



OBSERVATIONAL TESTS OF FUNDAMENTAL PHYSICS FROM GRAVITATIONAL WAVE DETECTIONS

Von der Fakultät für Mathematik und Physik
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
Dr. rer. nat.

genehmigte Dissertation von

M.Sc. Julian Westerweck

2023

Referent:

Dr. Frank Ohme

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut) Hannover
Leibniz Universität Hannover

Korreferenten:

Prof. Dr. Badri Krishnan

Radboud Universiteit Nijmegen

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut) Hannover
Leibniz Universität Hannover

Prof. Dr. Alex B. Nielsen

Universitetet i Stavanger

Tag der Promotion:

24.06.2022

Promotionskommission:

Prof. Dr. Gunther Seckmeyer
Leibniz Universität Hannover

Prof. Dr. Badri Krishnan
Radboud Universiteit Nijmegen
Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut) Hannover
Leibniz Universität Hannover

Prof. Dr. Vitor Cardoso
Instituto Superior Técnico da Universidade de Lisboa

Prof. Dr. Alex B. Nielsen
Universitetet i Stavanger

Prof. Dr. Bruce Allen
Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut) Hannover
Leibniz Universität Hannover

Abstract

With the detection of the signal GW150914 from the collision of two black holes in 2015, observational gravitational wave physics has begun. Many more signals have since been recorded, and new detections are now becoming routine. These observations offer a new window to probe fundamental physics in thus far inaccessible regimes of strong gravity, such as in the regions near black hole horizons. The work presented here pursues this through two approaches, studying predicted signals of either black holes of general relativity, or of proposed alternative objects without horizons.

A binary black hole collision creates a single perturbed black hole, which settles to its final state through the ringdown gravitational wave emission. The ringdown consists of a spectrum of modes, which the no-hair theorem in General Relativity predicts to be determined entirely by the black hole mass and angular momentum. Measurement of multiple modes allows to test this prediction but is challenging due to the weak and short-lived nature of the ringdown signal.

Two studies are presented on the feasibility of such tests using current and near-future detector sensitivities. Large populations of simulated ringdown signals are constructed based on observational models of the binary black hole population. Bayesian parameter estimation techniques are applied to these signals to place bounds on deviations from the no-hair prediction. Detections leading to stringent bounds are unlikely to occur for current instruments but can be found during a few years of operation at their planned future sensitivities. The prospects improve when extending the analysis to combine data from multiple detections into a single bound on deviations. At the sensitivity planned for the next observation run of current instruments, the detections from one year of data can be combined into stringent bounds. Solutions are provided to limitations uncovered for this type of study.

In a further study, strong evidence is found for the presence of a subdominant mode in the data of the event GW190521. A new method is employed to allow the analysis of only the ringdown part of the signal, without contamination from outside the analysis window and preventing windowing artefacts and signal loss. Tests of the no-hair theorem are performed, yielding unexpectedly tight constraints on deviations.

Two phenomenologically distinct signals from horizonless compact objects are studied, both following after the primary signal which is otherwise unchanged compared to that of a black hole binary. One takes the form of repeated pulses after the ringdown, called gravitational wave echoes, while the other consists of a very long-lived damped sinusoid with a small amplitude.

Using a simplified waveform model for echoes, evidence for such signals in the data of several detections is evaluated. Previous results from the first search for these are replicated, and the methods tested thoroughly. Through improved estimation methods, low statistical significance is established for these results, yet the presence of such signals cannot be ruled out by the analysis. An independent Bayesian analysis is performed for the same waveform model, with results for each event either preferring the absence of echoes in the data or being consistent with it. Bounds on the echo amplitudes ruled out by the data are produced.

The long-lived mode signal for a broad class of horizonless objects is considered in a Bayesian analysis. Methods are developed to accommodate the long duration of the signal, and their performance is tested with simulated signals and off-source data. They are then applied to the data of the event GW150914, yielding stringent bounds on the deviations from the Kerr geometry exhibited by such objects.

Keywords: Gravitational waves, black holes, tests of gravity theory, data analysis

Zusammenfassung

Mit der Detektion des Signals GW150914 von der Kollision zweier schwarzer Löcher im Jahr 2015 begann die beobachtungs-basierte Gravitationswellenphysik. Viele weitere Signale wurden seither aufgezeichnet und neue Detektionen werden zur Routine. Diese Beobachtungen eröffnen einen neuen Weg, fundamentale Physik im bisher unzugänglichen Regime starker Gravitation zu untersuchen, zum Beispiel in der Umgebung der Horizonte schwarzer Löcher. Die hier präsentierten Studien verfolgen dies durch zwei Ansätze, indem sie entweder die vorhergesagten Signale schwarzer Löcher in der Allgemeinen Relativitätstheorie oder vorgeschlagener alternativer Objekte ohne Horizonte untersuchen.

Die Kollision zweier schwarzer Löcher erzeugt ein einzelnes gestörtes schwarzes Loch, welches durch Emission der Abkling-Gravitationswellen schließlich in einen ungestörten Zustand übergeht. Die Abkling-Strahlung besteht aus einem Spektrum von Moden, welche dem Keine-Haare-Theorem der Allgemeinen Relativitätstheorie nach gänzlich durch Masse und Drehimpuls des schwarzen Loches bestimmt werden. Die Messung mehrerer Moden ermöglicht die Prüfung dieser Vorhersage, ist jedoch wegen des schwachen und kurzlebigen Abklingsignals schwierig.

Zwei Studien zur Durchführbarkeit solcher Tests mithilfe aktuell und in naher Zukunft verfügbarer Detektor-Empfindlichkeiten werden dargelegt. Große Populationen simulierter Abklingsignale werden konstruiert, basierend auf beobachtungsgestützten Modellen der Population von Binärsystemen schwarzer Löcher. Bayessche Parameterabschätzung wird auf diese Signale angewendet, um Abweichungen von der Keine-Haare-Vorhersage zu beschränken. Detektionen, die zu strikter Begrenzung führen, sind mit aktuellen Instrumenten unwahrscheinlich, können aber innerhalb weniger Jahre des Betriebs mit ihren geplanten zukünftigen Empfindlichkeiten erreicht werden. Diese Aussichten verbessern sich, wenn Daten mehrerer Detektionen in der Begrenzung kombiniert werden. Mit der geplanten Empfindlichkeit aktueller Instrumente im nächsten Beobachtungslauf können die in einem Jahr gesammelten Daten zu strikten Begrenzungen kombiniert werden. Lösungen für die entdeckten Limitationen dieser Art Analyse werden vorgestellt.

In einer weiteren Studie wird starke Evidenz für die Existenz einer subdominanten Mode in den Daten des Signals GW190521 gefunden. Eine neue Methode wird eingesetzt, welche die Analyse des Abkling-Signals ermöglicht, ohne Kontamination von außerhalb des Analyse-Fensters, Artefakte oder Signalverlust zu verursachen. Tests des Keine-Haare-Theorems werden durchgeführt und liefern unerwartet strikte Beschränkungen für Abweichungen.

Zwei phänomenologisch verschiedene Signale horizontfreier kompakter Objekte werden untersucht. Beide folgen dem Primärsignal, das ansonsten gegenüber dem schwarzer Löcher unverändert ist. Eines besteht aus wiederholten Pulsen, als Gravitationswellen-Echos bezeichnet, während das zweite die Form einer langlebigen, gedämpften Sinuswelle geringer Amplitude hat.

Anhand eines vereinfachten Modells der Echo-Wellenform wird die Evidenz solcher Signale in den Daten mehrerer Detektionen bewertet. Frühere Ergebnisse der ersten Suche nach Echos werden repliziert und die Methoden ausführlich geprüft. Durch verbesserte Abschätzungsmethoden wird eine geringe statistische Signifikanz der Ergebnisse etabliert, allerdings kann die Anwesenheit solcher Signale nicht durch diese Untersuchung ausgeschlossen werden. Eine unabhängige Bayessche Analyse wird mit derselben Wellenform durchgeführt, wobei die Ergebnisse die Abwesenheit des Signals bevorzugen oder mit Rauschen vereinbar sind. Grenzen für die von den Daten ausgeschlossenen Amplituden der Echos werden gefunden.

Das Signal einer langlebigen Mode von einer großen Klasse horizontfreier Objekte wird in einer Bayesschen Analyse betrachtet. Methoden werden entwickelt, um die lange Dauer des Signals handhaben zu können, und ihre Leistungsfähigkeit wird an simulierten Signalen und signalfreien Daten getestet. Auf die Daten des Signals GW150914 angewendet, liefern sie strikte Beschränkungen für die Abweichungen solcher Objekte von der Kerr-Geometrie.

