of black-hole spectroscopy is already contributing to our knowledge of the gravitational-wave sky.

In this thesis we investigate the ringdown in two ways. Firstly, we develop our ringdown models by performing fits directly to state-of-the-art numerical-relativity simulations. Our analysis considers, for the first time, a large selection of black-hole binaries with misaligned spins. In this clean and noise-free regime we can understand which quasinormal modes are important to include in our models, and it also informs us on the choice of ringdown start time. These studies help to guide our expectations when we analyse real gravitational-wave data. Secondly, we develop a novel frequency-domain ringdown analysis method for use with real data, applying it to both simulated signals and to real gravitational-wave data. Our method solves many of the problems associated with conventional ringdown-only analyses. Applying the method to GW150914, we quantify the significance with which a ringdown overtone can be found in the data.

Acknowledgements

First and foremost I would like to thank my supervisor, Chris Moore, for all the guidance and support. I'll always be grateful for the time you have invested in me, and for your patience and understanding. I would also like to thank Alberto Vecchio for introducing me to gravitational-wave astronomy, and for being a caring head of a brilliant research group. It's been a pleasure to be a part of it.

A special thanks goes to all the PhD students (Aysha, Diganta, Tom, Alice, Riccardo, Cressida, Andris, Daria, Alex, Matt, Paige, Dan, Evan, Xinyue, Lucy and Natalie) for the help, friendship, pubs, parties and board games. You have helped make a very challenging endeavour a very enjoyable one, too.

I'm also grateful for the administrative staff in the group, with particular thanks to Jo Cox and David Stops for all the help.

I have been able to be a part of the LIGO Collaboration during my PhD, and I am thankful for the Testing-GR group for comments on my paper drafts. I am especially thankful to Max Isi and Will Farr for their invitation to the ringdown workshop in New York.

Thank you to Aaron and Marc for all the laughs, and for all the years of friendship.

And finally thank you to my mum and brother, who have only ever been supportive throughout my life; this is for you.

Contents

Abstract	iii
Acknowledgements	v
1 Introduction	1
1.1 Binary black-hole mergers	1
1.2 Black-hole ringdown	11
1.3 Quasinormal modes	14
1.3.1 Scalar field on a Schwarzschild background	16
1.3.2 Quasinormal modes from the geodesic correspondence	23
1.3.3 Limitations of quasinormal modes	29
1.4 Black-hole spectroscopy	31
2 Modelling the Ringdown from Precessing Black-hole Binaries	39
2.1 Introduction	39
2.2 Model for the spherical modes	42
2.3 Fitting implementation	46
2.3.1 Single-mode fit	46
2.3.2 Multimode fit	48
2.4 Aligned-spin systems	49

2.4.1	Mirror modes	56
2.5	Misaligned-spin systems	58
2.5.1	Mirror modes	66
2.5.2	Higher harmonics	69
2.6	Surrogates	71
2.7	Numerical relativity errors	73
2.8	Conclusions	75
3	Frequency-domain Analysis of Black-hole Ringdowns	77
3.1	Introduction	77
3.2	Methods	81
3.2.1	Time- and frequency-domain likelihoods	81
3.2.2	Marginalising over the inspiral-merger	84
3.2.3	Model	86
3.3	Injection study	89
3.3.1	Determining the ringdown start time	96
3.3.2	Detecting additional QNMs	100
3.4	GW150914-like injection	103
3.5	Conclusions	107
4	Searching for a Ringdown Overtone in GW150914	109
4.1	Introduction	109
4.2	Methods	113
4.2.1	Wavelet-ringdown model	113
4.2.2	Data and priors	115
4.2.3	Reweighting	116

4.3	Results	119
4.3.1	Presence of an overtone	120
4.3.2	The nature of the overtone	126
4.3.3	The effect of noise and sampling frequency	131
4.3.4	Posteriors on the ringdown start time	133
4.3.5	Injection study	134
4.3.6	Wavelet posteriors	135
4.3.7	Other results	137
4.4	Conclusions	139
5	Conclusions and Prospects	143

List of Figures

1.1	Schematic diagram of an interferometer	2
1.2	Time-domain gravitational-wave signal from a binary black-hole merger	8
1.3	Frequency-domain gravitational-wave signal from a binary black-hole merger	9
1.4	Gravitational-wave ringdown waveform	12
1.5	Regge-Wheeler potential	19
1.6	Kerr quasinormal-mode spectrum	32
1.7	Black-hole spectroscopy illustration	35
2.1	Mismatch of the overtone model fitted to SXS:BBH:0305	51
2.2	Recovery of SXS:BBH:0305 remnant properties using the overtone model	52
2.3	Remnant-property errors and mode amplitudes for different numbers of overtones fitted to SXS:BBH:0305	53
2.4	Remnant-property errors and mismatches for the overtone model fitted to aligned-spin SXS simulations	56
2.5	Differences in the times of peak strain amplitude and peak gravitational- wave luminosity for aligned- and misaligned-spin SXS simulations	62
2.6	Mismatch for the overtone model fitted to SXS:BBH:1856	63
2.7	Recovery of SXS:BBH:1856 remnant properties using the overtone model	64
2.8	Remnant-property errors and mismatches for the overtone model fitted to misaligned-spin SXS simulations	65

2.9	Mismatch for the mirror-mode model fitted to SXS:BBH:1856	67
2.10	Recovery of SXS:BBH:1856 remnant properties using the mirror-mode model	68
2.11	Remnant-property errors and mismatches for the harmonic model fitted to misaligned-spin SXS simulations	70
2.12	Comparison of remnant-property errors from ringdown fits to two surrogate models and a selection of SXS simulations	72
3.1	Illustration of the frequency-domain ringdown analysis method	80
3.2	Posteriors on the remnant mass and spin for a GW190521-like injection analysed with a fixed sky location and ringdown start time, for different numbers of wavelets	91
3.3	Posteriors on the remnant mass and spin for a GW190521-like injection analysed while marginalising over the sky location and ringdown start time, for different numbers of wavelets	93
3.4	Posteriors on the ringdown start time for the same analyses as in Fig. 3.3 .	94
3.5	Maximum-likelihood frequency-domain waveform from the analysis of Fig. 3.3	95
3.6	Posteriors on the remnant mass and spin for a GW190521-like injection analysed with different choices of ringdown start-time prior	97
3.7	Posteriors on the ringdown start time for the same analyses as in Fig. 3.6 .	98
3.8	Posteriors on the ringdown start time in the geocentric frame for a GW190521- like injection analysed with two- and three- detector networks	99
3.9	Posteriors on the remnant mass and spin and on the quasinormal-mode amplitudes for a GW190521-like injection analysed with and without an overtone	101

3.10	Posteriors on the overtone amplitude for GW190521-like injections at a selection of signal-to-noise ratios, for both time- and frequency-domain analyses	102
3.11	Posteriors on the remnant mass and spin for a GW150914-like injection analysed with a fixed sky location and ringdown start time, for different numbers of wavelets	105
3.12	Whitened waveform reconstructions for the same analyses as in Fig. 3.11	106
4.1	Priors on the ringdown start time used for the GW150914 analysis	118
4.2	Posteriors on the GW150914 remnant mass and spin for each choice of ringdown start-time prior	121
4.3	Posteriors on the GW150914 overtone amplitude and Bayes' factors in favour of the overtone model for each choice of ringdown start-time prior	122
4.4	Posteriors on the GW150914 overtone amplitude, as in Fig. 4.3, rescaled to a fixed reference time	125
4.5	Posteriors on the deviation from Kerr for the real part of the GW150914 overtone frequency for each choice of the ringdown start-time prior	127
4.6	Posteriors on the deviation from Kerr for the real part of the GW150914 overtone frequency and ringdown start time for each choice of the ringdown start-time prior	128
4.7	Posterior on the complex frequency of the second GW150914 quasinormal mode assuming the first is associated with the fundamental mode	129
4.8	Posteriors on the GW150914 overtone amplitude divided by the fundamental-mode amplitude for each choice of the ringdown start-time prior	130
4.9	Posteriors on the GW150914 ringdown start time, using a flat prior	133
4.10	Posteriors on the overtone amplitude from a GW150914-like injection study	135

4.11 Posteriors on selected wavelet parameters from the GW150914 overtone analysis for the preferred choice of ringdown start-time prior	136
4.12 Posterior on the sky location from the GW150914 overtone analysis for the preferred choice of ringdown start-time prior	137
4.13 Whitened waveform reconstructions from the GW150914 overtone analysis for the preferred choice of ringdown start-time prior	138

Contributions Overview

This thesis contains the following published works (Refs. [1–3]):

- [1] Eliot Finch and Christopher J. Moore. “Modeling the ringdown from precessing black hole binaries”. Published in *Phys. Rev. D* 103, 084048 (2021).
- [2] Eliot Finch and Christopher J. Moore. “Frequency-domain analysis of black-hole ringdowns”. Published in *Phys. Rev. D* 104, 123034 (2021).
- [3] Eliot Finch and Christopher J. Moore. “Searching for a ringdown overtone in GW150914”. Published in *Phys. Rev. D* 106, 043005 (2022).

Supplementary material and data releases for the above works are available at Refs. [4–6].

Chapter contributions and synopses

Chapter 1 is an introductory chapter, containing standard results and calculations. I wrote this chapter with input from my supervisor, and I created all figures.

In Chapter 1 we provide wider context to the study of black-hole ringdown. This includes gravitational waves, their production via binary black-hole mergers, and the formation of an initially perturbed remnant black hole. We discuss the basic morphology of the ringdown signal, and provide details on its constituent quasinormal modes. This includes some

basic calculations to build intuition. Finally, we provide motivation for the study of the ringdown in the form of tests of general relativity.

Chapter 2 is a reformatted version of Ref. [1]. The introduction has been re-worked to fit into the thesis format, and some minor edits have been made to the text to include new references and improve readability. The study was conceived with help from my supervisor and co-author, Christopher J. Moore. I performed all analyses, wrote the code (available at Ref. [7]), and created all figures. I wrote the majority of the text, with input from my supervisor.

The study of Chapter 2 was prompted by numerical work of Giesler et al., who found that the linear ringdown description could be extended to relatively early times if multiple quasinormal-mode overtones were included in the model. We generalised this work (which considered only aligned-spin systems) to precessing systems, testing the validity of the model across a wider parameter space. We find a much greater variation in the performance of ringdown fits than in the aligned-spin case. The inclusion of mirror modes and higher harmonics, along with overtones, improves the reliability of ringdown fits with an early start time (perhaps a sign that mode mixing in the ringdown is generally more important in precessing systems). However, there remain cases with poor performing fits. This suggests that it is not possible to reliably model the ringdown from as early a time as claimed in the work of Giesler et al.

Chapter 3 is a reformatted version of Ref. [2]. The introduction has been re-worked to fit into the thesis format, and some minor edits have been made to the text to include

new references and improve readability. The idea for the data-analysis method presented in this chapter was proposed by my supervisor, and jointly developed to its final form. I performed all analyses, wrote the code (available at Ref. [8]), and created all figures. I wrote the majority of the text, with input from my supervisor.

In Chapter 3 we propose a novel frequency-domain ringdown analysis method, and demonstrate its use with simulated gravitational-wave signals. Our approach avoids the issues of spectral leakage that would normally be expected (associated with the abrupt start of the ringdown) by modelling the inspiral and merger parts of the signal using a flexible sum of sine-Gaussian wavelets truncated at the onset of the ringdown; these effectively marginalise out the inspiral and merger. Performing the analysis in the frequency domain allows us to use standard (and by now well-established) Bayesian inference pipelines for gravitational-wave data as well as giving us the ability to readily search over the sky position and the ringdown start time, although we find that it is necessary to use an informative prior for the latter. We test our method by using it to analyse several simulated signals with varying signal-to-noise ratios injected into two- and three-detector networks. We find that our frequency-domain approach is generally able to place tighter constraints on the remnant black-hole mass and spin than a standard time-domain analysis.

Chapter 4 is a reformatted version of Ref. [3]. The introduction has been re-worked to fit into the thesis format, and some minor edits have been made to the text to improve readability. I conceived the study, and performed all analyses with input from my supervisor. With the exception of Figs. 4.7 and 4.12 (created by my supervisor) I created all figures, and wrote the majority of the text.

Prompted by disagreements in the literature, in Chapter 4 we reanalyse the GW150914 data searching for a quasinormal-mode overtone using our frequency-domain approach. Our analysis has several advantages compared to other analyses; in particular, the source sky position and the ringdown start time are marginalised over (as opposed to simply being fixed) as part of a Bayesian ringdown analysis. We find tentative evidence for an overtone in GW150914, but at a lower significance than reported elsewhere. Our preferred analysis, marginalising over the uncertainty in the time of peak strain amplitude, gives a posterior on the overtone amplitude peaked away from zero at $\sim 1.8\sigma$.

Chapter 1

Introduction

1.1 Binary black-hole mergers

At the time of writing the LIGO [9] and Virgo [10] detectors have completed three observing runs, accumulating 90 confident gravitational-wave (GW) signal candidates [11–14]. The signals originate from the mergers of compact objects. The majority of these are binary black holes (BBHs), and these are the systems of interest throughout this thesis. Signals have also been observed from at least one binary neutron star system (GW170817 [15] remains the only unambiguous candidate) and from neutron-star – black-hole (BH) systems (the most confident candidate being GW200115_042309 [16]). GW150914 [17] marked the first direct observation of GWs, as well as the first observation of a BBH merger.

A prediction of general relativity (GR), GWs are produced by accelerating masses. As transverse waves which act to expand and contract space, they can be measured via their influence on freely-falling test masses (which will move relative to each other due to geodesic deviation). The LIGO and Virgo GW detectors employ interferometry to measure the motion of test masses to extreme precision. In this setup the mirrors at the end of the interferometer arms act as the freely-falling masses (they are hung as pendulums to isolate them from the environment), and changes in their position can be determined via




FIGURE 1.1: A simple representation of an interferometer, such as those used by LIGO and Virgo to detect GWs. Light from a laser is separated by a beam-splitter and travels down the two interferometer arms. Mirrors at the end of the arms (which are isolated from the environment and act as test masses) reflect the light, before it is recombined at the beam-splitter and sent to a photodetector.

measurements of the light travel time along each arm. The detector output is a time series of a dimensionless scalar quantity called the strain, h , which quantifies the fractional change of the interferometer arm lengths (which can also be thought of as changes in the position of the mirrors).

Being metric perturbations, GWs themselves are not scalar quantities but tensors (h is merely the measurement we make when GWs are “projected” onto the interferometer). GWs are described by two dimensionless amplitudes, h_+ and h_\times , which are the “plus” and “cross” polarisations respectively. These two polarisations, which represent small perturbations to a flat background metric, are encapsulated in the GW strain tensor, h_{ij} . To understand the response of an interferometer to incoming GWs, consider a setup

(depicted in Fig. 1.1) where the interferometer arms are aligned with the x - and y -axes such that the beam-splitter is at the origin and the mirrors are at the coordinates $(L_x, 0)$ and $(0, L_y)$. Incident GWs will influence the propagation of light in the interferometer arms. In the proper detector frame, the GW acts to displace the mirrors. This is an intuitive picture, and the response of the mirrors to the passing GWs can be described in terms of Newtonian forces. However, GWs take on a very simple form (and become easier to work with) when working in the transverse-traceless (TT) frame. In this description coordinates are defined by freely-falling masses, meaning the interferometer mirrors do not move under the influence of passing GWs. Instead, the propagation of the light in the interferometer arms is modified via the GWs influence on the space-time interval. In the simple case of a GW with only plus polarisation incoming from the positive z -direction (i.e. into the page in Fig. 1.1), the space-time interval becomes

$$ds^2 = -c^2 dt^2 + [1 + h_+(t)] dx^2 + [1 - h_+(t)] dy^2 + dz^2, \quad (1.1)$$

where $h_+(t)$ is the GW amplitude. Light (here, the laser in the interferometer arms) follows null geodesics, $ds^2 = 0$, so for the light propagation in the x -arm we can write

$$dx = \pm c dt [1 + h_+(t)]^{-1/2} \approx \pm c dt \left[1 - \frac{1}{2} h_+(t) \right] \quad (1.2)$$

for small h_+ (as is the case in practice, see Eq. 1.17). The plus sign corresponds to light travelling from the beam-splitter to the mirror in the positive x -direction; integrating along the length of the arm (where t_0 is the time the light leaves the beam-splitter and t_1

is the arrival time at the mirror) we have

$$\int_0^{L_x} dx = \int_{t_0}^{t_1} c dt \left[1 - \frac{1}{2} h_+(t) \right] \quad (1.3)$$

$$\implies L_x = c(t_1 - t_0) - \frac{c}{2} \int_{t_0}^{t_1} dt h_+(t). \quad (1.4)$$

Similarly, for the return trip we have

$$L_x = c(t_2 - t_1) - \frac{c}{2} \int_{t_1}^{t_2} dt h_+(t), \quad (1.5)$$

where t_2 is the arrival time back at the beam-splitter. Summing Eqs. 1.4 and 1.5 we obtain

$$T_x = t_2 - t_0 = \frac{2L_x}{c} + \frac{1}{2} \int_{t_0}^{t_2} dt h_+(t), \quad (1.6)$$

so the GW acts to modify the usual round-trip travel time of $2L_x/c$. If we set $h_+(t) = h$ (that is, if we consider a GW of constant amplitude), we can write

$$T_x = \frac{2L_x}{c} \left(1 - \frac{1}{2} h \right)^{-1} \approx \frac{2L_x}{c} \left(1 + \frac{1}{2} h \right). \quad (1.7)$$

Repeating the above for the y -arm, we obtain

$$T_y = \frac{2L_y}{c} \left(1 - \frac{1}{2} h \right). \quad (1.8)$$

Since the GW alters the round-trip time in the y -arm differently, light that left the photodetector initially in-phase will arrive back at the photodetector out of phase. Setting

$L_x = L_y = L$, we can write the difference in arrival time as

$$\Delta T = T_x - T_y = \frac{2L}{c}h \implies \Delta\phi = \omega\Delta T = \frac{2\omega L}{c}h \quad (1.9)$$

where ω is the light angular frequency. This difference in phase, $\Delta\phi$, is how the GW influences the power measured at the photodetector.

Einstein [18] showed the GW emission from slowly changing and weakly gravitating sources is quadrupolar to leading order, and far from the source the GW strain tensor can be written

$$h_{ij} = \frac{2G}{c^4 r} \ddot{Q}_{ij}. \quad (1.10)$$

Here, r is the distance to the source, and Q_{ij} is the traceless mass quadrupole of the source:

$$Q_{ij} = \int d^3\mathbf{x} \rho(\mathbf{x}) \left(x_i x_j - \frac{1}{3} |\mathbf{x}|^2 \delta_{ij} \right), \quad (1.11)$$

where $\rho(\mathbf{x})$ is the mass density and we have used Cartesian coordinates $\mathbf{x} = (x_1, x_2, x_3) = (x, y, z)$ with some suitable choice of origin. Note that the above is valid when h_{ij} is chosen to be in the TT gauge.

It was also shown that the GW luminosity is given by

$$\dot{E}_{\text{GW}} = \frac{G}{5c^5} \langle \ddot{Q}_{ij} \ddot{Q}^{ij} \rangle, \quad (1.12)$$

where the average (denoted by the angled brackets) is to be taken over \sim a few GW wavelengths.

To perform some simple order-of-magnitude calculations we can introduce a characteristic mass, length and timescale of the radiating system: M , R , and T respectively. Starting with the GW luminosity, from dimensional arguments we approximate the quadrupole

moment as

$$\ddot{Q} \sim \frac{MR^2}{T^3} \sim \frac{Mv^3}{R} \quad (1.13)$$

where v is the characteristic velocity of the system. Substituting into Eq. 1.12, we can express the order-of-magnitude GW luminosity as

$$\dot{E}_{\text{GW}} \sim \frac{G}{5c^5} \left(\frac{Mv^3}{R} \right)^2 \sim \frac{c^5}{G} \left(\frac{r_s}{R} \right)^2 \left(\frac{v}{c} \right)^6 \quad (1.14)$$

where $r_s = 2GM/c^2$ is the Schwarzschild radius (note that we have dropped all constants of order 1 in Eq. 1.14). From Eq. 1.14 we immediately see that to maximise energy radiated as GWs we need compact ($R \sim r_s$) and relativistic ($v \sim c$) systems. Of course, we also need a time-varying quadrupole (Eq. 1.10); the mergers of compact objects, then, make ideal systems. We can similarly obtain the order-of-magnitude GW strain. Assuming that our source is indeed a binary system, we can use Kepler's third law to relate the characteristic length scale and frequency $f = 1/T$:

$$R^3 \sim \frac{GM}{f^2}. \quad (1.15)$$

Our expression for the strain, Eq. 1.10, then becomes

$$h \sim \frac{G}{c^4 r} \frac{MR^2}{T^2} \sim \frac{(GM)^{5/3} f^{2/3}}{c^4 r}. \quad (1.16)$$

The first GW event, GW150914, occurred roughly 440 Mpc from Earth and had a total binary mass of $72 M_\odot$. And, as we will see, current GW detectors are most sensitive to frequencies ~ 100 Hz. Expressing the GW strain in terms of these numbers gives

$$h \sim 10^{-21} \left(\frac{M}{72 M_\odot} \right)^{5/3} \left(\frac{f}{100 \text{ Hz}} \right)^{2/3} \left(\frac{440 \text{ Mpc}}{r} \right), \quad (1.17)$$

emphasising the just how sensitive the interferometers need to be.

Going beyond these simple calculations, it is well-understood how the dynamics of compact-object mergers translate to the measured strain, which allows us to learn about the signal origin. The GW signal from a BBH merger can be split into three stages: the inspiral, merger, and ringdown. The inspiral corresponds to the two BHs in the binary orbiting each other. Unlike in the Newtonian case, the generation of GWs causes the system to lose energy (Eq. 1.12) and for the orbit to shrink and spiral. This stage in the binary evolution, where the motion of the orbit is slow compared to the speed of light, can be modelled with corrections to the Newtonian case in a post-Newtonian framework [19]. Eventually, the BHs come together in the highly energetic and dynamical merger; this process requires the full nonlinear equations of GR, and is the subject of numerical relativity (NR) [20]. Immediately after merging, the final BH is highly distorted. As it equilibrates it produces GWs which roughly take the form of a damped sinusoid. This is analogous to the ringing a bell makes after it has been struck, and this final stage is called the ringdown. Just as the early inspiral can be studied by considering deviations to a Newtonian system, the ringdown can be studied by considering perturbations to a static BH spacetime [21, 22]. It is not known beforehand when this perturbative treatment of the final BH becomes valid (i.e., when the ringdown actually starts), and so care must be taken; this will be a recurring theme throughout this work.

The GW signal from these three stages is shown in Fig. 1.2. Plotted is the (noiseless) GW strain signal projected onto the LIGO Hanford detector, h , multiplied by the distance to the source, r , as a function of time. We use geometric units ($G = c = 1$) to express these quantities in units of the total binary mass, M . In geometric units we can convert between lengths, masses, and times with appropriate combinations of G and c ; this means time/M and rh/M can be made dimensionless quantities. Using GW150914 again as an




FIGURE 1.2: The GW signal from a BBH merger. This waveform was generated using the NRSur7dq4 surrogate, and projected onto the Hanford detector using parameters consistent with GW150914 (this includes choosing a sky location and polarisation angle).

example, say we had a system with total mass $M = 72 M_{\odot} \approx 1.43 \times 10^{32} \text{ kg}$ at a distance $r = 440 \text{ Mpc} \approx 1.36 \times 10^{25} \text{ m}$. Expressing M in seconds gives us the conversion factor for time:

$$M \times \frac{G}{c^3} \approx 1.43 \times 10^{32} \text{ kg} \times \frac{6.67 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}}{(3 \times 10^8 \text{ m s}^{-1})^3} \approx 3.6 \times 10^{-4} \text{ s}. \quad (1.18)$$

So, for these system properties, in Fig. 1.2 we have a ringdown duration of $\sim 50 M \approx 50 \times 3.6 \times 10^{-4} \text{ s} \approx 0.018 \text{ s}$. Calculating the dimensionless quantity M/r gives us the conversion for the GW amplitude. One way to do this is to express r in seconds by dividing by c , then use our previous result for M :

$$\frac{M \times G/c^3}{r/c} \approx \frac{3.6 \times 10^{-4} \text{ s}}{1.36 \times 10^{25} \text{ m} / 3 \times 10^8 \text{ m s}^{-1}} \approx 7.85 \times 10^{-21}. \quad (1.19)$$

This gives a peak strain amplitude in Fig. 1.2 of $\sim 0.2 \times 7.85 \times 10^{-21} \approx 1.60 \times 10^{-21}$, in agreement with our order-of-magnitude calculation in Eq. 1.17. This can also be compared with what was actually measured for GW150914 (see Fig. 1 in Ref. [17]), and we see these numbers are consistent.

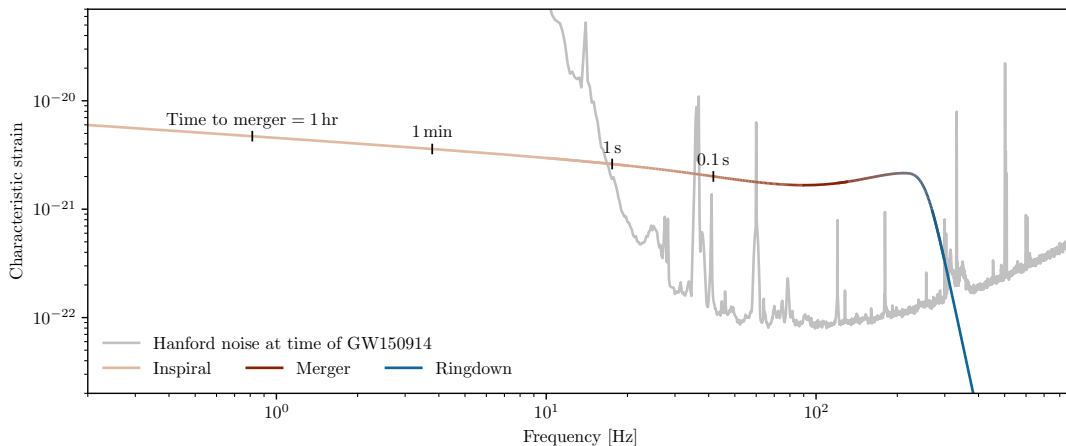


FIGURE 1.3: The frequency-domain GW signal from a BBH merger, shown with the Hanford sensitivity curve (both expressed through the characteristic strain). The signal was generated with IMRPHENOMD using parameters consistent with GW150914.

The waveform was generated using the NRSur7dq4 [23] surrogate waveform with zero spins, equal mass ratio, and zero inclination. Surrogates effectively “interpolate” the waveforms from full NR simulations, allowing quick waveform generation for any choice of system parameters (within the validity of the model). For NRSur7dq4 the simulations used to build the model came from the SXS catalog [24], which we use extensively in Chapter 2. Along with the intrinsic system parameters, to project onto the Hanford detector we also choose an event time, sky location and polarisation angle consistent with GW150914.

The above discussion disregards the sensitivity of the detector to different frequencies (i.e. the detector noise). It is standard to treat the noise in the LIGO and Virgo detectors as stationary and Gaussian [25]. Stationarity means that the noise covariance matrix is diagonal in the frequency domain (meaning there is no correlation between frequency bins), and so we can describe the noise with a power spectral density (PSD) $S_n(f)$. In each frequency bin we model the noise as having a random phase and an amplitude drawn from a Gaussian distribution with standard deviation $\sqrt{S_n(f)}$. We show the Hanford PSD in Fig. 1.3, along with a frequency-domain BBH GW signal (both expressed through the

characteristic strain, see Eqs. 1.20 and 1.21). Since we are now comparing with a PSD, we fix the BBH total mass and distance to the example numbers used above ($M = 72 M_{\odot}$, $r = 440$ Mpc). Roughly speaking, the GW signal from a BBH merger increases in frequency over time. The majority of the inspiral signal is at too low a frequency to be detected, with the GW signal only entering the detector band (~ 20 to ~ 1000 Hz) just before merger. Although the binary can spend billions of years inspiralling [26, 27], we observe only the last few tenths of a second of the merger. We choose to plot the signal and noise curve in terms of the characteristic strain [28], given by

$$h_c(f) = 2f \left| \tilde{h}(f) \right| \quad (1.20)$$

for the GW signal (where \tilde{h} is the frequency-domain GW in a given detector), and

$$h_n(f) = \sqrt{f S_n(f)} \quad (1.21)$$

for the noise curve. The characteristic strain has the property that the signal-to-noise ratio (SNR) ρ is given by

$$\rho^2 = \int_{-\infty}^{\infty} d \log f \left[\frac{h_c(f)}{h_n(f)} \right]^2. \quad (1.22)$$

Evaluating for the above figure we get $\rho \sim 17$, which is consistent with the Hanford SNR for GW150914 (assuming equal SNRs in Hanford and Livingston, a network SNR of 26 [13] gives a single detector SNR of ~ 18).

Also shown in the figure is the time to merger at select frequencies of the inspiral. This is to emphasise how the frequency of the binary evolves over time, spending the vast majority of its life in the inspiral with a slowly shrinking orbit. The leading-order time to

merger, τ_{merge} , as a function of GW frequency is given by

$$\tau_{\text{merge}} = \frac{5}{256} (\pi f)^{-\frac{8}{3}} \left(\frac{G\mathcal{M}}{c^3} \right)^{-\frac{5}{3}}, \quad (1.23)$$

where \mathcal{M} is a combination of the two masses in the binary known as the chirp mass,

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}. \quad (1.24)$$

For GW150914, we have $m_1 = 39 M_\odot$ and $m_2 = 33 M_\odot$. In Eq. 1.23 we also explicitly show the factors of G and c for clarity (in geometric units these would be set to 1, and it would be implied that the chirp mass should be expressed in units of time). This expression for the time to merger is valid in the quasi-circular orbit regime; it can be derived by equating the time derivative of the orbital energy of a binary system to (negative) the GW luminosity (Eq. 1.12), where the quadrupolar moment is most easily calculated using the equivalent one-body problem with a reduced mass.

The signal in Fig. 1.3 was generated with the IMRPHENOMD [29] waveform model, as implemented in the RIPPLE [30] Python package. Phenom models produce approximate waveforms using closed-form analytic expressions in the frequency-domain, making evaluation quick and suitable for GW searches and parameter estimation. The Hanford PSD is estimated from 1024s of off-source data at the time of GW150914 using a Welch periodogram [31].

1.2 Black-hole ringdown

The endpoint of a BBH, the ringdown signal is produced by the remnant BH settling to its stationary state. Just as a bell or drum has a characteristic sound, so too does a BH;




FIGURE 1.4: The ringdown waveform from a BBH, fitted with a simple damped sinusoid. At early times this simple description breaks down. The waveform is a NR simulation from the SXS catalog.

associated with the BH is a unique spectrum of oscillatory modes, determined purely by its properties (namely, for astrophysical BHs, a mass and spin). The characteristic oscillations of the remnant BH are called quasinormal modes (QNMs), so-called because, unlike normal modes, they decay over time (for reviews on the subject, see Refs. [32–35]). The QNM frequencies are complex, $\omega = 2\pi f - i/\tau$, with the real part f giving the oscillation frequency and the reciprocal of the imaginary part τ giving the damping time. The QNM spectrum is the subject of Sections 1.3 and 1.4. The ringdown signal consists of a sum of QNMs, each excited a different amount depending on the initial configuration of the

binary and how they merged; the excitation of different ringdown modes is among the areas of investigation in Chapter 2.

Fig. 1.4 shows (in blue) the ringdown portion of the waveform from a NR simulation, SXS:BBH:0305 [36]. This simulation has properties consistent with GW150914, and will be among the NR simulations studied in Chapter 2. NR waveforms are decomposed into spherical-harmonic modes indexed by ℓ and m (see Section 2.2 and Eq. 2.1 for more details), and here we plot the real part of the dominant $\ell = m = 2$ mode. The radial dependence of GW waveforms goes predominantly as r^{-1} , so the SXS catalog provides the strain multiplied by r (just as what was plotted in Fig. 1.2). For clarity, when dealing with NR simulations (as we do throughout Chapter 2), we will drop the r factor and have the r^{-1} scaling implied.

Overplotted is a simple model for the h_{22} mode: a exponentially damped sinusoid with complex frequency $\omega = 2\pi f - i/\tau$ and amplitude $C = Ae^{i\phi}$. Taking the real part we have

$$\text{Re}[Ce^{-i\omega(t-t_0)}] = A \cos[2\pi f(t - t_0) - \phi]e^{-(t-t_0)/\tau}. \quad (1.25)$$

That such a simple model describes the final stages of such a complicated system is a remarkable result, and is part of the power of studying the ringdown. A faded red line traces the damped sinusoid to earlier times, where it starts to lose validity. This is to be expected: at earlier times we approach the nonlinear and strong-gravity merger, where BH perturbation theory (which is where this simple damped-sinusoid description comes from) starts to break down. Note, however, that our model here only includes a single term (i.e., one ringdown QNM); with additional QNMs we may be able to get a better fit to the NR waveform, or to describe the waveform at earlier times (this idea is explored at length in Chapter 2). A particular subset of QNMs, known as “overtones”, will be of

particular interest throughout the thesis; as well as featuring in the numerical studies of Chapter 2, they will be a target in the analysis of GW data in Chapters 3 and 4.

A key goal in GW astronomy is the identification of QNMs in the ringdown signal. We will see in the following sections that QNMs carry information about the remnant BH, meaning that measurements of QNM frequencies give us a way of inferring the remnant BH properties independently of the rest of the GW signal. This forms the basis of important tests of GR, and is discussed further in Section 1.4. The testing-GR companion papers for the second [37] and third [38] GW event catalogs featured searches for QNMs; of the detected events, results were reported for 22. This can be taken as a rough guide for how many events have at least one measurable ringdown mode. On top of this, tentative evidence was reported for the identification of an additional ringdown mode in a few of the loudest events, including GW150914. However, as will be seen, the identification of subdominant ringdown modes is subtle. This is particularly true for one of the most massive BBH mergers observed so far, GW190521 [39], which has a total source frame mass of $\sim 150 M_{\odot}$. Its large mass means this event enters the detector band only very near merger, and so is ringdown dominated. This makes it a promising target for QNM searches, and along with GW150914 it is an event that will be referenced throughout this thesis.

1.3 Quasinormal modes

The characteristic vibrational modes of dissipative systems are known as QNMs. As stated above, these differ from usual normal modes because they decay over time. Although not limited to BHs (any real-world physical systems which are subject to damping will exhibit decaying modes), BH spacetimes are an interesting case because even idealised systems

are intrinsically dissipative.

Gravitational perturbations of the Schwarzschild geometry [40] were first studied by Regge and Wheeler [41], and this work was extended by Zerilli to a more general class of perturbations [42, 43]. Employing the perturbation techniques developed by Regge and Wheeler, Vishveshwara [44] performed numerical studies involving the scattering of GWs off a Schwarzschild BH. It was found that the late-time GW waveform consisted of damped sinusoids, the form of which carried information about the BH mass. Further numerical work by Press [45], studying the evolution of perturbations to the Schwarzschild geometry, identified the damped sinusoids as the “free oscillation of a black hole”. This work was also the first instance of describing the vibration of a BH as a “quasi-normal mode”.

Equations governing the perturbations of the Kerr metric [46] were found by Teukolsky [47, 48]. Describing rotating BHs, these are expected to be the most general class of astrophysical BH and will be the focus throughout this thesis. A remnant Kerr BH has “no hair” [49], meaning it is fully described by only a final mass and an angular momentum (which we will express via a dimensionless spin parameter). The same is true of the spectrum of Kerr QNM frequencies, which are also functions of only the mass and spin. This is how QNMs carry information about the remnant BH, and this fact forms the basis of the GR tests discussed further in Section 1.4.

The QNM frequencies can be calculated within the framework of linearised gravity, treating the gravitational field in the vicinity of the remnant as a small (linear) perturbation of the Kerr metric. Therefore, the QNM description of the GW signal is only expected to be valid at sufficiently late times, when the nonlinearities from the merger have largely decayed away. In the following subsections we discuss the calculation and physical picture of QNMs further to help build some intuition.

1.3.1 Scalar field on a Schwarzschild background

To help build some intuition regarding the origin of the QNMs, we will perform a demonstrative calculation involving a massless scalar field, $\psi(t, r, \theta, \phi)$, on a Schwarzschild background. This will lead to equations reminiscent of the Zerilli and Regge-Wheeler equations, without the complication of tensor spherical harmonics that comes with the full gravitational-perturbation treatment.

The metric tensor, in Schwarzschild coordinates, is

$$g_{\mu\nu} = \begin{pmatrix} -\left(1 - \frac{2M}{r}\right) & 0 & 0 & 0 \\ 0 & \left(1 - \frac{2M}{r}\right)^{-1} & 0 & 0 \\ 0 & 0 & r^2 & 0 \\ 0 & 0 & 0 & r^2 \sin^2 \theta \end{pmatrix}, \quad (1.26)$$

where now M is the mass of the BH (and not the total mass of the binary, as before). The relevant massless wave equation, qualitatively similar to equations describing GWs and electromagnetic waves (and used here as a toy model for the GW perturbations of a BH), is the Klein-Gordon equation

$$\nabla_\mu \nabla^\mu \psi = 0. \quad (1.27)$$

Using the fact that a covariant derivative reduces to the partial derivative on scalars, we can write this as

$$\nabla_\mu \nabla^\mu \psi = \frac{1}{\sqrt{-g}} \partial_\mu (\sqrt{-g} g^{\mu\nu} \partial_\nu \psi) = 0 \quad (1.28)$$

where g is the determinant of the metric tensor, $g^{\mu\nu}$ is the inverse of the metric tensor, and we have also used $\Gamma_{\mu\nu}^\mu = \partial_\nu \ln \sqrt{-g} = (-g)^{-1/2} \partial_\nu \sqrt{-g}$. For Schwarzschild we have

$\sqrt{-g} = r^2 \sin \theta$. Evaluating, we get

$$-\left(1 - \frac{2M}{r}\right)^{-1} \frac{\partial^2 \psi}{\partial t^2} + \frac{1}{r^2} \frac{\partial}{\partial r} \left[r^2 \left(1 - \frac{2M}{r}\right) \frac{\partial \psi}{\partial r} \right] + \frac{1}{r^2} \left[\frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial \psi}{\partial \theta} \right) + \frac{1}{\sin^2 \theta} \frac{\partial^2 \psi}{\partial \phi^2} \right] = 0. \quad (1.29)$$

Recognising that the scalar spherical harmonics, $Y_{\ell m}(\theta, \phi)$, are the eigenfunctions of the angular part:

$$\frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial Y_{\ell m}}{\partial \theta} \right) + \frac{1}{\sin^2 \theta} \frac{\partial^2 Y_{\ell m}}{\partial \phi^2} = -\ell(\ell + 1) Y_{\ell m}, \quad (1.30)$$

we will attempt a spherical-harmonic decomposition of the scalar field. With the expectation that the radial dependence of the field will go as $1/r$ (and also anticipating a change in radial coordinate) we write the scalar field as

$$\psi(t, r, \theta, \phi) = \frac{1}{r} \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \psi_{\ell m}(t, r) Y_{\ell m}(\theta, \phi). \quad (1.31)$$

Substituting, we are left with an equation for the radial part:

$$-\frac{\partial^2 \psi_{\ell m}}{\partial t^2} + \frac{1}{r} \left(1 - \frac{2M}{r}\right) \frac{\partial}{\partial r} \left[r^2 \left(1 - \frac{2M}{r}\right) \frac{\partial}{\partial r} \left(\frac{\psi_{\ell m}}{r} \right) \right] - \left(1 - \frac{2M}{r}\right) \left(\frac{\ell(\ell + 1)}{r^2} \right) \psi_{\ell m} = 0. \quad (1.32)$$

To proceed we introduce the tortoise coordinate, r_* ,

$$r_* = r + 2M \ln \left(\frac{r}{2M} - 1 \right), \quad (1.33)$$

which has the property that

$$\frac{dr_*}{dr} = \left(1 - \frac{2M}{r}\right)^{-1}. \quad (1.34)$$

Note that as r approaches the Schwarzschild radius ($r \rightarrow 2M$), the tortoise coordinate $r_* \rightarrow -\infty$. Making this change of variables, we arrive at

$$\frac{\partial^2 \psi_{\ell m}(t, r)}{\partial r_*^2} - \frac{\partial^2 \psi_{\ell m}(t, r)}{\partial t^2} - V_\ell(r) \psi_{\ell m}(t, r) = 0, \quad (1.35)$$

where

$$V_\ell(r) = \left(1 - \frac{2M}{r}\right) \left(\frac{\ell(\ell+1)}{r^2} + \frac{2M}{r^3}\right) \quad (1.36)$$

is an effective potential. Performing a Fourier transform,

$$\tilde{\psi}_{\ell m}(\omega, r) = \int_{-\infty}^{\infty} dt \psi(t, r) e^{-2\pi i f t}, \quad (1.37)$$

we can bring Eq. 1.35 into the form of a one-dimensional Schrödinger equation

$$\frac{\partial^2 \tilde{\psi}_{\ell m}(\omega, r)}{\partial r_*^2} + [\omega^2 - V_\ell(r)] \tilde{\psi}_{\ell m}(\omega, r) = 0. \quad (1.38)$$

Gravitational perturbations obey an equation of the same form (i.e. the Regge-Wheeler and Zerilli equations), and in fact the effective potential can be written in the unified form

$$V_\ell(r) = \left(1 - \frac{2M}{r}\right) \left(\frac{\ell(\ell+1)}{r^2} + \frac{(1-s^2)2M}{r^3}\right) \quad (1.39)$$

with $s = 0, -1$ and -2 for scalar, electrical, and gravitational perturbations respectively.