Schlagworte: Gravitationswellen, Schwarze Löcher, Gravitationstheorie-Tests, Datenanalyse

Contents

1	Introduction	11
1.1	Introduction	11
1.2	Chapters and contributions	11
1.3	General relativity and gravitational waves	13
1.3.1	Sources and types of signals	20
1.4	Gravitational wave detectors	23
1.5	Data analysis, Bayesian inference, and parameter estimation	27
1.5.1	Data model	28
1.5.2	Sampling parameter space	32
1.5.3	Model selection	33
1.6	Quasi-normal modes	35
1.6.1	No-hair theorem	42
1.7	Echoes	43
1.8	Results in context	47
2	Prospects for black hole spectroscopy in the next decade	51
2.1	Introduction	51
2.2	Bayesian framework	52
2.2.1	The likelihood function	53
2.2.2	The ringdown model	53
2.3	Populations	54
2.4	Analysis and results	55
2.4.1	Rates of measurable subdominant modes	55
2.4.2	Resolvable subdominant modes for testing GR	57
2.5	Conclusions	58
2.6	Appendix	60
3	Hierarchical black hole spectroscopy	63
3.1	Introduction	63
3.2	Analysis	64
3.2.1	Combining posteriors	64
3.2.2	Populations and analysis framework	66
3.3	Results	69
3.4	Conclusions	78
4	Black hole spectroscopy with GW190521	81
4.1	Introduction	81
4.2	Multimode agnostic search	82
4.3	Consistency with the Kerr solution	84
4.4	Discussion and conclusions	87
4.5	Supplemental	89
4.6	Materials and Methods	89
4.6.1	The ringdown signal model	89

4.6.2	Computational analysis methods	90
5	Statistical significance of evidence for gravitational wave echoes	99
5.1	Introduction	99
5.2	ADA’s Model and search procedure	100
5.3	General remarks	102
5.4	Validation of the matched-filter analysis	103
5.5	Prior ranges and template spacing	104
5.6	Extending the background estimation	107
5.7	Conclusions	117
6	Bayesian parameter estimation for gravitational wave echoes	121
6.1	Introduction	121
6.2	Methodology and analysis pipeline	123
6.3	Injections based on GW150914	124
6.4	Events in the first observing run	130
6.5	Discussion and conclusions	132
6.6	Appendix	133
6.6.1	Fiducial IMR waveform parameters	133
7	Constraining the Kerr-geometry of GW150914 with horizonless modes	135
7.1	Introduction	135
7.2	Theoretical framework	137
7.3	Signal model	138
7.4	Results	141
7.5	Conclusion and outlook	143
7.6	Appendix	144
7.6.1	Justifying the assumption of full reflection	144
7.6.2	Lower bound for ϵ	145
7.6.3	Data analysis details	146
7.6.4	Validation with noise and simulated signals	147
8	Conclusions and outlook	151
9	Acknowledgements	175
10	Curriculum Vitae	177

1 Introduction

1.1 Introduction

The theory of general relativity very successfully describes the phenomenon of gravity. It has been rigorously tested in the regimes available to observation and experiment. Among its predictions is the existence of exotic phenomena, such as black holes as objects in the strong field regime and the subtle effects of gravitational waves. Despite being two of the first solutions to the Einstein field equations found, the experimental accessibility of both has proven difficult.

These very strong and very weak gravitational effects meet in the form of gravitational-wave observations. While sources of gravitational waves are in principle ubiquitous, the majority emits only undetectably faint signals. The gravitational wave emission increases for large masses and accelerations, making binary systems of massive, compact objects in tight orbits a promising candidate. Such compact objects known to exist are, by increasing compactness, white dwarfs, neutron stars and black holes. Even these sources emit gravitational waves that for their respective typical distances result in minute effects on Earth. There, they would displace test masses by only the 10^{-21} -st part of their separation distance, corresponding to about a thousandth part of the diameter of a proton for masses that are separated by kilometres.

Nevertheless, detectors have been constructed that first observed a gravitational wave directly only a few years ago and are now routinely measuring gravitational-wave signals from the collisions of compact objects.

These detections not only allow to study the properties of gravitational waves themselves, but to observe otherwise dark objects. While black holes are studied through the orbits of nearby stars or radiation emitted from their accretion disk, gravitational waves allow their observation when those components are absent or undetectable. Additionally, gravitational waves may allow probing regions otherwise inaccessible, such as the vicinity of the black hole's event horizon.

Individual black holes may be studied through their gravitational wave emission when they are perturbed. As the collision of sufficiently massive compact objects creates a highly perturbed black hole, the late stage of a binary collision's signal consists of this emission. In many aspects, black holes are remarkably simple objects, which is reflected in the simple relations found in the signal of a perturbed black hole. Measuring these relations allows constraining potential deviations from the signal predicted in general relativity. Similarly, some theories introduce modifications to the description of black holes or propose alternative objects, leading to modifications in the expected signal that can be targeted in a search. Constraining deviations from the expected signal and determining the support for such modifications in the data helps to select the most promising theories for further study.

1.2 Chapters and contributions

The work presented in this thesis has been carried out in collaboration with several colleagues, and most of it has been published in the literature or is available as pre-print. This section presents a short overview of the contents of the respective chapters and details the individual contributions.

1 Introduction

Chapter 1 provides a general introduction to the basics of gravitational wave observation. The background of the signals studied in the following chapters is discussed, in particular the quasi-normal modes of perturbed black holes and late-time signals of proposed alternative compact objects. Finally, the concepts and tools for the analysis of the detector data are introduced, focusing on techniques to estimate the parameter values of given models and the preference for these models as supported by the data. The chapter was written by the author.

Chapter 2 is an adaptation of the work published as [1] and deals with the prospects of observational tests of the black hole no-hair theorem. We construct populations of simulated signals based on observational models of the binary black hole merger population and perform Bayesian parameter estimation on the signals. Tests of the no-hair prediction are carried out on signals where a multi-mode ringdown spectrum is found. Repeating this analysis for different current and future detector sensitivities, we estimate the rates of events suited to place stringent constraints on deviations from the mode relations expected in general relativity.

This work was led by the author, together with M. Cabero. The author has implemented the methods required to apply the Bayesian inference to the desired large populations of events. The populations were constructed and the parameter estimation executed by the author and M. Cabero. They led the discussion and interpretation of the results and the writing of the publication, each with participation from the other authors of [1]. S. Kumar calculated the rate of merger events occurring in the desired co-moving volumes.

The work presented in Chapter 3 has not yet been published. It continues the study of prospects for no-hair tests by extending the previous analysis to incorporate methods combining information from multiple events into a single measurement. These methods combine the results from the parameter estimation analysis of individual detections into a single probability distribution for deviations from the predictions of general relativity. We discover fundamental limitations of this type of analysis and provide strategies to prevent the recovery of biased results. Performing the analysis again on populations of simulated signals, we estimate the results likely obtainable through these methods using one year of data with the detector sensitivity planned for the next observing run.

The author implemented the tool used to combine results from multiple detections and performed the parameter estimation analysis. The performance of the new analysis and the limiting bias were characterised by the author. The authors of [1] participated in discussion and interpretation of the results. The chapter was written by the author.

Chapter 4 is an adaptation of the work in the pre-print [2], which presents strong evidence for the presence of multiple quasi-normal modes in the ringdown signal of the gravitational wave event GW190521. This evidence is obtained through an agnostic and a targeted approach, either assuming the expected relations between the modes' parameters or allowing them to vary freely. Each approach is applied to multiple times around the reported merger time of the event to account for uncertainties. A new method to exclude data before the start of the ringdown from the analysis is presented and applied, simultaneously preventing the introduction of artefacts and the loss of signal. The measured modes are then used for an independent estimate of the binary's mass-ratio and to perform no-hair tests.

This work was led by C. Capano, who performed the majority of the parameter estimation runs, while the author and M. Cabero performed subsets. C. Capano, S. Kastha and the author implemented and tested the method for data exclusion. The accompanying data release was prepared by C. Capano, M. Cabero and the author, and they produced the visualisation of the results with J. Abedi and S. Kastha. All authors of [2] contributed to discussions and interpretation of the results and to the writing of the paper.

Chapter 5 is an adaptation of the work published as [3] and deals with the statistical interpretation of results found previously in the first search for gravitational wave echo signals of alter-

native horizonless compact objects. An independent implementation of the employed methods is constructed and characterised in detail, uncovering features directly affecting the conclusions drawn from the search results. The original results are replicated, and a new estimate of their statistical significance is given, using extensive studies of simulated and real detector noise. The statistical significance is found to be low, while at the same time the presence of such signals cannot be ruled out confidently through these methods.

This work was led by the author, who implemented and tested the methods and carried out the analyses, with guidance from A. Nielsen and O. Birnholtz. All authors of [3] participated in the discussion of the results and editing of the publication. The writing of the publication was again led by the author with contributions from A. Nielsen and O. Birnholtz.

Chapter 6 is an adaptation of the work published as [4] and provides a Bayesian analysis searching for the same echo signals as discussed in the previous chapter. An independent assessment is given for the support for such models in the data and for their estimated parameters. Simulated signals are used to characterise the performance of the analysis. Through this, the model selection results are validated, and bounds are placed on the amplitudes of potential signals.

This work was led by A. Nielsen, who carried out the parameter estimation analysis. The author provided the echo waveform model and with C. Capano implemented improvements and increased its efficiency. Integration into the analysis pipeline was implemented by A. Nielsen, C. Capano and the author. The author and O. Birnholtz contributed to the discussion and interpretation of the results, and to the writing of the publication, which were led by A. Nielsen and C. Capano.

Chapter 7 is an adaptation of the work in the pre-print [5] and presents a search for modified quasi-normal mode signals from proposed horizonless compact objects. A new analysis method is developed and employed to accommodate this long-lived signal. Applying it to the data of GW150914, stringent bounds are placed on the deviation from the Kerr geometry for a class of the proposed objects.

This work was led by the author, who implemented and tested the new methods and carried out the data analysis. C. Capano provided guidance in the development of the methods, while Y. Sherf led the adaptation of the model of the horizonless object. Y. Sherf and the author led writing the publication, while all authors of [5] contributed to the text, as well as to discussions and interpretation of the results.

Chapter 8 provides some overall conclusions from the work presented in this thesis.

1.3 General relativity and gravitational waves

In general relativity, spacetime is modelled as a 4-dimensional Lorentzian manifold (M, g) and events are identified with points in this manifold. The metric tensor field g in Einstein's Theory of Gravity obeys the Einstein field equations,

$$G_{\mu\nu} = 8\pi T_{\mu\nu}, \quad (1.1)$$

where $G_{\mu\nu} = R_{\mu\nu} - \frac{R}{2}g_{\mu\nu}$ is the Einstein tensor and $T_{\mu\nu}$ is the stress-energy-momentum tensor. In this description, freely falling particles follow geodesic curves. For Eq. (1.1) we used geometric units, wherein the gravitational constant G and the speed of light c in vacuum are $G = c = 1$. In SI units, an additional factor G/c^4 appears on the right-hand side.

These rather compact statements are at the core of general relativity (GR), and this section shortly reviews the concepts and gives an overview of how they lead to the phenomena measured and analysed in the following chapters. They are detailed in many textbooks, and this introduction is based on several of them [6–11].

1 Introduction

A manifold M is a topological space that locally resembles Euclidean space in that it is locally homeomorphic to it. The corresponding homeomorphisms allow us to assign coordinates to the manifold, and several of these coordinate maps may be required to cover the manifold completely. A manifold is differentiable if it is covered by a collection of coordinate maps with differentiable transition functions between overlapping maps. The resulting differentiable structure allows us to work with vectors of the tangent space $T_p M$ and the one-forms of its dual, the cotangent space $T_p^* M$, at each point p of the manifold. Consequently, we can consider tensor spaces $(T_p)_n^m M$ as the tensor product of m copies of $T_p M$ and n copies of $T_p^* M$, or equivalently as multilinear mappings of m one-forms and n tangent vectors into the real numbers \mathbb{R} .

Tensor fields assign every $p \in M$ an element of the corresponding tensor space $(T_p)_n^m M$, of which vector fields and one-form fields are special cases. Choosing a basis of vector fields e_μ and one-form fields θ^ν , we write a tensor field as $T^\mu{}_\nu e_\mu \theta^\nu$, using here and in the following the convention of summing over indices appearing twice unless noted otherwise. The basis will often be omitted from equations, but the equations will be those between the tensors as geometric objects, not their components. We introduce the metric tensor field g on the differentiable manifold, a non-degenerate, symmetric tensor field, which maps two tangent vectors to \mathbb{R} . Finally, with the metric tensor having signature $(-, +, +, +)$, (M, g) is a Lorentzian manifold, modelling spacetime.

Deciphering the Einstein field equations, we find therein

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R g_{\mu\nu}, \quad (1.2)$$

$$R_{\mu\nu} = R^\alpha{}_{\mu\alpha\nu}, \quad (1.3)$$

$$R = R^\mu{}_\mu, \quad (1.4)$$

$$R^\alpha{}_{\beta\mu\nu} = \partial_\mu \Gamma^\alpha{}_{\nu\beta} - \partial_\nu \Gamma^\alpha{}_{\mu\beta} + \Gamma^\alpha{}_{\mu\gamma} \Gamma^\gamma{}_{\nu\beta} - \Gamma^\alpha{}_{\nu\gamma} \Gamma^\gamma{}_{\mu\beta}, \quad (1.5)$$

$$\Gamma^\alpha{}_{\mu\nu} = \frac{1}{2}g^{\alpha\beta} (\partial_\mu g_{\beta\nu} + \partial_\nu g_{\mu\beta} - \partial_\beta g_{\mu\nu}), \quad (1.6)$$

the Einstein tensor $G_{\mu\nu}$, Ricci tensor $R_{\mu\nu}$, Ricci scalar R , Riemann curvature tensor $R^\alpha{}_{\mu\beta\nu}$, and connection coefficients or Christoffel symbols $\Gamma^\alpha{}_{\mu\nu}$ on the left-hand side, and the energy-stress-momentum tensor $T_{\mu\nu}$ on the right-hand side. Partial derivatives $\partial_\mu = \frac{\partial}{\partial x^\mu}$ are taken with respect to the coordinate functions x^μ . We use the convention of Greek spacetime indices taking values $0, 1, 2, 3$ and Latin spatial indices taking values $1, 2, 3$. With each tangent vector, the metric tensor $g_{\mu\nu}$ associates a one-form from the dual cotangent space. Similarly, the inverse of the metric g^{-1} , with components $g^{\mu\nu}$, defines a tangent vector for each one-form of the cotangent space. In component notation, this takes the form of raising and lowering indices,

$$g_{\mu\alpha} g^{\alpha\nu} = \delta_\mu^\nu \quad T_\mu = g_{\mu\nu} T^\nu \quad T^\mu = g^{\mu\nu} T_\nu, \quad (1.7)$$

where $\delta_\nu^\mu = 1$ for $\mu = \nu$ and $\delta_\nu^\mu = 0$ for $\mu \neq \nu$.

The effect of gravity on a freely falling particle, i.e. with no forces acting upon it, is described by the particle following a geodesic curve through spacetime. Its path $x^\mu(\lambda)$, parameterised by λ , then obeys

$$\frac{d^2 x^\mu}{d\lambda^2} + \Gamma^\mu{}_{\nu\alpha} \frac{dx^\nu}{d\lambda} \frac{dx^\alpha}{d\lambda} = 0. \quad (1.8)$$

If additional forces act on the particle, they appear as F^μ on the right-hand side.

The left-hand side of Eq. (1.1) describes the curvature of spacetime. On the right-hand side, the energy-stress-momentum tensor describes its matter, energy, and momentum content, and

their fluxes. Thus, the energy and momentum content of spacetime can be seen as determining its curvature, while its curvature determines the motion of particles. An additional cosmological constant term may be added on the left-hand side as $\Lambda g_{\mu\nu}$, which can also be formally absorbed into $T_{\mu\nu}$, and which is relevant for the observed accelerating expansion of the universe.

Taking into account its components, we see that the Einstein field equations represent a coupled system of nonlinear partial differential equations for the metric $g_{\mu\nu}$. Finding a metric that is a solution for the equations is thus difficult, and the known solutions rely on simplifications or approximations. A simple analytic solution is the Minkowski-metric [12–14], a vacuum with $T_{\mu\nu} = 0$, describing flat spacetime. Important analytic solutions are the vacuum solutions describing black holes, the Schwarzschild metric [15] for non-rotating, and the Kerr metric [16] for rotating black holes. The Friedmann-Lemaître-Robertson-Walker metric [17–23] describes a homogeneous and isotropic universe and is relevant in cosmology. Numerical Relativity uses numerical methods to solve the Einstein field equations for more complex cases where no analytic solution is known, and can yield approximate solutions of high accuracy. Finally, using the approximation of only small deviations from a flat metric allows us to find solutions to the linearised equation, importantly leading to a wave solution.

The components of the metric tensor are commonly shown by writing the line element $ds^2 = g_{\mu\nu} dx^\mu dx^\nu$, from which the individual components can be read off. The line element of the Minkowski metric is given by

$$ds^2 = -dt^2 + dx^2 + dy^2 + dz^2. \quad (1.9)$$

The Schwarzschild solution is an asymptotically flat, static, spherically symmetric, vacuum solution. It is used as an approximation to describe a non-rotating astrophysical black hole.

The line element of the Schwarzschild metric in Schwarzschild coordinates is given by

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \left(1 - \frac{2M}{r}\right)^{-1} dr^2 + r^2 (d\theta^2 + \sin^2 \theta d\phi^2). \quad (1.10)$$

The Reissner-Nordström metric [24] is a generalisation to include electric charge, however, astrophysical black holes are expected to possess only negligible electric charge.

The Kerr metric is an asymptotically flat, stationary, axially-symmetric, vacuum solution. It is used as a more realistic approximation to describe astrophysical black holes, which are expected to possess angular momentum. The line element of the Kerr metric in Boyer-Lindquist coordinates [25] is given by

$$ds^2 = -\left(1 - \frac{2Mr}{\Sigma}\right) dt^2 - \frac{4Mr}{\Sigma} a \sin^2 \theta d\phi dt + \frac{\Sigma}{\Delta} dr^2 + \Sigma d\theta^2 + \left((r^2 + a^2) \sin^2 \theta + \frac{2Mr}{\Sigma} a^2 \sin^4 \theta\right) d\phi^2, \quad (1.11)$$

with $\Sigma = r^2 + a^2 \cos^2 \theta$, and $\Delta = r^2 + a^2 - 2Mr = (r - r_+)(r - r_-)$, where $r_{\pm} = M \pm \sqrt{M^2 - a^2}$. The angular momentum J and mass M appear in the spin parameter $a = \frac{J}{Mc}$, and we will often refer to the dimensionless spin parameter $\chi = \frac{a}{M}$. The generalisation of the Kerr solution to include electric charge is the Kerr-Newman metric [26, 27].