We plot $V_\ell(r)$ in Fig. 1.5 for the scalar and gravitational perturbation cases, and for a few different values of ℓ . Note that the minimum value of ℓ permitted is given by $|s|$.

When studying gravitational perturbations, we write the metric as

$$g'_{\mu\nu} = g_{\mu\nu} + h_{\mu\nu}, \quad (1.40)$$




FIGURE 1.5: The Regge-Wheeler potential as a function of the tortoise coordinate r_* (with the BH horizon at $r_* = -\infty$), for scalar and gravitational perturbations and for a selection of ℓ . In this subsection M refers to the mass of the Schwarzschild BH, and not the total binary mass.

where $g_{\mu\nu}$ is the unperturbed BH spacetime (i.e., for Schwarzschild, Eq. 1.26) and $h_{\mu\nu}$ is a perturbation. An analogous procedure to the above is followed, and it is worth noting that the angular dependence is now described by spin-weighted functions that are generalisations of the usual spherical harmonics. In the case of Schwarzschild, these are the spin-weighted spherical harmonics, ${}_s Y_{\ell m}(\theta, \phi)$. These are also the functions commonly used as the basis in the expansion of NR waveforms; see the discussion around Eq. 2.1 for more details. Spin-weighted functions differ from their usual scalar counterparts via their transformation properties, which is discussed in more detail in Sec. 2.5. In the case of Kerr, a set of functions called the spin-weighted spheroidal harmonics, ${}_s S_{\ell m}(\theta, \phi; \gamma)$ are required [48]. The quantity γ is called the spheroidicity or oblateness, and it depends on the BH spin (see the discussion around Eq. 2.2). The spheroidal harmonics are solutions

of the equation

$$\left[\frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial}{\partial \theta} \right) - \frac{(m + s \cos \theta)^2}{\sin^2 \theta} + (\gamma \cos \theta - s)^2 - s(s - 1) + {}_s A_{\ell m} \right] {}_s S_{\ell m} = 0, \quad (1.41)$$

where ${}_s A_{\ell m}(\gamma)$ is known as the angular separation constant. In the limit $\gamma \rightarrow 0$ (corresponding to the Schwarzschild limit), the solutions to Eq. 1.41 are the spin-weighted spherical harmonics and we have

$${}_s A_{\ell m} = \ell(\ell + 1) - s(s + 1). \quad (1.42)$$

If we then set $s = 0$, we recover Eq. 1.30 (since we have a ϕ -dependence of the form $e^{im\phi}$) and the solutions become the usual spherical harmonics.

Returning to the effective potential, Eq. 1.39, we note that the Zerilli equation has a slightly different form (the Zerilli equation describes polar, or even-parity, perturbations, whereas the Regge-Wheeler equations describes axial, or odd-parity, perturbations). However, it has been shown that the QNM spectrum resulting from the Regge-Wheeler and Zerilli equations are identical [50].

Given Eq. 1.38 and the effective potential, we must specify boundary conditions to find the QNM solutions. When solving for the normal modes on a string, or the energy levels of a quantum harmonic oscillator, we require the wavefunction to vanish at the boundaries (either at the string ends, or at $\pm\infty$ for the harmonic oscillator). The problem we're considering here is slightly different; the potential, depicted in Fig. 1.5, clearly does not admit bound states (it does not have any minima, and $V_\ell(r) > 0$ for all r). The implication is that we should look for plane wave solutions that are ingoing to the BH

horizon, and outgoing to infinity:

$$\begin{aligned}\tilde{\psi}_{\ell m}(\omega, r) &\sim e^{i\omega r_*} & (r_* \rightarrow -\infty) \\ \tilde{\psi}_{\ell m}(\omega, r) &\sim e^{-i\omega r_*} & (r_* \rightarrow \infty).\end{aligned}\tag{1.43}$$

This follows from the consideration that the field should radiate only inward at the horizon and only outward at spatial infinity.

This gives us a well-posed problem, and now all that is left is the computation of the QNMs. A variety of methods have been used [32, 51], with the first attempts consisting of the aforementioned time-domain evolutions of the Regge-Wheeler and Zerilli equations by Vishveshwara [44] and Press [45]. In principle the QNM frequencies can be extracted from the resulting waveform, but this approach does not return the complete spectrum (in practice only a subset of modes can be extracted). Chandrasekhar and Detweiler [52] employed a shooting method to directly integrate the wave equation in the frequency domain; this involves picking a value for the QNM frequency, integrating, and checking whether the boundary conditions are satisfied. This is an inefficient way of identifying QNMs, and this approach is also prone to numerical noise. Analytical methods were developed by Blome, Mashhoon and Ferrari [53–55], which consider the bound states of the inverted BH potential. Although not in general accurate, this approach offers physical insight which we touch upon in the next subsection. Motivated by the analogy between Eq. 1.38 and the Schrödinger equation, Schutz and Will [56] employed WKB methods to calculate a handful of fundamental-mode QNM frequencies. This approach has since been improved upon [57–59], and can give very accurate results for certain modes (but, again, breaks down in some limits). Finally, a numerical method utilising continued fractions (proposed by Leaver [60], and later improved upon by Nollert [61]) is known to be highly

accurate and fast. It is the method of choice in modern codes such as the QNM Python package [62] (which is used in work throughout this thesis). The QNM package actually employs a spectral method, developed in Ref. [63], alongside Leaver’s method to solve for the QNM frequencies and angular separation constants. In brief, when a separation of variables (similar to Eq. 1.31, but with spin-weighted spheroidal harmonics) is applied to the equation governing perturbations to the Kerr metric we obtain coupled radial and angular equations (the Teukolsky equations). The QNM package uses Leaver’s method to solve the radial equation, but the solutions for the angular equation (the spin-weighted spheroidal harmonics) are expanded in the basis of the spin-weighted spherical harmonics as

$${}_sS_{\ell m}(\theta, \phi; \gamma) = \sum_{\ell'=\ell_{\min}}^{\infty} \mu_{\ell' m \ell}(\gamma) {}_sY_{\ell' m}(\theta, \phi), \quad (1.44)$$

where $\ell_{\min} = \max(|m|, |s|)$. The coefficients, $\mu_{\ell' m \ell}(\gamma)$, are the spherical-spheroidal mixing coefficients [64] and will become important when modelling NR data (which is decomposed into spin-weighted spherical harmonics) as a sum of QNMs (whose natural basis is the spin-weighted spheroidal harmonics); see Section 2.3 for more details. When the spheroidal harmonics are expanded in this way (and the upper limit on the sum is replaced by some finite ℓ_{\max}), the angular equation becomes an eigenvalue problem with ${}_sA_{\ell m}(\gamma)$ as the eigenvalues and $\mu_{\ell' m \ell}(\gamma)$ as the eigenvectors. So, the spectral approach has the advantage of returning the spherical-spheroidal mixing coefficients as part of the solution.

In summary, when enforcing the above boundary conditions, only discrete values of ω satisfy Eq. 1.38; these are the QNMs. In the case of a Kerr BH, we denote the QNMs $\omega_{\ell mn}$. They are indexed by three numbers: the usual angular indices ℓ and m , and an additional “overtone” index n . We present the solutions in Section 1.4 and in Fig. 1.6. The full computation of the spectrum is beyond the scope of this work. We instead consider

a simple calculation of QNMs using the geodesic correspondence in order to build some physical intuition.

1.3.2 Quasinormal modes from the geodesic correspondence

First pointed out by Goebel [65], there exists a relation between BH QNMs and null geodesics around the BH spacetime. Having since been developed further [53–55, 66, 67], this approach provides a physical insight to QNMs; that is, QNMs can be interpreted as GWs slowly leaking out of a light-ring orbit around the BH. Although only valid for $\ell \gg 1$ (known as the eikonal, or geometrical optics, limit), this correspondence greatly simplifies the calculations of QNMs since it only depends on the background metric. Consequently, it offers a way to compute QNMs in beyond GR theories (for example, as was done in Ref. [68]), or for computing the QNM spectrum for charged (Kerr-Newman [69]) BHs [70, 71] which can be challenging with the previously discussed methods (but, see also Ref. [72] where a more sophisticated analysis was done).

The key result is that the QNM frequencies for a Kerr BH in the eikonal limit can be written as

$$\omega_{\ell mn} = 2\pi f_{\ell mn} - \frac{i}{\tau_{\ell mn}} = \left(\ell + \frac{1}{2}\right)\Omega_{\theta} + m\Omega_{\text{pre}} - i\left(n + \frac{1}{2}\right)\gamma_L. \quad (1.45)$$

Here, Ω_{θ} is the frequency of small oscillations in the polar direction of a perturbed geodesic. The precessional frequency of the orbital plane, Ω_{pre} , is given by $\Omega_{\text{pre}} = \Omega_{\phi} - \Omega_{\theta}$, where Ω_{ϕ} is the orbital frequency of the light ring. Finally, γ_L is the Lyapunov exponent of the light ring; this can be thought of as a measure of the stability of the orbit. More precisely, it characterises the rate at which the cross section of a congruence of null geodesics on the circular photon orbit increases under radial perturbations. Crucially,

these quantities are all determined by the metric. Below we take the case of a stationary and axisymmetric spacetime (i.e. the spacetime associated with a rotating BH) and show how one can calculate the relevant quantities. For this demonstrative calculation we set $\ell = m$, associated with prograde equatorial motion, to simplify the calculations. Note that, with the assumption $\ell = m \gg 1$, we can write the real part of the QNM frequency as

$$\left(\ell + \frac{1}{2}\right)\Omega_\theta + m\Omega_{\text{pre}} \sim \ell\Omega_\theta + \ell(\Omega_\phi - \Omega_\theta) = \ell\Omega_\phi, \quad (1.46)$$

which aligns with the interpretation of QNMs originating as GWs in light ring orbits.

First, we need the metric associated with a stationary and axisymmetric spacetime. The stationary and axisymmetric character requires that the metric coefficients be independent of t and ϕ , so that $g_{\mu\nu} = g_{\mu\nu}(r, \theta)$. Note that we haven't imposed equatorial motion ($\theta = \pi/2$) yet. We also require that the spacetime is invariant to the simultaneous inversion of the time t and the angle ϕ (i.e. to the transformation $t \rightarrow -t$, $\phi \rightarrow -\phi$). This has the physical meaning that the spacetime we are considering is associated with a rotating body. This invariance requires

$$g_{tr} = g_{t\theta} = g_{\phi r} = g_{\phi\theta} = 0. \quad (1.47)$$

Then we have

$$ds^2 = g_{tt}dt^2 + 2g_{t\phi}dtd\phi + g_{\phi\phi}d\phi^2 + [g_{rr}dr^2 + 2g_{r\theta}drd\theta + g_{\theta\theta}d\theta^2]. \quad (1.48)$$

It can be shown [73] that the term in square brackets can be brought to the diagonal form $g_{r'r'}dr'^2 + g_{\theta'\theta'}d\theta'^2$ by a change of coordinates $r' = r'(r, \theta)$ and $\theta' = \theta'(r, \theta)$. Renaming our

variables by removing the primes, this gives

$$ds^2 = g_{tt}dt^2 + g_{rr}dr^2 + g_{\theta\theta}d\theta^2 + g_{\phi\phi}d\phi^2 + 2g_{t\phi}dtd\phi. \quad (1.49)$$

We can find geodesic curves $x^\mu(\lambda)$ by extremising the action $S = \int d\lambda \mathcal{L}$ where the Lagrangian is given by

$$\begin{aligned} \mathcal{L} &= \frac{1}{2}g_{\mu\nu}\dot{x}^\mu\dot{x}^\nu \\ &= \frac{1}{2}\left(g_{tt}\dot{t}^2 + g_{rr}\dot{r}^2 + g_{\theta\theta}\dot{\theta}^2 + g_{\phi\phi}\dot{\phi}^2 + 2g_{t\phi}\dot{t}\dot{\phi}\right), \end{aligned} \quad (1.50)$$

and a dot denotes a derivative with respect to the affine parameter λ along the curve. We proceed with the Euler-Lagrange (EL) equations,

$$\frac{d}{d\lambda}\left(\frac{\partial\mathcal{L}}{\partial\dot{x}^\mu}\right) = \frac{\partial\mathcal{L}}{\partial x^\mu}. \quad (1.51)$$

Firstly, the stationarity of our spacetime leads to a constant of motion (the energy):

$$\frac{\partial\mathcal{L}}{\partial\dot{t}} = -E \implies g_{tt}\dot{t} + g_{t\phi}\dot{\phi} = -E. \quad (1.52)$$

Similarly, from the axisymmetry of the spacetime we have a second constant of motion (the angular momentum):

$$\frac{\partial\mathcal{L}}{\partial\dot{\phi}} = L \implies g_{\phi\phi}\dot{\phi} + g_{t\phi}\dot{t} = L. \quad (1.53)$$

From Eqs. 1.52 and 1.53 we can solve for the two components of the four-velocity \dot{t} and $\dot{\phi}$ to give

$$\dot{t} = E \frac{g_{\phi\phi} + g_{t\phi}\hat{L}}{(g_{t\phi})^2 - g_{tt}g_{\phi\phi}} \quad (1.54)$$

$$\dot{\phi} = E \frac{g_{t\phi} + g_{tt}\hat{L}}{g_{tt}g_{\phi\phi} - (g_{t\phi})^2} \quad (1.55)$$

where $\hat{L} = L/E$ is the specific angular momentum. We are also free to rescale our affine parameter $\lambda \rightarrow E\lambda$ to remove E from the above expressions. Using the above, the azimuthal orbital frequency can be calculated:

$$\Omega_\phi = \frac{d\phi}{dt} = \frac{\dot{\phi}}{\dot{t}} = -\frac{g_{t\phi} + g_{tt}\hat{L}}{g_{\phi\phi} + g_{t\phi}\hat{L}}. \quad (1.56)$$

In general, for the remaining two coordinates (r and θ) we can use the EL equations to find the second-order differential geodesic equations. However, now we impose the simplification of equatorial motion mentioned previously: $\theta = \pi/2$. This implies $\dot{\theta} = 0$ in the case of reflective symmetry about the equator. This allows us to get a first order equation for \dot{r} by considering the four-velocity for null geodesics:

$$g_{\mu\nu}\dot{x}^\mu\dot{x}^\nu = 0 \implies \dot{r}^2 = V_{\text{eff}}(r; \hat{L}), \quad (1.57)$$

where

$$\begin{aligned} V_{\text{eff}}(r; \hat{L}) &= \frac{-g_{tt}\dot{t}^2 - g_{\phi\phi}\dot{\phi}^2 - 2g_{t\phi}\dot{t}\dot{\phi}}{g_{rr}} \\ &= \frac{g_{tt}\hat{L}^2 + 2g_{t\phi}\hat{L} + g_{\phi\phi}}{g_{rr}(g_{t\phi}^2 - g_{tt}g_{\phi\phi})}, \end{aligned} \quad (1.58)$$

and in the final line we have used Eqs. 1.54 and 1.55 to eliminate \dot{t} and $\dot{\phi}$. In Eq. 1.58 the $g_{\mu\nu}$ metric coefficients are to be evaluated on the equatorial plane $\theta = \pi/2$ and so are only functions of r . A light ring is a circular null geodesic orbit. The radius, r_* , and angular momentum, \hat{L}_* , of such an orbit must satisfy $V_{\text{eff}} = V'_{\text{eff}} = 0$, where a prime denotes a radial derivative with respect to r . The first condition yields

$$\hat{L}_*(r) = \frac{-g_{t\phi} \pm \sqrt{g_{t\phi}^2 - g_{tt}g_{\phi\phi}}}{g_{tt}}, \quad (1.59)$$

while the second gives an implicit formula for r_* :

$$V'_{\text{eff}}(r_*; \hat{L}_*(r_*)) = 0. \quad (1.60)$$

In general we solve Eq. 1.60 numerically to obtain the light ring radius. For prograde equatorial orbits in the Kerr metric this will have a single root, but will depend on the spin chosen for the spacetime.

Having found the equations of the light ring, we now consider neighbouring geodesics (i.e. perturbations to the light ring) to determine Ω_θ and the Lyapunov exponent. First, consider polar motion. The EL equation for θ is

$$g_{\theta\theta}\ddot{\theta} + \left(\frac{\partial g_{\theta\theta}}{\partial \theta} \dot{\theta} + \frac{\partial g_{\theta\theta}}{\partial r} \dot{r} \right) \dot{\theta} = \frac{1}{2} \left(\frac{\partial g_{tt}}{\partial \theta} \dot{t}^2 + \frac{\partial g_{rr}}{\partial \theta} \dot{r}^2 + \frac{\partial g_{\theta\theta}}{\partial \theta} \dot{\theta}^2 + \frac{\partial g_{\phi\phi}}{\partial \theta} \dot{\phi}^2 + 2 \frac{\partial g_{t\phi}}{\partial \theta} \dot{t} \dot{\phi} \right). \quad (1.61)$$

We consider small perturbations in the θ direction about the light ring. That is, we set $r = r_*$ and $\theta = \pi/2 + \delta\theta$. Discarding terms $\mathcal{O}(\delta\theta^2)$ (and with \dot{t} and $\dot{\phi}$ given by Eqs. 1.54 and 1.55 respectively) we obtain

$$g_{\theta\theta}(r_*, \pi/2) \delta\ddot{\theta} = \frac{1}{2} \left(\frac{\partial^2 g_{tt}(r_*, \pi/2)}{\partial \theta^2} \dot{t}^2 + \frac{\partial^2 g_{\phi\phi}(r_*, \pi/2)}{\partial \theta^2} \dot{\phi}^2 + 2 \frac{\partial^2 g_{t\phi}(r_*, \pi/2)}{\partial \theta^2} \dot{t} \dot{\phi} \right) \delta\theta. \quad (1.62)$$

which describes simple harmonic motion, $\delta\ddot{\theta} = -\tilde{\Omega}_\theta^2\delta\theta$, where the constant $\tilde{\Omega}_\theta$ is the frequency of the oscillations with respect to the parameter λ . The frequency of the oscillations with respect to coordinate time t is given by

$$\begin{aligned}\Omega_\theta &= \frac{\tilde{\Omega}_\theta}{\dot{t}} = \frac{1}{\dot{t}} \sqrt{-\frac{1}{2g_{\theta\theta}} \left(\frac{\partial^2 g_{tt}}{\partial \theta^2} \dot{t}^2 + 2 \frac{\partial^2 g_{t\phi}}{\partial \theta^2} \dot{t} \dot{\phi} + \frac{\partial^2 g_{\phi\phi}}{\partial \theta^2} \dot{\phi}^2 \right)} \\ &= \sqrt{-\frac{1}{2g_{\theta\theta}} \left(\frac{\partial^2 g_{tt}}{\partial \theta^2} + 2 \frac{\partial^2 g_{t\phi}}{\partial \theta^2} \Omega_\phi + \frac{\partial^2 g_{\phi\phi}}{\partial \theta^2} \Omega_\phi^2 \right)},\end{aligned}\quad (1.63)$$

where all quantities on the right hand side are to be evaluated at the light ring. In the case of the Kerr metric $\Omega_\theta^2 > 0$ and the light ring is stable in the polar direction.

Now consider motion in the radial direction ($\dot{r} \neq 0$). Differentiating Eq. 1.57 with respect to λ gives

$$\ddot{r} = \frac{1}{2} \frac{\dot{V}_{\text{eff}}(r)}{\dot{r}} = \frac{1}{2} V'_{\text{eff}}(r). \quad (1.64)$$

Consider small perturbations $r = r_* + \delta r$ with $\theta = \pi/2$ and discarding $\mathcal{O}(\delta r^2)$ terms gives

$$\delta\ddot{r} = \frac{1}{2} V''_{\text{eff}}(r_*) \delta r. \quad (1.65)$$

Looking for periodic solutions, the frequency of the radial oscillations (with respect to coordinate time) is given by

$$\Omega_r = \sqrt{-\frac{V''_{\text{eff}}(r_*)}{2\dot{t}^2}}. \quad (1.66)$$

In the case of the Kerr metric the light ring orbit has $\Omega_r^2 < 0$ and is unstable in the radial direction; we can define the Lyapunov exponent as

$$\gamma_L = -i\Omega_r = \sqrt{\frac{V''_{\text{eff}}(r_*)}{2\dot{t}^2}} \quad (1.67)$$

and the corresponding decay timescale from Eq. 1.45 is

$$\tau_{\ell mn} = \left(n + \frac{1}{2}\right)^{-1} \sqrt{\frac{2t^2}{V_{\text{eff}}''(r_*)}}. \quad (1.68)$$

We conclude this section with a discussion of some the problems associated with QNMs. In Section 1.4 we will discuss the full (and accurately computed) QNM spectrum further, and the tests of GR that can be done with the observation of QNMs.

1.3.3 Limitations of quasinormal modes

Although QNMs offer a powerful tool to learn about the remnant BH and test our theories (see Section 1.4), there are some subtleties regarding the QNM spectrum and the ringdown signal that should be pointed out.

It is known that QNMs, in contrast to normal modes, do not form a complete basis [33, 74–77], meaning that there are other contributions to the ringdown waveform beyond the QNMs. In particular, it is known that at late times the ringdown consists of a power-law tail [78–80] (see also Ref. [81], and references therein, for a more recent discussion about our knowledge of tails), and nonlinear effects such as GW memory [82] may also be present. Although at a much lower amplitude than the QNM ringing and not resolvable at current detector sensitivities, at sufficient SNRs these effects may need to be considered to avoid biases. This idea was touched upon in Ref. [83], where it was found that their ability to constrain remnant parameters from the ringdown signal actually degraded at high SNRs (late-time tails were suggested as an explanation for this, but it is not clear if they were responsible in this study). See also Ref. [84] for a study on how neglecting these beyond-QNM contributions could bias future studies.

Motivated by work done in Refs. [74, 75], Nollert [85] attempted to make the QNM spectrum complete via modifications to the effective potential (for example, in the case of Schwarzschild, the Regge-Wheeler potential depicted in Fig. 1.5). However, it was found that the QNM spectrum is highly unstable to changes in the potential (no matter how small). That is, perturbations to the potential (which do not significantly change the physical ringdown waveform) can lead to very different QNM spectra; this brings the physical meaning of the QNMs into question. This fact was commented on again in Nollert and Price [76], later revisited by Daghigh et al. [86], and additional recent work is ongoing [87–91].

Putting these subtleties aside, we know we can calculate the QNM spectrum (as described throughout this section, and see Fig. 1.6 below) and based on the GW observations made so far it seems that there is physical significance to the QNMs. However, the excitations of each mode (that is, the amplitudes of each QNM in the ringdown waveform) are harder to predict. In principle the QNM amplitudes are a function of the initial conditions of the binary; the perturbation of the remnant BH depends on the way the binary merged. Unfortunately this mapping is non-trivial, particularly in the case of precessing BBHs. The QNMs we include in our models and target in the data must then be motivated by numerical studies. We contribute to this in Chapter 2, where we also discuss this issue in more detail.

Similarly to the mode amplitudes, a priori it is not known what the start time of the ringdown is. Even with clean, noise-free numerical simulations it is not clear if the start time can be determined unambiguously. This is an issue which will be present throughout the thesis, and in the data-analysis context (Chapters 3 and 4) we propose a method to alleviate this concern.

Finally, on the subject of data analysis, the fact that ringdown models are fundamentally

discontinuous leads to problems; GW analysis is usually done in the frequency domain, but Fourier transforms of a discontinuous model leads to spectral leakage. This is discussed further in Chapters 3 and 4, where we also propose a solution.

1.4 Black-hole spectroscopy

The full Kerr QNM spectrum is nowadays readily available, with modern codes (such as the QNM package [62]) making use of a version of Leaver’s method [60] to compute them accurately and quickly. We show the $\ell = 2, n = 0$; $\ell = 3, n = 0$; and $\ell = 2, n = 1$ branches of the Kerr spectrum in Fig. 1.6 (of course, the full spectrum would extend to infinity in both ℓ and n , but the branches shown here include the QNMs expected to be of most interest at current detector sensitivities).

There are different conventions used in the literature to label the modes; we aim to clarify these with Fig. 1.6. Firstly, we see the QNMs appear to come in pairs: those with positive real part, and those with negative real part. These are the “regular” and “mirror” (sometimes called “twin” or “conjugate”) QNMs respectively. Denoting the mirror QNM frequency with a prime, we can relate $\omega'_{\ell mn}$ to the regular QNMs as follows [92]:

$$\begin{aligned} f'_{\ell mn} &= -f_{\ell -mn}, & \tau'_{\ell mn} &= \tau_{\ell -mn} \\ \implies \omega'_{\ell mn} &= -\omega_{\ell -mn}^*. \end{aligned} \tag{1.69}$$

So, for example, when we refer to the “(2, −2, 0) mirror mode”, we are referring to the mode that is the “mirror image” of the regular (2, 2, 0) mode in Fig. 1.6. This way of referring to the modes allows us to unambiguously refer to every mode in the spectrum, including the $m = 0$ modes, and will be the preferred method in this work.

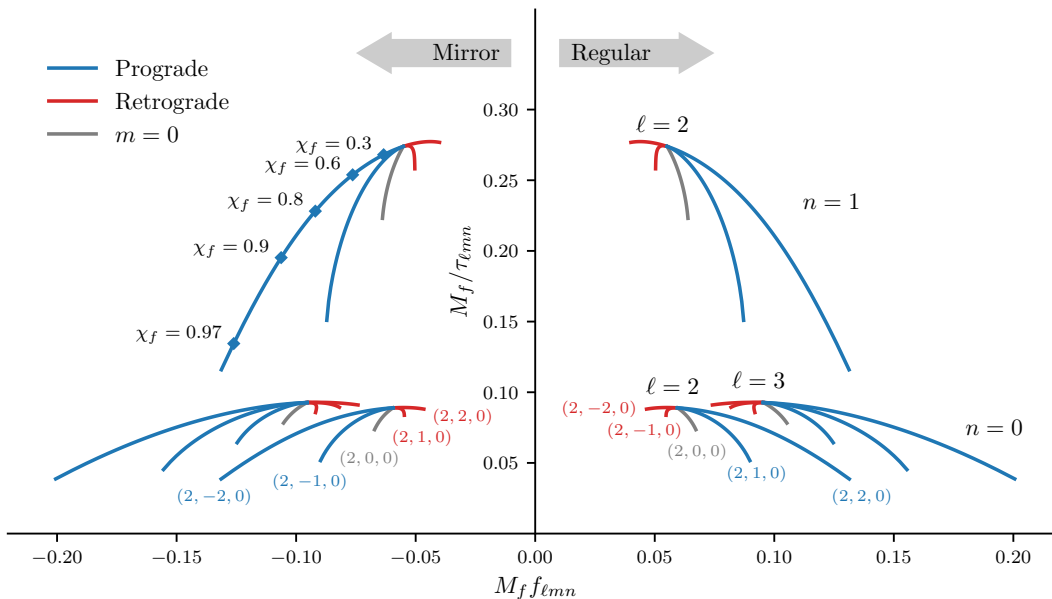


FIGURE 1.6: Selected modes of the Kerr QNM spectrum. BH QNM frequencies are conventionally represented as complex numbers, with the real part giving the angular frequency of the mode and the imaginary part giving (minus) the inverse of the damping time: $\omega_{\ell mn} = 2\pi f_{\ell mn} - i/\tau_{\ell mn}$. Here we plot $f_{\ell mn}$ and $1/\tau_{\ell mn}$, each scaled by the remnant BH mass M_f to make a dimensionless quantity. The spectrum of a Kerr BH also depends on the dimensionless BH spin magnitude, χ_f ; as the spin is increased from zero, branches with different m grow from the points of the Schwarzschild spectrum.

We can, alternatively, split the spectrum into “prograde” and “retrograde” modes. This description has the advantage of having a clear physical interpretation; prograde (retrograde) modes are those that are corotating (counterrotating) with the final BH spin. This can also be expressed as prograde modes satisfying $\text{sgn}(\text{Re}[\omega_{\ell mn}]) = \text{sgn}(m)$, and retrograde modes satisfying $\text{sgn}(\text{Re}[\omega_{\ell mn}]) = -\text{sgn}(m)$. For binaries where the individual BHs have low spin (or aligned spins that rotate in the same sense as the orbit) it is also expected that the retrograde modes will be suppressed compared to the prograde modes, which is simply a result of the geometry of merger (this has also been shown by studies of numerical simulations, for example in Refs. [92–94]). This assumption is less clear for precessing systems, and is investigated in Chapter 2. Note that this way of classifying the

modes breaks down for $m = 0$, leading to some ambiguity. This isn't a problem in current data analysis, which focuses on prograde $\ell = |m|$ modes, but since we also perform fits to numerical simulations in this thesis we opt to use the regular-mirror classification.

In terms of spherical harmonics it is known the $\ell = |m| = 2$ family of modes dominate the GW strain, and so we expect the same in the ringdown. Since the overtones decay more quickly (i.e. τ decreases) with increasing n , at late times the signal will be dominated by the fundamental $n = 0$ modes. Therefore, the most prominent QNM in the ringdown is expected to be the $(2, \pm 2, 0)$ prograde mode; the observational challenge is usually to detect the presence of other, subdominant modes.

As previously mentioned, the remnant Kerr BH is fully described by only a final mass and a dimensionless final spin [49]. We now denote these quantities M_f and $\chi_f = |\chi_f|$, where the f denotes “final” to avoid confusion with the binary properties. Consequently, the Kerr QNM spectrum is also a function of only the remnant mass and spin. In Fig. 1.6 we can see the mass enters as a scaling on the axes, and the spin gives the position along each branch (we show select spin values along the $(2, -2, 1)$ mirror-mode branch). The Schwarzschild spectrum, which depends only on the BH mass and has no m dependence, can be recovered by taking the branch point of each (ℓ, n) group.

This property of Kerr BHs (that they are described by only two parameters in GR) is known as the no-hair theorem, and provides various opportunities for testing GR. For example, consistency tests can be performed between the inspiral and ringdown [95–98]; this involves estimating the remnant properties from an inspiral-only analysis, and separately from a ringdown-only analysis. In the case of the ringdown, the remnant properties are given directly. For the inspiral, the binary properties must be projected forward using knowledge from NR simulations (see, for example, Ref. [99]). This test can be posed in an alternative way by using the BH areas; Hawking's area theorem states that the total

(classical) BH horizon area cannot decrease over time [100]. Put another way, in a BBH merger, the total area of the two initial BHs must be smaller than the area of the remnant. As the BH area is a simple function of the mass and spin, this test offers an alternative way of testing consistency between measurements of the inspiral and merger (and it does not rely on projecting inspiral parameters forward to merger). See Refs. [101, 102] for examples of applying this test.

Additionally, we can test the no-hair theorem directly; by simply measuring the QNM frequencies in a ringdown signal we can perform BH spectroscopy. In analogy with using spectral lines to identify atomic elements, in BH spectroscopy we use the QNM frequencies to identify BHs and their properties. Initial studies by Echeverria [103] and Finn [104] considered the precision with which the mass and spin of the remnant BH could be measured with a single QNM. It was in a work by Dreyer et al. [105] that the importance of measuring multiple QNMs was emphasised, since this leads to tests of the Kerr metric in GR (discussed further below). This multimode formalism was developed significantly by Berti et al. [92, 106], and there have since been multiple studies on alternative test methods [107, 108], combining events to place tighter constraints [109–112], and BH spectroscopy prospects with current and future GW detectors [113–118].

Fig. 1.7 (inspired by a similar figure in Ref. [105]) demonstrates the idea behind BH spectroscopy. On the left panel we show four branches of the Kerr spectrum; the $(2, 2, 0)$, $(2, 2, 1)$, $(3, 3, 0)$ and $(3, 3, 1)$ modes. The $(2, 2, 1)$ and $(3, 3, 0)$ modes are promising targets for a QNM measurement beyond the fundamental mode, and currently they are the only modes for which there is possible evidence in GW observations [119, 120] (these claims are, however, disputed; see Chapters 3 and 4 for further discussion). In BH spectroscopy, we imagine directly measuring a frequency and damping time in the GW data; say, f_* and τ_* . In the $M_f f_{\ell mn} - M_f / \tau_{\ell mn}$ space, this measurement corresponds to a straight line




FIGURE 1.7: BH spectroscopy (inspired by a similar figure in Ref. [105]): a frequency – damping-time measurement corresponds to a straight line in the left panel. Three such measurements are shown, and their intersections with the Kerr spectrum are indicated. Each measurement is associated with a particular marker shape, and each Kerr branch is associated with a particular marker colour. Each intersection can be converted to a remnant BH mass and spin, shown on the right panel. Only one mass and spin is consistent with all three measurements; these are the true BH properties.

intersecting with the origin; this simply comes from the fact that our scale is not fixed, since we express everything in terms of M_f . Fixing the value of M_f would collapse the line to a single point, (f_*, τ_*) . The gradient of the line is given by

$$\frac{1}{f_* \tau_*} = \frac{\pi}{Q_*}, \quad (1.70)$$

where $Q_* = \pi f_* \tau_*$ is the quality factor of the measured mode. We show three such measurements, and mark where they intersect the Kerr spectrum. Each intersection corresponds to a particular remnant BH mass, M_f , and spin, χ_f . The distance of the intersection along the Kerr branch gives the spin. The mass is found by, for example, dividing the value of $M_f f_{lmn}$ at the intersection by the measured frequency f_* . Since each measurement line will generically intersect with multiple Kerr branches, they return

multiple (M_f, χ_f) values. However, assuming the BH is indeed described by the Kerr metric, then there will be a single (M_f, χ_f) combination consistent with all the measurements (indicated by the horizontal and vertical lines in the right panel). If no such combination exists, then what is observed is either not an isolated BH, or it is not described by the Kerr metric in GR.

The above describes the key idea behind BH spectroscopy, but in practice the test is performed slightly differently. Firstly, in reality it may not be practical to measure multiple QNM frequencies directly due to parameter degeneracies that come with summing functionally equivalent QNMs. And secondly, it can be tricky to turn multiple frequency–damping-time measurements into a quantitative statement regarding compatibility with the GR Kerr spectrum. Instead, we can target the loudest QNMs in the data and parameterise all the QNM frequencies with a single mass and spin (according to the Kerr spectrum). That is,

$$\omega_{\ell mn} = \omega_{\ell mn}(M_f, \chi_f). \quad (1.71)$$

This reduces the dimensions of the problem, and also directly measures the physically relevant properties of the remnant BH. With the expectation that what we measure will not be greatly different from the GR Kerr prediction, we can measure any disagreement by allowing a subset of the QNMs in our model to deviate from Kerr in the following way:

$$f = f_{\ell mn}(M_f, \chi_f)e^{\delta f_{\ell mn}} \quad \tau = \tau_{\ell mn}(M_f, \chi_f)e^{\delta \tau_{\ell mn}}. \quad (1.72)$$

We emphasise that not all the modes in the model should be allowed to vary, since then the analysis becomes equivalent to a direct search of frequencies and damping times. The new parameters, $\delta f_{\ell mn}$ and $\delta \tau_{\ell mn}$, will have measurements consistent with zero if the data is well described by Kerr. Additionally, the uncertainty on the measurements of the δ

parameters give us a natural test confidence. This method of testing the no-hair theorem is described further in Ref. [121], and is employed in Section 4.3.2 for our analysis of GW150914.

Chapter 2

Modelling the Ringdown from Precessing Black-hole Binaries

2.1 Introduction

In 2019 Giesler et al. [122] demonstrated that using QNMs with $n > 0$ (that is, ringdown “overtones”) could push the validity of the linear ringdown model to times as early as the peak of the GW strain. Their work involved fitting ringdown models to a selection of aligned-spin SXS simulations, and in this chapter we extend their work to misaligned-spin (i.e. precessing) systems.

The 2019 study sparked many other papers which involve fitting ringdown models with overtones to NR simulations [94, 123–130] (including the present work), but it should be noted that fears of over-fitting and the physical validity of the overtones were also present. Even at the time of writing there is not a consensus on this point [131, 132], a central issue being that of the ringdown start time.

A prerequisite for any ringdown analysis is a suitable choice for the start time, t_0 , of the ringdown. Starting too early risks obtaining biased measurements, because a GW signal contaminated with nonlinearities from the merger cannot be described by a model based

solely on QNMs. On the other hand, due to the exponential decay of the ringdown, starting too late leaves too little SNR to make useful measurements. In ringdown studies, typically the start time is given in reference to the maxima of some time-dependant quantity. This could be the modulus of the $(2, 2)$ mode of the strain or the Ψ_4 Weyl scalar, or the total GW luminosity; these quantities peak at times that typically differ by a few tens of M (see Ref. [133] for a discussion of some possible choices for the ringdown start time). Often, to avoid concerns of fitting to the nonlinear merger, the ringdown start time is chosen to be $10M$ to $20M$ after the peak of these reference quantities. The work of Giesler et al. [122], which itself builds on previous studies of fitting ringdown models to NR simulations [84, 93, 107, 133–136], found that by including up to seven overtones the ringdown analysis can be started as early as the peak of the $(2, 2)$ mode strain. This might be considered a surprising result; the signal peak is expected to occur when the remnant BH (to the extent that it yet even makes sense to consider it as such) is most highly distorted and linear perturbation theory is not expected to be valid. The failure of this intuition was investigated in Ref. [137] which suggests much of the nonlinearity is trapped behind a forming common apparent horizon and never makes it out to future null infinity in the form of GWs. Further support came from a study on the overtone excitation factors [138], which quantify the ease of excitation of the modes, and it was found that higher overtones are relatively easy to excite. Even more surprising, Dhani [126] extended this approach via the inclusion of mirror modes along with overtones (thereby doubling the number of QNMs) and it was found that it was possible to start the ringdown analysis even earlier (up to $10M$ before the peak). Clearly it is not surprising that a model with so many free parameters is able to fit the GW signal well; the important point is that it is able to do so without obtaining biased values for the final mass and spin.

The previous studies mentioned have only considered aligned-spin BBH systems,

although we note that Ref. [136] performed some limited analyses on precessing simulations. We also note that some work on precessing systems has been done in the extreme mass ratio limit, see Refs. [139–141]. It is well known that misalignment between the orbital angular momentum and the spins of the component BHs cause the orbit to precess during the inspiral phase of the evolution, leading to qualitatively different GW signals at early times (see, e.g., Ref. [142]). It is less clear what effect, if any, misaligned component spins would have on the late-time ringdown signal which is generally associated with the remnant BH. The primary aim of this chapter is to address this question by systematically extending the analyses of Refs. [122, 126] to a large number of precessing BBH simulations from the SXS catalog [24]. We find that for BBH systems with misaligned spins, and that exhibit precession during their inspiral phase, a model consisting only of overtones (with or without mirror modes) cannot be reliably applied from the peak amplitude of the $(2, 2)$ strain. A more conservative ringdown start time corresponding to the peak of the total energy flux (i.e., the GW luminosity) improves reliability, but we still see significant variation in performance across different simulations. The introduction of a higher harmonic (QNMs with $\ell > 2$) to the overtone model helps to reduce this variation, hinting at the importance of mode mixing.

Previous studies have focused on using full NR simulations to test ringdown models. In this chapter we also briefly investigate the use of surrogates, which provide an opportunity to test models over a continuous parameter space. We find caution should be taken, particularly for surrogates of precessing systems, due to errors in the surrogate waveforms.

In Section 2.2 we write down the most general form of the ringdown model which will be used throughout this chapter. The method used to fit the ringdown model to the SXS simulations is then explained in Section 2.3. In Section 2.4 we reproduce some important results from Refs. [122, 126], which are later compared with those for precessing systems

in Section 2.5. With precessing systems, it is necessary to perform a frame rotation to account for the fact that the spin of the remnant BH will not be aligned with the initial coordinate axes used to set up the simulation; the procedure for doing this is also discussed in Section 2.5. In Section 2.6 we comment on the use of NR surrogates to test ringdown models, and in Section 2.7 we discuss the estimation of numerical errors present in the NR simulations. Finally, concluding remarks are presented in Section 2.8. Throughout, we use units in which $G = c = 1$.

2.2 Model for the spherical modes

NR expands the GW strain in the basis of (spin-weighted) spherical harmonics

$$h = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h_{\ell m}(t) {}_{-2}Y_{\ell m}(\Omega), \quad (2.1)$$

where Ω is used as shorthand for the angles θ, ϕ . By convention, the NR frame is uniquely fixed by requiring that initially the two component BHs are located on the x -axis and the orbital angular momentum, \mathbf{L} , points along the z -axis. The $h_{\ell m}(t)$ coefficients are referred to as the spherical-harmonic modes of the GW signal. The $\ell = |m| = 2$ modes are typically largest, while the remaining “higher modes” are generally subdominant. The output of an NR simulation usually includes the first few modes (e.g. $\ell \leq 8$) with the asymptotic radial dependence scaled out. The spherical-harmonic modes are defined with respect to a particular frame at infinity, chosen such that the centre-of-mass of the system is at rest at some initial time. Note, however, that this still leaves freedom to perform an overall rotation (as will become important when we discuss precessing systems).

At late times ($t \geq t_0$, where t_0 is to be determined), perturbation theory expands the GW strain in the basis of the (spin-weighted) spheroidal harmonics

$$h = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} \sum_{n=0}^{\infty} \left[C_{\ell mn} e^{-i\omega_{\ell mn}(t-t_0)} {}_{-2}S_{\ell mn}(\Omega) + C'_{\ell mn} e^{-i\omega'_{\ell mn}(t-t_0)} {}_{-2}S'_{\ell mn}(\Omega) \right]. \quad (2.2)$$

Here, $C_{\ell mn}$ are complex coefficients (containing an amplitude and a phase), $\omega_{\ell mn} = 2\pi f_{\ell mn} - i/\tau_{\ell mn}$ are the complex QNM frequencies (which are functions of the remnant BH mass M_f and spin χ_f), and ${}_{-2}S_{\ell mn}(\Omega) = {}_{-2}S_{\ell m}(\Omega; a\omega_{\ell mn})$ are the spheroidal harmonics. The spheroidal harmonics are functions of the spheroidicity $\gamma = a\omega_{\ell mn}$, where $a = M_f\chi_f$ is the Kerr parameter. The primes denote the mirror modes, which satisfy $\text{Re}[\omega'_{\ell mn}] = 2\pi f'_{\ell mn} < 0$. As discussed in Section 1.4, these mirror modes are related to the regular modes via $\omega'_{\ell mn} = -\omega_{\ell -mn}^*$. The prime on the spheroidal harmonic enters in the spheroidicity: ${}_{-2}S'_{\ell mn}(\Omega) = {}_{-2}S_{\ell m}(\Omega; a\omega'_{\ell mn})$.

It is important to note that Eq. 2.2 is valid in a frame in which the remnant BH is at rest, with its spin vector pointing along the positive z -direction (such a frame is unique up to an unimportant rotation about the z -axis). This ringdown frame is only the same as the NR frame for aligned-spin BBH systems (it is possible that systems with large component spins in the negative z -direction will exhibit a “spin flip”, where the final spin also points in the negative z -direction; in these cases the two frames will only differ by a sign). For misaligned-spin systems the remnant spin can point in essentially any direction and the NR and ringdown frames are misaligned; in these instances, as explained in Section 2.5, we need to rotate the frame of the NR simulation to bring it into the ringdown frame where we can then apply Eq. 2.2. The ringdown frame will also be moving with respect to the NR frame as a result of the recoil, or kick, from the anisotropic emission of GWs near merger. The effects of the kick are neglected here; it is assumed that the NR and

ringdown frames are related by a rotation.

Assuming that the remnant BH spin vector is aligned with the z -axis in the NR frame (that is, the required rotation has been applied to the NR spherical-harmonic modes), we can equate Eqs. 2.1 and 2.2 to get

$$\sum_{\ell m} h_{\ell m}(t) {}_{-2}Y_{\ell m}(\Omega) = \sum_{\ell mn} \left[C_{\ell mn} e^{-i\omega_{\ell mn} t} {}_{-2}S_{\ell mn}(\Omega) + C'_{\ell mn} e^{-i\omega'_{\ell mn} t} {}_{-2}S'_{\ell mn}(\Omega) \right], \quad (2.3)$$

where we have dropped the limits on the sums for clarity. We can then extract $h_{\ell m}$ using spherical-harmonic orthogonality:

$$h_{\ell' m'}(t) = \sum_{\ell mn} \left[C_{\ell mn} e^{-i\omega_{\ell mn} t} \left(\int_{\Omega} d\Omega {}_{-2}S_{\ell mn}(\Omega) {}_{-2}Y_{\ell' m'}^*(\Omega) \right) + C'_{\ell mn} e^{-i\omega'_{\ell mn} t} \left(\int_{\Omega} d\Omega {}_{-2}S'_{\ell mn}(\Omega) {}_{-2}Y_{\ell' m'}^*(\Omega) \right) \right]. \quad (2.4)$$

Following the convention of Ref. [62], the first integral is the spherical-spheroidal mixing coefficient:

$$\int_{\Omega} d\Omega {}_{-2}S_{\ell mn}(\Omega) {}_{-2}Y_{\ell' m'}^*(\Omega) = \mu_{\ell' m' \ell n} \delta_{m' m}. \quad (2.5)$$

To evaluate the second integral we first rewrite the primed spheroidal harmonic as

$$\begin{aligned} {}_{-2}S'_{\ell mn}(\Omega) &= {}_{-2}S_{\ell m}(\theta, \phi; a\omega'_{\ell mn}) \\ &= {}_{-2}S_{\ell m}(\theta, \phi; -a\omega_{\ell -mn}^*) \\ &= (-1)^{\ell} {}_{-2}S_{\ell -m}^*(\pi - \theta, \phi; a\omega_{\ell -mn}) \end{aligned} \quad (2.6)$$

where the last line follows from Eqs. 48b and 48c of Ref. [63]. Next, we use the symmetries of the spherical harmonics to write

$$-{}_2Y_{\ell m}^*(\theta, \phi) = (-1)^\ell {}_2Y_{\ell -m}(\pi - \theta, \phi). \quad (2.7)$$

Using Eqs. 2.6 and 2.7 we can rewrite the second integral of Eq. 2.4 as

$$\begin{aligned} \int_{\Omega} d\Omega {}_2S'_{\ell mn}(\Omega) {}_2Y_{\ell' m'}^*(\Omega) &= \\ &= \int_{\Omega} d\Omega (-1)^\ell {}_2S_{\ell -m}^*(\pi - \theta, \phi; a\omega_{\ell -mn}) (-1)^{\ell'} {}_2Y_{\ell' -m'}(\pi - \theta, \phi) \\ &= (-1)^{l+l'} \int_{\Omega} d\Omega {}_2S_{\ell -m}^*(\pi - \theta, \phi; a\omega_{\ell -mn}) {}_2Y_{\ell' -m'}(\pi - \theta, \phi) \\ &= (-1)^{l+l'} \mu_{\ell' -m' \ell n}^* \delta_{m' m}. \end{aligned} \quad (2.8)$$

Substituting for both of the integrals of Eq. 2.4 we get

$$\begin{aligned} h_{\ell' m'}(t) &= \sum_{\ell mn} \left[C_{\ell mn} e^{-i\omega_{\ell mn} t} \mu_{\ell' m' \ell n} \delta_{m' m} + C'_{\ell mn} e^{-i\omega'_{\ell mn} t} (-1)^{l+l'} \mu_{\ell' -m' \ell n}^* \delta_{m' m} \right] \\ &= \sum_{\ell n} \left[C_{\ell m' n} e^{-i\omega_{\ell m' n} t} \mu_{\ell' m' \ell n} + C'_{\ell m' n} e^{-i\omega'_{\ell m' n} t} (-1)^{l+l'} \mu_{\ell' -m' \ell n}^* \right] \\ &= \sum_{\ell n} \left[C_{\ell m' n} e^{-i\omega_{\ell m' n} t} \mu_{\ell' m' \ell n} + C'_{\ell m' n} e^{-i\omega'_{\ell m' n} t} \mu'_{\ell' m' \ell n} \right], \end{aligned} \quad (2.9)$$

where $\mu'_{\ell' m' \ell n} = (-1)^{l+l'} \mu_{\ell' -m' \ell n}^*$. Eq. 2.9 tells us how a given spherical-harmonic mode (as provided by NR simulations) can be expressed in terms of QNMs. It reveals that each spherical-harmonic mode has contributions from every QNM of the same m , weighted by the spherical-spheroidal mixing coefficients; this is an effect known as mode mixing [64].

2.3 Fitting implementation

Given some spherical-harmonic modes $h_{\ell m}$, we can turn Eq. 2.9 into a least-squares fitting problem to find the best-fit complex coefficients $C_{\ell mn}$ and $C'_{\ell mn}$ (this assumes the complex frequencies and mixing coefficients are also given). First, we write Eq. 2.9 as a matrix equation. Note that in the above we have separated the regular and mirror modes to show explicitly how to deal with the mirror modes (i.e. how to obtain their frequencies and mixing coefficients from the regular frequencies and mixing coefficients). We will now drop this distinction for clarity.

2.3.1 Single-mode fit

For simplicity, first consider the case when we want to model a single spherical mode (for example, the h_{22} mode). We write

$$\mathbf{h}_{\ell m} = \mathbf{a}_{\ell m} \cdot \mathbf{C}, \quad (2.10)$$

where $\mathbf{h}_{\ell m} = (h_{\ell m}(t_0), h_{\ell m}(t_1), \dots, h_{\ell m}(t_{K-1}))$ is the waveform data discretely sampled at a total of K times labelled by t_k . In general t_0 may not exist on the default array of simulation times, so care must be taken (for example, we can interpolate the simulation data and evaluate on a new grid of times, or we could use the first value after t_0 , or the closest value to t_0).

The matrix $\mathbf{a}_{\ell m}$ is where the choice of QNM content in our model enters. It has the form

$$\mathbf{a}_{\ell m} = \begin{pmatrix} e^{-i\omega_0(t_0-t_0)}\mu_{\ell m,0} & e^{-i\omega_1(t_0-t_0)}\mu_{\ell m,1} & \dots & e^{-i\omega_{J-1}(t_0-t_0)}\mu_{\ell m,J-1} \\ e^{-i\omega_0(t_1-t_0)}\mu_{\ell m,0} & e^{-i\omega_1(t_1-t_0)}\mu_{\ell m,1} & \dots & e^{-i\omega_{J-1}(t_1-t_0)}\mu_{\ell m,J-1} \\ \vdots & \vdots & \ddots & \vdots \\ e^{-i\omega_0(t_{K-1}-t_0)}\mu_{\ell m,0} & e^{-i\omega_1(t_{K-1}-t_0)}\mu_{\ell m,1} & \dots & e^{-i\omega_{J-1}(t_{K-1}-t_0)}\mu_{\ell m,J-1} \end{pmatrix}, \quad (2.11)$$

where we have suppressed the three QNM labels and instead labelled each QNM by a single number. There are a total of J QNMs included in the model. So, $\mathbf{a}_{\ell m}$ is a matrix of shape (K, J) . To reiterate, the QNM content in our model can consist of any regular or mirror modes, as long as the correct frequencies and mixing coefficients are used in the above matrix (we have just used ω_j and $\mu_{\ell m,j}$ as generic terms).

The vector \mathbf{C} contains our complex amplitudes, of which there are a total of J (one for each QNM included in the model). This has the form

$$\mathbf{C} = (C_0, C_1, \dots, C_{J-1})^T. \quad (2.12)$$

When we perform the matrix multiplication of Eq. 2.10 we are multiplying a matrix of shape (K, J) by a vector of length J , so we are left with a vector of length K . Written in this form, we can easily apply least-squares solvers to invert the equation for the vector of coefficients. This minimises the Euclidean 2-norm $\|\mathbf{h}_{\ell m} - \mathbf{a}_{\ell m} \cdot \mathbf{C}\|$ (i.e. the sum of the squares of the fit residuals). In this chapter we will only be dealing with single-mode fits as described here. However, for completeness, below we describe how multimode fits can be implemented (available in the code developed for this work [7]). We also note that the multimode-fit formalism is used in Fig. 4.8 to predict QNM amplitudes.

2.3.2 Multimode fit

Due to mode mixing, a given QNM contributes to all spherical-harmonic modes with the same m . This means we can perform a fit to multiple spherical-harmonic modes with a single set of shared QNM amplitudes \mathbf{C} . We will approach this by joining $\mathbf{h}_{\ell m}$ vectors together for each (ℓ, m) we want to include in the fit (to effectively form a single time series), and similarly by “stacking” $\mathbf{a}_{\ell m}$ matrices on top of each other. The single-mode matrix equation, Eq. 2.10, becomes

$$\mathbf{h} = \mathbf{a} \cdot \mathbf{C}, \quad (2.13)$$

where

$$\mathbf{h} = \begin{bmatrix} \mathbf{h}_0 & \mathbf{h}_1 & \cdots & \mathbf{h}_{I-1} \end{bmatrix} \quad (2.14)$$

and I is the number of spherical-harmonic modes to include in the fit (we have suppressed the two spherical-harmonic indices for clarity). So, \mathbf{h} is a vector of length $I \times K$. Similarly

$$\mathbf{a} = \begin{bmatrix} \mathbf{a}_0 \\ \mathbf{a}_1 \\ \vdots \\ \mathbf{a}_{I-1} \end{bmatrix}, \quad (2.15)$$

which has shape $(I \times K, J)$. We see that when we multiply this new coefficient matrix by the vector \mathbf{C} (length J) in Eq. 2.13 we recover a vector of correct length $I \times K$. The quantity we’re minimising is now

$$\|\mathbf{h} - \mathbf{a} \cdot \mathbf{C}\| = \sqrt{\sum_{\ell m} \|\mathbf{h}_{\ell m} - \mathbf{a}_{\ell m} \cdot \mathbf{C}\|^2}. \quad (2.16)$$

2.4 Aligned-spin systems

Following Giesler et al. [122], the spherical-harmonic modes of the ringdown signal can be modelled by writing each as a sum of N overtones:

$$h_{\ell m}^N(t) = \sum_{n=0}^N C_{\ell mn} e^{-i\omega_{\ell mn}(t-t_0)}, \quad \text{for } t \geq t_0. \quad (2.17)$$

This *overtone* model is a restriction of the sum in Eq. 2.9, where overlaps between different harmonic ℓ indices (mode mixing) as well as mirror modes are neglected. As in Ref. [122], we model each spherical-harmonic mode individually as a sum of QNMs. In Ref. [122], the efficacy of this model for $\ell = m = 2$ was demonstrated by performing least squares fits to the h_{22} mode for a selection of aligned-spin SXS simulations. The authors note that this was also verified for other values of (ℓ, m) .

The overtone model in Eq. 2.17 contains $2(N + 1)$ free parameters in the complex amplitudes, $C_{\ell mn}$, plus the two parameters, M_f and χ_f , that determine the $\omega_{\ell mn}$ frequencies. All of these parameters depend on the properties of the progenitor binary, but we do not study these dependencies here.

Our fitting algorithm finds the amplitudes $C_{\ell mn}$ that minimise the sum-of-the-squares of the fit residuals. We find it convenient to treat the remnant property parameters M_f and χ_f differently from the excitation amplitudes. If we also want to minimise over the remnant properties (as opposed to fixing them to the true values given by NR) then first a discrete 2-dimensional numerical grid of values for M_f and χ_f is constructed. At each point on this grid, we consider varying only the complex amplitudes $C_{\ell mn}$. Eq. 2.10 turns this minimisation problem into a linear algebra problem that can be efficiently solved with, for example, `numpy.linalg.lstsq` [143]. Finally, the point of the grid with the lowest overall value for the sum-of-the-squares of the fit residuals is chosen.

Once the least-squares fit to the data has been obtained, the quality of the fit is quantified via the mismatch and the error on the remnant parameters. The mismatch between signals h_1 and h_2 is defined as

$$\mathcal{M} = 1 - \frac{\text{Re}[\langle h_1|h_2 \rangle]}{\sqrt{\langle h_1|h_1 \rangle \langle h_2|h_2 \rangle}}, \quad (2.18)$$

where we use the following complex inner product [76]

$$\langle h_1|h_2 \rangle = \int_{t_0}^T h_1(t)h_2^*(t) dt. \quad (2.19)$$

We integrate from the ringdown start time, t_0 , to an upper limit T chosen such that the whole ringdown is captured (we use $T = t_0 + 100M$). When fitting models with very small mismatches, the finite accuracy of the NR simulations must be considered; this is discussed in Section 2.7. As noted in Ref. [122], a small mismatch is not sufficient by itself to justify the model. The overtone model contains more parameters as N is increased, and it is necessary to check for over-fitting. To address this, we check to see if the remnant BH properties are correctly recovered by the model. The combined error on the remnant mass and spin is quantified by [122]

$$\epsilon = \sqrt{\left(\frac{\delta M_f}{M}\right)^2 + (\delta\chi_f)^2}, \quad (2.20)$$

where $\delta M_f = M_{\text{bestfit}} - M_f$, and $\delta\chi_f = \chi_{\text{bestfit}} - \chi_f$. The best-fit values are those which minimise the mismatch, while the true values are taken from the metadata for the SXS simulation. A ringdown model can be said to perform well if it yields small values for both \mathcal{M} and ϵ .

Following Ref. [122], we now apply these ideas to the simulation SXS:BBH:0305 [36].




FIGURE 2.1: Mismatch as a function of ringdown start time for the overtone model (Eq. 2.17) when fitting to the h_{22} mode of the NR simulation SXS:BBH:0305. When using only the fundamental $\ell = m = 2$, $n = 0$ QNM the start time that gives the lowest mismatch with the NR data is well after the merger (the rising mismatch at late times is a numerical artefact). However, reproducing the results from Ref. [122], we find that by including $N = 7$ overtones the GW signal can be fitted using QNMs starting from as early as the peak strain. We also show (in light grey) the mismatch curves obtained when including up to 20 overtones. The dashed grey curve shows the estimate of the error in the underlying NR simulation and is described in Section 2.7.

This simulation has source parameters consistent with GW150914 and was originally chosen to demonstrate the success of the overtone model. Fig. 2.1 shows the mismatch values obtained with the overtone model when using the true values of M_f and χ_f . With $N = 7$ (that is, eight QNMs = the fundamental mode + seven overtones) the h_{22} mode can be fitted all the way back to the time of its peak amplitude, $t_{\text{peak}}^{h_{22}}$, while still achieving the smallest possible mismatch. Using a smaller number of overtones requires a later choice for the start time to achieve the smallest possible mismatch. A larger number of overtones can also be used (shown with the light grey lines in the figure), and low-mismatch fits as early as $\sim 10M$ before the time of peak amplitude can be achieved. We show mismatch curves with up to 20 overtones, by which point the inclusion of extra modes does not significantly help. Ref. [122] refer to the values of the QNM amplitudes to justify stopping




FIGURE 2.2: Recovery of the SXS:BBH:0305 remnant properties when fitting the overtone model (Eq. 2.17) to the h_{22} mode from the time of its peak strain. The heat map shows the mismatch for the fit with $N = 7$ overtones, which shows a pronounced minimum close ($\epsilon = 3.4 \times 10^{-4}$) to the true remnant parameters (indicated by the horizontal and vertical lines). The sequence of crosses shows the locations of the minima for fits performed with different values of N , all using the same start time (the cross colours correspond to the colours used in Fig. 2.1; crosses for $N = 5$ and 6 are omitted to avoid crowding the plot, but they converge towards the true remnant parameters). If we choose a different ringdown start time for each N corresponding to the mismatch minima in Fig. 2.1, we do see a reduction in ϵ for the lower N models, however the $N = 7$ model with $t_0 = t_{\text{peak}}^{h_{22}}$ remains the best performing model.

at $N = 7$; when performing fits at the time of peak strain the $n = 4$ mode has the largest amplitude, with the amplitude of higher overtones decaying rapidly. We perform a brief study of the overtone amplitudes in SXS:BBH:0305 (see Fig. 2.3), and we also refer the reader to Ref. [129] for a more in-depth study (and for the $(2, 2, 8)$ QNM frequency data used in Fig. 2.3).

In addition to giving a small ($\sim 10^{-6}$) mismatch, the $N = 7$ overtone model, with $t_0 = t_{\text{peak}}^{h_{22}}$, also achieves this minimum mismatch with the correct values for the remnant properties; this is shown by the heat map in Fig. 2.2 where the values of M_f and χ_f are now allowed to vary. We find, for the $N = 7$ model, a remnant error $\epsilon = 3.4 \times 10^{-4}$. Importantly, this is larger than the NR error on the remnant properties (which is estimated

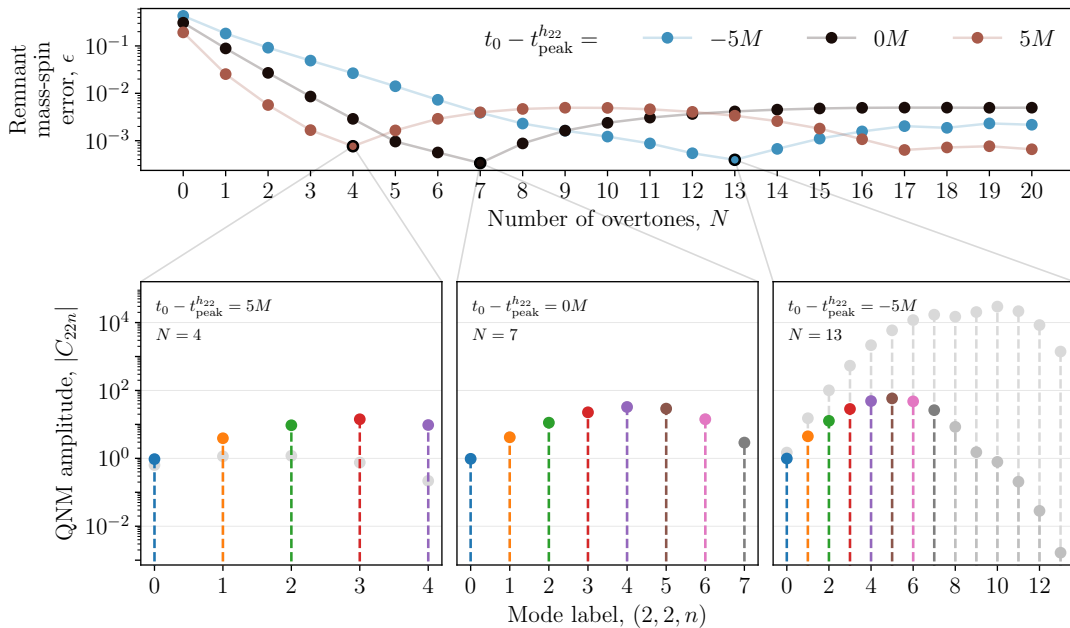


FIGURE 2.3: *Top*: The remnant mass-spin error (ϵ) from an overtone model fit to SXS:BBH:0305, for different numbers of overtones in the model (N) and for three different ringdown start times (line colours). For each start time there is a choice of N which gives a minimum in the curve (indicated with a black circle). *Bottom row*: The absolute value of each best-fit QNM amplitude from a fit at the minimum of the curve from the top panel (indicated by the connecting lines). We report the amplitudes rescaled to what they would be at the time $t_{\text{peak}}^{h_{22}}$. For the left and right panels, where the fit is performed $5M$ after and before that time respectively, the unscaled amplitudes are shown in light grey.

to be $\epsilon_{\text{NR}} = 2.1 \times 10^{-5}$, see Section 2.7 for details). This confirms that this is really the true scale of the bias in the inferred remnant parameters when using the overtone model, and not just the numerical noise floor in the NR simulation. Using a smaller number of overtones and starting the ringdown as early as $t_{\text{peak}}^{h_{22}}$ gives inferior results with the minimum in the mismatch being biased away from the true parameters. The results in Figs. 2.1 and 2.2 show that the overtone model performs well for SXS:BBH:0305 (i.e. yields small \mathcal{M} and ϵ) even when starting the ringdown as early as the peak in the strain.

In the top panel of Fig. 2.3 we show the values of ϵ obtained with different numbers of overtones (up to $N = 20$) and for three different ringdown start times (at the time of the

h_{22} amplitude peak, and $5M$ before/after the peak time). We see that, when $t_0 = h_{\text{peak}}^{h_{22}}$ (black line), $N = 7$ overtones does the best job at recovering the remnant properties. As expected, lower numbers of overtones perform worse (this is already shown in Fig. 2.2 by the coloured markers). Interestingly, larger numbers of overtones also perform worse. We speculate that this is due to over-fitting; as shown in Fig. 2.1, using more than seven overtones at this start time does not improve the mismatch significantly, and so the extra free parameters in the model are not necessary. We have found that using more QNMs than needed in the fits can lead to unstable behaviour (for example, in terms of the QNM amplitudes, which start to vary significantly as more overtones are added). Starting at a later time, $5M$ after the peak (red line), we see a similar behaviour but with $N = 4$ overtones giving the lowest ϵ (specifically, the first minimum in the ϵ vs N curve occurs at $N = 4$; the curve eventually reaches lower values of ϵ at high N , but this could be a non-physical result of over-fitting). This is consistent with what we see in Fig. 2.1, where the mismatch curve levels-out with four overtones when starting $\sim 5M$ after the peak (and the above argument regarding over-fitting with additional overtones still holds). Starting at the earlier time of $5M$ before the peak (blue line), we find that $N = 13$ overtones recovers a value of ϵ which is comparable to the result when seven overtones are fitted from the peak amplitude.

This hints at the interesting possibility of using even more overtones and starting the fit at even earlier times, but we note that care should be taken. Firstly, as shown in Fig. 2.1, as we use more overtones we are reaching mismatches further below the estimated waveform accuracy (dashed grey line) and so we are at risk of over-fitting. And secondly, we see some instability in the overtone amplitudes (which can be an indicator of whether these modes are physical or not). We demonstrate this with the bottom three panels of Fig. 2.3, which show the QNM amplitudes obtained with the lowest- ϵ fit at each of the

three start times. Being exponentially decaying modes, the value of the amplitude obtained in the fit depends strongly on the value of t_0 used. However, we know the expected decay time $\tau_{\ell mn}$ of each mode, and so to make a fair comparison of mode amplitudes we rescale them to a reference time (we perform the same procedure in Chapter 4 for the overtone amplitude, see Fig. 4.4). We choose to show the amplitudes rescaled to their value at the time of peak h_{22} strain, such that the amplitudes in the middle panel are unchanged. For the other two panels, in light-grey we show the unscaled amplitude values obtained from the least-squares fit. If the overtones were physical, we would expect the recovered amplitude to be stable with different choices of start time. Indeed, there is good agreement for the amplitudes up to and including the $n = 3$ mode. But, this agreement is less clear for the higher overtones. For example, the recovered amplitude of the $(2, 2, 7)$ mode is a factor of ~ 9 larger when performing the fit $5M$ before the peak (right panel) vs at the time of the peak (middle panel). See Ref. [129] for a more in-depth study of going beyond the $n = 7$ mode, where a selection of different NR simulations were also considered. Their conclusions are broadly in agreement with what is shown here; i.e., a larger number of overtones can be shown to fit the data well and recover the remnant properties, but the mode amplitudes show some instability. For simplicity, in the rest of this chapter we limit ourselves to seven overtones, which also aids comparison with previous work.

In order to see how robust the conclusions drawn from SXS:BBH:0305 are in general, the calculations of ϵ and \mathcal{M} were repeated for a wider selection of SXS simulations. Following Ref. [122], we consider only aligned-spin simulations with initial spin magnitudes $|\chi_{1,2}| = \chi_{1,2} < 0.8$, and mass ratios $q < 8$. We also require that the z -component of χ_f is greater than zero, which eliminates the “spin flip” systems. The simulations were chosen in the ID range SXS:BBH:1412 to SXS:BBH:1513, as these cover a range of initial spin magnitudes and mass ratios. After applying these cuts, this left 85 spin-aligned SXS

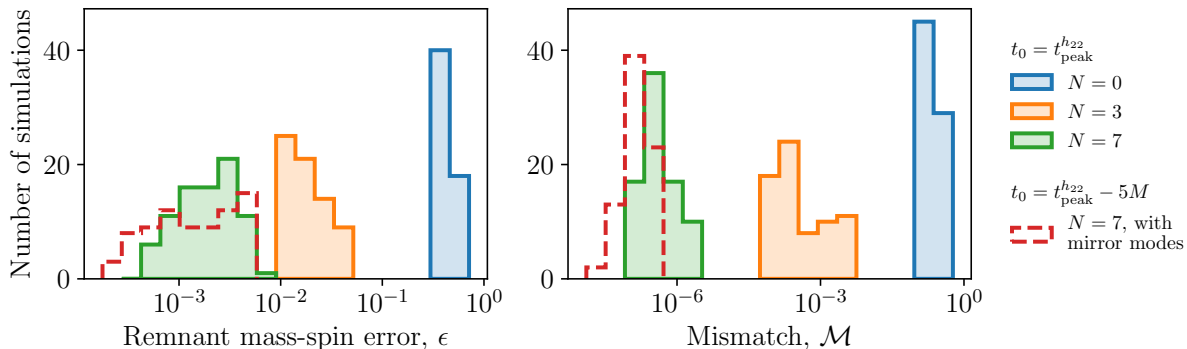


FIGURE 2.4: *Left*: histograms of the mass-spin remnant error ϵ from an overtone-model fit to 85 aligned-spin SXS simulations for several different overtone numbers N . *Right*: histograms of the mismatch from a fit with the true remnant mass and spin parameters, with the same overtone models and SXS simulations as in the left histogram. The solid histograms show results from fits performed starting at the peak of the h_{22} mode with N overtones of the fundamental $\ell = m = 2$ mode. The red dashed line shows results from a $N = 7$ model that also includes mirror modes (see Section 2.4.1) and was fitted with a ringdown starting $5M$ before the peak of the strain.

simulations in our test set. For each simulation, fits were performed using the overtone model with $N = 0, 3,$ and 7 and with a start time of $t_0 = t_{\text{peak}}^{h_{22}}$. The results are shown in Fig. 2.4. We see distributions similar to those in Fig. 3 of Ref. [122]. The inclusion of additional overtones systematically shifts the entirety of both the ϵ and \mathcal{M} histograms to smaller values. We note, as it will become important later, that the worst cases in these histograms improve, along with the median values. This demonstrates that, when using the overtone model on systems with aligned spins, the ringdown reliably starts as early as the peak in the h_{22} mode of the strain.

2.4.1 Mirror modes

For a given ℓ, m and n , the equations governing QNM frequencies allow two solutions: one, $\omega_{\ell mn} = 2\pi f_{\ell mn} - i/\tau_{\ell mn}$, with a positive real part; and another, $\omega'_{\ell mn} = 2\pi f'_{\ell mn} - i/\tau'_{\ell mn}$, with negative real part [92, 126]. The frequencies of the mirror modes $\omega'_{\ell mn}$ are related to the regular modes $\omega_{\ell mn}$ by Eq. 1.69.

A new ringdown model which explicitly includes the mirror modes can be written as

$$h_{\ell m}^{N, \text{mirror}}(t) = \sum_{n=0}^N \left[C_{\ell mn} e^{-i\omega_{\ell mn}(t-t_0)} + C'_{\ell mn} e^{-i\omega'_{\ell mn}(t-t_0)} \right] \quad \text{for } t \geq t_0. \quad (2.21)$$

This *mirror-mode* model is an extension of the overtone model in Eq. 2.17; if $C'_{\ell mn} = 0$ the mirror modes aren't excited and we recover the previous overtone model. This model has twice as many free parameters as the overtone model; $4(N+1)$ in the complex amplitudes, plus the two remnant parameters M_f, χ_f . The mirror-mode model is still a restriction of the full sum in Eq. 2.9 as overlaps between modes with different ℓ indices (i.e. mode mixing) are still not included. Substituting for $\omega'_{\ell mn}$ using the conjugate symmetry property in Eq. 1.69, we can rewrite the mirror-mode model in the form

$$h_{\ell m}^{N, \text{mirror}}(t) = \sum_{n=0}^N \left[C_{\ell mn} e^{-i\omega_{\ell mn}(t-t_0)} + C'_{\ell mn} e^{i\omega_{\ell - mn}^*(t-t_0)} \right] \quad \text{for } t \geq t_0. \quad (2.22)$$

This is how the model was implemented in practice.

As was shown by Dhani [126], the inclusion of mirror modes can improve the ringdown modelling of aligned-spin systems. In particular, the ringdown can be considered to start even earlier in the waveform, whilst still recovering the correct remnant properties. We confirm this here by repeating the above analysis for the same set of spin-aligned SXS simulations, but now using the mirror-mode model in Eq. 2.22 with $N = 7$ and an earlier choice for the ringdown start time, $t_0 = t_{\text{peak}}^{h_{22}} - 5M$. Although Ref. [126] demonstrated the mirror-mode model starting $10M$ before the peak in the h_{22} strain, we adopt a more conservative choice of $5M$. The results are shown in Fig. 2.4 using a dashed line. The addition of mirror modes gives a small improvement in the mismatch, but this is to be expected with the increased number of parameters. However, the ϵ histogram shows that the overall performance of the mirror-mode model is comparable to that of the $N = 7$

overtone model, despite the use of an earlier start time.

2.5 Misaligned-spin systems

The analyses in Section 2.4, and analyses in previous studies, was limited to BBH systems with component spins that are aligned with the orbital angular momentum, \mathbf{L} . This is a potentially serious limitation as misaligned spins are expected to be a generic feature of astrophysical BBHs. Misaligned spins generally lead to precession of the orbital plane during the inspiral phase of the evolution and a richer phenomenology in the GW signals [142]. There is strong evidence for precession in the GW events observed so far when looking at the population level [144, 145], and tentative evidence in individual events. This includes high-mass event GW190521 [39, 146], the asymmetric-mass event GW190412 [147], and recently there have been claims of precession in the GWTC-3 event GW200129_065458 [148] (but we note that there are potential data-quality issues associated with this event [149]). In this section we investigate the effect of precession on the modelling of the ringdown by repeating analyses like those in Section 2.4, but now on precessing NR simulations.

As already mentioned, the remnant BH will not have a spin vector aligned with the z -axis in the NR frame if it has undergone precession. The ringdown models we have been using are only valid in the frame with the remnant BH spin vector aligned with the z -axis. Therefore, before applying the ringdown models to precessing NR data, we need to rotate the NR data into the suitable frame.

The direction from which a GW source is viewed affects the observed signal (e.g. you see circularly/linearly polarised GWs with a larger/smaller amplitude when viewing parallel/perpendicular to \mathbf{L}). These differences in the GW signals also manifest themselves

at the level of individual spherical-harmonic modes as amplitude modulations. The frame in which the expansion (Eq. 2.1) is performed affects the values of the spherical-harmonic modes. The idea is to re-expand the NR data in the “ringdown frame” as follows:

$$h'(t, \Omega') = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h'_{\ell m}(t) {}_{-2}Y_{\ell m}(\Omega'), \quad (2.23)$$

where the prime on Ω indicates we are using new coordinates where the remnant spin is aligned with the z -axis. The coefficients in this expansion, $h'_{\ell m}$, are what we now fit our ringdown models to. In particular, we will focus on modelling the $\ell = m = 2$ spherical harmonic mode in the ringdown frame, h'_{22} .

For aligned-spin systems, the $h_{2\pm 2}$ are usually the dominant modes in the sum in Eq. 2.1. This is related to the fact that the GW signal amplitude is largest when viewed along the direction of the orbital angular momentum: \mathbf{L} or $-\mathbf{L}$. For misaligned-spin systems undergoing precession, other modes become important. This in turn is related to the constantly changing direction of the orbital angular momentum, $\mathbf{L}(t)$. Changing into the non-inertial, coprecessing frame in which \mathbf{L} always points along the z -direction has been found to account for most precessional effects and makes the precessing waveform remarkably similar to a non-precessing one. This transformation into the coprecessing frame has been successfully used to help model the full inspiral-merger-ringdown waveforms for precessing systems [150, 151] in the context of phenomenological [152–154], effective-one-body [155, 156] and NR surrogate [23, 157, 158] modelling. There is an analogy with the approach taken here for the modelling of the ringdown.

In order to simplify the task, we choose to work in a frame adapted to final spin angular momentum of the remnant, χ_f . So, in our case, the rotation required to get into this frame is not time dependant and our chosen frame is therefore inertial. This approach does not

account for changes to the ringdown frame throughout the duration of the ringdown; the system is still radiating angular momentum even after the time of peak strain, meaning the direction of the remnant spin vector will move throughout the ringdown. Although this effect is small, it may become important for the mismatches achieved in this work; a typical ringdown mismatch between a waveform with the final spin vector aligned along the z -axis (the transformation used in this work) and a waveform with the spin at the time of peak strain aligned along z (a typical choice for the ringdown start time) is $\sim 10^{-4}$. That the system is still radiating angular momentum and energy throughout the ringdown also raises the question of whether applying a fixed QNM spectrum given by the final Kerr BH makes sense. It may be that the frequencies and damping times should be time dependant, perhaps in response to the changing mass and spin of the oscillating remnant BH. Preliminary studies suggest that a time-dependant spectrum obtained by integrating the energy and angular momentum of the system backwards does not improve the fits; at earlier times the frequency of the signal decreases, but the spin magnitude increases and results in higher QNM frequencies (the increasing mass at earlier times does act to decrease the QNM frequencies, which may help the fits). The approach we take here avoids these issues and addresses the most important aspect of the final BH frame, and we leave these other points to future work.

To obtain an expression for $h'_{\ell m}$ in terms of the $h_{\ell m}$ provided by NR, we invert Eq. 2.23 to obtain

$$h'_{\ell m}(t) = \int_{\Omega'} h'(t, \Omega') {}_{-2}Y_{\ell m}^*(\Omega') d\Omega'. \quad (2.24)$$

For spin-weighted fields there is a subtlety that $h'(t, \Omega') \neq h(t, \Omega)$, but instead

$$h'(t, \Omega') = h(t, \Omega) e^{-is\gamma} \quad (2.25)$$

where s is the spin weight ($s = -2$ in our case), and γ is some angle (for example, a rotation about the third degree of freedom we have when going into a new coordinate system). Substituting into our expression for $h'_{\ell m}$ we get

$$h'_{\ell m}(t) = \int_{\Omega'} h(t, \Omega) e^{2i\gamma} {}_{-2}Y_{\ell m}^*(\Omega') d\Omega'. \quad (2.26)$$

Under a rotation, a spin-weighted spherical harmonic transforms into a linear combination of spin-weighted spherical harmonics of the same ℓ but different m [159]:

$${}_{-2}Y_{\ell m}(\Omega') = \sum_{m'=-\ell}^{\ell} [D_{m'm}^{\ell}(\mathbf{R})]^* {}_{-2}Y_{\ell m'}(\Omega) e^{2i\gamma} \quad (2.27)$$

where $D_{m'm}^{\ell}(\mathbf{R})$ is the Wigner D matrix for a rotation \mathbf{R} of the basis. Substituting into the expression for $h'_{\ell m}$ we get

$$\begin{aligned} h'_{\ell m}(t) &= \int_{\Omega'} h(t, \Omega) e^{2i\gamma} \left[\sum_{m'=-\ell}^{\ell} [D_{m'm}^{\ell}(\mathbf{R})]^* {}_{-2}Y_{\ell m'}(\Omega) e^{2i\gamma} \right]^* d\Omega' \\ &= \sum_{m'=-\ell}^{\ell} D_{m'm}^{\ell}(\mathbf{R}) \int_{\Omega'} h(t, \Omega) {}_{-2}Y_{\ell m'}^*(\Omega) d\Omega' \\ &= \sum_{m'=-\ell}^{\ell} D_{m'm}^{\ell}(\mathbf{R}) h_{\ell m'}(t), \end{aligned} \quad (2.28)$$

where the γ terms have cancelled (note that it doesn't matter the final integral is over the primed coordinates, as we're integrating over the full sphere).

This tells us how we can express the more natural $h'_{\ell m}$ modes (with coordinates suited to the remnant BH) in terms of the SXS $h_{\ell m}$ modes. Each $h'_{\ell m}$ mode is a superposition of $h_{\ell m}$ modes with the same ℓ but different m . The rotation \mathbf{R} can be obtained from the direction of the remnant BH spin vector (which is provided as metadata for all SXS




FIGURE 2.5: Histogram of the differences between the two possible start times considered in Section 2.5: the peak of the (rotated) strain mode h'_{22} , and the peak of the GW energy flux. The normalised distribution of the differences between these times is shown both for the 85 aligned-spin systems used in Section 2.4 and for the 252 precessing simulations considered in Section 2.5. The peak of the flux almost always occurs later than the peak of the strain, making this a more conservative choice for the ringdown start time. We note that there is a much greater variation amongst the population with misaligned spins.

simulations). Specifically, \mathbf{R} is any rotation that maps the z -axis onto the final spin vector.

We now apply the overtone model to the ringdown of an example precessing simulation SXS:BBH:1856 [23]. This simulation (at the reference time) has a mass ratio of $q = 2.78$ and dimensionless spins $\chi_1 = (0.18, -0.54, -0.45)$ and $\chi_2 = (-0.12, -0.31, -0.031)$ on the heavier and lighter components respectively. This simulation was chosen because it exhibits strong precession effects visible as amplitude modulations in $h_{22}(t)$. The final spin vector is $\chi_f = (-0.03, -0.19, 0.42)$ and the rotated mode $h'_{22}(t)$ was computed using Eq. 2.28.