The Friedmann-Lemaître-Robertson-Walker metric describes a homogeneous and isotropic universe but is not restricted to a vacuum. Depending on whether its spatial geometry is spherical, flat, or hyperbolic, it reads

1 Introduction

$$ds^2 = -d\tau^2 + a^2(\tau) \begin{cases} d\psi^2 + \sin^2 \psi (d\theta^2 + \sin^2 \theta d\phi^2) \\ dx^2 + dy^2 + dz^2 \\ d\psi^2 + \sinh^2 \psi (d\theta^2 + \sin^2 \theta d\phi^2) \end{cases}. \quad (1.12)$$

Gravitational waves can be found as small perturbations of an otherwise flat Minkowski spacetime, which approximates weak gravitational fields far from sources. For this approximately flat metric, coordinates can be found so that it can be written as

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.13)$$

where $\eta_{\mu\nu}$ is the Minkowski metric and for all components of $h_{\mu\nu}$ we have $|h_{\mu\nu}| \ll 1$ everywhere in spacetime in these coordinates [6]. We now follow the derivation in [7] and will raise and lower indices with $\eta_{\mu\nu}$ in the following, treating $h_{\mu\nu}$ like a tensor in a flat spacetime, while restricting coordinate transformations to gauge- and background Lorentz-transformations. These gauge-transformations satisfy $x'^{\alpha} = x^{\alpha} + \xi^{\alpha}(x^{\alpha})$ with $|\partial_{\beta}\xi^{\alpha}| \ll 1$, while the background Lorentz transformations simply have constant components as would ordinary Lorentz transformations when $h_{\mu\nu}$ vanished. Then, $h_{\mu\nu}$ is treated like a tensor and we define the trace reverse of $h_{\mu\nu}$,

$$\bar{h}^{\mu\nu} = h^{\mu\nu} - \frac{1}{2}\eta^{\mu\nu}h, \quad (1.14)$$

where $h = h^{\mu}_{\mu}$ is the trace of $h_{\mu\nu}$. We can perform gauge transformations to use the Lorenz gauge¹, such that

$$\partial_{\mu}\bar{h}^{\mu\nu} = 0. \quad (1.15)$$

Discarding terms quadratic in $h_{\mu\nu}$ or its derivatives, as we assumed these to be small, we can then write the Einstein tensor as

$$G_{\mu\nu} = -\frac{1}{2}\eta^{\alpha\beta}\partial_{\alpha}\partial_{\beta}\bar{h}_{\mu\nu}. \quad (1.16)$$

The d'Alembert-operator $\square = \eta^{\mu\nu}\partial_{\mu}\partial_{\nu}$ appears in the Einstein tensor, and inserting into the Einstein field equations, we find the linearised field equations

$$\square\bar{h}_{\mu\nu} = -16\pi T_{\mu\nu}. \quad (1.17)$$

We find a system of inhomogeneous wave equations, or homogeneous wave equations for a vacuum, with $T_{\mu\nu} = 0$. This is solved by plane waves,

$$\bar{h}_{\mu\nu} = \Re \left\{ C_{\mu\nu} e^{i\eta_{\alpha\beta}k^{\alpha}x^{\beta}} \right\}, \quad (1.18)$$

where $k^0 = \omega$ is the angular frequency of the wave, k^a are the spatial components of the direction of propagation, and the complex $C_{\mu\nu}$ represents amplitude and phase. We take the real part \Re on the right-hand side of Equation (1.18) to describe the physical gravitational wave. In the solution, k^{α} is a null vector, $k^{\alpha}k_{\alpha} = 0$, and thus the propagation velocity is the speed of light, $c = 1$. Indeed, any solution of the homogeneous wave equation in the Lorenz gauge can be written as a superposition of plane waves [7, 8].

¹In analogy to the gauge of electromagnetism by L. Lorenz, but often written Lorentz gauge in GR texts. A special case of the more general de Donder or harmonic gauge [10].

The Lorenz gauge is not unique and still allows the gauge freedom of a change in coordinates $x'_\mu = x_\mu + \xi_\mu$, where $\square \xi_\mu = 0$. Through such a transformation, we now use the transverse-traceless gauge such that

$$C^\mu{}_\mu = 0 \quad \text{and} \quad C_{\mu\nu}U^\nu = 0 \quad (1.19)$$

$$\Rightarrow \bar{h}^{TT}_{\mu\nu} = h^{TT}_{\mu\nu}, \quad (1.20)$$

where U^ν is a four-velocity and the condition $C^\mu{}_\mu = 0$ means tracelessness. We can perform a background Lorentz transformation to coordinates such that U^ν is the time basis vector in these coordinates, $U^\nu = \delta_0^\nu$, i.e. the four-velocity of a particle at rest in the reference frame of the coordinates. Using Cartesian coordinates t, x, y, z , we additionally choose the spatial coordinates such that the wave propagates in the z -direction, so $(k^\mu) = (\omega, 0, 0, \omega)$, as $k^\mu k_\mu = 0$ implies $(k^0)^2 = k^a k_a$. This leads to

$$C^{TT}_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & C_+ & C_\times & 0 \\ 0 & C_\times & -C_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (1.21)$$

with only C_+, C_\times remaining as independent free parameters (and the form justifying the term transverse for this gauge, as $C^{TT}_{\mu z} = 0$). These degrees of freedom represent the two independent polarisations of gravitational waves, plus and cross.

We can write the two polarisation components in terms of the unit linear-polarisation tensors [8],

$$e_+^{ij} = e_x^i e_x^j - e_y^i e_y^j \quad (1.22)$$

$$e_\times^{ij} = e_x^i e_y^j + e_y^i e_x^j, \quad (1.23)$$

where e_x^i, e_y^i are coordinate-basis unit vectors for the x - and y -coordinate. A gravitational wave is linearly polarised with $+$ -polarisation for $(h^{ij})^{TT} \propto e_+^{ij}$ and \times -polarisation for $(h^{ij})^{TT} \propto e_\times^{ij}$. An arbitrarily polarised wave can then be expressed as the superposition of two linearly polarised waves

$$(h^{ij})^{TT}(t, z) = h_+(t, z)e_+^{ij} + h_\times(t, z)e_\times^{ij}, \quad (1.24)$$

where t, x, y, z are the coordinates of the coordinate system chosen above, and

$$h_+(t, z) = A_+ \cos(\omega(t - z) - \phi_+) \quad (1.25)$$

$$h_\times(t, z) = A_\times \cos(\omega(t - z) - \phi_\times). \quad (1.26)$$

Here, ω is the angular frequency, and the magnitude and phase of the complex $C_{+/\times} = |C_{+/\times}| \exp[i\varphi_{+/\times}]$ give the amplitudes $A_{+/\times} = |C_{+/\times}|$ and phases $\phi_{+/\times} = \varphi_{+/\times}$.

The effect of a gravitational wave on freely falling particles can be used to measure a passing gravitational wave, as the wave can affect the proper distance between such particles. We consider a particle in its rest frame and use the transverse-traceless gauge corresponding to its four-velocity U^μ . We now let the particle encounter a gravitational wave. The resting, freely falling particle obeys the geodesic equation, with its path parameterised by its proper time τ and determined by its initial four-velocity U^μ ,

$$\frac{d}{d\tau}U^\mu + \Gamma^\mu{}_{\nu\rho}U^\nu U^\rho = 0. \quad (1.27)$$

1 Introduction

The particle is initially at rest, $U^i = 0$, and the initial acceleration is

$$\frac{d}{d\tau}U^\mu = -\Gamma^\mu_{00} = -\frac{1}{2}\eta^{\mu\nu}(2\partial_0 h_{\nu 0} - \partial_\nu h_{00}). \quad (1.28)$$

The initial acceleration vanishes as in the transverse-traceless gauge the $h_{\alpha 0}$ vanish. Thus, the particle will initially remain at rest and by the same argument stay at rest in these coordinates indefinitely.