The overtone model in Eq. 2.17 was fitted to the rotated $\ell = m = 2$ mode of the strain, $h'_{22}(t)$, in the same way as was done for the aligned-spin systems in Section 2.4. There is some ambiguity in how to choose the ringdown start time t_0 in a way that gives as fair a comparison as possible with the non-precessing case. We cannot use the peak of the $h_{22}(t)$




FIGURE 2.6: Mismatch as a function of ringdown start time for the overtone model (Eq. 2.17) when fitting to the rotated h'_{22} mode of the NR simulation SXS:BBH:1856. When using $N = 7$ overtones, the lowest mismatch is achieved starting slightly ($\sim 10M$) before the peak in the GW energy flux. However, the minimum mismatch is ~ 100 times larger than that obtained for the example spin-aligned system SXS:BBH:0305 in Fig. 2.1. The dashed grey curve shows the estimate of the error in the underlying NR simulation and is described in Section 2.7.

strain, as was done in Section 2.4, as this mode suffers from precession induced amplitude modulations. One option would be to use instead the peak of the rotated strain mode $h'_{22}(t)$. However, we find that using the peak of the GW energy flux, \dot{E} , (which can be computed from the modes in either frame, see Eq. 3.8 in Ref. [160]) gives more consistent results between simulations. For example, some precessing configurations show a peak in the (rotated) strain relatively early in the signal, leading to poorer fits. The use of the peak in the energy flux is also a conservative choice in the sense that $t_{\text{peak}}^{\dot{E}} > t_{\text{peak}}^{h'_{22}}$ in almost all cases (see Fig. 2.5).

Fig. 2.6 shows how the mismatch varies for SXS:BBH:1856 as a function of ringdown start time, for different values of N in the overtone model Eq. 2.17. With each additional overtone, the minimum mismatch is reached at an earlier time (the same behaviour as was seen in Fig. 2.1). However, the values of the minimum mismatch are a factor of ~ 100 larger than those obtained in the aligned-spin case.




FIGURE 2.7: Recovery of the SXS:BBH:1856 remnant properties when fitting the overtone model (Eq. 2.17) to the rotated h'_{22} mode from the time of its peak energy flux. The heat map shows the mismatch for the fit with $N = 7$, while the crosses show the locations of the minima in the mismatch for fits performed with different values of N . The mismatch shows a much broader and less deep minimum than that seen for the spin-aligned system SXS:BBH:0305 in Fig. 2.2. The minimum in the mismatch is also biased away from the true remnant parameters with $\epsilon = 0.025$ for the $N = 7$ fit. The sequence of crosses for fits with different values of N also do not show the same convergent trend towards the true remnant parameters that was observed for SXS:BBH:0305 in Fig. 2.2.

The $N = 7$ model achieves a minimum mismatch $\sim 10M$ before the time of peak GW energy flux. This is fairly typical behaviour among the misaligned-spin SXS simulations considered. However, we note there is a much greater variety of possible behaviours for misaligned-spin systems than for the aligned-spin population. The greater variation amongst the misaligned-spin population has already been hinted at in Fig. 2.5, where the spread of start times is greater than in the aligned-spin cases.

The heat map of Fig. 2.7 shows the mismatch as a function of the remnant BH properties, for the $N = 7$ model. The coloured crosses indicate the mismatch minimum for different values of N . Comparing with Fig. 2.2, we see the mismatch minimum is less pronounced than the aligned-spin case, which is probably contributing to the larger value of ϵ (for $N = 7$ we find $\epsilon = 0.025$, which is much larger than the estimated numerical




FIGURE 2.8: *Left*: histograms of the mass-spin remnant error ϵ from an overtone model fit to the rotated h'_{22} modes of 252 misaligned-spin SXS simulations for several different overtone numbers N . *Right*: histograms of the mismatch from a fit with the true remnant mass and spin parameters, with the same overtone models and SXS simulations as in the left histogram. The solid histograms show results from fits performed starting at the peak of the energy flux with N overtones of the fundamental $\ell = m = 2$ mode. The red dashed line shows results from a $N = 7$ model that also includes mirror modes and was fitted with a ringdown starting $5M$ before the peak in the energy flux. These histograms should be compared with those in Fig. 2.4; we note that the effect of precession is to (i) significantly broaden the histograms (i.e. the quality of the fit is much more varied) and (ii) to significantly degrade the quality of the fit for some systems.

error $\epsilon_{\text{NR}} = 8.6 \times 10^{-5}$). In addition, the convergent behaviour with increasing N is not present. For $N \geq 1$, all mismatch minima appear randomly distributed around the true remnant properties. If we reproduce this figure with a earlier start time of $t_0 = t_{\text{peak}}^{\dot{E}} - 10M$ (motivated by the time of minimum mismatch for $N = 7$ in Fig. 2.6), the heat map remains unchanged, and the value of ϵ recovered for $N = 7$ is not significantly improved ($\epsilon = 0.013$). The earlier start time does cause the value of ϵ for $N \leq 3$ to increase significantly, which may be expected as we are now using a start time before those models reach a mismatch minimum.

Following Section 2.4, we now extend this analysis to a wider selection of SXS simulations to investigate the robustness (or lack thereof) of this behaviour. We consider only misaligned-spin simulations, chosen such that the angle between the initial spins, χ_θ , satisfies $\pi/16 < \chi_\theta < 15\pi/16$. We again require initial spin magnitudes $\chi_{1,2} < 0.8$ and

mass ratios $q < 8$. The 252 simulations were chosen in the ID range SXS:BBH:1643 to SXS:BBH:1899, as these cover a range of mass ratios and initial spin configurations.

The results are shown in Fig. 2.8. When compared to the $N = 0$ model, the addition of three overtones reduces the remnant error and mismatch. However, the inclusion of additional overtones does not change the ϵ histogram, and produces only a minor reduction in the mismatch. Comparing the $N = 7$ histogram for ϵ to that found in Fig. 2.4, we see that, on average, ϵ increases by a factor of ~ 10 and, in the worst cases, by a factor of ~ 20 (however, the overtone model does still perform similarly well for a small fraction of simulations). The histograms for ϵ reflect the behaviour of Fig. 2.7, where models with $N \geq 1$ don't show systematic improvements. It would be interesting to investigate whether the binary parameters correlate with ϵ , and if certain binary configurations are responsible for the largest remnant errors. We have performed preliminary studies which reveal no clear correlations of ϵ with either the amount of precession (quantified via χ_p [161]) or the recoil velocity. We defer a more detailed study of this question to future work.

It was checked if using an earlier start time of $t_{\text{peak}}^{\dot{E}} - 10M$ changed the recovered distribution on ϵ . This choice was motivated by the location of the mismatch minimum typically seen for misaligned-spin simulations (e.g. see Fig. 2.6). It was found the $N = 7$ model results did not significantly change. However, the $N = 3$ and $N = 0$ models performed worse. Finally, we also note that all of the histograms are wider than those in Fig. 2.4. This may be due to mirror modes and/or higher harmonics having a more important role for precessing systems (see below).

2.5.1 Mirror modes

We repeat the population analysis with the $N = 7$ mirror-mode model, again shifting the ringdown start time back by $5M$ to make a clear comparison to Fig. 2.4. The results are




FIGURE 2.9: Mismatch as a function of ringdown start time for the mirror-mode model (Eq. 2.22) when fitting to the rotated h'_{22} mode of the NR simulation SXS:BBH:1856. Comparing with Fig. 2.6, the locations of the mismatch minima are roughly unchanged in time, but the inclusion of mirror modes reduces the mismatch to values similar to those in Fig. 2.1. The dashed grey curve shows the estimate of the error in the underlying NR simulation and is described in Section 2.7.

shown by the red dashed lines in Fig. 2.8. The histogram for ϵ doesn't reach values as high as the overtone model (with worst-case values of $\epsilon \sim 0.04$ compared to the overtone model's ~ 0.2), but otherwise has a broadly similar distribution. However, there is a significant improvement on the recovered mismatch values. This is expected because of the large number of parameters. And, as discussed, this alone isn't enough to say the model is successful.

Inspecting individual simulations, we see that the inclusion of mirror modes can make the mismatch minima in the mass-spin plane more pronounced (advantageous, as it reduces uncertainty on ϵ). For example, Figs. 2.9 and 2.10 show how mirror mode fits perform for SXS:BBH:1856. We see significantly smaller mismatches, and a stronger mismatch peak around the true remnant properties. However, on average this does not translate to smaller values of ϵ for the $N = 7$ model (as can be seen from the red dashed histogram in Fig. 2.8). For SXS:BBH:1856, the $N = 7$ model gives $\epsilon = 0.014$, which is not a significant

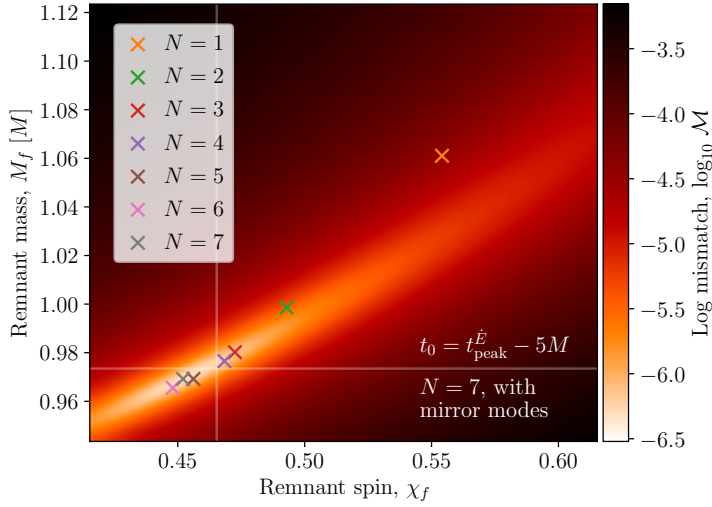


FIGURE 2.10: Recovery of the SXS:BBH:1856 remnant properties when fitting the mirror-mode model (Eq. 2.22) to the rotated h'_{22} mode from $5M$ before the time of its peak energy flux. The heat map shows the mismatch for the fit with $N = 7$, while the crosses show the locations of the minima in the mismatch for fits performed with different values of N ($N = 0$ lies outside the figure, and is not included for clarity). Comparing with Fig. 2.7, the inclusion of mirror modes sharpens the mismatch peak and achieves smaller mismatch values. However, when averaged across the population of precessing simulations, the mirror-mode model doesn't give smaller values for the remnant error (see dashed curve in Fig. 2.8). Here, $\epsilon = 0.014$ for the $N = 7$ model.

improvement.

To investigate whether the choice of ringdown start time could be contributing to the wider histograms seen in Fig. 2.8, the behaviour of the mismatch heat maps (e.g. Figs. 2.2, 2.7, 2.10) with varying start time was explored for selected SXS simulations. Animations of ringdown fits with varying start time can be found at Ref. [4]. For the aligned-spin simulation SXS:BBH:0305, we see that the location of the mismatch minimum in the mass-spin plane settles on the true remnant properties for a sufficiently late choice of the start time ($t_0 \geq t_{\text{peak}}^{h_{22}}$ for the $N = 7$ overtone model). In addition, the mismatch minimum stays centred on the true remnant properties until numerical noise takes over. For earlier choices of the start time, the $N = 7$ overtone model gives biased values for the final mass and spin, see Ref. [4]. Applying the $N = 7$ overtone model to the misaligned-spin

simulation SXS:BBH:1856, we see that the location of the mismatch minimum moves around the mass-spin plane as start time is varied. Even at late times, it never settles on the true remnant properties. The inclusion of mirror modes, as seen in Fig. 2.10, narrows the mismatch minimum. The movement of the mismatch minimum around the mass-spin plane is reduced as well, however it still doesn't settle on the location of the true remnant properties. This behaviour may explain some of the observed widening of the histograms, and perhaps hints something is missing from the ringdown model.

2.5.2 Higher harmonics

As demonstrated by Fig. 2.8, the overtone and mirror-mode models considered so far achieve median values for the remnant error $\epsilon \sim 0.01$, a factor of 10 or more higher than the aligned-spin fits of Fig. 2.4. In addition, the spread of ϵ values recovered is significantly larger, leading to values of ϵ up to ~ 0.1 . These models perform significantly worse in some cases for precessing systems than aligned-spin systems.

We now investigate whether the inclusion of higher harmonics (that is, QNMs with $\ell > 2$) can improve the fits to $h'_{22}(t)$. These higher harmonics were neglected by both the overtone (Eq. 2.17) and mirror-mode (Eq. 2.22) models. However, mode mixing occurs as a consequence of the different angular basis functions used in the waveform decompositions in Eqs. 2.1 and 2.2 and the fact that these basis functions are not mutually orthogonal [64]. The amount of mode mixing between the spherical mode ${}_{-2}Y_{\ell m}$ and the spheroidal mode ${}_{-2}S_{\ell mn}$ is determined by the remnant spin χ_f and the QNM frequency. This can be quantified by how much these functions fail to be orthogonal; i.e. by the spherical-spheroidal mixing coefficients (Eq. 2.5). A translational offset between the NR and ringdown frames (e.g. due to a kick) can also lead to mixing between m -modes [162];




FIGURE 2.11: *Left:* histograms of the mass-spin remnant error ϵ from harmonic model fits (Eq. 2.29) to the same 252 misaligned-spin SXS simulations used in Fig. 2.8. Shown (in dashed lines) are the $L = 3$ and $L = 4$ models with $N = 7$ overtones and mirror modes. Shown in green is the overtone model with $N = 7$ and $L = 2$ (no mirror modes); this is the same as the green histogram in Fig. 2.8 and is included here to aid comparison. We also show for comparison the $L = 3$ model without mirror modes (grey histogram). *Right:* histograms of the mismatch from a fit with the true remnant mass and spin parameters, with the same models and SXS simulations as in the left histogram. The harmonic model with mirror modes, which includes many free parameters, achieves small mismatches but without significant improvement in the remnant error. We note that the inclusion of $L = 4$ does not bring any additional improvements over $L = 3$.

this effect, expected to be small for astrophysical kick velocities (but, like the time-dependant direction of the remnant spin, may become important for the low mismatches achieved with many-mode models), is neglected here. To include the contribution from higher harmonics, we define a new ringdown model for the spherical-harmonic modes which now allows for a sum over different ℓ :

$$h_{\ell m}^{N, L, \text{mirror}}(t) = \sum_{n=0}^N \sum_{l=2}^L [C_{l m n} e^{-i\omega_{l m n}(t-t_0)} + C'_{l m n} e^{i\omega_{l m n}^*(t-t_0)}] \quad \text{for } t \geq t_0. \quad (2.29)$$

This *harmonic* model contains all of the allowed QNMs in Eq. 2.2, including the mirror modes and the overtones. This comes at the expense of a large number of free parameters; there are $4(N+1)(L-\ell+1)$ in the complex amplitudes, plus the two remnant parameters M_f, χ_f that determine the complex QNM frequencies.

Multiple variations of this harmonic model were trialled (varying N, L , and the

inclusion of mirror modes) on the same population of 252 misaligned-spin SXS simulations. Fig. 2.11 shows the chosen subset of results. All results shown include seven overtones, and include $L = 2$ (16 free parameters in the complex amplitudes), $L = 3$ without and with mirror modes (32 and 64 free parameters respectively), and $L = 4$ with mirror modes (96 free parameters). As before, we fit to the rotated $h'_{22}(t)$ spherical harmonic mode. To make a clear comparison with the previous models, we again use a ringdown start time corresponding to the peak of the GW energy flux.

The inclusion of higher harmonics with the mirror modes drastically improves the mismatch. A small mismatch is not surprising for a model with so many free parameters, and in some of these cases we are likely pushing beyond the limits of accuracy of the NR simulations. See Section 2.7 for a discussion of the numerical errors. There is a modest reduction in ϵ for some systems, and in particular we see less systems with $\epsilon > 0.01$. This hints at the importance of higher harmonics in some precessing systems. Despite this, we still see worst-case values of $\epsilon \sim 0.04$. We also note that higher harmonics alone (i.e. without mirror modes, shown with the grey histograms) leads to a greater spread of mismatches, and more systems with $\epsilon > 0.01$.

2.6 Surrogates

NR simulations are computationally expensive, and although the number of simulations available in public catalogs is growing they are still limited in their parameter space coverage. NR surrogate models [23, 158, 163, 164] would appear to be an attractive alternative. These models use reduced-order and surrogate modelling techniques to extend the results of a set of NR simulations smoothly across parameter space. The use of surrogates could, in principle, allow us to extend the results of this chapter to include




FIGURE 2.12: Comparison of the remnant error ϵ from two surrogate models and a selection of SXS simulations. All are zero initial spin. The fits were performed on the h_{22} mode with the $N = 7$ overtone model, Eq. 2.17, starting from the time of peak strain. The labels on each cross correspond to the SXS ID. The dashed line indicates where we are outside the training range of NRSur7dq4.

many more systems as well as allowing us to study how the excitations of the various QNMs vary during a smooth exploration of parameter space. However, care must be taken as the surrogate modelling necessarily introduces an additional source of error into the waveforms, on top of the errors originally in the NR waveforms themselves.

When attempting to fit QNM ringdown models with overtones to NRSur7dq4 [23] waveforms, it was found that incorrect values for M_f and χ_f were being recovered (particularly at high mass ratios). This being the case even for aligned-spin or non-spinning systems. Although the NRSur7dq4 waveforms do not provide the remnant properties, these can be obtained via NRSur7dq4Remnant [23] (it was found the problem did not lie with the values returned by NRSur7dq4Remnant but rather with the waveform surrogate).

To investigate the performance of NRSur7dq4 ringdown waveforms, a series of simulations with zero initial spin with increasing mass ratio q from 1 to 6 were used. The $N = 7$ overtone model (Eq. 2.17) was fitted to the $h_{22}(t)$ mode of each starting from

the peak strain (as in Section 2.4) and the remnant error ϵ (Eq. 2.20) was calculated for each. The results are shown in Fig. 2.12, along with the results for similar fits performed directly on 11 zero-spin SXS simulations at discrete values of the mass ratio. The fits to the NRSur7dq4 surrogate produce values for ϵ that are 1-2 orders of magnitude higher than for the equivalent SXS simulations. Also shown are the results from a similar analysis with the more restrictive aligned-spin surrogate NRHybSur3dq8 [99, 164]; this was found to be in close agreement with the SXS simulations.

Residuals and mismatches can also be computed between surrogate and NR waveforms (taking care to align the waveforms in both time and phase). For SXS:BBH:0168, the $q = 3$, zero-spin simulation used in Fig. 2.12, we find $\sim 2\%$ residuals in the ringdown when comparing to the NRSur7dq4 surrogate with the same parameters. This leads to a mismatch between the surrogate and SXS:BBH:0168 of 3.7×10^{-4} , when integrating over the ringdown. For comparison, we have a $\sim 10^{-6}$ mismatch between the ringdown model Eq. (2.17) and the SXS simulation. The relatively high mismatch between the NRSur7dq4 and SXS waveforms translates to the relatively high values of ϵ seen in Fig. 2.12.

It seems that the high-dimensional precessing surrogate NRsur7dq4 is not yet sufficiently accurate in the ringdown for the purposes of QNM overtone studies that, by virtue of their large number of free parameters, fit the ringdown with very small mismatches. By contrast, the lower-dimensional aligned-spin surrogate NRHybSur3dq8 does appear to be sufficiently accurate for such studies.

2.7 Numerical relativity errors

It is important to remember the finite accuracy of the NR simulations used in ringdown studies. This is particularly true when using models with many QNMs which, by their very

nature, use a large number of free parameters and regularly achieve very small ($\sim 10^{-6}$) mismatches. If care is not taken, we risk fitting our models to the numerical noise. In this section we describe the numerical checks performed on the two individual simulations used in this chapter: SXS:BBH:0305, and SXS:BBH:1856. In each case the numerical errors were estimated by comparing results obtained using data from the two highest resolutions (levels) available in the SXS catalog.

First, we quantify the numerical error in the mismatch. This was done by calculating the mismatch between the two highest NR resolutions from a time t_0 to a time $T = t_0 + 100M$, for a range of t_0 . For each start time, we optimally align the two waveforms in time (taking the absolute value in the mismatch automatically optimises the mismatch over phase). The alignment in time can be done by matching the time of peak strain, for example, or by numerically rolling the waveform to find the optimal time shift for each mismatch calculation. The results are shown by the grey dashed lines in the mismatch vs start time plots in Figs. 2.1 and 2.6 (duplicated in Fig. 2.9). Generally, we see numerical error estimates at or below the model mismatches, particularly at late times, indicating that we are not fitting to the numerical noise. The main exception is Fig. 2.9 where the mirror-mode model is applied to a precessing system. This is expected; precessing NR simulations, and those with high mass ratios are generally expected to have larger numerical errors. Additionally, the mirror mode and harmonic models have the highest numbers of free parameters making them more likely to reach the accuracy of the NR simulation.

Second, we investigate the numerical error on the remnant mass and spin. We quantify the numerical error with ϵ_{NR} , the Euclidean distance (Eq. 2.20) between the remnant properties reported in the two highest resolution levels of the NR simulation. The ϵ_{NR} values are reported in the previous sections. In all cases $\epsilon_{\text{NR}} < \epsilon$. This supports the

conclusions in this chapter and indicates they are likely to be robust against numerical noise in the underlying NR simulations used.

2.8 Conclusions

This chapter has made a first systematic attempt at using QNMs to model the ringdown of BHs formed from BBHs with misaligned component spins in the inspiral. Previously, for aligned-spin systems, it has been found that the ringdown can be modelled with low mismatch and low remnant errors using a model that includes overtones of the fundamental QNM [122]. For seven overtones, the ringdown can be reliably modelled from the peak of the $h_{22}(t)$ strain for a range of SXS simulations. Additionally, the inclusion of mirror modes can allow the ringdown to be modelled from even earlier times [126]. In this chapter, which generalised these studies to precessing systems, we find that while QNM models can reliably achieve small mismatches, in the worst cases the remnant errors are more than a factor of 10 higher. This is the case even when choosing to start the ringdown at the more conservative (i.e. later) peak in GW energy flux. The inclusion of higher harmonics reduces the remnant error in some cases, perhaps a sign that mode mixing in the ringdown is generally more important in precessing systems. However, in other cases, a bias remains in the recovered remnant properties. We conclude that it is not possible to reliably model the ringdown from the peak in the flux, or indeed from the peak in the strain.

We end by sounding a brief note of caution to any who attempt to construct a QNM model starting at or before the peak flux or strain. While such a model will work in some cases, it risks biased results in others. This risk is subtle because QNM models can give small mismatches even when they fail to adequately describe the remnant.

Chapter 3

Frequency-domain Analysis of Black-hole Ringdowns

3.1 Introduction

A key motivation of the NR ringdown studies discussed in the previous chapter is to understand what to expect from nature. This is the subject of the next two chapters: testing our expectations with real GW data.

As discussed in the previous chapter, theoretical studies using catalogs of BBH simulations suggest the ringdown typically starts early in the merger process, i.e. at or even slightly before the time of peak strain amplitude. This is only possible if the ringdown modelling includes overtones ($n \geq 1$), and possibly a combination of mirror modes and/or higher harmonics ($\ell \geq 3$). This early start time is good for the prospects of observing QNMs because it means the signal amplitude is still large when the ringdown starts and consequently the SNR in the ringdown is large.

Once QNMs have been correctly identified in an observed signal, they provide an exciting opportunity for testing GR, the Kerr metric hypothesis, and the no-hair theorem. The idea of using QNM frequencies for such tests predates the detection of GWs, and

so experimental QNM tests of GR started immediately with the first GW observation. Despite the difficulties discussed in Section 1.3.3, detecting and characterising QNMs is a key goal in GW astronomy. GW data analysis is almost universally performed in the frequency domain, where it is most natural to characterise the noise. To satisfy the periodicity requirements of the Fourier transform and avoid spectral leakage, the data is windowed before and after the GW signal. However, to perform a ringdown-only analysis, a sharp cut must be applied in the time domain to isolate the ringdown signal. Applying any window functions would suppress the ringdown signal, or (if the analysis data is extended to earlier times) the inspiral-merger would contaminate the analysis. See Ref. [121] for a discussion. This problem can be solved by instead working in the time domain, where the ringdown can be isolated cleanly from a chosen start time, t_0 , and the noise is instead characterised by a non-diagonal covariance matrix. However, as the particular segment of data to be analysed is chosen before the analysis begins, the ringdown start time (and consequently the source sky location) usually must be fixed in these analyses. We note that it is possible to extend the analysis data such that the ringdown start time and sky position can be varied (as was done in Ref. [165]), but a ringdown-only model may be biased towards early start times due to the presence of an inspiral in the signal. Another drawback of working in the time domain is that, as the noise covariance matrix is no longer diagonal, the computational cost of the likelihood is increased. The covariance matrix is constructed with the autocovariance function, which characterises the noise in the time domain. And, as discussed in Ref. [121], care must be taken when estimating the autocovariance to avoid corrupting the ringdown data. Therefore, there are additional subtleties in a time-domain analysis compared to a frequency-domain analysis.

This time-domain approach is not the only way to formulate a robust ringdown analysis. Ref. [120] proposed an alternative approach which is expressed in the frequency domain,

but uses a modified form of the likelihood (involving “in-painting” the data before the start of the ringdown in such a way as to remove its contribution to the likelihood). However, it has been shown that this method is formally equivalent to the standard time-domain approach [121] and suffers from the same limitations. It should also be noted that Ref. [120] applied their method to GW190521, where a higher harmonic was identified in the ringdown. This finding is apparently in contradiction with the earlier results of Ref. [37], which use a time-domain approach. That these two formally equivalent analyses [37, 120] can come to different conclusions regarding the QNM content of GW190521 highlights some of the difficulties that come with this type of analysis, where important choices (that can affect the result) for the ringdown start time have to be made and care must be taken with the noise estimation.

Another approach [166] is to make use of the full GW signal, and perform the analysis with augmented effective-one-body [167] waveforms. By calibrating QNM amplitudes with NR simulations, QNM frequencies can be measured. An additional QNM “frequency-filter” method has also been proposed [168], which targets particular QNMs and removes them from the data. How consistent the remaining data is with noise can be used to infer remnant properties without requiring any knowledge of the QNM amplitudes, and to estimate when the linear ringdown regime begins. These approaches are a valuable way of investigating the QNM content of GW data. But, although advantageous in some ways, the lack of measurements on QNM amplitude means there is one less way to quantify the presence of the modes. The latter QNM filter method also requires a fixed sky location in order to align data from multiple detectors.

In this chapter, we present a new approach to performing ringdown analyses in the frequency domain. We employ a flexible sum of truncated wavelets to model the inspiral-merger signal (inspired by BAYESWAVE [169, 170]) and QNMs to model the ringdown. The




FIGURE 3.1: This figure is intended to illustrate the idea behind our approach for analysing the ringdown in the frequency domain. The grey line in the bottom panel shows the whitened and band-passed Livingston time-series data around GW190521. For the purpose of illustration, we use the simplest version of our model where a single, truncated sine-Gaussian wavelet is used to model the inspiral-merger part of the signal and the fundamental $\ell = m = 2, n = 0$ QNM is used to model the ringdown. The ringdown start time, t_0 , is allowed to vary as part of a Bayesian analysis which also searches over different values of wavelet parameters, the source sky position, the QNM amplitude and phase, and over the remnant BH mass and spin. Full details of this analysis will be presented elsewhere. The top panel shows the maximum-likelihood waveform broken down into its wavelet (h^{IM} , orange) and ringdown (h^{R} , purple) parts as well as into plus (solid) and cross (dashed) polarisations. The discontinuity in our model can be clearly seen in the coloured lines. However, when these polarisations are combined and projected (see Eq. 3.16) onto the interferometer (black line) the result is nearly continuous; we emphasise that this continuity has not been imposed by the model but is rather “learnt” from the data. When the projected maximum-likelihood waveform is whitened according to the detector noise curve it becomes completely continuous, this is plotted as the central blue line in the bottom panel. In the bottom panel we also plot the uncertainty (90% credible region; blue shaded band) on the recovered signal.

general idea behind our approach is illustrated in Fig. 3.1. By working in the frequency domain, we can use the standard, and now very mature, GW data analysis pipelines (our analysis pipeline borrows from the public BILBY package [171]). Also, working with a complete description of the signal makes it trivial to search and marginalise over the sky position the ringdown start time, t_0 . We hope this approach can complement the

existing analysis methods. The code associated with this work is available at Ref. [8], and a complete set of posterior samples from this work are made available at Ref. [5].

The details of our method are described in Section 3.2, where we compare and contrast the time- and frequency-domain likelihood functions before introducing our frequency-domain approach. In Section 3.3 we present the results of a series of analyses on simulated GW signals where we test the performance of our approach and compare it with time-domain methods. We present our conclusions in Section 3.5. A complete set of posterior samples from this work are made available at Ref. [5].

3.2 Methods

This section describes the details of the proposed frequency-domain approach to the analysis of BH ringdowns. Section 3.2.1 describes the GW likelihood function and its implementation in both the time and frequency domains. Section 3.2.2 motivates our frequency-domain approach by describing an extreme limit in which it becomes equivalent to the standard, time-domain approach. Finally, Section 3.2.3 describes the combination of truncated wavelets and QNMs that comprise our waveform model.

3.2.1 Time- and frequency-domain likelihoods

Most GW data analysis is done in the frequency domain because, with the usual assumptions of stationary zero-mean Gaussian noise, the instrumental noise is fully described by the (one-sided) noise power spectral density (PSD), $S_n(f)$. In the literature, the most commonly encountered expression for the log-likelihood in one interferometer is the integral

$$\log \mathcal{L}(d|\vec{\theta}) = -2 \int_0^\infty df \frac{|\tilde{d}(f) - \tilde{h}(f; \vec{\theta})|^2}{S_n(f)} + \text{norm}, \quad (3.1)$$

where $d(t)$ is the observed data, and $h(t; \vec{\theta})$ is the signal model (projected onto the interferometer) described by parameters $\vec{\theta}$. The normalisation constant in the likelihood is unimportant for our purposes and will be dropped in all following equations. A tilde denotes the Fourier transform of a time series. Because the noise is uncorrelated between two well separated interferometers, the log-likelihood for a network is obtained by summing the independent contributions from each instrument.

In practice, it is necessary to work with discretely sampled time series; $d_j = d(t_j)$, where $t_j = j\delta t$ for $j = 0, 1, \dots, J - 1$, and where $1/\delta t$ is the sampling frequency. The discrete Fourier transform $\tilde{d}_k = \tilde{d}(f_k)$ is sampled at (positive) frequencies $f_k = k/(J\delta t)$ for $k = 0, 1, \dots, K - 1$, where $K = \lfloor (J + 2)/2 \rfloor$. The log-likelihood in terms of the discretely sampled frequency series is given by the following sum,

$$\log \mathcal{L}(d|\vec{\theta}) = \frac{-2}{J\delta t} \sum_k \frac{|\tilde{d}_k - \tilde{h}_k(\vec{\theta})|^2}{S_n(f_k)}, \quad (3.2)$$

which can be compared to Eq. 3.1. The noise PSD is usually estimated from off-source data using a Welch periodogram [31]. The frequency-domain expression for the log-likelihood involves a single sum; the noise covariance matrix is diagonal in the frequency domain. Although, for finite duration time series there can be small correlations between frequency bins [172].

The log-likelihood can also be expressed in the time domain via

$$\log \mathcal{L}(d|\vec{\theta}) = -\frac{1}{2} \sum_{jj'} \left[d_j - h_j(\vec{\theta}) \right] C_{jj'}^{-1} \left[d_{j'} - h_{j'}(\vec{\theta}) \right], \quad (3.3)$$

where $C_{jj'}$ is the noise covariance matrix. The time-domain expression for the log-likelihood involves a double sum over a dense covariance matrix which is computationally more costly to evaluate than the frequency-domain expression [$\mathcal{O}(J^2)$ as opposed to $\mathcal{O}(J)$].

Because the noise is assumed to be stationary, $C_{jj'}$ has the Toeplitz structure

$$C_{jj'} = \rho_{|j-j'|}, \quad (3.4)$$

where ρ_j is the noise autocovariance. This can also be estimated from off-source data using the following two-point expectation:

$$\rho_j = \frac{1}{\mathcal{J}} \sum_{j'=0}^{\mathcal{J}-1} n_{j'} n_{(j'+j)}. \quad (3.5)$$

Here, \mathcal{J} is the length of some off-source data segment, which is usually chosen to be longer than the analysis data (i.e. $\mathcal{J} \gg J$) [121]. It is also necessary to treat the ‘‘edges’’ of the data segment (i.e. where $j + j' > \mathcal{J} - 1$) either by zero-padding or imposing periodicity: $\rho_j = \rho_{\mathcal{J}-j}$. Although these different treatments result in an autocovariance that differs for large j , if \mathcal{J} is sufficiently large then the autocovariance will be consistent for $j < J$ (which is what enters the calculation of the likelihood).

The two expressions for the log-likelihood are equivalent. The noise autocovariance (which appears in the time-domain log-likelihood) is related to the PSD (in the frequency-domain log-likelihood) via a discrete Fourier transform (Wiener-Khinchin theorem), when imposing the circularity condition [121]:

$$\frac{1}{2} S_n(f_k) = \delta t \sum_j \rho_j \exp\left(\frac{-2\pi i j k}{J}\right). \quad (3.6)$$

We use the inverse of Eq. 3.6 to estimate the autocovariance, which comes with the requirement that the off-source segment length used in the PSD estimate is much longer than the analysis length [121].

The time-domain expression for the log-likelihood has hitherto been considered more

suitable for ringdown analyses. This is because in the time-domain expression no Fourier transform of the data or model is required, and so no periodicity has to be ensured. The abrupt start of ringdown models means they do not satisfy this periodicity condition, which leads to spectral leakage upon Fourier transforming. When performing Fourier transforms of the GW data, periodicity is ensured by applying window functions to taper the data. This makes it difficult to isolate the ringdown region of a GW signal in some data; a sharp cut at the ringdown start time would introduce a discontinuity, whereas a smooth window would either suppress the ringdown signal or risk contamination of the ringdown by including unwanted parts of the inspiral-merger (see Fig. 7 in Ref. [121] for an illustration of this). These problems are naturally avoided in the time domain. In the following sections we describe how these problems can also be overcome in a frequency-domain analysis.

3.2.2 Marginalising over the inspiral-merger

We now motivate our approach to analysing the ringdown by first discussing a special case in which it becomes formally equivalent to the standard time-domain approach.

Consider first the case of a single interferometer. The observed data is a discretely sampled time series:

$$\dots, d_{-2}, d_{-1}, d_0, d_1, d_2, \dots \quad (3.7)$$

We assume that the ringdown has been identified as starting at the time of d_0 . The GW signal has a large amplitude around d_0 , but decays to zero at early and late times.

The standard approach to analysing the ringdown is to cut the data at the start time where the signal amplitude is large and take only the data after that time (i.e.

d_0, d_1, d_2, \dots) and to model this using a superposition of QNMs described by parameters $\vec{\theta}$ (e.g. the remnant mass, spin, amplitudes and phases for each QNM). The likelihood is written as

$$\mathcal{L}(d_0, d_1, d_2, \dots | \vec{\theta}), \quad (3.8)$$

and, because we have cut the signal where the amplitude is large, this must be expressed in the time domain (Eq. 3.3) to avoid problems of spectral leakage.

We could instead extend our analysis segment backwards by starting at d_{-1} , and then analyse this longer data stream with a new model that treats the value of the signal at d_{-1} as being a completely free parameter. The model is otherwise unchanged at later times and is now described by parameters $(\hat{d}_{-1}, \vec{\theta})$. Note that the new model for d_{-1} is entirely unphysical and is also discontinuous in the sense that there is no requirement that the model takes similar values at times -1 and 0 . The likelihood for this new model

$$\mathcal{L}(d_{-1}, d_0, d_1, d_2, \dots | \hat{d}_{-1}, \vec{\theta}), \quad (3.9)$$

is also given by Eq. 3.3, only with a slightly larger covariance matrix. If we marginalise this with respect to the ‘‘inspiral-merger parameter’’ \hat{d}_{-1} (adopting a flat, improper prior on \hat{d}_{-1} ranging between $\pm\infty$) then we recover the original likelihood in Eq. 3.8, i.e.

$$\mathcal{L}(d_0, d_1, d_2, \dots | \vec{\theta}) = \int_{-\infty}^{\infty} d\hat{d}_{-1} \mathcal{L}(d_{-1}, d_0, d_1, d_2, \dots | \hat{d}_{-1}, \vec{\theta}). \quad (3.10)$$

This follows from well-known properties of the Gaussian distribution.

We can of course include more early-time data in a similar way. If we treat the value of the GW signal at each of the times d_{-2}, d_{-3}, \dots as free parameters and marginalise

over all of them then we recover the original likelihood in Eq. 3.8. The point of doing this is that it allows us to start the analysis at early times when the amplitude of the GW signal in the data is small. This means we can apply windowing, aka tapering, operations to the data without fear of suppressing the signal, and we can therefore transform the likelihood into the frequency domain without encountering spectral leakage problems.

The extension of this argument to multiple interferometers is straightforward. The values of the model at each early time in each interferometer must all be treated independently as free parameters and marginalised over in the manner of Eq. 3.10.

The point we wish to emphasise is that an analysis of only the ringdown data (usually done in the time domain to avoid problems of spectral leakage) is equivalent to an analysis of all the data (which can be done in the frequency domain using standard GW data analysis techniques) if the inspiral-merger part of the signal is suitably marginalised out. This requires the use of an unphysical and discontinuous model for the inspiral-merger signal. For the equivalence to be exact, the inspiral-merger model should include an extremely large number of free parameters (one for each time stamp in each interferometer), but we will argue below that in practice it is sufficient to use a smaller number of parameters provided a sufficiently flexible model is used.

This discussion motivates the model we describe in Section 3.2.3. Once the likelihood is expressed in the frequency domain, several extensions of the analysis (such as treating the ringdown start time as a parameter of the model; see Section 3.3.1) become natural.

3.2.3 Model

As it is clearly impractical to model the data at each time stamp as a free parameter, we propose to use a continuous (but very flexible) inspiral-merger model instead. We choose a sum of sine-Gaussian wavelets, which are then truncated at the ringdown start time

and attached to a ringdown QNM model. With this method, we aim to model the full inspiral-merger-ringdown signal.

The ringdown model is zero for early times, and after a start time t_0 takes the form

$$h^{\text{R}}(t) = h_+^{\text{R}}(t) - ih_{\times}^{\text{R}}(t) = \sum_{\ell mn} A_{\ell mn} e^{-i[\omega_{\ell mn}(t-t_0) - \phi_{\ell mn}]}, \quad t \geq t_0, \quad (3.11)$$

where the complex QNM frequencies $\omega_{\ell mn} = 2\pi f_{\ell mn} - i/\tau_{\ell mn}$ are functions of the remnant BH mass M_f and dimensionless spin magnitude χ_f . Here, $f_{\ell mn}$ is the oscillation frequency, and $\tau_{\ell mn}$ is the damping time. Each QNM is further described by an amplitude, $A_{\ell mn}$, and phase parameter, $\phi_{\ell mn}$. It is possible to analytically take the Fourier transform of this expression and thereby write the ringdown model in the frequency domain as

$$\tilde{h}^{\text{R}}(f) = \int_{-\infty}^{\infty} dt [h_+^{\text{R}}(t) - ih_{\times}^{\text{R}}(t)] e^{-2\pi i f t} = \sum_{\ell mn} \frac{A_{\ell mn} e^{-i[2\pi f t_0 - \phi_{\ell mn}]} }{i(\omega_{\ell mn} + 2\pi f)}. \quad (3.12)$$

We model the inspiral-merger part of the signal as a truncated sum of W wavelets. The inspiral-merger model is zero for late times, but before the start time t_0 takes the form

$$\begin{aligned} h^{\text{IM}}(t) &= h_+^{\text{IM}}(t) - ih_{\times}^{\text{IM}}(t) \\ &= \sum_{w=1}^W \mathcal{A}_w \exp \left[-2\pi i \nu_w (t - \eta_w) - \left(\frac{t - \eta_w}{\tau_w} \right)^2 + i\varphi_w \right], \quad t < t_0. \end{aligned} \quad (3.13)$$

The limit on time is equivalent to a model which is multiplied by a Heaviside step function, $H(t_0 - t)$. Each wavelet is described by five parameters: \mathcal{A}_w and φ_w are the wavelet amplitudes and phases, τ_w are the wavelet widths, ν_w are the wavelet frequencies, and η_w are the wavelet central times. Again, it is possible to analytically take the Fourier transform of this expression and thereby write the inspiral-merger model in the frequency

domain as

$$\begin{aligned} \tilde{h}^{\text{IM}}(f) &= \int_{-\infty}^{\infty} dt [h_+^{\text{IM}}(t) - ih_\times^{\text{IM}}(t)] e^{-2\pi i f t} \\ &= \sum_{w=1}^W \mathcal{A}_w \exp[-2\pi i f \eta_w - \pi^2 (f + \nu_w)^2 \tau_w^2 + i\varphi_w] \\ &\quad \times \frac{\sqrt{\pi}}{2} \tau_w \left(1 + \operatorname{erf} \left[\frac{t_0 - \eta_w}{\tau_w} + \pi i (f + \nu_w) \tau_w \right] \right). \end{aligned} \quad (3.14)$$

The full IMR model is simply given by

$$h(t) = h^{\text{IM}}(t) + h^{\text{R}}(t). \quad (3.15)$$

We refer the reader to Fig. 3.1, which provides an illustration of this model with a single QNM and a single wavelet ($W = 1$).

If N QNMs are used, then the ringdown part of the model is described by $2N + 2$ parameters. The inspiral-merger part of the model is described by $5W$ parameters (these numbers do not include t_0). Additionally, there are two sky position angles (α and δ) and a polarisation angle (ψ) which enter the detector response described below. Finally, the ringdown start time (t_0) can also be treated as a model parameter in our frequency-domain approach. For a typical choice of these parameters, our IMR model is discontinuous at t_0 . The inspiral-merger part of the model contains no information of the physics of the source and no attempt is made to enforce continuity between the two parts; this helps to decouple the ringdown inference from the inspiral-merger parts of the data and thereby ensure that we are really performing a ringdown analysis.

To complete the description of our model, the detector response is given by projecting the waveform polarisations onto each interferometer (IFO) with the appropriate antenna patterns, $F_{+,\times}^{\text{IFO}}$ (see, for example, Ref. [173]). For a given sky location and GW polarisation

angle the detector response for each IFO $\in \{H, L, V\}$ is given by

$$h^{\text{IFO}}(t) = F_+^{\text{IFO}}(\alpha, \delta, \psi) h_+(t + \Delta t_{\text{IFO}}) + F_\times^{\text{IFO}}(\alpha, \delta, \psi) h_\times(t + \Delta t_{\text{IFO}}). \quad (3.16)$$

Here, $\Delta t_{\text{IFO}}(\alpha, \delta)$ accounts for the different signal arrival times at the detectors and is also a function of sky location. By definition, $h_+(t) = \text{Re}\{h(t)\}$, and $h_\times(t) = -\text{Im}\{h(t)\}$. Note, the frequency-domain waveforms presented here are Fourier transforms of the complex polarisation sum $h_+(t) - ih_\times(t)$. This means the separation into the plus and cross polarisations is not as simple as for the time-domain waveforms. Instead, the property that $\tilde{h}_{+,\times}^*(-f) = \tilde{h}_{+,\times}(f)$ for a real time-series implies that

$$\tilde{h}_+(f) = \frac{\tilde{h}(f) + \tilde{h}^*(-f)}{2}, \quad \tilde{h}_\times(f) = -\frac{\tilde{h}(f) - \tilde{h}^*(-f)}{2i}. \quad (3.17)$$

3.3 Injection study

We use the NR surrogate NRHybSur3dq8 [164] to simulate the full inspiral, merger and ringdown signal from GW190521-like and GW150914-like sources. We use these two sources to test our frequency-domain approach on the analysis of the ringdown and compare the results with a standard, time-domain analysis. Results from the GW190521-like analyses are shown here while the results from the GW150914-like analyses are shown in Section 3.4.

For all the GW190521-like injections, the surrogate was initialised with a total mass of $271 M_\odot$ (all masses are given in the detector frame) and a mass ratio of 1.27. For simplicity, all of the component spins were set to zero and the inclination angle was also chosen to be zero (i.e. the source was injected “face-on”). The simulated sky location and GW polarisation angle were taken to be the maximum-likelihood values from the NRSur7dq4 analysis in Refs. [12, 174] ($\alpha = 0.164$, $\delta = -1.14$, $\psi = 2.38$). The distance to

the binary was chosen so that it gives a particular value of the optimal SNR in Livingston; this was usually chosen to be 15 (corresponding to a distance of 4016 Mpc) so that it would likely be possible to detect multiple QNMs (in particular overtones), but several smaller values are also considered in Section 3.3.2. We perform zero-noise injections (i.e. analysing simulated data with a noise realisation of zero) into a three-interferometer H-L-V LIGO-Virgo network (except in Section 3.3.1 where a two-interferometer H-L injection is performed for comparison) and use the average PSDs from the first three months of O3 (available at Ref. [175]).

Although an important advantage of our approach is that it allows for easy marginalisation over the source sky position and ringdown start time, we first apply it to the case where these are fixed. This allows us to compare our results more directly to those from a time-domain analysis. The sky position and polarisation angles α , δ and ψ are fixed to their injected values and the ringdown start time t_0 is fixed to be 12.7 ms ($\sim 10 M_f$ in geometric units) after the peak of the strain (this choice follows the analysis of the real GW190521 signal in Ref. [39]).

First, for reference, an analysis using the time-domain expression for the log-likelihood was performed on this injection searching for the fundamental QNM (i.e. $\ell = m = 2$ and $n = 0$). Only the ringdown data $t \geq t_0$ was analysed; the time series in each interferometer was shifted to geocenter time using the injected sky position, then cut to include 0.1 s of data from the start of the ringdown. For each interferometer, the PSD was converted into an autocovariance function using the inverse of the transformation in Eq. 3.6 and this was used to construct the covariance matrix with Eq. 3.4. We sample over the remnant mass (M_f , using a flat prior between $100 M_\odot$ and $400 M_\odot$), the dimensionless remnant spin (χ_f , using a flat prior between 0 and 0.99), the QNM phase (ϕ_{220} , using a flat, periodic prior between 0 and 2π), and the QNM amplitude (A_{220} , using a flat prior




FIGURE 3.2: Posteriors on the (detector frame) remnant mass and dimensionless spin for the GW190521-like injection using a single QNM and analysed with a fixed sky position, GW polarisation angle and ringdown start time. The main panel shows the 90% confidence contour while the side plots show the one-dimensional marginalised posteriors. The solid blue line shows the results of a time-domain (TD) analysis. The other dashed and dotted lines show the results of frequency-domain analyses using different numbers of wavelets, W . The vertical and horizontal lines indicate the true values.

between 0 and 5×10^{-21}). As modes beyond the fundamental will likely have amplitude posteriors consistent with zero, we use a flat prior (as opposed to a log-uniform prior) on the amplitudes; we have checked the choice of prior has minimal influence on the results. We emphasise that t_0 is fixed in this analysis; i.e. using a delta-function prior. The resultant posterior on the remnant parameters is shown in Fig. 3.2. The posterior is consistent with the true remnant properties indicated by the vertical and horizontal lines. The true values were obtained with the NRSur3dq8Remnant model [99] which, provided with the injection parameters, can estimate the remnant properties.

Second, the corresponding frequency-domain analyses using $W = 1, 2$ and 3 truncated wavelets were also performed on the same injection but now using 4 s of data centred on the signal. The same ringdown parameters and priors as before were used. In addition, we now sample over the wavelet amplitudes (\mathcal{A}_w , using a flat prior between 0 and 5×10^{-21}), phases (φ_w , using a flat, periodic prior between 0 and 2π), widths (τ_w , with flat priors between 5 ms and 100 ms, or in geometric units between $\sim 4 M_f$ and $\sim 80 M_f$) and frequencies (ν_w , with flat priors between 30 Hz and 100 Hz, or in geometric units between $\sim 0.04 M_f^{-1}$ and $\sim 0.13 M_f^{-1}$). The label-switching ambiguity among the wavelets was removed by enforcing the ordering $\mathcal{A}_w \leq \mathcal{A}_{w+1}$ via the hypertriangulation transformation described in Ref. [176]. We also sample over the wavelet central times (η_w) using a Gaussian prior with a width of 50 ms ($\sim 40 M_f$) centred on the ringdown start time; this choice was empirically found to be sufficiently flexible, whilst also encouraging the wavelets to accurately model the signal near the peak. Recall that, for the moment, we are fixing the parameters α , δ , ψ and t_0 . The resultant posteriors on the remnant parameters are shown in Fig. 3.2.

From Fig. 3.2 we see that our frequency-domain approach gives posteriors on the remnant properties that are consistent with both the true values and the time-domain analysis. We see slight variations in the results of the frequency-domain analyses depending on the number of wavelets used. This is a high-mass injection with a short inspiral-merger in-band, so it would be expected that a small number of wavelets would be sufficient. In all cases our frequency-domain approach yields slightly more precise measurements of the remnant properties than the time-domain approach. We speculate that this is because of some coupling between the ringdown and inspiral-merger parts of the model, which leads to information from the early-time data informing our measurement of the remnant properties. Indeed, the inspiral-merger model with finite W is only an approximation to the maximally flexible model described in Section 3.2.2. We demonstrate this in Fig. 3.5,




FIGURE 3.3: Similar to Fig. 3.2, posteriors on the remnant mass and dimensionless spin for the GW190521-like injection using a single QNM while marginalising over the sky position, polarisation angle and ringdown start time. Also shown in the grey shaded region is the result of a frequency-domain analysis that does not include any wavelets (i.e. $W = 0$); as expected, since this model has an abrupt discontinuity at t_0 , this analysis yields severely biased estimates of the remnant mass and spin. This $W = 0$ analysis is included here to highlight the important role the wavelets play in our frequency-domain approach.

which shows our model in the frequency domain.

It is also possible to visually check the performance of the frequency-domain approach by plotting the whitened waveform reconstructions. These reconstructions, which are the relevant time series that enter the frequency-domain log-likelihood, were found to be in excellent agreement with the injected data. Examples of such reconstructions are shown for the GW150914-like injection in Chapter 3.4.

We now turn to the case where the source sky position, polarisation angle and the ringdown start time are treated as free parameters in the frequency-domain analysis. We




FIGURE 3.4: The prior and posterior distributions on the ringdown start time in the Hanford frame for the same frequency-domain analyses as in Fig. 3.3. The prior (solid blue line) is a Gaussian centred 12.7 ms after the time of peak strain in Hanford, with a standard deviation of 1 ms. This prior has been chosen to be informative; the posterior distributions do not differ significantly from the prior. It is necessary to use such an informative prior because we find that the ringdown start time cannot be reliably inferred solely from the data (see discussion in Section 3.3.1). We observe a slight preference for an early start time when using a small number of wavelets; we speculate that this is due to the wavelet model being less flexible than the maximally flexible model described in Section 3.2.2.

use a uniform prior over the sphere of the sky and a flat, periodic prior on ψ between 0 and π . As the sky position is now allowed to vary in the analysis, the time delay between the different interferometers and the geocenter is not constant. Therefore, for the ringdown start time, we choose to place a Gaussian prior on the start time in one of the detectors where the ringdown is clearly visible (we choose Hanford). The Gaussian prior was centred on the fixed value used in the time-domain analysis and has a relatively narrow width of 1 ms ($\sim 0.8 M_f$). This choice of prior is quite restrictive (i.e. assuming good prior knowledge of t_0) and the reasons for this are discussed further in Section 3.3.1; although, we note here that this is still more flexible than the delta-function prior used above. The resultant posteriors on the remnant parameters are shown in Fig. 3.3, and we see that




FIGURE 3.5: The maximum-likelihood waveform in Livingston from the $W = 1$ analysis of Fig. 3.3. Shown are the separate wavelet (h^{IM} , Eq. 3.14) and ringdown (h^{R} , Eq. 3.12) parts of the model, along with sum (h). We choose to plot the characteristic strain, as in Fig. 1.3, with the Livingston O3 noise curve (which was used in the analysis) also plotted for reference.

the performance of our frequency-domain approach is not significantly degraded when the searching over α , δ , ψ and t_0 . Shown in Fig. 3.4 are the posteriors on the ringdown start time, t_0 , which are discussed in more detail in the next section.

In Fig. 3.5 we show the maximum-likelihood waveform from the $W = 1$ analysis, projected onto the Livingston detector. Complementary to the time-domain waveform shown in Fig. 3.1, this representation shows clearly the overlap that exists between the two parts of the model in the frequency domain. Although there is a clear ordering of frequencies, with the inspiral-merger occurring at lower frequencies than the ringdown (and so the wavelet power peaks at a lower frequency than the ringdown), in the frequency domain there is not a clear transition between the inspiral-merger and the ringdown. This may be responsible for the ringdown inference being coupled to the wavelet inference, resulting in the observed differences with the time-domain analyses.

Stochastic sampling was performed using the DYNESTY [177] implementation of the

nested sampling algorithm [178, 179]. For the sampling method, we used random walks with fixed proposals. Typically, the minimum number of steps used in the random walk was 2000, and the number of live points was 4000. We note that our frequency-domain ringdown analysis has many more parameters than a time-domain analysis due to the $5W$ parameters used in the wavelet sum. Posteriors on these inspiral-merger parameters are not presented here, but we note that as the number of wavelets is increased strong degeneracies develop among these parameters. This is expected, and is desirable in this context, as the wavelet part of the model is designed to be extremely flexible. These degeneracies in the inspiral-merger part of the model are not a problem for our present purpose as they do not inhibit our ability to measure the QNMs or the remnant properties. All posterior samples, including those for the wavelet parameters, are available at Ref. [5].

3.3.1 Determining the ringdown start time

The ringdown start time, t_0 , appears as a model parameter in our frequency-domain approach. This raises two interesting questions: what prior should be placed on t_0 , and how well can t_0 be measured from the data? The possibility of determining t_0 from the data is particularly enticing because the ringdown start time is theoretically uncertain and its choice is crucial for any ringdown analysis.

Unfortunately, we find that when using wide priors on t_0 , early ringdown start times are generally favoured and this leads to a bias in the recovered remnant mass and spin. This can be seen in Figs. 3.6 and 3.7, where the $W = 1$ analysis previously shown in Fig. 3.3 is repeated with increased values of the t_0 prior width. The posterior on t_0 is also affected by the number of wavelets used; as can be seen from the Fig. 3.4, larger values of W tend to favour later ringdown start times. We have repeated the analyses in Figs. 3.6 and 3.7 with larger values of W to see if this counteracts the preference for




FIGURE 3.6: Similar to Fig. 3.3, posteriors on the remnant mass and spin for the GW190521-like injection using a single wavelet ($W = 1$) and a single QNM, but using different priors on t_0 . The markers indicate the maximum-likelihood values. The dashed orange curve is identical to that in Fig. 3.3.

an early start time (and hence removes biases in the remnant parameters), however, this was found not to be the case. These calculations show that the posterior obtained on the parameter t_0 in our approach depends on the prior and on the number of wavelets used in the (unphysical) model of the inspiral-merger signal. Therefore, it does not seem to be possible to reliably measure the ringdown start time from the data alone. In the next chapter we draw similar conclusions from the analysis of GW150914; wide priors on the ringdown start time lead to some biases and also uninformative posteriors that are hard to interpret (see Fig. 4.9). It is for this reason that a narrow, informative, t_0 prior must be used in the analyses described in the previous section. Although not desirable, this is still an improvement over the fixed t_0 routinely used in most time-domain analyses.




FIGURE 3.7: Posteriors on the ringdown start time corresponding to Fig. 3.6. For each, the prior is a Gaussian centred on the vertical line, with widths given in the legend. It can be seen that a wider prior causes earlier ringdown start times to be favoured (this is the case even when additional wavelets are included). As a result of the earlier start time, a bias appears in the recovered remnant parameters and higher values of both M_f and χ_f are favoured.

Finally, we discuss the choice that was made in the previous section to place the prior on t_0 in the frame of one of the interferometers. Because the sky position is allowed to vary, using the geocenter time is inappropriate due to coupling with the sky position. The orange curve in Fig. 3.8 shows the posterior on t_0 transformed into the geocenter frame from the $W = 1$ frequency-domain analysis using the narrow 1 ms prior on the ringdown start time. Also shown in the blue-dashed line is a posterior from an identical injection into a two-interferometer H-L network. Due to the multimodal sky posterior, the posterior on t_0 in the geocenter frame can also be multimodal (this is present in the three-detector analysis to a smaller extent but is most clear in the two-detector analysis). This makes choosing a suitable prior for the ringdown start time more difficult in the geocenter frame. It is for this reason that for the analyses described above, the prior was specified in the frame of one of the detectors. The results in Fig. 3.8 also show that our frequency-domain

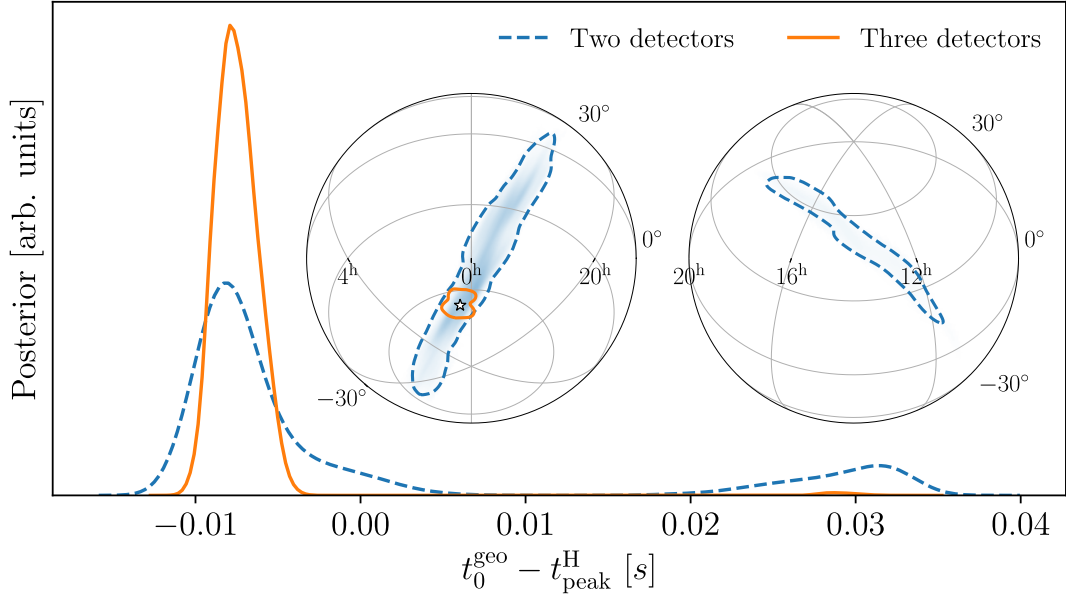


FIGURE 3.8: *Main panel:* posterior on the ringdown start time in the geocentric frame. The orange line corresponds to the $W = 1$ model applied to the three-detector network injection (this is the same analysis presented in Figs. 3.3 and 3.6, which was plotted with a dashed orange line). The dashed blue line corresponds to a similar analysis on a two-detector network injection (with Virgo removed). This is to motivate the choice to parameterise t_0 in the frame of a detector; in the geocentric frame, a multimodal structure appears as a result of different possible sky locations. This makes it harder to place a sensible prior. *Left inset plot:* the sky location posterior on the southern hemisphere (orthographic projection). This contains the injected source location (indicated by the star) which is correctly recovered with sky area $\sim 77 \text{ deg}^2$ (90% confidence) for the three-detector network, and $\sim 1800 \text{ deg}^2$ for the two-detector network. *Right inset plot:* the northern hemisphere of the sky contains a secondary mode when using the two-detector network, which correlates with t_0^{geo} . Both modes of the sky posterior are elongated along the circle of constant time delay between the two detectors.

approach yields a posterior on the source sky position as a by-product of the ringdown analysis. However, it should be stressed that this is not a ringdown-only result; the entire IMR model, including the unphysical wavelet part, is contributing to this sky localisation.

3.3.2 Detecting additional QNMs

A key goal in the analysis of BH ringdowns is the detection of additional QNMs beyond the fundamental $\ell = m = 2, n = 0$ mode. This has already been achieved; see, for example, Ref. [119] where the $\ell = m = 2, n = 1$ overtone was identified in GW150914 using a time-domain analysis. In this section we show, using our GW190521-like injection, that our frequency-domain approach is also able to identify additional QNMs.

As a first step towards testing our model we search for the $n = 1$ overtone of the fundamental QNM. It would also be possible to search for higher harmonics (e.g. modes with $\ell \geq 3$); however, the results of previous investigations on NR simulations (see, e.g. Refs. [1, 122, 124, 126]) suggest that overtones are generally more prominent than harmonics in the ringdown and are therefore a natural first target for any search.

We reanalyze the GW190521-like injection in the frequency domain using the $W = 1$ inspiral-merger model but this time including an overtone in the ringdown (the choice to use a single wavelet is motivated by the previous results; it is sufficient to model the inspiral-merger for this high-mass injection, and we see no significant improvements with additional wavelets). When using overtones, it is appropriate to start the ringdown analysis at an earlier time. For the frequency-domain analyses a Gaussian prior with a standard deviation of 1 ms centred at the time of the peak strain was used (this is 12.7 ms earlier than was used above). The results of this “2QNM” analysis are shown in Fig. 3.9, along with the fundamental only “1QNM” analysis for comparison. Posteriors are plotted for the QNM amplitudes and the remnant mass and spin parameters. We see that the overtone can be confidently detected with non-zero amplitude. The 2QNM analysis yields more precise measurements of the remnant mass and spin due to a combination of the earlier ringdown start time (which gives a larger ringdown SNR) and the improved ringdown model.




FIGURE 3.9: Posteriors on the QNM amplitudes and remnant mass and spin for one- and two-mode analyses (1QNM and 2QNM respectively) of the GW190521-like injection, performed in the frequency domain. The results in blue are for the recovery using two QNMs [the fundamental $(\ell, m, n) = (2, 2, 0)$ and its first overtone $(2, 2, 1)$] which are both detected with non-zero amplitudes using a t_0 prior centred on the time of the peak strain and with a width of 1 ms. Also shown in orange for comparison are the results using one QNM [the fundamental $(2, 2, 0)$ only] with a prior centred 12.7 ms after the peak, again with a width of 1 ms. The vertical and horizontal solid grey lines indicate the true values of the remnant mass and spin and the diagonal dashed grey line indicates $A_{220} = A_{221}$. The difference in the A_{220} amplitude between the two analyses can be explained by the different ringdown start times and the decay of the $(2, 2, 0)$ QNM. Over a time ~ 12.7 ms, we expect the A_{220} amplitude to decay by a factor $\sim \exp[-12.7 \text{ ms}/\tau_{220}] \sim 0.5$; this is shown in the shaded grey posterior in the top-left panel where the 2QNM posterior is used to predict the value of the amplitude at the later start time.




FIGURE 3.10: Overtone amplitude posteriors from a $W = 1$ frequency-domain (FD) analysis, where the fundamental $(\ell, m, n) = (2, 2, 0)$ QNM and its first overtone $(2, 2, 1)$ are included in the ringdown model. For comparison, the overtone amplitudes from a time-domain (TD) analysis are shown with the dashed lines. The injected SNR in Livingston is controlled by changing the injection luminosity distance: $D_L = \{4016.3, 5020.4, 6693.9, 10040.8\}$ Mpc for $\text{SNR} = \{15, 12, 9, 6\}$ respectively. The maximum-likelihood values scale as D_L^{-1} .

We now turn our attention to the resolvability of this additional QNM as a function of the injected SNR and compare the sensitivities of the time- and frequency-domain approaches. The GW190521-like source was re-injected at a series of lower SNRs: 12, 9, and 6 (in the Livingston detector). A $W = 1$ frequency-domain analysis was re-performed on this sequence of injections, along with a time-domain analysis for comparison. Following the use of an earlier ringdown start time in the frequency-domain analysis, for the time-domain analysis the ringdown start time was fixed to the peak of the strain. The posteriors on the amplitude A_{221} of the overtone are shown in Fig. 3.10. As the SNR is decreased, the overtone becomes increasingly difficult to detect and the posteriors become consistent with $A_{221} = 0$. This is the case for both the time- and frequency-domain analyses which give similar results. This suggests the time- and frequency-domain approaches are equally sensitive to additional QNMs.

Further evidence supporting this conclusion comes from the odds ratios (aka Bayes' factors) in favour of the overtone. The Bayes' factors $\mathcal{B}_{1\text{QNM}}^{2\text{QNM}}$ (computed with equal prior odds) in favour of the second QNM were computed from both the time- and frequency-domain analyses. In order to do this, we perform an additional set of analyses on the series of injections used in Fig. 3.10 with the same ringdown start time (fixed at the peak for the time-domain analysis, and a Gaussian prior centred on the peak in Hanford for the frequency-domain) but without the overtone included. We can then compute the evidence, $\mathcal{B}_{1\text{QNM}}^{2\text{QNM}}$, in favour of the 2QNM analysis (with an overtone) over the 1QNM analysis (fundamental mode only) keeping every other part of the analysis identical. The log-Bayes' factors for each of the different SNR injections are shown in Table. 3.1 where it can be seen that the time- and frequency-domain approaches are equally sensitive to the overtone mode.

SNR	15	12	9	6
TD	1.1	0.3	-0.3	-0.7
FD	1.1	0.4	-0.4	-0.6

TABLE 3.1: The log-Bayes' factors $\log_{10} \mathcal{B}_{1\text{QNM}}^{2\text{QNM}}$ in favour of an overtone for the series of GW190521-like injections at different SNRs, for both time-domain (TD) and $W = 1$ frequency-domain (FD) analyses. The uncertainties on these Bayes' factors are all $\pm(0.1 - 0.2)$, with errors on the evidences estimated from within a single nested sampling run.

3.4 GW150914-like injection

In this chapter we test the frequency-domain approach further by analysing a GW150914-like injection, which will set up the analysis in the final chapter.

The surrogate was initialised with a total mass of $72.2 M_{\odot}$ and a mass ratio of 1.16. As before, all of the component spins were set to zero for simplicity. The simulated sky location and GW polarisation angle were taken to be $\alpha = 1.95$, $\delta = -1.27$, and $\psi = 0.82$.

These are consistent with the GW150914 posterior and were chosen to match the values used in Ref. [119]. The distance to the binary was set to 471.4 Mpc, which gave an optimal SNR in Hanford of 25.

The inclination angle was chosen to be π (i.e the source is injected “face-off”) which is consistent with the GW150914 posterior. The source inclination affects the GW polarisation, and this is handled via the introduction of a “ellipticity parameter” ϵ which has the effect of transforming $h_{\times}(t) \rightarrow \epsilon h_{\times}(t)$. For the “face-on” injections in the main text $\epsilon = 1$ was used, while for the “face-off” injections considered here $\epsilon = -1$. A more general analysis would allow ϵ to vary as a free parameter, such as what was done in Ref. [121].

We perform zero-noise injections into the two-interferometer H-L LIGO network that was operating at the time of the first detection. We use the PSDs associated with the data surrounding GW150914 (available at Ref. [180]). These different parameters (particularly the lower total mass) result in a signal with a longer inspiral. This is an important test case for our model as there is a much larger fraction of the SNR in the inspiral-merger (compared to the GW190521-like injection) which has to be “marginalised out” in the analysis.

As was done initially for the GW190521-like injection, we fix the sky location and polarisation angle to the injected values to simplify the problem and aid comparison to the time-domain analysis. The ringdown start time is also fixed to 3 ms after the time of the peak strain ($\sim 10 M_f$ in geometric units).

Following the same procedure as in Section 3.3, a time-domain analysis was first carried out to recover the fundamental QNM. The prior on the remnant mass was adjusted to reflect the lower injected value (flat between $50 M_{\odot}$ and $100 M_{\odot}$), but otherwise the analysis was unchanged from the time-domain analyses described in the main text. The resultant remnant mass and spin posterior is shown by the blue solid line in Fig. 3.11.




FIGURE 3.11: Similar to Fig. 3.2, posteriors on the recovered remnant mass and spin for the GW150914-like injection using the fundamental QNM and a varying numbers of wavelets. Also shown for comparison is the result of a time-domain analysis (solid blue line).

Secondly, a series of frequency-domain analyses were carried out using an increasing number of wavelets. Results for $W = 1, 3,$ and 5 are shown in Fig. 3.11. We found a slightly more restrictive prior on the wavelet central times, η_w , was required to aid the inference; the Gaussian width was reduced to 10 ms ($\sim 30 M_f$). This encouraged the wavelets to fit near the ringdown start time, which is the part of the signal we are most interested in. The upper bound on the wavelet frequencies was also increased to 500 Hz ($\sim 0.17 M_f^{-1}$), as the lower binary mass means the merger-ringdown occurs at a higher frequency. We see that a single wavelet is not quite sufficient to avoid bias in the remnant parameters, which may be expected when working with a longer inspiral. As the number of wavelets is increased the bias disappears, and the remnant posteriors seem to converge to a solution that is stable against the inclusion of additional wavelets. As was the case for the

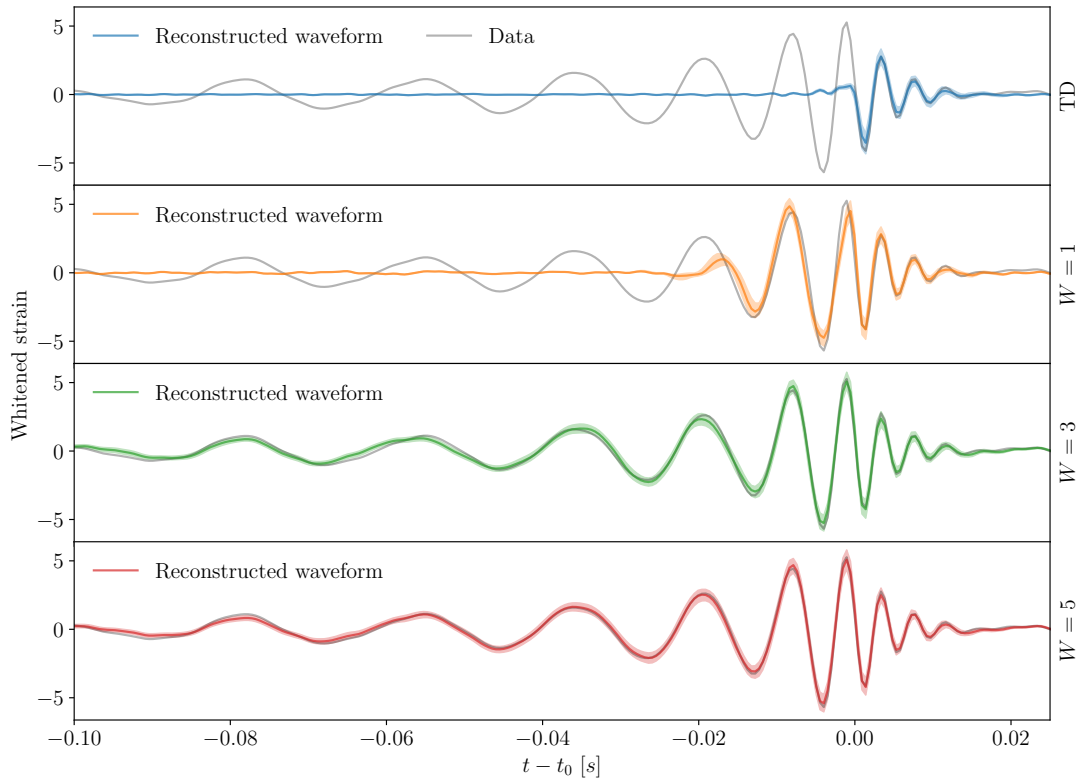


FIGURE 3.12: Whitened waveform reconstructions (in Hanford) corresponding to the results of Fig. 3.11. The top panel shows the waveform from a time-domain analysis. In the time-domain approach the data before the ringdown start time is excluded, but here the waveform is plotted for all times. This highlights one of the problems with using a ringdown-only model in the frequency domain: the abrupt start of the model leads to spectral leakage when Fourier transforming (visible as oscillations before the ringdown start time). The following panels show waveforms from the frequency-domain approach. Problems with spectral leakage are avoided, due to the wavelets smoothly connecting to the ringdown part of the model. Just a single wavelet fails to model the full GW150914-like inspiral-merger, which is to be expected because of its longer duration in-band. As more wavelets are included in the model, more of the inspiral-merger is captured by the model. The difference in the reconstruction for three and five wavelets is minimal, showing the model is converging on the signal.

GW190521-like analysis, we find the frequency-domain model achieves tighter constraints on the remnant parameters in comparison to the time-domain analysis.

Finally, in Fig. 3.12 we inspect the whitened waveform reconstructions for all four analyses shown in Fig. 3.11. We focus on the waveform in the Hanford detector. We take samples from the posterior of each run and use these to compute the projected

waveform $F_+^{\text{H}}h_+(t + \Delta t_{\text{H}}) + F_{\times}^{\text{H}}h_{\times}(t + \Delta t_{\text{H}})$ (see Eq. 3.16). This quantity is always discontinuous for both the time- and frequency-domain analyses. We then whiten this waveform (and the data) using the Hanford PSD. After whitening, the projected waveform is continuous. In Fig. 3.12 we plot the median and 5% – 95% credible region of the whitened waveform reconstructions. In a time-domain analysis no Fourier transforms are performed, and usually data before the ringdown start time is excluded, so the top panel is not representative of the actual analysis. Instead, we include it to here to highlight the problem of spectral leakage, which occurs when taking Fourier transforms of discontinuous models. A ringdown-only model (as shown in the top panel) has a discontinuity at the ringdown start time and this causes oscillations to appear before the start time in the whitened waveform. The following panels, which include wavelets to model the inspiral-merger signal, remove this discontinuity and prevent Fourier transform artifacts. A single wavelet is not sufficient to capture the full inspiral-merger signal, which likely causes the bias seen in Fig. 3.11. Increasing the number of wavelets makes the model flexible enough to model the inspiral-merger signal, and also to remove bias in the mass-spin posterior.

3.5 Conclusions

BH ringdown and QNMs are a key area of study in the burgeoning field of GW astronomy and are particularly important for testing GR. Ringdown analyses are usually performed in the time domain as this provides a natural way to work with discontinuous models and to apply sharp cuts to the data. However, in these analyses the ringdown start time and sky position usually have been fixed beforehand. The log-likelihood is also more computationally expensive than in the frequency domain.

We have presented a novel approach for analysing the ringdown in the frequency

domain. Our approach uses a flexible combination of sine-Gaussian wavelets, truncated at the start of the ringdown, to effectively marginalise over the inspiral and merger parts of the signal. The benefits of performing the analysis in the frequency domain include being able to easily vary the source sky position and ringdown start time model parameters as part of the analysis. As virtually all other GW data analysis is already performed in the frequency domain, a further benefit of our approach is that it allows us to utilise standard, and now very well-tested, GW analysis software packages for performing the Bayesian inference and also for estimating the noise properties.

We have tested our frequency-domain approach by analysing a series of NR surrogate injections and by comparing our results with those from a time-domain analysis. We find that our frequency-domain approach is equally sensitive to additional QNMs compared to the time-domain approach. However, we find that it generally yields more precise measurements of the remnant BH mass and spin parameters which we speculate is due to some small coupling with the inspiral and merger signal. We also paid particular attention to the choice of prior on t_0 ; although this appears as a model parameter in our approach it was found that, unfortunately, it was not possible to reliably determine it solely from the data.

In the next chapter we apply this method to real GW data. Specifically the first GW event, GW150914, in a search for a ringdown overtone.

Chapter 4

Searching for a Ringdown Overtone in GW150914

4.1 Introduction

The very first GW event, GW150914 [17], remains among the best candidates for studying the ringdown. This is a result of several factors, including its large SNR of $\rho \sim 24$ and its total mass of $M \sim 70 M_{\odot}$ which places the merger and ringdown in the center of the LIGO [9] sensitive frequency band at ~ 200 Hz. Additionally, GW150914 is by now the most well-studied GW event and therefore the signal and the properties of the noise in the surrounding data are extremely well understood.

The first tests of GR performed using GW150914 included an investigation of the ringdown [181]. The ringdown signal, after a fixed starting time t_0 , was modelled using a single damped sinusoid; the parameters of which were checked for consistency with the predicted least-damped QNM of the remnant BH. This first attempt at a ringdown analysis was performed using the standard Whittle frequency-domain log-likelihood [182], commonly used in GW data analysis. The ringdown was isolated by choosing a lower limit of ~ 130 Hz in the frequency integral, effectively cutting the data mid-signal. As

discussed in the previous chapter, this approach suffers from several shortcomings. Firstly the frequency-domain cut at ~ 130 Hz only approximately separates the ringdown from the early-time signal due to the breakdown of the stationary phase approximation near merger. Secondly the nonzero amplitude at the start of the signal model breaks the assumption of circularity for the Fourier transform, thereby introducing contamination in the form of spectral leakage. Therefore, this approach does not scale well to higher SNRs where noise will no longer dominate over the systematic errors introduced by the sharp frequency-domain cut. Despite these drawbacks, this approach was successfully used in Ref. [181] to identify the fundamental QNM in the GW150914 signal.

Since this initial attempt, several groups have developed new time-domain frameworks specifically for ringdown analyses [119, 120, 165]. The principle motivation for working in the time domain is that it is easy to impose sharp cuts on the data at specific times (without any spectral leakage) and to analyse only data after a chosen start time (see Ref. [121] for details of time-domain analysis methods). These approaches have also enabled going beyond the fundamental mode. Outside of the testing-GR catalog papers [37, 38], other groups have attempted to identify additional QNMs in the ringdown data and perform tests of the no-hair theorem. This includes claims of detection of the $(2, 2, 1)$ overtone in GW150914 by Isi et al. [119], and claims of detection of the $(3, 3, 0)$ harmonic in GW190521 by Capano et al. [120]. By allowing the QNM frequency of the secondary mode to deviate from the GR Kerr prediction, the above works found the measured spectrum to be in agreement with the no-hair hypothesis to within $\sim 20\%$ (68% credibility) and $\sim 1\%$ (90% credibility) respectively. More recently, the QNM filter approach of Ma et al. [168] was also applied to GW150914. They found results broadly in agreement with Isi et al., with a test of the no-hair hypothesis at a similar level [183, 184].

The detection of an overtone in GW150914 was made possible partly because, compared

to earlier studies, the authors chose to use an earlier start time for the ringdown; this was motivated by contemporary NR studies discussed in Chapter 2, that demonstrated that when overtones are included the ringdown can be considered to start as early as the time of peak strain amplitude.

However, a recent paper by Cotesta et al. [185] reanalysed the GW150914 signal using very similar methods and found no significant evidence for an overtone. It was also suggested that the earlier detection claims of Ref. [119] were noise dominated. (This prompted a response from Isi et al. [186] where they restated their claim to have detected an overtone in GW150914.) Ref. [187] also found weaker evidence for an overtone using an analysis method closer to that of Ref. [181]. There are similar disagreements regarding the claim of a higher harmonic in GW190521 made in Ref. [120], as these are in tension with the results of Ref. [38]. Amid this confusion, it is particularly concerning that the supposedly identical analyses in Refs. [119, 186], and [185] come to such different conclusions concerning which QNMs are in the data. Discrepancies of this sort risk jeopardising the science that can be done using future ringdown observations.

These discrepancies highlight some of the difficulties inherent in time-domain ringdown analysis, where important choices (that affect the results) for fixed quantities such as the ringdown start time have to be made and care must be taken with the noise covariance estimation. If ringdown studies are to be used to make precision measurements of BH properties or as reliable tests of GR we must first be able to make reliable and reproducible determinations of the QNM content. This is also not a problem that will be removed in the future with observations at higher SNR. Even if an event has a higher SNR that is sufficient for a clear detection of the first QNM overtone, the focus will then simply shift to trying to identify the next overtone (or else the next QNM harmonic) in the countably infinite ringdown sum [187].

To complement the time-domain analysis frameworks, in the last chapter we proposed a new method for ringdown analyses which works in the frequency domain. A flexible sum of sine-Gaussian wavelets, truncated at the ringdown start time, is used to effectively marginalise over the inspiral-merger (i.e. pre-ringdown) part of the signal. The model is completed by attaching this to the usual sum of QNMs which model the ringdown. No continuity is enforced between the two parts of the model in order to keep the ringdown inference independent from the rest of the signal. However, we find the continuity is effectively learned from the data, and any remaining discontinuities disappear entirely when the signal is “whitened” according to the instrumental noise. In a particular limit, this approach can be shown to be formally equivalent to the time-domain analyses described above. However, this frequency-domain approach can be generalised and offers several advantages over time-domain approaches: well-established GW data analysis methods and pipelines can be used (which are all built in the frequency domain), the inspiral-merger data informs the noise estimation at the start of the ringdown (improving parameter estimation accuracy), and the ringdown start time and the source sky position can be easily treated as free parameters and marginalised over as part of a Bayesian analysis (instead of being fixed). We note, however, that (as discussed in Chapter 3) a narrow and informative prior on the ringdown start time must be used. Reweighting techniques can be employed to investigate different ringdown start time prior choices computationally efficiently in post processing (see Section 4.2.3) obviating the need for the large number of analyses performed in Refs. [185, 186].

In this chapter the new frequency-domain method is applied to reanalysing the ringdown of GW150914 paying particular attention to the presence (or absence) of an overtone. We perform analyses with and without an overtone and investigate different choices of the ringdown start time. We also perform additional analyses with varying data sampling

frequencies and integration limits to verify the stability of our results. Finally, a mock injection study into real detector noise is also performed to further assess the significance of any overtone detection. Section 4.2 describes the signal model, the data, and the analysis methods used in this chapter. Section 4.3 presents our main results including posteriors on the remnant BH properties and overtone amplitude, and Bayes' factors for the overtone model. The results are discussed further in Section 4.4. Throughout this chapter we make use of natural units where $G = c = 1$.

All data products and plotting scripts used to make the figures in this chapter are made publicly available at Ref. [6].

4.2 Methods

This section briefly describes the frequency-domain method for analysing BH ringdowns introduced in Chapter 3: the wavelet-ringdown model is described in Section 4.2.1; the data, likelihood and priors are described in Section 4.2.2; and our approach for dealing with changes to the ringdown start time is described in Section 4.2.3.

4.2.1 Wavelet-ringdown model

Our model consists of two parts: one for early times before t_0 which is referred to here as the *inspiral-merger*, and another for the *ringdown* after the start time t_0 .