However, this is not a geometric statement but only refers to the coordinate distance. The proper distance between two particles can nonetheless be influenced by the gravitational wave. We can calculate the proper distance of two nearby particles at fixed coordinates, one at the coordinate origin p_0 , and one shifted from the origin by δx along the x -axis to p_1 . The path connecting the points, parameterised by λ , is $p^\mu \rightarrow (0, \lambda, 0, 0)$, and $v^\mu = \frac{dp^\mu}{d\lambda} \rightarrow (0, 1, 0, 0)$. Integrating along this path, we find the proper length

$$l_x = \int_{p_0}^{p_1} |g_{\mu\nu}v^\mu v^\nu|^{\frac{1}{2}} d\lambda \quad (1.29)$$

$$= \int_{p_0}^{p_1} |(\eta_{\mu\nu} + h_{\mu\nu})v^\mu v^\nu|^{\frac{1}{2}} d\lambda \quad (1.30)$$

$$= \int_{p_0}^{p_1} |(\eta_{xx} + h_{xx})|^{\frac{1}{2}} d\lambda \quad (1.31)$$

$$= \int_0^{\delta x} |(1 + h_{xx}^{TT})|^{\frac{1}{2}} dx \quad (1.32)$$

$$\approx |1 + h_{xx}^{TT}(x=0)|^{\frac{1}{2}} \delta x \quad (1.33)$$

$$\approx \left(1 + \frac{1}{2}h_{xx}^{TT}(x=0)\right) \delta x + \mathcal{O}(h^2). \quad (1.34)$$

As h_{xx}^{TT} is time-dependent, the proper distance between the particles changes over time as the wave passes. Repeating the calculation for particles along the y - and z -axes leads to an expression of the same form. As $h_{zz}^{TT} = 0$, the distance between particles separated in the propagation direction of the wave remains unaffected. From $h_{yy}^{TT} = -h_{xx}^{TT}$ we see for the x - and y -directions that the distance in one direction is stretched when it is contracted in the other direction. Thus distance along one direction is stretched and contracted periodically, and the same occurs in the perpendicular direction with a phase offset compared to the first direction.

The same calculation can be performed for particles placed on the diagonals between the x - and y -axes, where the relevant components of the metric tensor will be $h_{xy}^{TT} = h_{yx}^{TT}$. This leads to the same behaviour of periodical stretching and contraction along the two diagonals, offset by a difference in phase.

The effect of a passing gravitational wave on a ring of freely falling particles illustrates the two polarisations in Fig. 1.1. Both polarisations of the gravitational wave leave the proper distances along the z -direction unaffected. In the plane perpendicular to the propagation of the wave, distances are periodically stretched and contracted along one direction, and alternately along the direction perpendicular to the first. This behaviour is the same for both polarisations, with the two directions of stretching and contraction rotated by $\pi/4$ compared to the respective opposite polarisation.

We note the effect of a $+$ -polarised wave on two particles placed on the diagonals, e.g. $(x_1, y_1) = (a, a)$ and $(x_2, y_2) = (a, -a)$. These are displaced perpendicular to the diagonals, so the displacement has no radial component and their proper distance to the origin does not change initially. For small displacements, the proper distance to these particles will thus remain

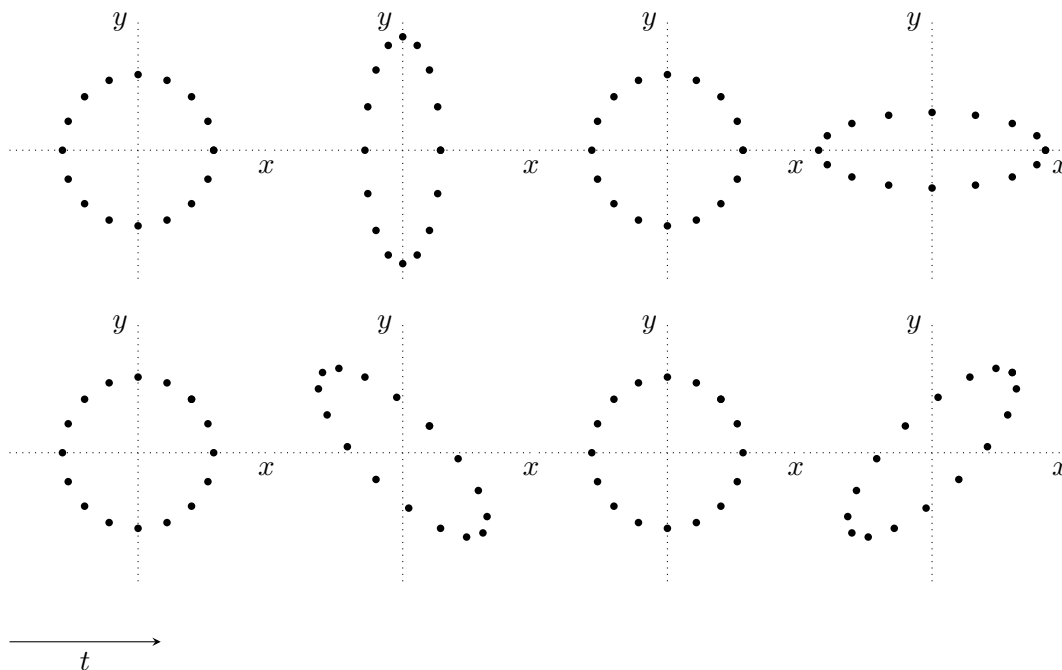


Figure 1.1: Effect of a gravitational wave on a ring of freely falling particles. The two rows of figures show the behaviour of a ring of particles in the x - y plane when encountering a gravitational wave propagating in the z -direction perpendicular to the plane. Depicted are the proper distances between the particles. The first row shows the effect of a wave with $+$ -polarisation, first stretching the circle in the y -direction, then returning to a circle, and then stretching it in x -direction, before the cycle repeats. The shape of these distortions suggests the $+$ -symbol for this polarisation. Similarly, the distortions in the second row for a \times -polarised wave are strongest along the diagonals, suggesting the \times -symbol.

approximately unchanged, and a detector based on particles placed in this manner could not detect a $+$ -polarised wave by measuring the difference in the change of distance of these particles. Conversely, a detector with particles on the x - and y -axes cannot detect a \times -polarised wave. This principle is encoded in the detector antenna-pattern functions as described in Section 1.4.

We notice that the change in proper distance is proportional to the initial separation of the particles, so placing the particles further apart increases the observable effect of the gravitational wave in absolute terms. The alternating stretching and contraction can be used to increase the observable effect of the wave by using three particles, one at the origin, and each one at coordinate-distance $\delta x = \delta y = L$ on the x - and y -axis, respectively. With the changes in proper length $\delta l_x = l_x - \delta x$, $\delta l_y = l_y - \delta y$ we then find

$$\frac{\delta l_x}{\delta x} - \frac{\delta l_y}{\delta y} = \frac{\delta l_x - \delta l_y}{L} = \frac{\Delta L}{L} = h_{xx}^{TT}. \quad (1.35)$$

Using such a set of particles and measuring the differential change in proper distance relative to its initial value allows us to directly measure gravitational waves.

The description for gravitational waves as found for a flat background metric can be used to approximate regions sufficiently distant to strong gravitational fields, i.e. far both from the source of the gravitational wave and from other massive bodies. In this local wave zone, the propagation of the waves is only negligibly influenced by the background curvature, and a multipole

1 Introduction

decomposition of the radiation field shows only quadrupolar or higher multipole contributions [8, 28]. A special case can be used to approximate the emission of a binary of compact objects, which are orbiting each other, but are well separated, as seen from large distances. If the source region containing the objects is small compared to the gravitational wavelength, and the motion of the source components is slow, we can express the gravitational wave emission using the reduced quadrupole moment of the source,

$$I_{ij}(t) = \int \rho(t, \vec{x}) \left(x_i x_j - \frac{1}{3} r^2 \delta_{ij} \right) d^3x. \quad (1.36)$$

Here, $\rho(t, \vec{x})$ is the mass density of the source, r is the distance to the centre of the source system, coinciding with the coordinate origin, and $\vec{x} = (x^1, x^2, x^3)$. In the transverse-traceless gauge, and at distance r from the source, the gravitational wave tensor to first order in the inverse distance then is

$$h_{ij}^{TT}(t) = \frac{2}{r} \ddot{I}_{ij}^{TT}(t - r). \quad (1.37)$$

In this, the second time derivative of the transverse-traceless part of the reduced quadrupole moment \ddot{I}_{ij}^{TT} appears, with the transverse projection perpendicular to the radial propagation direction of the wave [8, 29]. The dominant contribution to the emission from these sources thus comes from the second time derivative of the source mass-density's quadrupole moment, while the retarded time $t - r$ appears due to the travel time from the source to the considered point at distance r . We see that the emission described by Equation (1.37) vanishes for systems with spherically symmetric mass-density (as follows generally from Birkhoff's theorem [30, 31]), as well as for axisymmetric mass distributions with movement about the symmetry axis.

1.3.1 Sources and types of signals

Several astrophysical sources are expected to produce directly detectable gravitational wave signals [32]. Binary systems of very compact objects, such as black holes or neutron stars, emit gravitational waves that have already been detected [33–35]. Supernova explosions and rapidly rotating neutron stars with non-axisymmetric deformations are among further candidates that are being searched for [36–38].

The quadrupole formula shows that rapid changes in the mass distribution and large masses are favourable for the production of high amplitudes in the emitted gravitational waves. Binary systems of compact objects are thus a suitable candidate, as the objects are of high mass and, due to their small size compared to their mass, can be in tight orbits with high orbital velocities and short periods. Figure 1.2 shows a simple overview of the expected signal strengths of different sources and the sensitivities of different planned and existing detectors.

The gravitational wave emission of a compact binary system can be broadly divided into three phases. These phases of the binary evolution and the corresponding signal are illustrated in Figure 1.3, using the first observed gravitational wave signal GW150914 as an example. During the initial inspiral phase, the objects in the binary system are orbiting each other with a comparatively large separation and small orbital velocity, allowing a Keplerian description. Assuming circular orbits for the component objects, the orbital distances and velocities are approximately constant. This leads to a gravitational wave signal which is approximately monochromatic, consisting of a single sine wave at twice the orbital frequency. As the system emits energy through the gravitational waves, the orbital distance shrinks and the orbital frequency increases. Through this, both the frequency and the amplitude of the emitted radiation increase, leading to the characteristic chirp form of the signal [40]. As the distance between the objects decreases and their velocities grow, post-Newtonian corrections become relevant for an accurate description [41,




Figure 1.2: The expected characteristic strain for different types of sources is shown as a function of signal frequency, as well as the sensitivities of different existing and planned detectors. Ground-based detectors target the frequency range of compact binary systems and rotating neutron stars, while space-based detectors cover the range of supermassive compact binaries. The figure was produced using the `GWplotter` online tool, described in [39].

[42]. The gravitational waveforms can be found through such post-Newtonian calculations or the effective one-body formalism [43, 44].

When the two objects have approached sufficiently, the system enters the merger phase, during which the individual objects form a single one. This phase is highly dynamic, with the objects reaching relativistic velocities, and has to be described through the full Einstein equations [42, 45]. As the analytical approximations such as used for the inspiral phase become insufficient, modelling this phase relies on the methods of Numerical Relativity to solve the Einstein equations [46–48]. Gravitational waveforms can be extracted from numerical results, and more efficient phenomenological and interpolated waveform models can be constructed from them [49, 50].

The resulting single object is highly perturbed and settles down to its final state during the ringdown phase, wherein the perturbations decrease through further emission of gravitational waves. If the individual objects were neutron stars of sufficiently low masses, the final object may be a single neutron star [51]. For neutron stars of higher mass, or if at least one of the objects was a black hole, the final state will be a black hole approximated by the Kerr geometry. The late ringdown phase of the final black hole can again be modelled analytically through linear perturbations of the Kerr spacetime. The resulting signal consists of a superposition of exponentially damped sinusoids with frequencies and damping times characteristic of the black hole [52–54]. These are described in more detail in Section 1.6.

Waveforms for the complete evolution of a binary can be constructed by combining results




Figure 1.3: Figure 2 from the publication of the first detection of gravitational waves from a binary black hole system, GW150914, [33]. The upper panel shows a schematic representation of the phases in the evolution of the binary system, with the corresponding gravitational wave strain signal. The lower panel shows the evolution of the orbital separation and velocity of the black holes in the binary.

for the individual phases, and banks of these waveform templates are then used in the matched-filtering search for signals [55]. The distinction between these phases is not unique, but serves as a guideline regarding the relevant physical effects, models, and methods at different stages of the binary evolution.

Individual rotating neutron stars present a second type of source if their mass density is not symmetric about the axis of rotation. Such asymmetries may be due to deformations such as millimetre-high mountains on the surface, supported by the crust or magnetic fields, or due to interior fluid oscillations or precession of the neutron star [56–58]. The resulting continuous wave signal is approximately monochromatic and persists over long times, slowly shifting towards lower frequencies as the neutron star’s rotation slows down due to the energy lost through the gravitational wave (and electromagnetic) emission. For pulsars, the rotation frequency of the neutron star can be well-constrained from electromagnetic observations, allowing for gravitational wave searches targeting a narrow range of parameters and improving the sensitivity of the search [59–61]. All-sky searches for signals from unknown rotating neutron stars are also performed with an accordingly much larger parameter space [37, 38].

Supernova explosions, when sufficiently deviating from spherical symmetry, are another potential source. Searches for unmodelled bursts of power coherent between detectors are used to search for this type of source [62].

These signals can be broadly divided into being of transient or continuous nature, while this distinction depends on the detector under consideration. While binary systems of compact objects emit a continuous signal during the early inspiral, only the frequencies emitted during

the late inspiral, merger, and ringdown are within the most sensitive range of the Earth-based detectors. For these detectors, they are thus transient sources with the observable signal having a short duration. Burst signals are transient due to the similarly transient nature of their sources, while rotating neutron stars can emit continuous radiation at frequencies observable for ground-based detectors.

Detections of gravitational wave signals from compact binaries yield information on the populations of these objects [63, 64]. They can be used to learn about the behaviour of matter under extreme conditions, to test predictions of general relativity, and to gain insight into fundamental physics in the strong-gravity regime. The signal from the collision of neutron stars carries information about the behaviour of matter under extremely high pressures and in strong gravity, and can be accompanied by an electromagnetic counterpart signal, as was the case for the first such detection [34]. Vice versa, it allows for tests of general relativity in the presence of matter [65]. Detections of continuous waves from rotating neutron stars can similarly be used to learn about their structure and composition, and for further tests of GR [66–68]. Signals from black hole collisions provide an opportunity to test fundamental physics and general relativity in the regime of even stronger gravity [69–72].

1.4 Gravitational wave detectors

The most prominent gravitational wave detectors currently operating are the two advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) detectors in the US, advanced Virgo in Italy, and KAGRA in Japan, and data of the first three is analysed in this thesis.

The effect of a passing gravitational wave on freely falling particles can be used to detect it by measuring the proper distance between the particles. As the change of the proper distance is very small, it is measured interferometrically, with the current detectors using a Michelson interferometer design [73–75]. The particles or test masses are represented by the mirrors at the ends of the interferometer’s arms. An unequal change in the proper distance along the arms’ optical paths then changes the interference pattern observed at the interferometer’s output. The quantity measured at the output photodetector is the power at a specific point in the interference pattern, which will change when the pattern shifts between constructive and destructive interference. As the difference in proper distance and thus light travel time along the detector arms changes, so does the phase difference between the two beams and the interference pattern.

The geometry of the Michelson interferometer with its perpendicular arms is beneficial as it maximises the measurable effect for a gravitational wave’s polarisation-component aligned with the arms. We see this from Equation (1.35), where the difference between the changes in proper distances along the arms ΔL appears. Defining the dimensionless strain $h(t)$ introduced in the detector by a gravitational wave of arbitrary polarisation and orientation, we rewrite this equation for ΔL . For a detector with arm lengths $L_x = L_y = L$ we find

$$\begin{aligned} h(t) &= \frac{\delta l_x}{L_x} - \frac{\delta l_y}{L_y} = \frac{1}{2} (e_x^i e_x^j - e_y^i e_y^j) h_{ij}^{TT} \\ \Rightarrow \Delta L &= \delta l_x - \delta l_y = h(t)L, \end{aligned} \tag{1.38}$$

where e_x^i, e_y^i are spatial unit-vectors along the detector arms [56]. The dimensionless strain amplitude $h(t)$ is the quantity we desire to measure. We use here the approximation that the gravitational wave’s wavelength λ is much larger than the arm length, $\lambda \gg L$, considering the metric tensor to be approximately constant during the laser beam’s propagation through the interferometer [7, 76]. We also see that larger arm lengths increase the magnitude of the measurable change in length, so using an interferometer with long arms is beneficial. In addition, the detectors can use a Fabry-Pérot cavity in each of the arms to increase the effective length

1 Introduction

of the optical path [73, 74]. The relative phase between the two arms' beams at the output then depends on $h(t)$, the mean number of round trips in the cavities b , and light frequency ν_0 [77–79],

$$\Delta\phi(t) = 2\pi b\nu_0 Lh(t). \quad (1.39)$$

The change in length introduced in the detector by a gravitational wave depends on the detector's geometry, and the wave's polarisation and direction. We have already seen this partially from the effect of a wave on a ring of freely falling particles. A gravitational wave consisting of plus- and cross-polarised components $h_+(t), h_\times(t)$, causes a strain in the detector described using the detector's antenna pattern functions F_+, F_\times ,

$$h(t) = F_+(\theta, \phi, \Psi)h_+(t) + F_\times(\theta, \phi, \Psi)h_\times(t). \quad (1.40)$$

Here, $\theta \in [0, \pi), \phi \in [0, 2\pi)$ are the spherical polar coordinates of the source in the frame of the detector. They correspond to the Cartesian coordinate system in which the x - and y -axes are oriented along the arms of the detector and the z -direction points upwards. In addition to this detector frame, we use the geocentric, radiation, and source frames explained in the following.