First, we describe the ringdown part of the model. After a ringdown start time t_0 , which is itself a parameter in the model, the model takes the form

$$h^{\text{R}}(t) = h_{+}^{\text{R}}(t) + ih_{\times}^{\text{R}}(t) = \sum_{n=0}^N A_n e^{-i[\omega_{22n}(t-t_0) + \phi_n]}, \quad t \geq t_0. \quad (4.1)$$

Because our focus in this chapter is on the presence of an overtone, we fix the angular indices to $\ell = m = 2$ and vary only the number of QNM overtones, N , in the model (N is always taken to be either 0 or 1 in this chapter). Note that the form of this equation differs slightly from Eq.3.11 in Chapter 3. This is because the source inclination angle is fixed to be “face-off” (i.e. $\iota = \pi$). In the notation of (for example) Refs. [1, 126, 130], this is equivalent to using the $\ell = -m = 2$ mirror modes. Or, in notation of Ref. [121], using an ellipticity of $\epsilon = -1$. The complex QNM frequencies, $\omega_{\ell mn} = 2\pi f_{\ell mn} - i/\tau_{\ell mn}$, are functions of the remnant BH mass M_f (detector frame) and dimensionless spin χ_f . Additionally, each QNM is further described by an amplitude, A_n , and a phase, ϕ_n .

Second, we describe the inspiral-merger part of the model. This is modelled as a truncated sum of W wavelets. At early times the model takes the form

$$\begin{aligned} h^{\text{IM}}(t) &= h_+^{\text{IM}}(t) + ih_{\times}^{\text{IM}}(t) \\ &= \sum_{w=1}^W \mathcal{A}_w \exp \left[-2\pi i \nu_w (t - \eta_w) - \left(\frac{t - \eta_w}{\tau_w} \right)^2 - i\varphi_w \right], \quad t < t_0. \end{aligned} \quad (4.2)$$

Again, the minor differences in sign conventions compared to Chapter 3 come from fixing the inclination angle to be face-off. The wavelets are each described by five parameters: \mathcal{A}_w and φ_w are the wavelet amplitudes and phases, τ_w are the wavelet widths, ν_w are the wavelet frequencies, and η_w are the wavelet central times. In this chapter we use $W = 3$ (three wavelets) in our model. This number was empirically found to be sufficient (see Section 3.4, where the number of wavelets was varied for a GW150914-like injection).

The full signal model is given by discontinuously joining the inspiral-merger to the ringdown at t_0 ,

$$h(t) = h^{\text{IM}}(t) + h^{\text{R}}(t). \quad (4.3)$$

And finally the detector response must be considered; we project the waveform polarisations

onto each interferometer (IFO) with the antenna patterns, $F_{+, \times}^{\text{IFO}}$. The detector response for each IFO $\in \{\text{H}, \text{L}\}$ is given by Eq. 3.16, which also accounts for the different signal arrival times at each detector. Throughout this chapter we quote times in the Hanford frame. So, in particular, t_0 refers to the ringdown start time in Hanford.

4.2.2 Data and priors

We use the GW150914 strain data sampled at 4096 Hz for both the Hanford and Livingston interferometers, which was obtained from Refs. [188, 189]. A total of 4096 s of data around the event was downloaded, from which the mean was subtracted (this is effectively equivalent to applying a ~ 1 Hz highpass filter). Pre-computed power spectral densities (PSDs) associated with GW150914 from the GWTC-1 release were used [180]. It has been verified our results are insensitive to the exact noise PSD used; for example, our results are unchanged when using a PSD estimated from a length of off-source data. The analysis data consists of 4 s of data centred on the event GPS time (1126259462.4 s), and a Tukey window with an alpha parameter of 0.2 was applied to this analysis data. The Bayesian analysis used the standard frequency-domain log-likelihood function (see, e.g., Eq. 3.1), with the limits of the frequency integration between 20 and 1000 Hz. The choices of sampling rate and upper limit of frequency integration are discussed further in Section 4.3.3.

All the model parameters described in Section 4.2 were sampled over as part of a Bayesian analysis. For the wavelet parameters, uniform priors are used for the amplitudes ($\mathcal{A}_w \in [0, 10^{-20}]$), phases ($\varphi_w \in [0, 2\pi]$), frequencies ($\nu_w \in [20, 200]$ Hz), and widths ($\tau_w \in [4, 80] \tilde{M}_f$, or equivalently $\sim [1.4, 27]$ ms). Here, $\tilde{M}_f = 68.779 M_\odot = 0.33875$ ms is a fixed point estimate of the final, detector-frame mass (obtained using the median value from Ref. [11]) and should not be confused with the varying model parameter M_f . The label-switching ambiguity among the wavelets was removed by enforcing the ordering

$\nu_w \leq \nu_{w+1}$ via the hypertriangulation transformation described in Ref. [176]. We sample over the wavelet central times (η_w) using a Gaussian prior in the Hanford frame with a width of $50 \tilde{M}_f$ (~ 17 ms) centred on $t_{\text{ref}} = 1126259462.423$ s. This choice was found to be sufficiently flexible, whilst at the same time encouraging the wavelets to accurately model the signal near the peak (see the discussion in Section 3.4).

For the ringdown, uniform priors are used for the amplitudes ($A_n \in [0, 10^{-19}]$), phases ($\phi_n \in [0, 2\pi]$), remnant mass ($M_f \in [40, 100] M_\odot$), remnant spin ($\chi_f \in [0, 0.99]$) and ringdown start time ($t_0 - t_{\text{ref}} \in [-15, 15] \tilde{M}_f$, which in SI units corresponds to $\sim [-5.1, 5.1]$ ms). We use a uniform prior on t_0 so that the samples can be easily reweighted in post-processing (see Section 4.2.3). For the remaining parameters, we used a uniform prior over the sphere of the sky (parameterised using α and δ) for the source location and a flat, periodic prior on the polarisation angle ψ in the range 0 to π .

The nested sampling [179] algorithm as implemented in DYNesty [177] was used to sample the posterior with 4000 live points and using the random walk sampling method with a walk length parameter of 2000.

4.2.3 Reweighting

An ever present issue in ringdown analyses is the choice of ringdown start time, t_0 , and this choice is closely related to the issue of the presence of an overtone. To address this issue, previous time-domain analyses [119, 185, 186] perform large numbers of Bayesian analysis runs with different choices of start time.

One key conceptual benefit of the frequency-domain approach of Chapter 3 is that the ringdown start time enters as a parameter of the model and can therefore be easily marginalised over, instead of simply being fixed (although, see Ref. [165] where the

ringdown start time was varied in a time-domain analysis). However, it is necessary to choose an informative (narrow) prior for the parameter t_0 .

A related computational benefit of our approach is that we can do a single Bayesian analysis run with a broad uniform prior on t_0 . We can then explore different, narrower priors by reweighting the results in post processing. This is an example of importance sampling (see, for example, Ref. [190]) and is the approach adopted here. This removes the need to perform the large number of runs used to explore the effect of varying the ringdown start time when performing time-domain ringdown analyses.

Given a model that depends on parameters $\vec{\theta}$, a likelihood $\mathcal{L}(\text{data}|\vec{\theta})$, and a prior $\pi(\vec{\theta})$, nested sampling can be used to draw a large number of samples $\vec{\theta}_i$ from the posterior, which is given by Bayes' theorem $P(\vec{\theta}|\text{data}) \propto \mathcal{L}(\text{data}|\vec{\theta})\pi(\vec{\theta})$. Samples from the posterior have associated weights w_i (samples may often be equally weighted with $w_i = 1$, but we do not require this to be the case). Such samples can be used to approximate integrals via a Monte-Carlo sum; $\int d\vec{\theta} P(\vec{\theta}|\text{data})f(\vec{\theta}) = \sum_i w_i f(\vec{\theta}_i)/W$, where $W = \sum_i w_i$. If we choose a new prior $\hat{\pi}(\vec{\theta})$, then the Bayesian posterior is given instead by $\hat{P}(\vec{\theta}|\text{data}) \propto \mathcal{L}(\text{data}|\vec{\theta})\hat{\pi}(\vec{\theta})$. We can define the new weights via

$$\hat{w}_i = w_i \frac{\hat{\pi}(\vec{\theta}_i)}{\pi(\vec{\theta}_i)}. \quad (4.4)$$

The same samples can then be used to approximate integrals of the form $\int d\vec{\theta} \hat{P}(\vec{\theta}|\text{data})f(\vec{\theta})$ via the Monte-Carlo sum $\sum_i \hat{w}_i f(\vec{\theta}_i)/\hat{W}$, where $\hat{W} = \sum_i \hat{w}_i$.

It is also possible to reweight the Bayesian evidence for the new choice of prior. In a GW context this approach has been used previously for inference with higher-order modes [191]. The Bayesian evidence (i.e. the normalisation denominator in Bayes' theorem) under the original prior is given by $Z = P(\text{data}) = \int d\vec{\theta} \mathcal{L}(\text{data}|\vec{\theta})\pi(\vec{\theta})$. The Bayesian

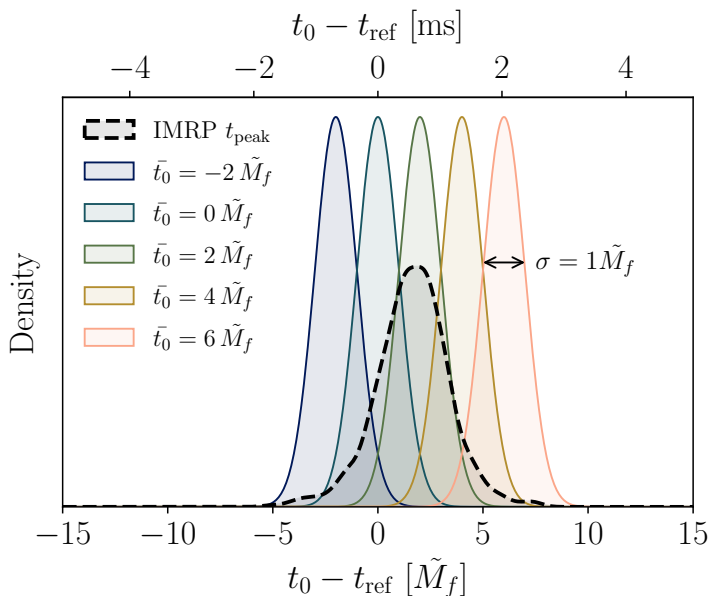


FIGURE 4.1: Our ringdown inference is run initially using a flat, uniform prior on the ringdown start time, t_0 , over the plot range $\pm 15\tilde{M}_f$ relative to t_{ref} (Hanford frame). In post processing, the posterior samples can be reweighted to a different choice of prior on t_0 (see Section 4.2.3). The different prior choices used in this chapter are shown in this figure. We use a sequence of narrow Gaussian priors (with different means \bar{t}_0 defined relative to t_{ref} and fixed standard deviation of $1\tilde{M}_f$) as well as using the posterior on the time of peak strain from a full IMR analysis as a prior.

evidence under the new prior, $\hat{\pi}(\vec{\theta})$, is $\hat{Z} = \int d\vec{\theta} \mathcal{L}(\text{data}|\vec{\theta}) \hat{\pi}(\vec{\theta})$. Using the reweighted samples to approximate the integral, it can be shown that the new evidence is given by

$$\hat{Z} = Z \frac{\hat{W}}{W}. \quad (4.5)$$

The process of reweighting to the new, target prior reduces the effective number of posterior samples available. For this not to be a problem, we require the original prior to have significant support across the target prior. Here, we reweight on just a single parameter, the ringdown start time t_0 . As described above, we use a uniform prior on t_0 as the original prior, π , in our analyses. For the target prior we use a variety of different choices, this removes the need for performing a large number of runs with different start

times. Our prior choices are plotted in Fig. 4.1. Narrow Gaussians centred at different start times are used to explore the start-time dependence on the results, and we use the notation \bar{t}_0 to indicate the mean of the Gaussian relative to t_{ref} . For more details on the t_0 reweighting, see Section 4.3.4.

We also use the posterior on t_{peak} from a full inspiral-merger-ringdown (IMR) analysis from Ref. [186], obtained with the IMRPHENOMPv2 (IMRP) waveform model [152], as another prior on t_0 . Our aim in doing this is to marginalise over our uncertainty on the ringdown start time, given that it occurs at the time of peak strain. This is motivated by the observations of Refs. [1, 94, 122–126, 128–130] described above, which show that generically the ringdown can be considered to start at around this time.

4.3 Results

There are several ways to investigate and quantify the evidence for additional QNMs in the ringdown. Section 4.3.1 contains the results of a series of analyses designed to study the presence of a possible overtone in GW150914. Section 4.3.2 contains the results of a series of analyses designed to test whether or not what has been detected really is an overtone and is not the accumulation of other effects. Section 4.3.3 describes further checks on the stability of the results, and Section 4.3.7 contains some additional results that further demonstrate the capabilities of the frequency-domain approach to ringdown analysis.

Throughout Secs. 4.3.1 and 4.3.2, we compare our results with those in Refs. [185] and [186]. This is done in the hope of helping to resolve the controversy over the evidence for a ringdown overtone in GW150914. However, it should be stressed that our results are produced using a very different method and care should therefore be taken in making direct comparisons. Although the frequency-domain analysis is formally equivalent to the

time-domain analysis in a particular limit (as discussed in the introduction, and in more detail in Section 3.2.2) we do not take this limit in a practical analysis. Furthermore, the frequency-domain analysis is further generalised with respect to the time-domain analysis in that it marginalises over parameters such as the sky position and ringdown start time (which are fixed in the analyses of Refs. [185, 186]). Results from our frequency-domain analyses should therefore not be expected to agree perfectly with those from previous time-domain analyses.

4.3.1 Presence of an overtone

In order to investigate the presence of an overtone in the GW150914 ringdown, we initially perform two analyses using the model described in Section 4.2.1: one analysis uses only the fundamental QNM ($N = 0$) and the other includes the first overtone ($N = 1$). Aside from the inclusion of the overtone in the ringdown (which introduces two additional parameters: an amplitude and a phase), these two analyses are otherwise identical.

In Fig. 4.2 we plot the posterior distributions on the remnant BH mass, M_f , and dimensionless spin, χ_f , for both of these analyses. Results are shown for the different choices of the prior on the ringdown start time shown in Fig. 4.1 (these results were obtained by reweighting the samples obtained with a flat prior using the approach described in Section 4.2.3). The earliest start time ($\bar{t}_0 = -2\tilde{M}_f$) is omitted from the fundamental-only ($N = 0$) plot in the left-hand panel of Fig. 4.2 because of a low number of posterior samples at these time (see Section 4.3.4). Also shown for comparison are the much tighter constraints resulting from the full IMR analysis. These IMR posterior samples were obtained from Ref. [192], which (as detailed in Refs. [119, 186]) are obtained from applying fitting formulas to the samples available at Ref. [193]. When only the fundamental QNM is used ($N = 0$), and when the analysis is started at early times (e.g. $t_0 - t_{\text{ref}} \lesssim -2\tilde{M}_f$) our




FIGURE 4.2: Posterior distributions on the remnant mass, M_f , and dimensionless spin, χ_f , for different choices of t_0 prior (the colours and line styles correspond to those used in Fig. 4.1). *Left*: the results from the $(2, 2, 0)$ fundamental-mode-only analysis (i.e. $N = 0$). *Right*: the results from the overtone analysis including the $(2, 2, 0)$ and $(2, 2, 1)$ modes (i.e. $N = 1$). Each line corresponds to a different choice of t_0 prior. Coloured lines correspond to Gaussians with widths of $1\tilde{M}_f$ and means \tilde{t}_0 (see Fig. 4.1). The dashed black line corresponds to using the posterior on time of peak strain (from a full IMR analysis) as our prior, which marginalises over uncertainty on the time of peak strain. Also shown for reference (dotted line) is the posterior from a full IMR analysis. The main panel shows the 90% confidence contours while the side panels show the one-dimensional marginalised posteriors.

posteriors on the remnant parameters are biased towards high values of M_f and χ_f . This behaviour is expected; a single QNM is only able to model the ringdown signal starting well after the time of peak strain. Including an overtone ($N = 1$) allows the ringdown analysis to start at earlier times, as can be seen by the removal of the bias in the right panel. This improvement is suggestive that the data supports the inclusion of an overtone. Additionally, using an earlier ringdown start time increases the SNR in the ringdown and reduces the posterior width; this effect can be seen in both the $N = 0$ and $N = 1$ analyses.

Our results in Fig. 4.2 can be compared to the corresponding results of the time-domain

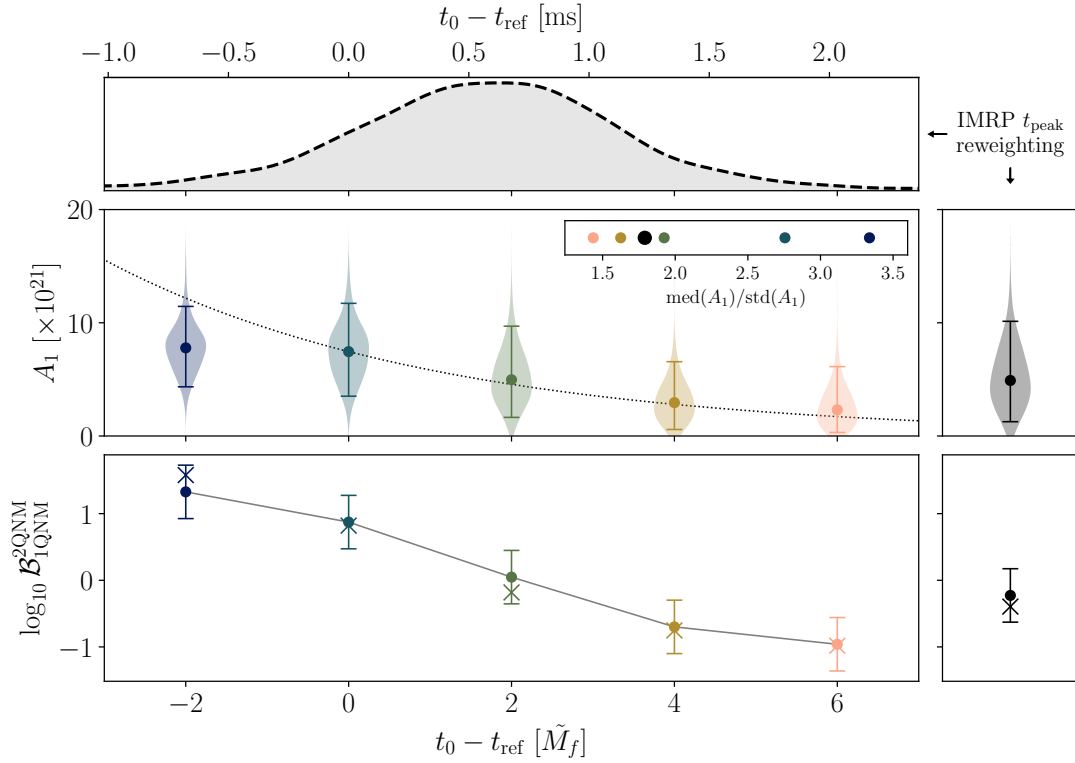


FIGURE 4.3: Posteriors on the overtone amplitude, and Bayes’ factors in favour of the overtone model for different choices of t_0 prior (the colours and line styles correspond to those used in Fig. 4.1). *Top*: posterior on the time of peak strain in the Hanford frame, from a IMRPHENOMPv2 analysis, as in Fig. 4.1 (and originally from Ref. [186]). *Middle*: overtone amplitude posteriors for different choices of t_0 prior. The left panel corresponds to Gaussian priors with standard deviation $1\tilde{M}_f$, centred at the time they are plotted. The dotted line indicates the expected exponential decay of the A_1 mode; this is included merely to guide the eye and was produced using the median mass and spin values from the full IMR analysis and the median value of A_1 from the $\bar{t}_0 = 0$ prior. The right panel corresponds to using the IMRPHENOMPv2 time of peak strain as a prior. For earlier start times the posteriors on the amplitude are peaked further away from zero; this is quantified in the inset plot where the ratio of the median to the standard deviation of the A_1 posterior is plotted. *Bottom*: the Bayes’ factor in favour of the overtone model for each prior choice; circles with error bars show the Bayes’ factor calculated from nested sampling (with errors estimated by the sampler) while the crosses show the results calculated using the Savage-Dickey density ratio.

analyses shown in Fig. 1 from Cotesta et al. [185] and Figs. 4 and 5 from Isi & Farr [186].

In general terms, there is broad agreement between all three sets of results. In particular, all three sets of authors find that the overtone analyses ($N = 1$) always gives results that

are more consistent with the IMR result and get increasingly broader for later choices of the ringdown start time. All sets of authors also find that for the fundamental-only analysis ($N = 0$) starting at early times (i.e. $t_0 - t_{\text{ref}} \lesssim 0$) leads to posteriors that are inconsistent with the IMR result. However, there are subtle differences between the various results. Our results with $N = 0$ and early start times gives posteriors biased to large values of M_f and χ_f ; this is also seen in Ref. [186], but not in Ref. [185] (where the posterior consistently reaches lower values of χ_f). Our results with $N = 0$ and late start times (i.e. $t_0 - t_{\text{ref}} \gtrsim 4\tilde{M}_f$) are partially consistent with the IMR results; this is also seen in Ref. [185], but not in Ref. [186] who never find consistency with the IMR result for any choice of start time. Finally, when including the overtone ($N = 1$) and starting at late times, Ref. [185] find results that are consistent with $\chi_f = 0$ (i.e. a Schwarzschild BH) at 90% confidence, in stark disagreement with Ref. [186] who find $\chi_f \gtrsim 0.2$. Our results are in better agreement with those of Ref. [186].

In the middle panel of Fig. 4.3 we investigate our $N = 1$ overtone analysis further by plotting the one-dimensional marginalised posteriors for the amplitude, A_1 , of the QNM overtone. An amplitude posterior peaked away from zero has been suggested (particularly by Ref. [119]) as one good indication for the presence of an overtone in the data. As expected, the QNM overtone decays quickly and when starting at later times we find a small value for the amplitude. The degree to which the A_1 posterior is peaked away from zero can be quantified using the ratio between the median and standard deviation; this is plotted in the inset of the middle panel of Fig. 4.3. For values of \bar{t}_0 between $-2\tilde{M}_f$ and $+6\tilde{M}_f$, we find posteriors on A_1 that are peaked away from zero at between 1.44 and 3.34σ . If we reweight using the IMRP t_{peak} prior, we find a posterior peaked away from zero at 1.79σ .

Our results in the middle panel of Fig. 4.3 can be compared to the corresponding results

of the time-domain analyses shown in Fig. 1 of Ref. [186] and Fig. 2 of Ref. [185]. All three sets of authors find values of A_1 that are smaller at later times, consistent with the expected exponential decay of the overtone, but they disagree on the absolute value of the amplitude and the significance with which a zero amplitude can be excluded. Refs. [119, 186] find the largest values; they report a posterior peaked 3.6σ away from zero. Ref. [185] finds much smaller values which are consistent with zero for many choices of start time. These analyses use essentially the same method and should therefore agree exactly. Our result, produced using a different method, lies somewhere in between; we do find nonzero values are preferred for a range of start times, but only with a modest significance of $\sim 1.79\sigma$ for our preferred IMRP t_{peak} prior which we consider to be the best description of our uncertainty on the ringdown start time.

The comparison of our results with those of Refs. [119, 185, 186] is complicated by the fact that we use subtly different definitions for the amplitude. The time-domain analyses naturally define the mode amplitudes at a fixed time, usually t_0 . Our frequency-domain analysis also defines the mode amplitudes at t_0 , but this start time is then varied as part of the analysis, blurring the exact time at which the amplitude is defined. This is a fairly small effect for the narrow Gaussian priors, but more significant for the wider IMRP t_{peak} prior. We can correct for this effect by rescaling all the overtone amplitudes to any fixed reference time (here we use t_{ref}) using the known decay rate for the QNMs;

$$A_{1,\text{ref}} = A_1 \exp\left(\frac{t_0 - t_{\text{ref}}}{\tau_{221}(M_f, \chi_f)}\right), \quad (4.6)$$

where $\tau_{221}(M_f, \chi_f)$ is the exponential decay time of the $(2, 2, 1)$ QNM and is a function of the remnant mass and spin. This rescaling can be done for any QNM and the resulting amplitude parameters $A_{\ell mn, \text{ref}}$ are more directly comparable with the amplitudes used in




FIGURE 4.4: Posteriors on the overtone amplitude from our $N = 1$ overtone analysis, rescaled to a fixed reference time of t_{ref} . The rescaling does not significantly affect the significance with which the posteriors are peaked away from zero. The colours and line styles indicate the prior used on t_0 and correspond to those used in Fig. 4.1.

time-domain analyses. Posteriors on $A_{1,\text{ref}}$ are shown in Fig. 4.4.

In the bottom panel of Fig. 4.3 we plot the Bayes' factors between the fundamental only ($N = 0$) and overtone ($N = 1$) analyses. This is defined as $\mathcal{B}_{1\text{QNM}}^{2\text{QNM}} = Z_{N=1}/Z_{N=0}$. The Bayes' factor has been suggested (particularly by Ref. [185]) as another good way for quantifying the support for an overtone in the data. The Bayes' factor was computed in two different ways. Firstly, DYNesty was used to calculate the evidences $Z_{N=0}$ and $Z_{N=1}$ for both of the analyses described above, and these were reweighted to the desired t_0 prior using Eq. 4.5. Nested sampling also returns an estimate for the error on the evidences, and these are used to plot the error bars in Fig. 4.3. Secondly, exploiting the fact that the $N = 0$ model is nested within the $N = 1$ model, the Bayes' factors were computed using the posterior on A_1 from the $N = 1$ analysis to find the Savage-Dickey density ratio [194].

Our results in the bottom panel of Fig. 4.3 can be compared to the corresponding results of the time-domain analyses shown in Fig. 7 of Ref. [186] and Fig. 2 of Ref. [185].

Ref. [185] computes the Bayes' factors using the ratio of evidences evaluated with nested sampling, whereas Ref. [186] computes Bayes' factors using Savage-Dickey density ratios. All sets of authors find Bayes' factors that decrease for later ringdown start times, although they disagree on the exact value. Ref. [186] finds the strongest log-Bayes' factor of ~ 1.7 at $t_0 - t_{\text{ref}} \sim 0$. Ref. [185] finds a slightly negative log-Bayes' factor starting at this time. Again, our result lies somewhere in between, we find a moderate log-Bayes' factor of ~ 1.0 when marginalising over a narrow prior on t_0 centred at this time. If we instead marginalise over the time of peak strain using the broader IMRP t_{peak} prior, the log-Bayes' factor is slightly negative. However, as discussed in Section 4.4 below, we consider the actual values of the Bayes factors to be less important than their trend with varying start time.

4.3.2 The nature of the overtone

The results of the previous section show that there is tentative evidence for something beyond the fundamental $(2, 2, 0)$ mode in the GW150914 data. In the previous section it was assumed that this is the $(2, 2, 1)$ QNM overtone; this is motivated by our expectations from NR experiments (see, for example, Ref. [122]). In this section, we address this assumption by measuring the frequency and amplitude of the QNM overtone and comparing with the expectations from GR.

Fig. 4.5 shows the results of a third ringdown analysis that also includes two QNMs. In this analysis the complex frequency of the second QNM is allowed to deviate from the Kerr overtone value. This differs from the $N = 1$ overtone analysis described above, where the frequency of the overtone was fixed by the remnant mass and spin to the Kerr value, $\omega_{221} = 2\pi f_{221} - i/\tau_{221}$. Recovering a value of δf consistent with zero has been suggested (particularly by Ref. [186]) as further evidence for the presence of an overtone; otherwise, it might be expected that the extra parameters would fit to the noise and would not recover




FIGURE 4.5: Posteriors on the deviation from the Kerr value for the real part of the overtone frequency. The colours and line styles distinguish different choices for the t_0 prior and correspond to those used in Fig. 4.1. The mode frequency is given by $f_{221}^{\text{Kerr}} \exp(\delta f)$, so that $\delta f = 0$ is the expected result for the Kerr metric. For all choices of t_0 prior the data is consistent with $\delta f = 0$.

the Kerr value. We use the parameterisation from Ref. [186]; the complex frequency of the second QNM is now $\omega_{221} = 2\pi f - i/\tau$, where $f = f_{221} \exp(\delta f)$ and $\tau = \tau_{221} \exp(\delta\tau)$. This introduces the two new dimensionless parameters δf and $\delta\tau$ into the model, for which we use uniform priors in the range $[-0.5, 0.5]$. The $\delta\tau$ parameter is not well constrained, therefore we focus on δf . We find posteriors on δf consistent with zero for all choices of t_0 prior with standard deviations ~ 0.2 . This is consistent with what was found in Ref. [119] and can be viewed as a test of the no-hair theorem at the $\sim 20\%$ level.

Our results in Fig. 4.5 can be compared with Fig. 2 of Ref. [186]. Our preferred run, using the IMRP t_{peak} prior on t_0 , is broadly consistent with that result. However, there is less noticeable broadening of our posterior for later choices of the ringdown start time (as was found by Ref. [186]). A broadening would be expected if an overtone was present, as both the overtone amplitude and ringdown SNR decay with later ringdown start times. Since in our analysis we obtain a posterior on the ringdown start time, we can also plot




FIGURE 4.6: Two-dimensional posteriors on the deviation from Kerr parameter (as in Fig. 4.5) and ringdown start time. The 90% contours are plotted, with the colours and line styles indicating the t_0 prior and correspond to those used in Fig. 4.1.

two-dimensional posteriors of the δf parameter and t_0 ; this is shown in Fig. 4.6. With the exception of the earliest $\bar{t}_0 = -2\tilde{M}_f$ posterior, with this representation it is easier to make out the slight broadening of the posterior at later times. The breakdown of this trend for the earliest t_0 deserves more attention, perhaps indicating sampling issues or unexpected model behaviour.

To investigate the consistency with Kerr further, we use the results of the ringdown analysis where the frequency of the overtone is allowed to vary freely to address another important question. If the data does indeed contain a second QNM, can we determine which mode it is? Theoretical studies of NR simulations suggest that the $(\ell, |m|, n) = (2, 2, 1)$ will be the next most prominent, especially for early start times [122]. In Fig. 4.7 we plot the posterior on the dimensionless complex frequency (allowing both the real and imaginary parts to vary freely) of the second QNM, ωM_f for the preferred IMRP t_{peak} reweighting. This plot uses the value for M_f calculated from the complex frequency of the first QNM, assuming this is the expected $(2, 2, 0)$ fundamental mode of Kerr. We find that we can confidently conclude that the second mode is an overtone ($n \geq 1$) but




FIGURE 4.7: The posterior on the dimensionless complex frequency of the second QNM (50% and 90% regions), assuming the first is the fundamental $(\ell, |m|, n) = (2, 2, 0)$ mode and for the preferred IMRP t_{peak} reweighting. Lines indicate the Kerr frequencies parameterised by the remnant spin; dots and crosses indicate points with $\chi_f = 0.7$ and 0 respectively. Lines are coloured according to their ℓ and n indices and the m index increases left to right in each set. The frequency of the second QNM is consistent with the expected $(2, 2, 1)$ overtone, but also with several other modes. However, all fundamental modes (those with $n = 0$) are excluded.

that it is not possible to say from the data alone exactly which overtone. For example, the modes $(2, 2, 1)$ and $(2, 1, 1)$ are both equally compatible with the data. In general, when searching for additional QNMs it is necessary to be guided by our prior expectations regarding which modes are expected to be excited with the highest amplitudes.

We now turn our attention to the measured amplitude A_1 and whether this matches the theoretical expectations for the $(2, 2, 1)$ overtone. For convenience, we choose to work with the amplitude ratio A_1/A_0 which eliminates factors common to all modes, such as the distance to the source. Because the two QNMs decay exponentially at different rates, the amplitude ratio depends strongly on the chosen ringdown start time. Our two-dimensional posteriors on the amplitude ratio and ringdown start time are plotted in Fig. 4.8. As expected we find that the amplitude ratio decreases for later start times, and the error on the amplitude ratio increases for later start times because of the decreasing SNR in the

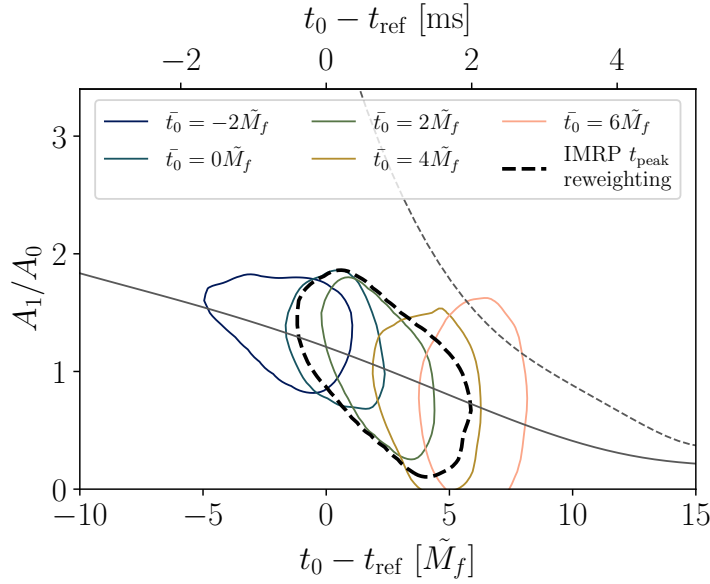


FIGURE 4.8: Posteriors on the amplitude ratio A_1/A_0 from our $N = 1$ overtone analysis. The 90% contours are plotted, with the colours and line styles indicating the t_0 prior and correspond to those used in Fig. 4.1. The solid grey curve shows the results of a two-QNM fit to the NR simulation SXS:BBH:0305 which has parameters consistent with GW150914. The dashed grey curve shows the results of a multi-QNM fit to SXS:BBH:0305 which follows closely the expected exponential decay rate for the amplitude ratio.

ringdown.

In order to check whether this is consistent with the theoretical expectation for an overtone we compare with fits to the NR simulation SXS:BBH:0305 [36] which has parameters consistent with GW150914. Fixing the remnant mass and spin to the values reported in the simulation metadata, we perform QNM least-squares fits to this simulation for a range of ringdown start times using the method developed in Chapter 2 (see Section 2.3 for details on the fitting implementation, and Ref. [7] for the code). Results are shown in Fig. 4.8 for two such fits. Firstly, we performed a two-QNM fit intended to mimic the analysis of the real GW150914 data described in Section 4.3.1 above. In this analysis the ${}_{-2}Y_{22}$ spherical harmonic mode of the simulation is modelled as a sum of the $(2, 2, 0)$ and $(2, 2, 1)$ QNMs and the amplitude ratio is recorded. The results from this two-QNM fit

agree very well with what is seen in the real data giving us further confidence that there is nothing unexpected present in the data and that our results are not unduly affected by noise fluctuations (see also the discussion in Section 4.3.3). Secondly, we perform a full multi-QNM fit to all the spherical harmonic modes (up to and including $\ell = 8$) with a ringdown model that includes all QNMs (including both prograde and retrograde modes) up to $\ell = 8$ and $n = 7$ (1232 QNMs in total). The ratio of the amplitudes of the $(2, 2, 0)$ and $(2, 2, 1)$ prograde modes from this fit behaves very differently; the ratio follows very closely an exponential time evolution which can be understood in terms of the difference between the two QNM decay times.

The fact that the two-QNM analysis gives a very different amplitude ratio compared to the full multi-QNM analysis for ringdown start times near the peak strain is related to the extreme destructive interference observed in the QNM overtone fits of Refs. [1, 94, 122–126, 128–130] with large values of N . This shows that the amplitude A_1 recovered from a two-QNM analysis is not purely the amplitude of the first overtone but also includes significant contributions from higher overtones and other harmonics. However, absorbing these contributions into the first overtone introduces a systematic bias in the remnant properties that is smaller than the statistical uncertainty; this can be seen in, for example, Fig. 4.2 and Section IV C of Ref. [122]. For this reason, it still makes sense to describe the results of the two-QNM analysis as a measurement of the overtone, even though there are undoubtedly other contributions present in the signal.

4.3.3 The effect of noise and sampling frequency

One of the key claims made in Ref. [185] was the overtone detection was highly sensitive to noise fluctuations. This was disputed by Ref. [186]. In order to address this issue, we performed a noise injection study mirroring closely what was done in Ref. [185]. The

results of this injection study are presented in Section 4.3.5. As expected, the results of injecting into different noise realisations show some scatter. However, this scatter is not larger than expected and we are unable to reproduce the claim in Ref. [185] with our (very different) analysis method.

It has been suggested [195] that the results of ringdown analyses, in particular those including overtones, might be sensitive to aliasing effects when using downsampled strain data due to the reduced Nyquist frequency. QNM overtones, $(\ell, m, n \geq 1)$, have roughly the same real part of the frequency as the corresponding fundamental, $(\ell, m, 0)$, mode, but they have a shorter damping time. See, for example, Fig. 4.7. This means that if a single, isolated, mode is viewed in Fourier space, the power spectrum is broader and contains significant power at higher frequencies.

Early ringdown studies, including those in Refs. [119, 181, 186], generally used strain data that had been downsampled to 2 kHz. This was done for convenience and computational speed and was not originally anticipated to be a problem because the merger of GW150914 occurs at ~ 200 Hz, safely below the Nyquist frequency.

In this chapter the 4 kHz data is used for all the analyses in the main text. Additionally, a frequency-domain log-likelihood with an upper integration limit of $f_{\text{high}} = 1000$ Hz was used. Our method is very different from the time-domain analyses, and we do not expect our results to be sensitive to small changes in these choices. To check that this is the case we have repeated the $N = 1$ overtone analysis using the 16 kHz sampled data (obtained from Ref. [188]) and we find no significant changes in our results. Using reweighting techniques (this time applied to the likelihood) we have also investigated the effect of changing the upper limit of integration in the likelihood. By re-evaluating the $N = 1$ posterior chain on a likelihood with $f_{\text{high}} = 1500$ Hz and 2000 Hz, and then reweighting, we again find no significant changes in our results.

4.3.4 Posteriors on the ringdown start time

As described in Section 4.2.3, we initially perform Bayesian inference on the ringdown using a broad, flat prior on the ringdown start time parameter t_0 . The posteriors on t_0 from both the $N = 0$ and $N = 1$ analyses are shown in Fig. 4.9. We do not consider these posteriors to be physically meaningful results because they were obtained with a prior that does not correctly describe our state of knowledge about when the ringdown should start. These results are produced merely as an intermediate step in our analysis, before the reweighting was applied, and are shown here only to further illustrate the reweighting procedure described in Section 4.2.3.

In order for the subsequent reweighting step to be accurate, it is necessary for the posterior chains (particularly for the $N = 1$ overtone analysis) to contain samples across




FIGURE 4.9: Posteriors on the Hanford ringdown start time obtained from our initial analysis using a flat prior over the range shown in the plot. Results are shown for the fundamental only ($N = 0$) and overtone ($N = 1$) analyses. Vertical coloured lines show the locations of the means \bar{t}_0 of the narrow Gaussian priors used for the subsequent reweighting (see Fig. 4.1). The $N = 1$ posterior has ample support across the entire range of interest, as required for the reweighting to remain accurate. The $N = 0$ posterior has enough support everywhere except the $\bar{t}_0 = -2\tilde{M}_f$ prior.

the range of start times that we consider. For this reason, the DYNESTY sampler settings described in Section 4.2.2 were chosen to ensure a large number of posterior samples were produced; we obtained 203697 and 218882 posterior samples from the $N = 0$ and $N = 1$ analyses respectively. This is sufficient for the reweighting to remain accurate everywhere except for the earliest start time in the $N = 0$ analysis. This is the reason why this result is omitted from Fig. 4.2.

4.3.5 Injection study

Closely following the injection study performed in Ref. [185], we inject GW150914-like signals in the instrumental noise surrounding the true GW150914 event and reperform our overtone analysis ($N = 1$).

The $\ell = 2$ spin-weighted spherical harmonic of the NR simulation SXS:BBH:0305 [36] was used as the mock signal, scaled to a total mass of $72 M_{\odot}$ and injected with a face-off orientation at a luminosity distance of 410 Mpc. The sky position was taken to be $\alpha = 1.95$ rad, $\delta = -1.27$ rad. This signal was injected into the data surrounding GW150914, such that the peak of the absolute value of the strain occurred at times $[-20, -15, -10, 5, 15, 20, 25, 30, 35, 40]$ s relative to t_{ref} . These choices ensure the mock signal does not overlap with the real event. Additionally, a zero-noise injection was performed for comparison.

We performed the frequency-domain ringdown analysis on these mock datasets using the same setup as was used for the real data and as described in Section 4.2.2. This includes using the same PSD in the likelihood for all datasets. We plot the resulting posteriors on the overtone amplitudes in Fig. 4.10. As with the real data, prior reweighting (see Section 4.2.3) has been used to show results for different choices of the ringdown start time prior. We also investigated the Bayes' factors and found the same declining trend.