The polarisation angle Ψ relates the detector and radiation frame. It is the angle between the plane of constant $\phi = \phi_s$ and the Cartesian x -direction of the radiation frame [80, 81], where ϕ_s is the ϕ -component of the line-of-sight vector towards the source. The radiation frame is that described in Section 1.3, with the z -axis pointing along the direction of propagation, in this case from the source towards the Earth. It is given in the transverse-traceless gauge, with a freely falling particle as origin, and the x - and y -axes pointing along the directions of stretching and contraction for the $+$ -polarisation, as shown in Section 1.3. The antenna pattern functions for an L-shaped interferometric detector with a right angle between the arms then are [80, 81]

$$F_+(\theta, \phi, \Psi) = \frac{1}{2} (1 + \cos^2 \theta) \cos 2\phi \cos 2\Psi - \cos \theta \sin 2\phi \sin 2\Psi \quad (1.41)$$

$$F_\times(\theta, \phi, \Psi) = \frac{1}{2} (1 + \cos^2 \theta) \cos 2\phi \sin 2\Psi + \cos \theta \sin 2\phi \sin 2\Psi. \quad (1.42)$$

Due to the rotation of the Earth, the angles θ, ϕ, Ψ as given in the detector frame are time-dependent. For transient signals from compact binary systems, however, this time-dependency will typically be negligible for the short duration during which the signal is observable.

The geocentric frame or equatorial coordinate system is independent of the considered detector [82, 83], and centred on the Earth. The sky-location of an observed source is given in this frame by the right ascension $\alpha \in [0, 2\pi)$ and declination $\delta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, and the polarisation angle Ψ can then be constructed similarly as in the detector frame. While centred on the Earth, these coordinates are not co-rotating with it. Instead, they project the Earth's equator onto the visible sky and use its intersection with the ecliptic of the solar system as a reference point to define spherical coordinates on the sky.

Finally, the source frame is fixed to the source system and coordinates adapted to conveniently describe its properties may be chosen, often aligning one of the basis vectors with the angular momentum vector of the system and having the origin coincide with the centre of mass. The inclination angle ι relates the source frame to the radiation and detector frames. It is the angle between the vector pointing from the source towards the Earth and the angular momentum vector of the source system, meaning the orbital angular momentum for a binary and the total angular momentum for an individual black hole [84]. The gravitational wave signal seen by an observer depends on this angle due to the angular dependency of the multipolar expansion, as described for the quasi-normal mode signal from a perturbed black hole in Section 1.6. Together,




Figure 1.4: Antenna patterns for a detector with arms pointing into the x - and y -directions, for waves with polarisation angle $\Psi = 0$. For each direction, the absolute values of the antenna pattern functions F_+ , F_\times and their combination $F_+^2 + F_\times^2$ are shown as the distance of the surface from the origin. Additionally, the same quantity is mapped to the colour of the surface, with lighter colours showing larger values.

ι , Ψ , α , and δ relate the source and radiation frames to the geocentric frame, and in turn to the detector frames via the known positions and orientations of the detectors on the Earth.

For distant sources, the signal may be significantly redshifted due to the expansion of the universe, which has to be taken into account when comparing quantities in the source and detector frames. For example, the mass measured in the detector frame differs from the source frame mass by a factor $(1 + z)$, where z is the redshift corresponding to the distance of the source [85].

As mentioned earlier, the LIGO and Virgo detectors are based on a Michelson interferometer installed in a system of vacuum tubes, with additional Fabry-Pérot cavities along the arms to increase the effective path length of the arms. These also increase the laser power circulating in the arms, as do signal- and power-recycling mirrors placed at the photodetector and laser source. These mirrors reflect the light coming toward the photodetector and that returning to the laser source back into the interferometer, thus forming a cavity with the interferometer itself. This further increases the laser power and thus sensitivity, and allows to manipulate the most sensitive frequency band of the detector. The test-mass mirrors are mounted on a series of pendulum suspensions to isolate them against external motion transmitted through the ground. Active control elements such as actuators allow to manipulate the motion of the mirrors and further compensate unwanted influences. Many components of the detector design are thus intended to prevent or remove undesired effects.

The desired observable is the dimensionless strain amplitude $h(t)$, but many additional effects influence the motion of the test masses and the observable output signal, limiting the sensitivity of the detectors as noise. Most noise sources are well understood and can be modelled, and many additional sensors are used to monitor these in auxiliary data channels. In Figure 1.5, the noise resulting from individual sources is shown for Advanced LIGO both for the predictions from the design study and as measured during the second observing run.

The following are examples of noise sources, explained in [73, 78]. Quantum noise consists of radiation pressure noise and shot noise. Radiation pressure noise is caused by the transfer of momentum from the photons of the light beams to the mirrors and setting them into motion. As the number of photons interacting with the mirror fluctuates, so does the transferred momentum, introducing uncertainty in the mirrors' positions. Shot noise results from variations




Figure 1.5: Each panel shows the strain equivalent of individual detector noise sources as a function of frequency. The left panel lists predictions from the Advanced LIGO design study, Figure 2 of [73]. In the right panel, the noise contributions measured during Advanced LIGO’s second observing run are shown for the two detectors H1 and L1, Figure 3 of [86].

in the arrival rate of photons at the photodetector, introducing fluctuations in the interference pattern being measured. Thermal effects introduce noise through several mechanisms, such as the Brownian motion of the atoms in the coating of the test-mass mirrors. Fluctuations in the local temperature of the suspensions result in random motion of the test masses. In addition, the optical properties of the test masses and their coatings are temperature dependent, leading to small fluctuations also in these. Seismic noise is due to the motion of the ground at the detector site being transferred to the mirrors. This is highly damped through the suspension system at higher frequencies but becomes dominant at low frequencies below about 11 Hz. Gravity gradient noise is due to density fluctuations in the Earth near the detectors, caused by seismic waves, which lead to varying gravitational forces on the test masses. The vacuum in the beam tubes is not perfect, and the residual gas causes noise through the interaction of the beam with the remaining molecules, changing the refractive index of the beam path. Technical noise encompasses various additional noise sources, such as frequency and amplitude fluctuations of the laser sources, and noise from actuators in the control loops used to manipulate the mirror position and orientation. A number of peaks, narrow in frequency, are present in the noise curve, called lines. These have various sources, such as resonances of the mirror suspensions, the AC power grid frequency and its harmonics, and artificial lines at known frequencies generated for calibration [87–89].

A further ground-based detector currently operating is the technology testbed GEO600 [75, 90, 91], while a site for the construction of the LIGO India detector has been selected [92]. Third-generation detectors are planned or proposed, such as the Einstein Telescope or Cosmic Explorer, utilising cryogenically cooled components, underground placement and even longer interferometer arms [93–95]. Space-based detectors such as LISA are currently in development, which offer the advantage of allowing far larger arm lengths and the removal of seismic noise sources [96]. At the same time, the space-based nature introduces constraints such as imprac-

tical maintenance and weight limits, while adding other noise sources such as attitude control thrusters.

In addition to different noise sources, the arm length of the different detector types determines in which frequency band they are most sensitive. For long arm lengths, the approximation of the test masses being initially separated by a small distance does not hold. The sensitivity curve for space-based detectors is then calculated depending on the lengths of the arms [97], with an estimate of the sensitivity for LISA and expected sources shown again in Figure 1.2. We also note that if the detector's arm length is sufficiently larger than the wavelength of the gravitational wave, the wave could pass it undetected as the effects of stretching and contraction along the light path may yield no net length-change for the beam's round trip.

1.5 Data analysis, Bayesian inference, and parameter estimation

We are interested in analysing the data collected by the gravitational wave detectors to quantify the support it lends to various models. In particular, we often want to decide whether a specific signal is present in the data and, if so, what are the values of the parameters used to describe the signal. The methods we use are those of Bayesian inference, the basics of which are reviewed here.

We write the probability of a statement A being true given that statement B is true as $P(A|B)$. For the probability of both A and B being simultaneously true, given that statement C is true, we can write the product rule [98],

$$P(A, B|C) = P(A|B, C)P(B|C) = P(B|A, C)P(A|C). \quad (1.43)$$

Writing now \bar{A} for A being false, we have the sum rule, stating that the sum of the probabilities of all possibilities should be 1,

$$P(A|B) + P(\bar{A}|B) = 1. \quad (1.44)$$

We can rearrange the product rule to find what is known as Bayes' theorem,

$$P(H|D, X) = P(H|X) \frac{P(D|H, X)}{P(D|X)}, \quad (1.45)$$

where we have replaced A, B, C with H, D, X , conventional for the following meaning. We consider a hypothesis H to be tested, given some prior knowledge or assumption X and the new data D to be used to update our belief about the hypothesis. We call $P(H|D, X)$ the posterior probability and $P(H|X)$ the prior probability. For fixed data D , we call $P(D|H, X)$ the likelihood and $P(D|X)$ the evidence. Here we will primarily be interested in two applications of this, parameter estimation and model selection.

In parameter estimation, we want to use given data to assign probabilities to the values of a considered model's parameters. Our hypothesis now is that a parameter θ of the model H has a certain value. Here, H means our model assumption, encompassing the description of the data, both signal and noise, and incorporating our prior knowledge. The model's parameters are often continuous, with real numbers as values, and the probability for them taking a specific value would thus vanish. So instead, we consider probability densities p , where for the parameter θ of unknown value and the model H we have

$$\int p(\theta|H) d\theta = 1. \quad (1.46)$$

1 Introduction

For each interval in parameter space, p assigns a probability for the unknown true value of the parameter θ lying in that interval. Bayes' theorem then reads [99]

$$p(\theta|d, H) = p(\theta|H) \frac{p(d|\theta, H)}{p(d|H)}, \quad (1.47)$$

and in this context, the evidence $p(d|H)$ can be treated as a normalisation constant. It depends only on the data and model, both of which we consider fixed when estimating the parameters. Our parameter estimation will be based on Equation (1.47), where we calculate the probability for some value of the parameter based on an expected prior probability distribution $p(\theta|H)$ and the likelihood $p(d|\theta, H)$ of the data to be found assuming a model and this specific value of the parameter.