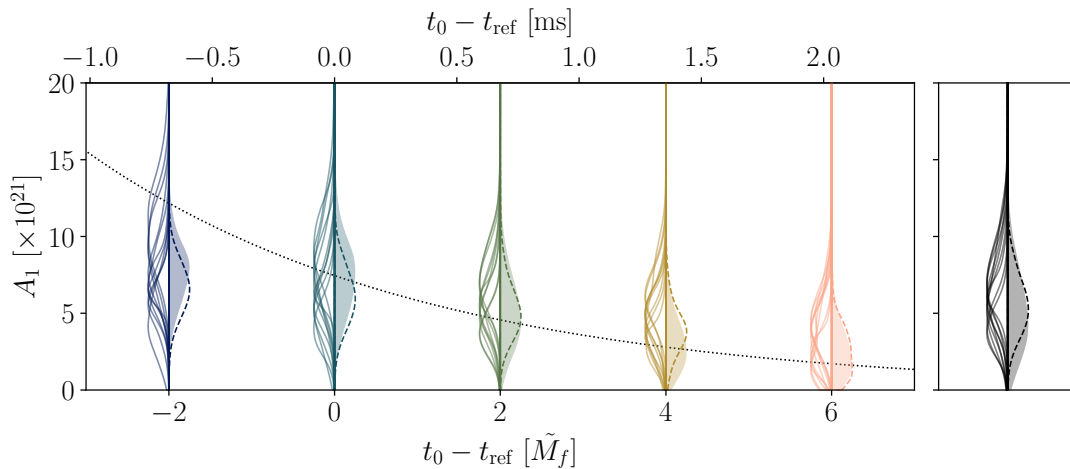


FIGURE 4.10: This is similar to the middle panel of Fig. 4.3 in the main text, but shows the posteriors on the overtone amplitude from the noise injection study. The different violin plots are for the different priors on the ringdown start time and the colours are the same as those used in Figs. 4.1 and 4.3. On the right-hand side of each set of violin plots, the filled posterior shows the result obtained using the real GW150914 data (this is the same as what is plotted in Fig. 4.3). On the left-hand side are all the posteriors from the injection campaign, which indicate the spread in results due to different noise realisations. Finally, the dashed lines on the right-hand side are the posteriors from the zero-noise injection. This plot is intended to be compared to Fig. 2 of Ref. [185], and Fig. 6 of Ref. [186].

As expected, different noise realisations introduce some scatter into the results and we observe a spread in the locations of the maximum posterior values for the overtone amplitudes. However, this spread is consistent with the width of the posterior. The analysis described in the main text found only tentative evidence for the overtone, but there is no indication that this is overly affected by noise fluctuations.

4.3.6 Wavelet posteriors

The frequency-domain ringdown analysis method described in Chapter 3 and used here marginalises over the early-time inspiral-merger signal using a flexible combination of sine-Gaussian wavelets (see Eq. 4.2). In the GW150914 analyses presented in this chapter $W = 3$ wavelets were used. This choice was found empirically to be large enough to model




FIGURE 4.11: Posteriors on selected parameters for the $W = 3$ wavelets used in the $N = 1$ overtone analysis, reweighted using the IMRP t_{peak} prior on the ringdown start time. *Left:* the wavelet central times, η_w . *Middle:* the wavelet widths, τ_w . *Right:* the wavelet frequencies, ν_w . The index runs over values $w = 1, 2$ and 3 , where the numbering of the wavelets is chosen to enforce the ordering $\nu_w < \nu_{w+1}$. All plots use SI units on the upper x -axis and natural units on the lower x -axis.

the inspiral-merger signal without biasing the ringdown inference. We have also verified that no strong correlations are observed between the wavelet and QNM parameters, and that further increasing the number of wavelets does not significantly affect the results for physically meaningful parameters (such as remnant parameters M_f and χ_f). These tests are described further in Section 3.4.

The whitened strain posterior on the sum of these wavelets, together with the QNMs, can be seen in the early-time signal in Fig. 4.13 where the fit to the data is seen to be excellent. The wavelet parameters themselves are not physical; the wavelets are being used here solely to marginalise out the inspiral-merger. Nevertheless, in this section we show some additional posterior plots on the wavelet parameters, see Fig. 4.11. As expected, in order to describe the “chirping” inspiral signal, the wavelets naturally order themselves with their amplitudes and frequencies increasing with time.

4.3.7 Other results

One important benefit of the frequency-domain approach to ringdown data analysis introduced in Chapter 3 and used here is that it naturally allows us to search (and hence to numerically marginalise) over source sky position and ringdown start time. This should be contrasted with the treatment of these parameters in most time-domain analyses where these parameters are fixed, potentially biasing the results. (Although it is technically possible to search over the sky in a time-domain analysis [121, 165], this is rarely done in practice.) To emphasise this, we plot the posterior on the sky location of GW150914 from our $N = 1$ overtone analysis reweighted to the IMRP t_{peak} prior on the ringdown start time. This can be compared with the publicly available LIGO sky posterior for GW150914 obtained using the samples from Ref. [196]. This is shown in Fig. 4.12. As discussed in

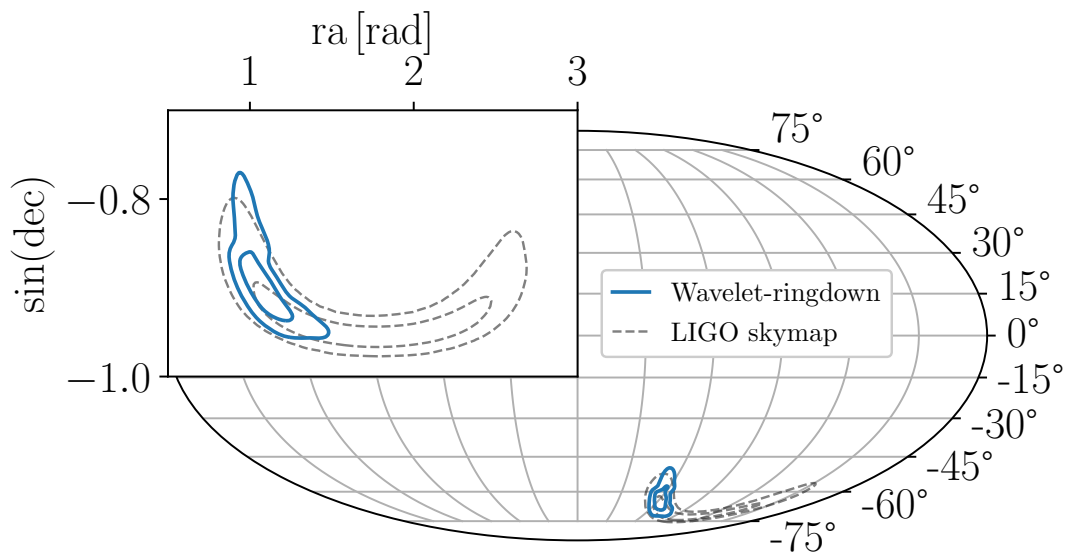


FIGURE 4.12: Posterior on the source sky position using geocentric coordinates in Mollweide projection. Shown in blue is the results from the $N = 1$ overtone analysis using the IMRP t_{peak} reweighting for the prior on the ringdown start time. The LIGO skymap for this event is shown by the dashed black line for comparison. The inset plot shows a zoomed-in map plotted using right ascension and the sine of the declination. In both cases, 50% and 90% contours are plotted.




FIGURE 4.13: Posterior on the reconstructed whitened waveform. Shown in grey is the strain data from both LIGO interferometers (*top*: Hanford, *bottom*: Livingston) whitened according to the noise amplitude spectral density in the detector and bandpass filtered between 32 and 512 Hz for clarity. Shown in blue is the waveform reconstruction from the $N = 1$ overtone analysis with the IMRP t_{peak} reweighting for the prior on the ringdown start time. The blue lines and shaded regions indicate the median and the 90% credible interval. The signal is plotted as a function of time from t_{ref} using both SI and natural units on the upper and lower x -axis respectively.

Chapter 3 (see the discussion around Fig. 3.8), it should be emphasised that this sky posterior is not a ringdown-only result because much of the information is also coming from the wavelets used to model the inspiral-merger portion of the signal. Our result is consistent with the LIGO skymap, but we note that our posterior does not include as much of the sky. In particular, our result only has support in a section of the LIGO arc (which itself corresponds to a section of a constant-time-delay ring). This may be due to a sampling issue, or it is possible that the choice to fix the inclination is responsible. We choose to model the signal as being face-off, and is partly motivated by aiding comparison with previous analyses (which also fix the inclination in this manner). Since the inclination is degenerate with amplitude, this choice is also conservative in the sense that it minimises value of the recovered overtone amplitude (a high-amplitude off-axis system will give the

same signal as a on-axis lower-amplitude system). The LIGO analysis is consistent with a face-off inclination, hence this choice in the analysis, but also shows some support at off-axis values. The coupling of the inclination to the GW amplitude and relative weights of the plus and cross polarizations may lead to a wider support on the sky.

Because the inspiral and merger parts of the signal are being modelled using truncated wavelets as part of the frequency-domain ringdown analysis, this allows us to plot a full waveform reconstruction from our results. This reconstruction is shown in Fig. 4.13 for our $N = 1$ overtone analysis reweighted to the IMRP t_{peak} prior. The full waveform model used in our analysis is discontinuous at t_0 . However, as discussed in Chapter 3, the whitened waveform reconstruction plotted here is smooth; this is a result of marginalising over the location of the discontinuity at t_0 , the waveform model “learning” the continuity from the data, and the whitening process (i.e., the PSD) used to make the figure. This waveform reconstruction uses the posterior on all of the model parameters, including those for the wavelets; more details on these parameters are given in Section 4.3.6.

4.4 Conclusions

The main motivation for this work comes from the ongoing discussion in the literature about whether a ringdown overtone can be confidently detected in the GW150914 data. In particular, the detection claim made in Ref. [119] was disputed by Ref. [185] where a nearly identical time-domain analysis was reperformed (see also the reply Ref. [186]). Applying the frequency-domain ringdown analysis originally presented in Chapter 3, we contribute to this discussion with a thorough reanalysis of the GW150914 data. This includes performing analyses with and without an overtone while considering different ringdown start times, as well as performing a noise injection study and studying the effects

of different data sampling rates and frequency integration limits on our results. Although the method used here differs significantly from previous time-domain analyses, we present our results in a way that makes it as easy as possible to compare with earlier work. In conclusion, we do find tentative evidence for a ringdown overtone, but not at the high level of significance originally claimed in Ref. [119].

In order to be more quantitative, it is first necessary to be able to say clearly what it even means to “detect a overtone”. Although intuitively obvious, it is not clear how to make this notion precise (this issue has previously been discussed in Ref. [186]). Several approaches have been suggested: looking to see if including the overtone improves the posterior on the remnant parameters (see Fig. 4.2); looking at the posterior on the overtone amplitude for a range of start times (see middle panel of Fig. 4.3); computing the Bayes’ factor in favour of an overtone (see bottom panel of Fig. 4.3); and allowing the frequency of the second QNM to vary freely to see if the data prefers, or at least is consistent with, the expected Kerr value (see Figs. 4.5, 4.6, and 4.7). Although these are not all independent from one another, they all help shed light on which QNMs are present. The results of all of these tests can also be compared to results from a noise injection study.

As well as not being completely independent of each other, none of these tests are, by themselves, sufficient to justify a claim of a detection. For example, one issue that has been raised is that the Bayes’ factor can be made to take any value with a suitable adjustment to the prior range. There are also conceptual problems regarding what it means to compare two models, neither of which is expected to fully describe the data. Here we are comparing the fundamental-only mode model (with a single QNM) to the overtone model (with two QNMs) when our firm prior belief is that the true signal should contain an infinite number of QNMs plus additional corrections (e.g. from nonlinearities in the merger, tails, and memory effects).

From the above discussion, it is clear that ringdown analyses are rather subtle. We think our frequency-domain method has some important advantages over what has been done before. For example, it marginalises over the ringdown start time and sky position which is preferable to fixing these parameters (which potentially introduces systematic biases). Ideally, we should also marginalise over the uncertainties in the noise power spectral density (see, e.g., Ref. [170]) and detector calibration (see, e.g., Ref. [197]) as part of a ringdown analysis. The ability to do this is, in principle, another benefit of the frequency-domain analysis approach used here as this can be done using techniques that are standard in the field.

Since the time of writing, the spin-weighted spheroidal harmonics have been implemented in the waveform model (via the `PYKERR` [198] package). This will allow the inclination to be marginalised over, which will be important for analysis of other events with less well constrained inclination (and getting closer towards the goal of marginalising over all uncertainty). Going forwards, the unexpected results at an early ringdown start time from the Kerr-deviation tests (Figs. 4.5 and 4.6) warrants further investigation. This type of analysis was not considered in the injection studies of Chapter 3, and a controlled test would be helpful to see if the model behaves as expected for this sort of test.

We stress that while our results have been compared with those of previous time-domain studies, our frequency-domain method is rather different and therefore we do not expect to find perfect agreement. In contrast, the results of Refs. [119, 185, 186] are produced using essentially identical methods and should therefore be expected to agree exactly. The reason for the disagreement that is seen there is currently unknown and the subject of an ongoing investigation by both sets of authors. It is vitally important for QNM science that all results are reproducible. To that end we have made all our data products and plotting scripts publicly available at Ref. [6].

If QNMs are going to fulfill their promise for testing GR, fundamental physics and the Kerr metric hypothesis, then the community must be able to agree on standards for what it means to detect them and to be able to robustly quantify their significance. This field is still very young, and that there is already significant controversy regarding the QNM content of GW150914 and GW190521 is concerning, and we risk the situation becoming more confused with many more suitable events expected in O4. And, as discussed in the introduction, this is a conceptual issue that will not be resolved with more observations, even at higher SNRs. This issue needs input from the whole community; however, we suggest that (as a minimum) future claims of an overtone detection are accompanied by the investigations in Fig. 4.2, both panels of Fig. 4.3 and Fig. 4.5. That is, posteriors on the remnant properties with and without the overtone, posteriors on the overtone amplitude, a study of the Bayes' factor trends for different start times, and posteriors on deviations from Kerr when the overtone frequency is allowed to vary.

Chapter 5

Conclusions and Prospects

In this thesis we have investigated the ringdown in two ways; fits of ringdown models to noiseless NR simulations inform us about what to look for in GW observations, and analysis of GW data tests our expectations. For the former, we performed (for the first time) a systematic analysis of QNM fits to precessing BBH mergers. It was shown that overtones alone are not sufficient to model the ringdown at early times in this extended parameter space, and even the inclusion of mirror modes and mode mixing (although improving the model in many instances) fail in some cases. Although beyond the limits of current detector sensitivities, these studies are valuable for understanding the nature of the ringdown, and also for testing our waveform models. We also showed that a current state-of-the-art surrogate waveform, NRSur7dq4, is not sufficiently accurate in the ringdown for studies of this sort.

Regarding the nature of the ringdown, there remain many unanswered questions in ringdown modelling. It is still unclear if ringdown overtones are actually physical, despite convincing evidence for their effectiveness in fitting ringdown waveforms and improving constraints on remnant properties. Their tendency to destructively interfere with each other, resulting in inflated amplitudes, and their rapid decay naturally raises suspicions of over-fitting. Where, exactly, the nonlinearities from the merger go (if the linear ringdown

model can indeed be extended to early times) is also the subject of ongoing study [137].

In principle, the QNM amplitudes encode information about how the remnant BH was perturbed (i.e., the initial binary configuration). Mappings between aligned-spin binary properties and QNM amplitudes have been found [93], but a similar relationship for precessing systems remains elusive. This is closely related to the issue of ringdown start time, for which choosing a reliable value becomes even more complicated in the precessing case. With the code [7] developed for the work in Chapter 2, studies of QNM amplitudes is a natural target. Since the publication of Ref. [1] the code has been developed to incorporate multimode ringdown fitting, and also has the capability to do fits of quadratic QNM frequencies. A second-order effect, quadratic QNMs can be thought of as QNMs sourced by their first-order counterparts. These “nonlinear” QNMs have now been identified in NR waveforms [199, 200], hinting at their importance in waveform modelling.

But, it is now clear that QNM studies are being limited by the accuracy of state-of-the-art NR waveforms. Typically, NR waveforms extrapolate the signal to large distances from the source. However, this is known to be sub-optimal for the ringdown. In particular, it is known that the extrapolation technique does not capture GW memory effects, and post-processing must be performed to ensure the remnant BH is in a suitable frame (for example, to perform frame rotations and to remove the remnant kick). In future work it will be important to utilise CCE waveforms [130], which allow for more control over subtle choices of the waveform frame. This machinery provides opportunities for studying in detail topics such as nonlinear QNMs, a non-inertial ringdown frame, and GW memory, and is readily compatible with our code.

Turning to the analysis of real data, the ringdown presents some unique challenges. These include low SNR signals, uncertainties with mode content and start time, and that it is more natural to work in the time domain (due to the ringdown’s abrupt start) instead

of the more commonly used frequency domain. Differences in analysis choices have already led to disagreements in the literature regarding the QNM content present in the data; in particular, the presence of an overtone in GW150914, and the presence of a higher harmonic in GW190521. To resolve this uncertainty we have developed a frequency-domain ringdown analysis pipeline. This approach is introduced in Chapter 3, and the code is available at Ref. [8]. By employing sine-Gaussian wavelets to model the pre-ringdown signal we can effectively marginalise it out, leaving a ringdown-only analysis. With this approach we can move back into the frequency domain and make use of well-established analysis methods.

The method was then employed to quantify evidence for a GW150914 overtone in Chapter 4, contributing to an ongoing discussion in the literature by marginalising over multiple uncertainties to report a single significance which was not possible before. There is also scope for further studies, in particular for GW190521, and the method can be readily applied to other GW events with any choice of QNM content. With a frequency-domain likelihood we can marginalise over uncertainty in the noise curve, and with closed-form waveform expressions there is also scope for speeding up sampling (for example, by using the waveform derivatives).

My studies at Birmingham have coincided with the birth of GW astronomy, and the field is rapidly evolving. From the first measurement of a GW signal in 2015, we now have a catalog of nearly 100. And with the fourth LIGO-Virgo observing run due to start soon, we expect hundreds of additional events over the next few years. This abundance of data leaves plenty of opportunities for the application of tools and techniques I have developed during the PhD, providing greater insight into the ringdown and the nature of gravity.

Bibliography

- [1] Eliot Finch and Christopher J. Moore. “Modeling the ringdown from precessing black hole binaries”. *Phys. Rev. D* 103.8 (2021), p. 084048. DOI: [10.1103/PhysRevD.103.084048](https://doi.org/10.1103/PhysRevD.103.084048). arXiv: [2102.07794](https://arxiv.org/abs/2102.07794) [gr-qc].
- [2] Eliot Finch and Christopher J. Moore. “Frequency-domain analysis of black-hole ringdowns”. *Phys. Rev. D* 104.12 (2021), p. 123034. DOI: [10.1103/PhysRevD.104.123034](https://doi.org/10.1103/PhysRevD.104.123034). arXiv: [2108.09344](https://arxiv.org/abs/2108.09344) [gr-qc].
- [3] Eliot Finch and Christopher J. Moore. “Searching for a ringdown overtone in GW150914”. *Phys. Rev. D* 106.4 (2022), p. 043005. DOI: [10.1103/PhysRevD.106.043005](https://doi.org/10.1103/PhysRevD.106.043005). arXiv: [2205.07809](https://arxiv.org/abs/2205.07809) [gr-qc].
- [4] Eliot Finch and Christopher J. Moore. *Supplementary material for "Modelling the Ringdown from Precessing Black Hole Binaries"*. Zenodo, Feb. 2021. DOI: [10.5281/zenodo.4538193](https://doi.org/10.5281/zenodo.4538193). URL: <https://doi.org/10.5281/zenodo.4538193>.
- [5] Eliot Finch and Christopher J. Moore. *Posterior samples for "Frequency-Domain Analysis of Black-Hole Ringdowns"*. Zenodo, Oct. 2021. DOI: [10.5281/zenodo.5569667](https://doi.org/10.5281/zenodo.5569667). URL: <https://doi.org/10.5281/zenodo.5569667>.
- [6] Eliot Finch and Christopher J. Moore. *Posterior samples for "Searching for a Ringdown Overtone in GW150914"*. Zenodo, May 2022. DOI: [10.5281/zenodo.6554512](https://doi.org/10.5281/zenodo.6554512). URL: <https://doi.org/10.5281/zenodo.6554512>.
- [7] Eliot Finch and Christopher J. Moore. *qnmfits*. URL: <https://github.com/eliotfinch/qnmfits>.
- [8] Eliot Finch and Christopher J. Moore. *fdringdown*. URL: <https://github.com/eliotfinch/fdringdown>.
- [9] J. Aasi et al. “Advanced LIGO”. *Class. Quant. Grav.* 32 (2015), p. 074001. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001). arXiv: [1411.4547](https://arxiv.org/abs/1411.4547) [gr-qc].
- [10] F. Acernese et al. “Advanced Virgo: a second-generation interferometric gravitational wave detector”. *Class. Quant. Grav.* 32.2 (2015), p. 024001. DOI: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001). arXiv: [1408.3978](https://arxiv.org/abs/1408.3978) [gr-qc].
- [11] 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 (2019), p. 031040. DOI: [10.1103/PhysRevX.9.031040](https://doi.org/10.1103/PhysRevX.9.031040). arXiv: [1811.12907](https://arxiv.org/abs/1811.12907) [astro-ph.HE].

- [12] R. Abbott et al. “GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run”. *Phys. Rev. X* 11 (2021), p. 021053. DOI: [10.1103/PhysRevX.11.021053](https://doi.org/10.1103/PhysRevX.11.021053). arXiv: [2010.14527](https://arxiv.org/abs/2010.14527) [gr-qc].
- [13] R. Abbott et al. “GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run” (Aug. 2021). arXiv: [2108.01045](https://arxiv.org/abs/2108.01045) [gr-qc].
- [14] R. Abbott et al. “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run”. (Nov. 2021). arXiv: [2111.03606](https://arxiv.org/abs/2111.03606) [gr-qc].
- [15] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. *Phys. Rev. Lett.* 119.16 (2017), p. 161101. DOI: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101). arXiv: [1710.05832](https://arxiv.org/abs/1710.05832) [gr-qc].
- [16] R. Abbott et al. “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences”. *Astrophys. J. Lett.* 915.1 (2021), p. L5. DOI: [10.3847/2041-8213/ac082e](https://doi.org/10.3847/2041-8213/ac082e). arXiv: [2106.15163](https://arxiv.org/abs/2106.15163) [astro-ph.HE].
- [17] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. *Phys. Rev. Lett.* 116.6 (2016), p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). arXiv: [1602.03837](https://arxiv.org/abs/1602.03837) [gr-qc].
- [18] Albert Einstein. “Über Gravitationswellen”. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1918 (1918), pp. 154–167.
- [19] Luc Blanchet. “Gravitational Radiation from Post-Newtonian Sources and Inspiral Compact Binaries”. *Living Rev. Rel.* 17 (2014), p. 2. DOI: [10.12942/lrr-2014-2](https://doi.org/10.12942/lrr-2014-2). arXiv: [1310.1528](https://arxiv.org/abs/1310.1528) [gr-qc].
- [20] Matthew D. Duez and Yosef Zlochower. “Numerical Relativity of Compact Binaries in the 21st Century”. *Rept. Prog. Phys.* 82.1 (2019), p. 016902. DOI: [10.1088/1361-6633/aadb16](https://doi.org/10.1088/1361-6633/aadb16). arXiv: [1808.06011](https://arxiv.org/abs/1808.06011) [gr-qc].
- [21] Misao Sasaki and Hideyuki Tagoshi. “Analytic black hole perturbation approach to gravitational radiation”. *Living Rev. Rel.* 6 (2003), p. 6. DOI: [10.12942/lrr-2003-6](https://doi.org/10.12942/lrr-2003-6). arXiv: [gr-qc/0306120](https://arxiv.org/abs/gr-qc/0306120).
- [22] Adam Pound and Barry Wardell. “Black hole perturbation theory and gravitational self-force” (Jan. 2021). DOI: [10.1007/978-981-15-4702-7_38-1](https://doi.org/10.1007/978-981-15-4702-7_38-1). arXiv: [2101.04592](https://arxiv.org/abs/2101.04592) [gr-qc].
- [23] Vijay Varma et al. “Surrogate models for precessing binary black hole simulations with unequal masses”. *Phys. Rev. Research.* 1 (2019), p. 033015. DOI: [10.1103/PhysRevResearch.1.033015](https://doi.org/10.1103/PhysRevResearch.1.033015). arXiv: [1905.09300](https://arxiv.org/abs/1905.09300) [gr-qc].
- [24] Michael Boyle et al. “The SXS Collaboration catalog of binary black hole simulations”. *Class. Quant. Grav.* 36.19 (2019), p. 195006. DOI: [10.1088/1361-6382/ab34e2](https://doi.org/10.1088/1361-6382/ab34e2). arXiv: [1904.04831](https://arxiv.org/abs/1904.04831) [gr-qc].

- [25] Benjamin P Abbott et al. “A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals”. *Class. Quant. Grav.* 37.5 (2020), p. 055002. DOI: [10.1088/1361-6382/ab685e](https://doi.org/10.1088/1361-6382/ab685e). arXiv: [1908.11170](https://arxiv.org/abs/1908.11170) [gr-qc].
- [26] Krzysztof Belczynski et al. “The first gravitational-wave source from the isolated evolution of two 40-100 Msun stars”. *Nature* 534 (2016), p. 512. DOI: [10.1038/nature18322](https://doi.org/10.1038/nature18322). arXiv: [1602.04531](https://arxiv.org/abs/1602.04531) [astro-ph.HE].
- [27] Carl L. Rodriguez et al. “Dynamical Formation of the GW150914 Binary Black Hole”. *Astrophys. J. Lett.* 824.1 (2016), p. L8. DOI: [10.3847/2041-8205/824/1/L8](https://doi.org/10.3847/2041-8205/824/1/L8). arXiv: [1604.04254](https://arxiv.org/abs/1604.04254) [astro-ph.HE].
- [28] C. J. Moore, R. H. Cole, and C. P. L. Berry. “Gravitational-wave sensitivity curves”. *Class. Quant. Grav.* 32.1 (2015), p. 015014. DOI: [10.1088/0264-9381/32/1/015014](https://doi.org/10.1088/0264-9381/32/1/015014). arXiv: [1408.0740](https://arxiv.org/abs/1408.0740) [gr-qc].
- [29] Sebastian Khan et al. “Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era”. *Phys. Rev. D* 93.4 (2016), p. 044007. DOI: [10.1103/PhysRevD.93.044007](https://doi.org/10.1103/PhysRevD.93.044007). arXiv: [1508.07253](https://arxiv.org/abs/1508.07253) [gr-qc].
- [30] Thomas D. P. Edwards et al. “ripple: Differentiable and Hardware-Accelerated Waveforms for Gravitational Wave Data Analysis” (Feb. 2023). arXiv: [2302.05329](https://arxiv.org/abs/2302.05329) [astro-ph.IM].
- [31] P. Welch. “The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms”. *IEEE Transactions on Audio and Electroacoustics* 15.2 (1967), pp. 70–73. DOI: [10.1109/TAU.1967.1161901](https://doi.org/10.1109/TAU.1967.1161901).
- [32] Kostas D. Kokkotas and Bernd G. Schmidt. “Quasinormal modes of stars and black holes”. *Living Rev. Rel.* 2 (1999), p. 2. DOI: [10.12942/lrr-1999-2](https://doi.org/10.12942/lrr-1999-2). arXiv: [gr-qc/9909058](https://arxiv.org/abs/gr-qc/9909058).
- [33] Hans-Peter Nollert. “TOPICAL REVIEW: Quasinormal modes: the characteristic ‘sound’ of black holes and neutron stars”. *Class. Quant. Grav.* 16 (1999), R159–R216. DOI: [10.1088/0264-9381/16/12/201](https://doi.org/10.1088/0264-9381/16/12/201).
- [34] Valeria Ferrari and Leonardo Gualtieri. “Quasi-Normal Modes and Gravitational Wave Astronomy”. *Gen. Rel. Grav.* 40 (2008), pp. 945–970. DOI: [10.1007/s10714-007-0585-1](https://doi.org/10.1007/s10714-007-0585-1). arXiv: [0709.0657](https://arxiv.org/abs/0709.0657) [gr-qc].
- [35] Emanuele Berti, Vitor Cardoso, and Andrei O. Starinets. “Quasinormal modes of black holes and black branes”. *Class. Quant. Grav.* 26 (2009), p. 163001. DOI: [10.1088/0264-9381/26/16/163001](https://doi.org/10.1088/0264-9381/26/16/163001). arXiv: [0905.2975](https://arxiv.org/abs/0905.2975) [gr-qc].
- [36] Geoffrey Lovelace et al. “Modeling the source of GW150914 with targeted numerical-relativity simulations”. *Class. Quant. Grav.* 33.24 (2016), p. 244002. DOI: [10.1088/0264-9381/33/24/244002](https://doi.org/10.1088/0264-9381/33/24/244002). arXiv: [1607.05377](https://arxiv.org/abs/1607.05377) [gr-qc].

- [37] R. Abbott et al. “Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog”. *Phys. Rev. D* 103.12 (2021), p. 122002. DOI: [10.1103/PhysRevD.103.122002](https://doi.org/10.1103/PhysRevD.103.122002). arXiv: [2010.14529](https://arxiv.org/abs/2010.14529) [gr-qc].
- [38] R. Abbott et al. “Tests of General Relativity with GWTC-3”. (Dec. 2021). arXiv: [2112.06861](https://arxiv.org/abs/2112.06861) [gr-qc].
- [39] R. Abbott et al. “GW190521: A Binary Black Hole Merger with a Total Mass of $150M_{\odot}$ ”. *Phys. Rev. Lett.* 125.10 (2020), p. 101102. DOI: [10.1103/PhysRevLett.125.101102](https://doi.org/10.1103/PhysRevLett.125.101102). arXiv: [2009.01075](https://arxiv.org/abs/2009.01075) [gr-qc].
- [40] Karl Schwarzschild. “On the gravitational field of a mass point according to Einstein’s theory”. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1916 (1916), pp. 189–196. arXiv: [physics/9905030](https://arxiv.org/abs/physics/9905030).
- [41] Tullio Regge and John A. Wheeler. “Stability of a Schwarzschild singularity”. *Phys. Rev.* 108 (1957), pp. 1063–1069. DOI: [10.1103/PhysRev.108.1063](https://doi.org/10.1103/PhysRev.108.1063).
- [42] Frank J. Zerilli. “Effective potential for even parity Regge-Wheeler gravitational perturbation equations”. *Phys. Rev. Lett.* 24 (1970), pp. 737–738. DOI: [10.1103/PhysRevLett.24.737](https://doi.org/10.1103/PhysRevLett.24.737).
- [43] F. J. Zerilli. “Gravitational field of a particle falling in a schwarzschild geometry analyzed in tensor harmonics”. *Phys. Rev. D* 2 (1970), pp. 2141–2160. DOI: [10.1103/PhysRevD.2.2141](https://doi.org/10.1103/PhysRevD.2.2141).
- [44] C. V. Vishveshwara. “Scattering of Gravitational Radiation by a Schwarzschild Black-hole”. *Nature* 227 (1970), pp. 936–938. DOI: [10.1038/227936a0](https://doi.org/10.1038/227936a0).
- [45] William H. Press. “Long Wave Trains of Gravitational Waves from a Vibrating Black Hole”. *Astrophys. J. Lett.* 170 (1971), pp. L105–L108. DOI: [10.1086/180849](https://doi.org/10.1086/180849).
- [46] Roy P. Kerr. “Gravitational field of a spinning mass as an example of algebraically special metrics”. *Phys. Rev. Lett.* 11 (1963), pp. 237–238. DOI: [10.1103/PhysRevLett.11.237](https://doi.org/10.1103/PhysRevLett.11.237).
- [47] S. A. Teukolsky. “Rotating black holes - separable wave equations for gravitational and electromagnetic perturbations”. *Phys. Rev. Lett.* 29 (1972), pp. 1114–1118. DOI: [10.1103/PhysRevLett.29.1114](https://doi.org/10.1103/PhysRevLett.29.1114).
- [48] Saul A. Teukolsky. “Perturbations of a rotating black hole. 1. Fundamental equations for gravitational electromagnetic and neutrino field perturbations”. *Astrophys. J.* 185 (1973), pp. 635–647. DOI: [10.1086/152444](https://doi.org/10.1086/152444).
- [49] B. Carter. “Axisymmetric Black Hole Has Only Two Degrees of Freedom”. *Phys. Rev. Lett.* 26 (1971), pp. 331–333. DOI: [10.1103/PhysRevLett.26.331](https://doi.org/10.1103/PhysRevLett.26.331).
- [50] Subrahmanyan Chandrasekhar. “On the equations governing the perturbations of the Schwarzschild black hole”. *Proc. Roy. Soc. Lond. A* 343.1634 (1975), pp. 289–298. DOI: [10.1098/rspa.1975.0066](https://doi.org/10.1098/rspa.1975.0066).

- [51] Emanuele Berti. “Black hole quasinormal modes: Hints of quantum gravity?” *Conf. Proc. C* 0405132 (2004), pp. 145–186. arXiv: [gr-qc/0411025](#).
- [52] S. Chandrasekhar and Steven L. Detweiler. “The quasi-normal modes of the Schwarzschild black hole”. *Proc. Roy. Soc. Lond. A* 344 (1975), pp. 441–452. DOI: [10.1098/rspa.1975.0112](#).
- [53] Hans-Joachim Blome and Bahram Mashhoon. “Quasi-normal oscillations of a schwarzschild black hole”. *Physics Letters A* 100.5 (1984), pp. 231–234. ISSN: 0375-9601. DOI: [10.1016/0375-9601\(84\)90769-2](#).
- [54] Valeria Ferrari and Bahram Mashhoon. “Oscillations of a Black Hole”. *Phys. Rev. Lett.* 52.16 (1984), p. 1361. DOI: [10.1103/PhysRevLett.52.1361](#).
- [55] Valeria Ferrari and Bahram Mashhoon. “New approach to the quasinormal modes of a black hole”. *Phys. Rev. D* 30 (1984), pp. 295–304. DOI: [10.1103/PhysRevD.30.295](#).
- [56] Bernard F. Schutz and Clifford M. Will. “BLACK HOLE NORMAL MODES: A SEMIANALYTIC APPROACH”. *Astrophys. J. Lett.* 291 (1985), pp. L33–L36. DOI: [10.1086/184453](#).
- [57] Sai Iyer and Clifford M. Will. “Black Hole Normal Modes: A WKB Approach. 1. Foundations and Application of a Higher Order WKB Analysis of Potential Barrier Scattering”. *Phys. Rev. D* 35 (1987), p. 3621. DOI: [10.1103/PhysRevD.35.3621](#).
- [58] Jerzy Matyjasek and Michał Opala. “Quasinormal modes of black holes. The improved semianalytic approach”. *Phys. Rev. D* 96.2 (2017), p. 024011. DOI: [10.1103/PhysRevD.96.024011](#). arXiv: [1704.00361 \[gr-qc\]](#).
- [59] R. A. Konoplya, A. Zhidenko, and A. F. Zinhailo. “Higher order WKB formula for quasinormal modes and grey-body factors: recipes for quick and accurate calculations”. *Class. Quant. Grav.* 36 (2019), p. 155002. DOI: [10.1088/1361-6382/ab2e25](#). arXiv: [1904.10333 \[gr-qc\]](#).
- [60] E. W. Leaver. “An Analytic representation for the quasi normal modes of Kerr black holes”. *Proc. Roy. Soc. Lond. A* 402 (1985), pp. 285–298. DOI: [10.1098/rspa.1985.0119](#).
- [61] Hans-Peter Nollert. “Quasinormal modes of Schwarzschild black holes: The determination of quasinormal frequencies with very large imaginary parts”. *Phys. Rev. D* 47 (1993), pp. 5253–5258. DOI: [10.1103/PhysRevD.47.5253](#).
- [62] Leo C. Stein. “qnm: A Python package for calculating Kerr quasinormal modes, separation constants, and spherical-spheroidal mixing coefficients”. *J. Open Source Softw.* 4.42 (2019), p. 1683. DOI: [10.21105/joss.01683](#). arXiv: [1908.10377 \[gr-qc\]](#).

- [63] Gregory B. Cook and Maxim Zaluskiy. “Gravitational perturbations of the Kerr geometry: High-accuracy study”. *Phys. Rev. D* 90.12 (2014), p. 124021. DOI: [10.1103/PhysRevD.90.124021](https://doi.org/10.1103/PhysRevD.90.124021). arXiv: [1410.7698](https://arxiv.org/abs/1410.7698) [gr-qc].
- [64] Emanuele Berti and Antoine Klein. “Mixing of spherical and spheroidal modes in perturbed Kerr black holes”. *Phys. Rev. D* 90.6 (2014), p. 064012. DOI: [10.1103/PhysRevD.90.064012](https://doi.org/10.1103/PhysRevD.90.064012). arXiv: [1408.1860](https://arxiv.org/abs/1408.1860) [gr-qc].
- [65] C. J. Goebel. “Comments on the “vibrations” of a Black Hole.” *Astrophys. J. Lett.* 172 (1972), p. L95. DOI: [10.1086/180898](https://doi.org/10.1086/180898).
- [66] Vitor Cardoso et al. “Geodesic stability, Lyapunov exponents and quasinormal modes”. *Phys. Rev. D* 79.6 (2009), p. 064016. DOI: [10.1103/PhysRevD.79.064016](https://doi.org/10.1103/PhysRevD.79.064016). arXiv: [0812.1806](https://arxiv.org/abs/0812.1806) [hep-th].
- [67] Huan Yang et al. “Quasinormal-mode spectrum of Kerr black holes and its geometric interpretation”. *Phys. Rev. D* 86 (2012), p. 104006. DOI: [10.1103/PhysRevD.86.104006](https://doi.org/10.1103/PhysRevD.86.104006). arXiv: [1207.4253](https://arxiv.org/abs/1207.4253) [gr-qc].
- [68] Jose Luis Blázquez-Salcedo et al. “Perturbed black holes in Einstein-dilaton-Gauss-Bonnet gravity: Stability, ringdown, and gravitational-wave emission”. *Phys. Rev. D* 94.10 (2016), p. 104024. DOI: [10.1103/PhysRevD.94.104024](https://doi.org/10.1103/PhysRevD.94.104024). arXiv: [1609.01286](https://arxiv.org/abs/1609.01286) [gr-qc].
- [69] E T. Newman et al. “Metric of a Rotating, Charged Mass”. *J. Math. Phys.* 6 (1965), pp. 918–919. DOI: [10.1063/1.1704351](https://doi.org/10.1063/1.1704351).
- [70] Vitor Cardoso et al. “Black holes and gravitational waves in models of minicharged dark matter”. *JCAP* 05 (2016). [Erratum: *JCAP* 04, E01 (2020)], p. 054. DOI: [10.1088/1475-7516/2016/05/054](https://doi.org/10.1088/1475-7516/2016/05/054). arXiv: [1604.07845](https://arxiv.org/abs/1604.07845) [hep-ph].
- [71] Hai-Tian Wang et al. “Quasinormal modes of the Kerr-Newman black hole: GW150914 and fundamental physics implications”. *Phys. Rev. D* 104.10 (2021), p. 104063. DOI: [10.1103/PhysRevD.104.104063](https://doi.org/10.1103/PhysRevD.104.104063). arXiv: [2104.07594](https://arxiv.org/abs/2104.07594) [gr-qc].
- [72] Gregorio Carullo et al. “Constraints on Kerr-Newman black holes from merger-ringdown gravitational-wave observations”. *Phys. Rev. D* 105.6 (2022), p. 062009. DOI: [10.1103/PhysRevD.105.062009](https://doi.org/10.1103/PhysRevD.105.062009). arXiv: [2109.13961](https://arxiv.org/abs/2109.13961) [gr-qc].
- [73] Subrahmanyan Chandrasekhar. *The mathematical theory of black holes*. 1985. ISBN: 978-0-19-850370-5.
- [74] E. S. C. Ching et al. “Quasinormal mode expansion for linearized waves in gravitational system”. *Phys. Rev. Lett.* 74 (1995), pp. 4588–4591. DOI: [10.1103/PhysRevLett.74.4588](https://doi.org/10.1103/PhysRevLett.74.4588). arXiv: [gr-qc/9408043](https://arxiv.org/abs/gr-qc/9408043).
- [75] E. S. C. Ching et al. “Wave propagation in gravitational systems: Completeness of quasinormal modes”. *Phys. Rev. D* 54 (1996), pp. 3778–3791. DOI: [10.1103/PhysRevD.54.3778](https://doi.org/10.1103/PhysRevD.54.3778). arXiv: [gr-qc/9507034](https://arxiv.org/abs/gr-qc/9507034).

- [76] Hans-Peter Nollert and Richard H. Price. “Quantifying excitations of quasinormal mode systems”. *J. Math. Phys.* 40 (1999), pp. 980–1010. DOI: [10.1063/1.532698](https://doi.org/10.1063/1.532698). arXiv: [gr-qc/9810074](https://arxiv.org/abs/gr-qc/9810074).
- [77] E. S. C. Ching et al. “Quasinormal-mode expansion for waves in open systems”. *Rev. Mod. Phys.* 70 (1998), p. 1545. DOI: [10.1103/RevModPhys.70.1545](https://doi.org/10.1103/RevModPhys.70.1545). arXiv: [gr-qc/9904017](https://arxiv.org/abs/gr-qc/9904017).
- [78] Richard H. Price. “Nonspherical perturbations of relativistic gravitational collapse. 1. Scalar and gravitational perturbations”. *Phys. Rev. D* 5 (1972), pp. 2419–2438. DOI: [10.1103/PhysRevD.5.2419](https://doi.org/10.1103/PhysRevD.5.2419).
- [79] Edward W. Leaver. “Spectral decomposition of the perturbation response of the Schwarzschild geometry”. *Phys. Rev. D* 34 (1986), pp. 384–408. DOI: [10.1103/PhysRevD.34.384](https://doi.org/10.1103/PhysRevD.34.384).
- [80] E. S. C. Ching et al. “Wave propagation in gravitational systems: Late time behavior”. *Phys. Rev. D* 52 (1995), pp. 2118–2132. DOI: [10.1103/PhysRevD.52.2118](https://doi.org/10.1103/PhysRevD.52.2118). arXiv: [gr-qc/9507035](https://arxiv.org/abs/gr-qc/9507035).
- [81] Reinaldo J. Gleiser, Richard H. Price, and Jorge Pullin. “Late time tails in the Kerr spacetime”. *Class. Quant. Grav.* 25 (2008), p. 072001. DOI: [10.1088/0264-9381/25/7/072001](https://doi.org/10.1088/0264-9381/25/7/072001). arXiv: [0710.4183 \[gr-qc\]](https://arxiv.org/abs/0710.4183).
- [82] Marc Favata. “The gravitational-wave memory effect”. *Class. Quant. Grav.* 27 (2010). Ed. by Zsuzsa Marka and Szabolcs Marka, p. 084036. DOI: [10.1088/0264-9381/27/8/084036](https://doi.org/10.1088/0264-9381/27/8/084036). arXiv: [1003.3486 \[gr-qc\]](https://arxiv.org/abs/1003.3486).
- [83] Eric Thrane, Paul D. Lasky, and Yuri Levin. “Challenges testing the no-hair theorem with gravitational waves”. *Phys. Rev. D* 96.10 (2017), p. 102004. DOI: [10.1103/PhysRevD.96.102004](https://doi.org/10.1103/PhysRevD.96.102004). arXiv: [1706.05152 \[gr-qc\]](https://arxiv.org/abs/1706.05152).
- [84] Vishal Baibhav et al. “Black Hole Spectroscopy: Systematic Errors and Ringdown Energy Estimates”. *Phys. Rev. D* 97.4 (2018), p. 044048. DOI: [10.1103/PhysRevD.97.044048](https://doi.org/10.1103/PhysRevD.97.044048). arXiv: [1710.02156 \[gr-qc\]](https://arxiv.org/abs/1710.02156).
- [85] Hans-Peter Nollert. “About the significance of quasinormal modes of black holes”. *Phys. Rev. D* 53 (1996), pp. 4397–4402. DOI: [10.1103/PhysRevD.53.4397](https://doi.org/10.1103/PhysRevD.53.4397). arXiv: [gr-qc/9602032](https://arxiv.org/abs/gr-qc/9602032).
- [86] Ramin G. Daghigh, Michael D. Green, and Jodin C. Morey. “Significance of Black Hole Quasinormal Modes: A Closer Look”. *Phys. Rev. D* 101.10 (2020), p. 104009. DOI: [10.1103/PhysRevD.101.104009](https://doi.org/10.1103/PhysRevD.101.104009). arXiv: [2002.07251 \[gr-qc\]](https://arxiv.org/abs/2002.07251).
- [87] José Luis Jaramillo, Rodrigo Panosso Macedo, and Lamis Al Sheikh. “Pseudospectrum and Black Hole Quasinormal Mode Instability”. *Phys. Rev. X* 11.3 (2021), p. 031003. DOI: [10.1103/PhysRevX.11.031003](https://doi.org/10.1103/PhysRevX.11.031003). arXiv: [2004.06434 \[gr-qc\]](https://arxiv.org/abs/2004.06434).

- [88] Wei-Liang Qian et al. “Asymptotical quasinormal mode spectrum for piecewise approximate effective potential”. *Phys. Rev. D* 103.2 (2021), p. 024019. DOI: [10.1103/PhysRevD.103.024019](https://doi.org/10.1103/PhysRevD.103.024019). arXiv: [2009.11627](https://arxiv.org/abs/2009.11627) [gr-qc].
- [89] José Luis Jaramillo, Rodrigo Panosso Macedo, and Lamis Al Sheikh. “Gravitational Wave Signatures of Black Hole Quasinormal Mode Instability”. *Phys. Rev. Lett.* 128.21 (2022), p. 211102. DOI: [10.1103/PhysRevLett.128.211102](https://doi.org/10.1103/PhysRevLett.128.211102). arXiv: [2105.03451](https://arxiv.org/abs/2105.03451) [gr-qc].
- [90] Mark Ho-Yeuk Cheung et al. “Destabilizing the Fundamental Mode of Black Holes: The Elephant and the Flea”. *Phys. Rev. Lett.* 128.11 (2022), p. 111103. DOI: [10.1103/PhysRevLett.128.111103](https://doi.org/10.1103/PhysRevLett.128.111103). arXiv: [2111.05415](https://arxiv.org/abs/2111.05415) [gr-qc].
- [91] Emanuele Berti et al. “Stability of the fundamental quasinormal mode in time-domain observations against small perturbations”. *Phys. Rev. D* 106.8 (2022), p. 084011. DOI: [10.1103/PhysRevD.106.084011](https://doi.org/10.1103/PhysRevD.106.084011). arXiv: [2205.08547](https://arxiv.org/abs/2205.08547) [gr-qc].
- [92] Emanuele Berti, Vitor Cardoso, and Clifford M. Will. “On gravitational-wave spectroscopy of massive black holes with the space interferometer LISA”. *Phys. Rev. D* 73 (2006), p. 064030. DOI: [10.1103/PhysRevD.73.064030](https://doi.org/10.1103/PhysRevD.73.064030). arXiv: [gr-qc/0512160](https://arxiv.org/abs/gr-qc/0512160).
- [93] Lionel London, Deirdre Shoemaker, and James Healy. “Modeling ringdown: Beyond the fundamental quasinormal modes”. *Phys. Rev. D* 90.12 (2014). [Erratum: *Phys.Rev.D* 94, 069902 (2016)], p. 124032. DOI: [10.1103/PhysRevD.90.124032](https://doi.org/10.1103/PhysRevD.90.124032). arXiv: [1404.3197](https://arxiv.org/abs/1404.3197) [gr-qc].
- [94] Xisco Jiménez Forteza et al. “Spectroscopy of binary black hole ringdown using overtones and angular modes”. *Phys. Rev. D* 102.4 (2020), p. 044053. DOI: [10.1103/PhysRevD.102.044053](https://doi.org/10.1103/PhysRevD.102.044053). arXiv: [2005.03260](https://arxiv.org/abs/2005.03260) [gr-qc].
- [95] Scott A. Hughes and Kristen Menou. “Golden binaries for LISA: Robust probes of strong-field gravity”. *Astrophys. J.* 623 (2005), pp. 689–699. DOI: [10.1086/428826](https://doi.org/10.1086/428826). arXiv: [astro-ph/0410148](https://arxiv.org/abs/astro-ph/0410148).
- [96] Hiroyuki Nakano, Takahiro Tanaka, and Takashi Nakamura. “Possible golden events for ringdown gravitational waves”. *Phys. Rev. D* 92.6 (2015), p. 064003. DOI: [10.1103/PhysRevD.92.064003](https://doi.org/10.1103/PhysRevD.92.064003). arXiv: [1506.00560](https://arxiv.org/abs/1506.00560) [astro-ph.HE].
- [97] Abhirup Ghosh et al. “Testing general relativity using golden black-hole binaries”. *Phys. Rev. D* 94.2 (2016), p. 021101. DOI: [10.1103/PhysRevD.94.021101](https://doi.org/10.1103/PhysRevD.94.021101). arXiv: [1602.02453](https://arxiv.org/abs/1602.02453) [gr-qc].
- [98] Abhirup Ghosh et al. “Testing general relativity using gravitational wave signals from the inspiral, merger and ringdown of binary black holes”. *Class. Quant. Grav.* 35.1 (2018), p. 014002. DOI: [10.1088/1361-6382/aa972e](https://doi.org/10.1088/1361-6382/aa972e). arXiv: [1704.06784](https://arxiv.org/abs/1704.06784) [gr-qc].

- [99] Vijay Varma et al. “High-accuracy mass, spin, and recoil predictions of generic black-hole merger remnants”. *Phys. Rev. Lett.* 122.1 (2019), p. 011101. DOI: [10.1103/PhysRevLett.122.011101](https://doi.org/10.1103/PhysRevLett.122.011101). arXiv: [1809.09125](https://arxiv.org/abs/1809.09125) [gr-qc].
- [100] S. W. Hawking. “Gravitational radiation from colliding black holes”. *Phys. Rev. Lett.* 26 (1971), pp. 1344–1346. DOI: [10.1103/PhysRevLett.26.1344](https://doi.org/10.1103/PhysRevLett.26.1344).
- [101] Miriam Cabero et al. “Observational tests of the black hole area increase law”. *Phys. Rev. D* 97.12 (2018), p. 124069. DOI: [10.1103/PhysRevD.97.124069](https://doi.org/10.1103/PhysRevD.97.124069). arXiv: [1711.09073](https://arxiv.org/abs/1711.09073) [gr-qc].
- [102] Maximiliano Isi et al. “Testing the Black-Hole Area Law with GW150914”. *Phys. Rev. Lett.* 127.1 (2021), p. 011103. DOI: [10.1103/PhysRevLett.127.011103](https://doi.org/10.1103/PhysRevLett.127.011103). arXiv: [2012.04486](https://arxiv.org/abs/2012.04486) [gr-qc].
- [103] F. Echeverria. “Gravitational Wave Measurements of the Mass and Angular Momentum of a Black Hole”. *Phys. Rev. D* 40 (1989), pp. 3194–3203. DOI: [10.1103/PhysRevD.40.3194](https://doi.org/10.1103/PhysRevD.40.3194).
- [104] Lee S. Finn. “Detection, measurement and gravitational radiation”. *Phys. Rev. D* 46 (1992), pp. 5236–5249. DOI: [10.1103/PhysRevD.46.5236](https://doi.org/10.1103/PhysRevD.46.5236). arXiv: [gr-qc/9209010](https://arxiv.org/abs/gr-qc/9209010).
- [105] Olaf Dreyer et al. “Black hole spectroscopy: Testing general relativity through gravitational wave observations”. *Class. Quant. Grav.* 21 (2004), pp. 787–804. DOI: [10.1088/0264-9381/21/4/003](https://doi.org/10.1088/0264-9381/21/4/003). arXiv: [gr-qc/0309007](https://arxiv.org/abs/gr-qc/0309007).
- [106] Emanuele Berti et al. “Matched-filtering and parameter estimation of ringdown waveforms”. *Phys. Rev. D* 76 (2007), p. 104044. DOI: [10.1103/PhysRevD.76.104044](https://doi.org/10.1103/PhysRevD.76.104044). arXiv: [0707.1202](https://arxiv.org/abs/0707.1202) [gr-qc].
- [107] Ioannis Kamaretsos et al. “Black-hole hair loss: learning about binary progenitors from ringdown signals”. *Phys. Rev. D* 85 (2012), p. 024018. DOI: [10.1103/PhysRevD.85.024018](https://doi.org/10.1103/PhysRevD.85.024018). arXiv: [1107.0854](https://arxiv.org/abs/1107.0854) [gr-qc].
- [108] S. Gossan, J. Veitch, and B. S. Sathyaprakash. “Bayesian model selection for testing the no-hair theorem with black hole ringdowns”. *Phys. Rev. D* 85 (2012), p. 124056. DOI: [10.1103/PhysRevD.85.124056](https://doi.org/10.1103/PhysRevD.85.124056). arXiv: [1111.5819](https://arxiv.org/abs/1111.5819) [gr-qc].
- [109] J. Meidam et al. “Testing the no-hair theorem with black hole ringdowns using TIGER”. *Phys. Rev. D* 90.6 (2014), p. 064009. DOI: [10.1103/PhysRevD.90.064009](https://doi.org/10.1103/PhysRevD.90.064009). arXiv: [1406.3201](https://arxiv.org/abs/1406.3201) [gr-qc].
- [110] Huan Yang et al. “Black hole spectroscopy with coherent mode stacking”. *Phys. Rev. Lett.* 118.16 (2017), p. 161101. DOI: [10.1103/PhysRevLett.118.161101](https://doi.org/10.1103/PhysRevLett.118.161101). arXiv: [1701.05808](https://arxiv.org/abs/1701.05808) [gr-qc].
- [111] C. F. Da Silva Costa et al. “Detection of (2,2) quasinormal mode from a population of black holes with a constructive summation method”. *Phys. Rev. D* 98.2 (2018), p. 024052. DOI: [10.1103/PhysRevD.98.024052](https://doi.org/10.1103/PhysRevD.98.024052). arXiv: [1711.00551](https://arxiv.org/abs/1711.00551) [gr-qc].

- [112] Gregorio Carullo et al. “Empirical tests of the black hole no-hair conjecture using gravitational-wave observations”. *Phys. Rev. D* 98.10 (2018), p. 104020. DOI: [10.1103/PhysRevD.98.104020](https://doi.org/10.1103/PhysRevD.98.104020). arXiv: [1805.04760](https://arxiv.org/abs/1805.04760) [gr-qc].
- [113] Emanuele Berti et al. “Spectroscopy of Kerr black holes with Earth- and space-based interferometers”. *Phys. Rev. Lett.* 117.10 (2016), p. 101102. DOI: [10.1103/PhysRevLett.117.101102](https://doi.org/10.1103/PhysRevLett.117.101102). arXiv: [1605.09286](https://arxiv.org/abs/1605.09286) [gr-qc].
- [114] Swetha Bhagwat, Duncan A. Brown, and Stefan W. Ballmer. “Spectroscopic analysis of stellar mass black-hole mergers in our local universe with ground-based gravitational wave detectors”. *Phys. Rev. D* 94.8 (2016). [Erratum: *Phys.Rev.D* 95, 069906 (2017)], p. 084024. DOI: [10.1103/PhysRevD.94.084024](https://doi.org/10.1103/PhysRevD.94.084024). arXiv: [1607.07845](https://arxiv.org/abs/1607.07845) [gr-qc].
- [115] Andrea Maselli, Kostas Kokkotas, and Pablo Laguna. “Observing binary black hole ringdowns by advanced gravitational wave detectors”. *Phys. Rev. D* 95.10 (2017), p. 104026. DOI: [10.1103/PhysRevD.95.104026](https://doi.org/10.1103/PhysRevD.95.104026). arXiv: [1702.01110](https://arxiv.org/abs/1702.01110) [gr-qc].
- [116] Vishal Baibhav and Emanuele Berti. “Multimode black hole spectroscopy”. *Phys. Rev. D* 99.2 (2019), p. 024005. DOI: [10.1103/PhysRevD.99.024005](https://doi.org/10.1103/PhysRevD.99.024005). arXiv: [1809.03500](https://arxiv.org/abs/1809.03500) [gr-qc].
- [117] Swetha Bhagwat et al. “Detectability of the subdominant mode in a binary black hole ringdown”. *Phys. Rev. D* 102.2 (2020), p. 024023. DOI: [10.1103/PhysRevD.102.024023](https://doi.org/10.1103/PhysRevD.102.024023). arXiv: [1910.13203](https://arxiv.org/abs/1910.13203) [gr-qc].
- [118] Miriam Cabero et al. “Black hole spectroscopy in the next decade”. *Phys. Rev. D* 101.6 (2020), p. 064044. DOI: [10.1103/PhysRevD.101.064044](https://doi.org/10.1103/PhysRevD.101.064044). arXiv: [1911.01361](https://arxiv.org/abs/1911.01361) [gr-qc].
- [119] Maximiliano Isi et al. “Testing the no-hair theorem with GW150914”. *Phys. Rev. Lett.* 123.11 (2019), p. 111102. DOI: [10.1103/PhysRevLett.123.111102](https://doi.org/10.1103/PhysRevLett.123.111102). arXiv: [1905.00869](https://arxiv.org/abs/1905.00869) [gr-qc].
- [120] Collin D. Capano et al. “Observation of a multimode quasi-normal spectrum from a perturbed black hole” (May 2021). arXiv: [2105.05238](https://arxiv.org/abs/2105.05238) [gr-qc].
- [121] Maximiliano Isi and Will M. Farr. “Analyzing black-hole ringdowns” (July 2021). arXiv: [2107.05609](https://arxiv.org/abs/2107.05609) [gr-qc].
- [122] Matthew Giesler et al. “Black Hole Ringdown: The Importance of Overtones”. *Phys. Rev. X* 9.4 (2019), p. 041060. DOI: [10.1103/PhysRevX.9.041060](https://doi.org/10.1103/PhysRevX.9.041060). arXiv: [1903.08284](https://arxiv.org/abs/1903.08284) [gr-qc].
- [123] Swetha Bhagwat et al. “Ringdown overtones, black hole spectroscopy, and no-hair theorem tests”. *Phys. Rev. D* 101.4 (2020), p. 044033. DOI: [10.1103/PhysRevD.101.044033](https://doi.org/10.1103/PhysRevD.101.044033). arXiv: [1910.08708](https://arxiv.org/abs/1910.08708) [gr-qc].

- [124] Iara Ota and Cecilia Chirenti. “Overtones or higher harmonics? Prospects for testing the no-hair theorem with gravitational wave detections”. *Phys. Rev. D* 101.10 (2020), p. 104005. DOI: [10.1103/PhysRevD.101.104005](https://doi.org/10.1103/PhysRevD.101.104005). arXiv: [1911.00440](https://arxiv.org/abs/1911.00440) [gr-qc].
- [125] Gregory B. Cook. “Aspects of multimode Kerr ringdown fitting”. *Phys. Rev. D* 102.2 (2020), p. 024027. DOI: [10.1103/PhysRevD.102.024027](https://doi.org/10.1103/PhysRevD.102.024027). arXiv: [2004.08347](https://arxiv.org/abs/2004.08347) [gr-qc].
- [126] Arnab Dhani. “Importance of mirror modes in binary black hole ringdown waveform”. *Phys. Rev. D* 103.10 (2021), p. 104048. DOI: [10.1103/PhysRevD.103.104048](https://doi.org/10.1103/PhysRevD.103.104048). arXiv: [2010.08602](https://arxiv.org/abs/2010.08602) [gr-qc].
- [127] Pierre Mourier et al. “Quasinormal modes and their overtones at the common horizon in a binary black hole merger”. *Phys. Rev. D* 103.4 (2021), p. 044054. DOI: [10.1103/PhysRevD.103.044054](https://doi.org/10.1103/PhysRevD.103.044054). arXiv: [2010.15186](https://arxiv.org/abs/2010.15186) [gr-qc].
- [128] Arnab Dhani and B. S. Sathyaprakash. “Overtones, mirror modes, and mode-mixing in binary black hole mergers” (July 2021). arXiv: [2107.14195](https://arxiv.org/abs/2107.14195) [gr-qc].
- [129] Xisco Jiménez Forteza and Pierre Mourier. “High-overtone fits to numerical relativity ringdowns: beyond the dismissed $n=8$ special tone” (July 2021). arXiv: [2107.11829](https://arxiv.org/abs/2107.11829) [gr-qc].
- [130] Lorena Magaña Zertuche et al. “High Precision Ringdown Modeling: Multimode Fits and BMS Frames”. *Phys. Rev. D* 105.10 (2022), p. 104015. DOI: [10.1103/PhysRevD.105.104015](https://doi.org/10.1103/PhysRevD.105.104015). arXiv: [2110.15922](https://arxiv.org/abs/2110.15922) [gr-qc].
- [131] Vishal Baibhav et al. “Agnostic black hole spectroscopy: quasinormal mode content of numerical relativity waveforms and limits of validity of linear perturbation theory” (Feb. 2023). arXiv: [2302.03050](https://arxiv.org/abs/2302.03050) [gr-qc].
- [132] Peter James Nee, Sebastian H. Völkel, and Harald P. Pfeiffer. “To ring or not to ring, the tale of black hole quasi-normal modes” (Feb. 2023). arXiv: [2302.06634](https://arxiv.org/abs/2302.06634) [gr-qc].
- [133] Emanuele Berti et al. “Inspirals, merger and ringdown of unequal mass black hole binaries: A Multipolar analysis”. *Phys. Rev. D* 76 (2007), p. 064034. DOI: [10.1103/PhysRevD.76.064034](https://doi.org/10.1103/PhysRevD.76.064034). arXiv: [gr-qc/0703053](https://arxiv.org/abs/gr-qc/0703053).
- [134] Ernst Nils Dorband et al. “A Numerical study of the quasinormal mode excitation of Kerr black holes”. *Phys. Rev. D* 74 (2006), p. 084028. DOI: [10.1103/PhysRevD.74.084028](https://doi.org/10.1103/PhysRevD.74.084028). arXiv: [gr-qc/0608091](https://arxiv.org/abs/gr-qc/0608091).
- [135] Alessandra Buonanno, Gregory B. Cook, and Frans Pretorius. “Inspirals, merger and ring-down of equal-mass black-hole binaries”. *Phys. Rev. D* 75 (2007), p. 124018. DOI: [10.1103/PhysRevD.75.124018](https://doi.org/10.1103/PhysRevD.75.124018). arXiv: [gr-qc/0610122](https://arxiv.org/abs/gr-qc/0610122).
- [136] Ioannis Kamaretsos, Mark Hannam, and B. Sathyaprakash. “Is black-hole ringdown a memory of its progenitor?” *Phys. Rev. Lett.* 109 (2012), p. 141102. DOI: [10.1103/PhysRevLett.109.141102](https://doi.org/10.1103/PhysRevLett.109.141102). arXiv: [1207.0399](https://arxiv.org/abs/1207.0399) [gr-qc].

- [137] Maria Okounkova. “Revisiting non-linearity in binary black hole mergers” (Apr. 2020). arXiv: [2004.00671 \[gr-qc\]](#).
- [138] Naritaka Oshita. “Ease of excitation of black hole ringing: Quantifying the importance of overtones by the excitation factors”. *Phys. Rev. D* 104.12 (2021), p. 124032. DOI: [10.1103/PhysRevD.104.124032](#). arXiv: [2109.09757 \[gr-qc\]](#).
- [139] Scott A. Hughes et al. “Learning about black hole binaries from their ringdown spectra”. *Phys. Rev. Lett.* 123.16 (2019), p. 161101. DOI: [10.1103/PhysRevLett.123.161101](#). arXiv: [1901.05900 \[gr-qc\]](#).
- [140] Halston Lim et al. “Exciting black hole modes via misaligned coalescences: II. The mode content of late-time coalescence waveforms”. *Phys. Rev. D* 100.8 (2019), p. 084032. DOI: [10.1103/PhysRevD.100.084032](#). arXiv: [1901.05902 \[gr-qc\]](#).
- [141] Halston Lim, Gaurav Khanna, and Scott A. Hughes. “Measuring quasinormal mode amplitudes with misaligned binary black hole ringdowns”. *Phys. Rev. D* 105.12 (2022), p. 124030. DOI: [10.1103/PhysRevD.105.124030](#). arXiv: [2204.06007 \[gr-qc\]](#).
- [142] Theocharis A. Apostolatos et al. “Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries”. *Phys. Rev. D* 49 (1994), pp. 6274–6297. DOI: [10.1103/PhysRevD.49.6274](#).
- [143] Charles R. Harris et al. “Array programming with NumPy”. *Nature* 585.7825 (2020), pp. 357–362. DOI: [10.1038/s41586-020-2649-2](#). arXiv: [2006.10256 \[cs.MS\]](#).
- [144] R. Abbott et al. “Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog”. *Astrophys. J. Lett.* 913.1 (2021), p. L7. DOI: [10.3847/2041-8213/abe949](#). arXiv: [2010.14533 \[astro-ph.HE\]](#).
- [145] R. Abbott et al. “The population of merging compact binaries inferred using gravitational waves through GWTC-3” (Nov. 2021). arXiv: [2111.03634 \[astro-ph.HE\]](#).
- [146] R. Abbott et al. “Properties and Astrophysical Implications of the 150 M_{\odot} Binary Black Hole Merger GW190521”. *Astrophys. J. Lett.* 900.1 (2020), p. L13. DOI: [10.3847/2041-8213/aba493](#). arXiv: [2009.01190 \[astro-ph.HE\]](#).
- [147] R. Abbott et al. “GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses”. *Phys. Rev. D* 102.4 (2020), p. 043015. DOI: [10.1103/PhysRevD.102.043015](#). arXiv: [2004.08342 \[astro-ph.HE\]](#).
- [148] Mark Hannam et al. “General-relativistic precession in a black-hole binary”. *Nature* 610.7933 (2022), pp. 652–655. DOI: [10.1038/s41586-022-05212-z](#). arXiv: [2112.11300 \[gr-qc\]](#).
- [149] Ethan Payne et al. “Curious case of GW200129: Interplay between spin-precession inference and data-quality issues”. *Phys. Rev. D* 106.10 (2022), p. 104017. DOI: [10.1103/PhysRevD.106.104017](#). arXiv: [2206.11932 \[gr-qc\]](#).

- [150] Patricia Schmidt et al. “Tracking the precession of compact binaries from their gravitational-wave signal”. *Phys. Rev. D* 84 (2011), p. 024046. DOI: [10.1103/PhysRevD.84.024046](https://doi.org/10.1103/PhysRevD.84.024046). arXiv: [1012.2879](https://arxiv.org/abs/1012.2879) [gr-qc].
- [151] 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. D* 86 (2012), p. 104063. DOI: [10.1103/PhysRevD.86.104063](https://doi.org/10.1103/PhysRevD.86.104063). arXiv: [1207.3088](https://arxiv.org/abs/1207.3088) [gr-qc].
- [152] Mark Hannam et al. “Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms”. *Phys. Rev. Lett.* 113.15 (2014), p. 151101. DOI: [10.1103/PhysRevLett.113.151101](https://doi.org/10.1103/PhysRevLett.113.151101). arXiv: [1308.3271](https://arxiv.org/abs/1308.3271) [gr-qc].
- [153] Sebastian Khan et al. “Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects”. *Phys. Rev. D* 100.2 (2019), p. 024059. DOI: [10.1103/PhysRevD.100.024059](https://doi.org/10.1103/PhysRevD.100.024059). arXiv: [1809.10113](https://arxiv.org/abs/1809.10113) [gr-qc].
- [154] Geraint Pratten et al. “Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes”. *Phys. Rev. D* 103.10 (2021), p. 104056. DOI: [10.1103/PhysRevD.103.104056](https://doi.org/10.1103/PhysRevD.103.104056). arXiv: [2004.06503](https://arxiv.org/abs/2004.06503) [gr-qc].
- [155] Yi Pan et al. “Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism”. *Phys. Rev. D* 89.8 (2014), p. 084006. DOI: [10.1103/PhysRevD.89.084006](https://doi.org/10.1103/PhysRevD.89.084006). arXiv: [1307.6232](https://arxiv.org/abs/1307.6232) [gr-qc].
- [156] Serguei Ossokine et al. “Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation”. *Phys. Rev. D* 102.4 (2020), p. 044055. DOI: [10.1103/PhysRevD.102.044055](https://doi.org/10.1103/PhysRevD.102.044055). arXiv: [2004.09442](https://arxiv.org/abs/2004.09442) [gr-qc].
- [157] Jonathan Blackman et al. “A Surrogate Model of Gravitational Waveforms from Numerical Relativity Simulations of Precessing Binary Black Hole Mergers”. *Phys. Rev. D* 95.10 (2017), p. 104023. DOI: [10.1103/PhysRevD.95.104023](https://doi.org/10.1103/PhysRevD.95.104023). arXiv: [1701.00550](https://arxiv.org/abs/1701.00550) [gr-qc].
- [158] Jonathan Blackman et al. “Numerical relativity waveform surrogate model for generically precessing binary black hole mergers”. *Phys. Rev. D* 96.2 (2017), p. 024058. DOI: [10.1103/PhysRevD.96.024058](https://doi.org/10.1103/PhysRevD.96.024058). arXiv: [1705.07089](https://arxiv.org/abs/1705.07089) [gr-qc].
- [159] Michael Boyle. “Angular velocity of gravitational radiation from precessing binaries and the corotating frame”. *Phys. Rev. D* 87.10 (2013), p. 104006. DOI: [10.1103/PhysRevD.87.104006](https://doi.org/10.1103/PhysRevD.87.104006). arXiv: [1302.2919](https://arxiv.org/abs/1302.2919) [gr-qc].
- [160] Milton Ruiz et al. “Multipole expansions for energy and momenta carried by gravitational waves”. *Gen. Rel. Grav.* 40 (2008), p. 2467. DOI: [10.1007/s10714-007-0570-8](https://doi.org/10.1007/s10714-007-0570-8). arXiv: [0707.4654](https://arxiv.org/abs/0707.4654) [gr-qc].

- [161] 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”. *Phys. Rev. D* 91.2 (2015), p. 024043. DOI: [10.1103/PhysRevD.91.024043](https://doi.org/10.1103/PhysRevD.91.024043). arXiv: [1408.1810](https://arxiv.org/abs/1408.1810) [gr-qc].
- [162] Michael Boyle. “Transformations of asymptotic gravitational-wave data”. *Phys. Rev. D* 93.8 (2016), p. 084031. DOI: [10.1103/PhysRevD.93.084031](https://doi.org/10.1103/PhysRevD.93.084031). arXiv: [1509.00862](https://arxiv.org/abs/1509.00862) [gr-qc].
- [163] Jonathan Blackman et al. “Fast and Accurate Prediction of Numerical Relativity Waveforms from Binary Black Hole Coalescences Using Surrogate Models”. *Phys. Rev. Lett.* 115.12 (2015), p. 121102. DOI: [10.1103/PhysRevLett.115.121102](https://doi.org/10.1103/PhysRevLett.115.121102). arXiv: [1502.07758](https://arxiv.org/abs/1502.07758) [gr-qc].
- [164] Vijay Varma et al. “Surrogate model of hybridized numerical relativity binary black hole waveforms”. *Phys. Rev. D* 99.6 (2019), p. 064045. DOI: [10.1103/PhysRevD.99.064045](https://doi.org/10.1103/PhysRevD.99.064045). arXiv: [1812.07865](https://arxiv.org/abs/1812.07865) [gr-qc].
- [165] Gregorio Carullo, Walter Del Pozzo, and John Veitch. “Observational Black Hole Spectroscopy: A time-domain multimode analysis of GW150914”. *Phys. Rev. D* 99.12 (2019). [Erratum: *Phys.Rev.D* 100, 089903 (2019)], p. 123029. DOI: [10.1103/PhysRevD.99.123029](https://doi.org/10.1103/PhysRevD.99.123029). arXiv: [1902.07527](https://arxiv.org/abs/1902.07527) [gr-qc].
- [166] Richard Brito, Alessandra Buonanno, and Vivien Raymond. “Black-hole Spectroscopy by Making Full Use of Gravitational-Wave Modeling”. *Phys. Rev. D* 98.8 (2018), p. 084038. DOI: [10.1103/PhysRevD.98.084038](https://doi.org/10.1103/PhysRevD.98.084038). arXiv: [1805.00293](https://arxiv.org/abs/1805.00293) [gr-qc].
- [167] A. Buonanno and T. Damour. “Effective one-body approach to general relativistic two-body dynamics”. *Phys. Rev. D* 59 (1999), p. 084006. DOI: [10.1103/PhysRevD.59.084006](https://doi.org/10.1103/PhysRevD.59.084006). arXiv: [gr-qc/9811091](https://arxiv.org/abs/gr-qc/9811091).
- [168] Sizheng Ma et al. “Quasinormal-mode filters: A new approach to analyze the gravitational-wave ringdown of binary black-hole mergers”. *Phys. Rev. D* 106.8 (2022), p. 084036. DOI: [10.1103/PhysRevD.106.084036](https://doi.org/10.1103/PhysRevD.106.084036). arXiv: [2207.10870](https://arxiv.org/abs/2207.10870) [gr-qc].
- [169] Neil J. Cornish and Tyson B. Littenberg. “BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches”. *Class. Quant. Grav.* 32.13 (2015), p. 135012. DOI: [10.1088/0264-9381/32/13/135012](https://doi.org/10.1088/0264-9381/32/13/135012). arXiv: [1410.3835](https://arxiv.org/abs/1410.3835) [gr-qc].
- [170] Neil J. Cornish et al. “BayesWave analysis pipeline in the era of gravitational wave observations”. *Phys. Rev. D* 103.4 (2021), p. 044006. DOI: [10.1103/PhysRevD.103.044006](https://doi.org/10.1103/PhysRevD.103.044006). arXiv: [2011.09494](https://arxiv.org/abs/2011.09494) [gr-qc].

- [171] Gregory Ashton et al. “BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy”. *Astrophys. J. Suppl.* 241.2 (2019), p. 27. DOI: [10.3847/1538-4365/ab06fc](https://doi.org/10.3847/1538-4365/ab06fc). arXiv: [1811.02042](https://arxiv.org/abs/1811.02042) [astro-ph.IM].
- [172] Colm Talbot et al. “Inference with finite time series: Observing the gravitational Universe through windows” (June 2021). arXiv: [2106.13785](https://arxiv.org/abs/2106.13785) [astro-ph.IM].
- [173] Warren G. Anderson et al. “An Excess power statistic for detection of burst sources of gravitational radiation”. *Phys. Rev. D* 63 (2001), p. 042003. DOI: [10.1103/PhysRevD.63.042003](https://doi.org/10.1103/PhysRevD.63.042003). arXiv: [gr-qc/0008066](https://arxiv.org/abs/gr-qc/0008066).
- [174] The LIGO Scientific and Virgo Collaborations. *GWTC-2 Data Release: Parameter Estimation Samples and Skymaps*. URL: <https://dcc.ligo.org/LIGO-P2000223/public>.
- [175] The LIGO Scientific and Virgo Collaborations. *Noise curves used for Simulations in the update of the Observing Scenarios Paper*. URL: <https://dcc.ligo.org/LIGO-T2000012/public>.
- [176] Riccardo Buscicchio et al. “Label switching problem in Bayesian analysis for gravitational wave astronomy”. *Phys. Rev. D* 100.8 (2019), p. 084041. DOI: [10.1103/PhysRevD.100.084041](https://doi.org/10.1103/PhysRevD.100.084041). arXiv: [1907.11631](https://arxiv.org/abs/1907.11631) [astro-ph.IM].
- [177] Joshua S. Speagle. “dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences”. *Mon. Not. Roy. Astron. Soc.* 493.3 (2020), pp. 3132–3158. DOI: [10.1093/mnras/staa278](https://doi.org/10.1093/mnras/staa278). arXiv: [1904.02180](https://arxiv.org/abs/1904.02180) [astro-ph.IM].
- [178] John Skilling. “Nested Sampling”. *AIP Conference Proceedings* 735.1 (2004), pp. 395–405. DOI: [10.1063/1.1835238](https://doi.org/10.1063/1.1835238). eprint: <https://aip.scitation.org/doi/pdf/10.1063/1.1835238>. URL: <https://aip.scitation.org/doi/abs/10.1063/1.1835238>.
- [179] John Skilling. “Nested sampling for general Bayesian computation”. *Bayesian Analysis* 1.4 (2006), pp. 833–859. DOI: [10.1214/06-BA127](https://doi.org/10.1214/06-BA127).
- [180] The LIGO Scientific and Virgo Collaborations. *Power Spectral Densities (PSD) release for GWTC-1*. URL: <https://dcc.ligo.org/LIGO-P1900011/public>.
- [181] B. P. Abbott et al. “Tests of general relativity with GW150914”. *Phys. Rev. Lett.* 116.22 (2016). [Erratum: *Phys.Rev.Lett.* 121, 129902 (2018)], p. 221101. DOI: [10.1103/PhysRevLett.116.221101](https://doi.org/10.1103/PhysRevLett.116.221101). arXiv: [1602.03841](https://arxiv.org/abs/1602.03841) [gr-qc].
- [182] P. Whittle. “Curve and Periodogram Smoothing”. *Journal of the Royal Statistical Society. Series B (Methodological)* 19.1 (1957), pp. 38–63. ISSN: 00359246. URL: <http://www.jstor.org/stable/2983994> (visited on 01/17/2023).
- [183] Sizheng Ma, Ling Sun, and Yanbei Chen. “Using rational filters to uncover the first ringdown overtone in GW150914”. *Phys. Rev. D* 107.8 (2023), p. 084010. DOI: [10.1103/PhysRevD.107.084010](https://doi.org/10.1103/PhysRevD.107.084010). arXiv: [2301.06639](https://arxiv.org/abs/2301.06639) [gr-qc].

- [184] Sizheng Ma, Ling Sun, and Yanbei Chen. “Black Hole Spectroscopy by Mode Cleaning”. *Phys. Rev. Lett.* 130.14 (2023), p. 141401. DOI: [10.1103/PhysRevLett.130.141401](https://doi.org/10.1103/PhysRevLett.130.141401). arXiv: [2301.06705](https://arxiv.org/abs/2301.06705) [gr-qc].
- [185] Roberto Cotesta et al. “On the detection of ringdown overtones in GW150914”. (Jan. 2022). arXiv: [2201.00822](https://arxiv.org/abs/2201.00822) [gr-qc].
- [186] Maximiliano Isi and Will M. Farr. “Revisiting the ringdown of GW150914”. (Feb. 2022). arXiv: [2202.02941](https://arxiv.org/abs/2202.02941) [gr-qc].
- [187] Juan Calderón Bustillo, Paul D. Lasky, and Eric Thrane. “Black-hole spectroscopy, the no-hair theorem, and GW150914: Kerr versus Occam”. *Phys. Rev. D* 103.2 (2021), p. 024041. DOI: [10.1103/PhysRevD.103.024041](https://doi.org/10.1103/PhysRevD.103.024041). arXiv: [2010.01857](https://arxiv.org/abs/2010.01857) [gr-qc].
- [188] The LIGO Scientific and Virgo Collaborations. *Gravitational Wave Open Science Center*. URL: <https://www.gw-openscience.org/>.
- [189] Rich Abbott et al. “Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo”. *SoftwareX* 13 (2021), p. 100658. DOI: [10.1016/j.softx.2021.100658](https://doi.org/10.1016/j.softx.2021.100658). arXiv: [1912.11716](https://arxiv.org/abs/1912.11716) [gr-qc].
- [190] Christian Robert and George Casella. *Monte Carlo statistical methods*. Springer texts in statistics. Springer, 2013. ISBN: 038798707X.
- [191] Ethan Payne, Colm Talbot, and Eric Thrane. “Higher order gravitational-wave modes with likelihood reweighting”. *Phys. Rev. D* 100.12 (2019), p. 123017. DOI: [10.1103/PhysRevD.100.123017](https://doi.org/10.1103/PhysRevD.100.123017). arXiv: [1905.05477](https://arxiv.org/abs/1905.05477) [astro-ph.IM].
- [192] Maximiliano Isi and Will M. Farr. *Revisiting the ringdown of GW150914*. Version v1. Zenodo, Feb. 2022. DOI: [10.5281/zenodo.5965772](https://doi.org/10.5281/zenodo.5965772). URL: <https://doi.org/10.5281/zenodo.5965772>.
- [193] The LIGO Scientific and Virgo Collaborations. *Parameter estimation sample release for GWTC-1*. URL: <https://dcc.ligo.org/LIGO-P1800370/public>.
- [194] James M. Dickey. “The Weighted Likelihood Ratio, Linear Hypotheses on Normal Location Parameters”. *The Annals of Mathematical Statistics* 42.1 (1971), pp. 204–223. ISSN: 00034851. URL: <http://www.jstor.org/stable/2958475> (visited on 01/17/2023).
- [195] Will M. Farr and Maximiliano Isi. *private communication*. 2022.
- [196] The LIGO Scientific and Virgo Collaborations. *Data Release for “Population properties of compact objects from the second LIGO-Virgo Gravitational-Wave Transient Catalog”*. URL: <https://dcc.ligo.org/LIGO-P2000434/public>.
- [197] Craig Cahillane et al. “Calibration uncertainty for Advanced LIGO’s first and second observing runs”. *Phys. Rev. D* 96.10 (2017), p. 102001. DOI: [10.1103/PhysRevD.96.102001](https://doi.org/10.1103/PhysRevD.96.102001). arXiv: [1708.03023](https://arxiv.org/abs/1708.03023) [astro-ph.IM].

-
- [198] Collin Capano. *pykerr*. URL: <https://github.com/cdcapano/pykerr>.
- [199] Mark Ho-Yeuk Cheung et al. “Nonlinear Effects in Black Hole Ringdown”. *Phys. Rev. Lett.* 130.8 (2023), p. 081401. DOI: [10.1103/PhysRevLett.130.081401](https://doi.org/10.1103/PhysRevLett.130.081401). arXiv: [2208.07374](https://arxiv.org/abs/2208.07374) [gr-qc].
- [200] Keefe Mitman et al. “Nonlinearities in Black Hole Ringdowns”. *Phys. Rev. Lett.* 130.8 (2023), p. 081402. DOI: [10.1103/PhysRevLett.130.081402](https://doi.org/10.1103/PhysRevLett.130.081402). arXiv: [2208.07380](https://arxiv.org/abs/2208.07380) [gr-qc].