Typically, our model will depend on several parameters θ_i , and we want to find the joint probability density for all parameters $\boldsymbol{\theta} = \{\theta_i\}$. If the joint probability density is known, an individual parameter's probability density is found by marginalising over the remaining parameters, e.g. for θ_1 ,

$$p(\theta_1|d, H) = \int p(\boldsymbol{\theta}|d, H) d\theta_2 \dots d\theta_n. \quad (1.48)$$

As the probability density function does not assign a probability to individual points in parameter space, we instead quote properties of the distribution to characterise results. The $x\%$ credible interval is simply an interval in the one-dimensional space of a specific parameter which yields a probability of $x\%$ when integrating the probability density. This may be interpreted as the interval containing the true value of the parameter with $x\%$ probability. Commonly quoted are 90% credible intervals, such that integrating the probability density over the interval will yield 90% probability. This still allows freedom in the choice of the interval. Typically, we will choose the interval such that the probabilities of values above and below the quoted boundaries are equal. We will additionally quote the median value of the distribution, which is the value such that the probabilities of the intervals above and below it are equal and thus each 50%. We then write the result as A_{-B}^{+C} , where A is the median value and $A - B$ and $A + C$ are the lower and upper bounds of the 90% credible interval [99]. When we expect a parameter to have a value at the edge of the parameter space, for example when constraining the amplitude of a signal that may not be present in the data, we can also use a one-sided credible interval to quote constraints on the possible values. In this case, the credible interval starts either at the lowest or the highest possible values, such that the remainder of the distribution will be entirely above or below its range, respectively. When considering a two-dimensional joint distribution for two parameters, we can similarly show the contour bounding a credible region, where again integration over the enclosed region yields a probability of $x\%$. For this, we will choose the smallest region giving the desired probability, thus choosing the region with the highest density of the distribution.

1.5.1 Data model

To evaluate the probability for a parameter of a given model via Bayes' theorem, we still have to describe the likelihood $\mathcal{L} = p(d|\theta, H)$, i.e. the probability for the observed data d to be found given the model and its parameters. We have seen in Section 1.4 that the detectors measure the unitless strain amplitude introduced by gravitational waves, and that various other effects lead to the detector output responding in the same way as if there was a change in the strain. We model the time dependent data s as consisting of contributions from the gravitational wave signal, h , and from the noise sources, n , added linearly,

$$s(t) = n(t) + h(t). \quad (1.49)$$

The detector output is sampled at times t_i , such that the strain-amplitude data produced by the gravitational wave detectors is given as a series of discrete, real-valued samples s_i and corresponding times t_i .

The noise is modelled as a Gaussian stochastic process that is wide-sense stationary, with zero mean and a known variance, which is estimated from the measured power spectrum of the data [99, 100]. By stochastic process we mean a collection of random variables, in our case the detector output due to noise at the times t_i is a random variable. The expectation value or mean of a random variable X taking possible values x with probability density function $f(x)$ is $\mu(X) = \int x f(x) dx$. For a collection of random variables $\{X_{t_i}\}$, the autocovariance $K(t_1, t_2)$ and autocorrelation $C(t_1, t_2)$ of the random variables for two different t_i are

$$K(t_1, t_2) = \mu((X_{t_1} - \mu(X_{t_1}))(X_{t_2} - \mu(X_{t_2}))^*), \quad (1.50)$$

$$C(t_1, t_2) = \mu((X_{t_1})(X_{t_2})^*), \quad (1.51)$$

where the star $*$ means complex conjugation [101].

Wide-sense stationarity means that the mean μ of the noise process is constant over time, and that the autocovariance K of two samples depends only on their separation in time, thus the latter is true also for their autocorrelation C [101, 102],

$$\mu(t) = \mu(t + \tau), \quad \forall \tau \quad (1.52)$$

$$K(t_1, t_2) = K(t_2 - t_1, 0), \quad \forall t_1, t_2 \quad (1.53)$$

$$C(t_1, t_2) = C(t_2 - t_1, 0), \quad \forall t_1, t_2. \quad (1.54)$$

We can then write these as only depending on the time interval $\tau = t_2 - t_1$, $K(\tau)$ and $C(\tau)$.

We use the Wiener-Khinchin theorem to relate the noise power spectral density $S_n(f)$ to the autocorrelation function $C(\tau)$ of the noise n through

$$S_n(f) = \int_{-\infty}^{\infty} C(\tau) e^{-i2\pi f\tau} d\tau \quad (1.55)$$

$$= \lim_{T \rightarrow \infty} \frac{1}{T} \mu(|\tilde{n}(f)|^2). \quad (1.56)$$

We may consider either the first or the second line as the definition of the power spectral density (PSD). In our applications, we consider discrete, finite sets of data samples and use the discrete Fourier transform, while noting subtleties for the continuous case [103, 104]. Here, the data n is understood to vanish outside a finite time interval of duration T , and \tilde{n} is the Fourier transform of n . The second line has the intuitive interpretation of the PSD as the average power of the signal in each frequency interval. By power here we mean the squared absolute value of the data, which need not correspond to physical power, as we see in our case of a unitless strain signal.

Given an estimate of the noise PSD, we can now construct the noise as a time-ordered set of random variables. Each has a Gaussian distribution for the probability density function, with mean $\mu = 0$, and a multivariate Gaussian distribution describes the entirety of the noise over time. The covariance matrix for the multivariate Gaussian distribution is given by the autocorrelation function, which we can find from the estimated noise PSD using Equation (1.55). We call the series of samples produced by one draw from this joint distribution a realisation of the noise, approximating the output samples expected from the detector in a given time interval.

Following [100], one finds the likelihood for a specific noise realisation n in the continuum limit for large numbers of samples with small separation in time to be

$$p(n) \propto \exp\left[-\frac{1}{2} \langle n|n \rangle\right], \quad (1.57)$$

1 Introduction

with the noise-weighted inner product $\langle \cdot | \cdot \rangle$,

$$\langle x | y \rangle = 4\Re \left(\int_0^\infty \frac{\tilde{x}(f)\tilde{y}^*(f)}{S_n(f)} df \right). \quad (1.58)$$

This inner product represents the matched filter for a known signal or template y in data x , which is the optimal filter for such a known signal in stationary Gaussian noise [105, 106]. The corresponding signal-to-noise ratio (SNR) ρ then is defined by

$$\rho^2(t) = \frac{|\langle s | h \rangle(t)|^2}{\langle h | h \rangle}, \quad (1.59)$$

where s is the data and h the template.

In practice, when calculating the inner product, we account for the power spectrum by applying a whitening filter to each argument, e.g. to a template and the data. Each is divided by the amplitude spectral density, the square-root of the power spectral density, in the frequency domain [107, 108]. Such whitened data can be used to make features of a present signal visible by suppressing dominant noise components, such as at low frequencies or from lines [109]. The finite sampling-rate of the data limits the integration to frequencies below the Nyquist-frequency, which is half of the sampling rate [86]. Above this frequency, aliasing occurs, and all our analyses based on the discrete Fourier-transform will be restricted to frequencies below the Nyquist frequency.

The actual detector noise is known not to be exactly Gaussian but can still be approximated by this model [107, 110]. Non-Gaussian noise features can be of transient or persistent nature, with some being of unknown origin. Transient features can often be treated individually for parameter estimation, for example by removing the affected data from the analysis if the feature is of sufficiently short duration. They can be problematic when estimating the significance of detections if they mimic signals sufficiently well to yield high values in the chosen detection-statistic. Long duration features that are limited to a narrow band in frequency, known as lines in the PSD, can similarly be treated individually by removing the affected frequencies from the data.

As we will use discretely sampled data, we review the construction of the noise likelihood for this case in more detail, following the derivation in [2]. The detector output is sampled uniformly in time, with a timestep Δt . Starting to observe at time t_0 , this gives an ordered set of samples s_i at times $t_i = t_0 + i\Delta t$ which we also write as a vector \mathbf{s} . For a total sampling time T we collect $N = \lceil T/\Delta t \rceil$ samples s_0, \dots, s_{N-1} at times t_0, \dots, t_{N-1} . Here, $\lceil a \rceil$ means the smallest integer greater or equal to a , while $\lfloor a \rfloor$ is the largest integer smaller or equal to a . We model the noise \mathbf{n} through a multivariate Gaussian distribution with covariance matrix \mathbf{C} . If the data consist only of the noise contributions, the likelihood $p(\mathbf{n}) = p(\mathbf{s}|\mathbf{n})$ for a specific noise realisation \mathbf{n} is that of a random draw from the distribution,

$$p(\mathbf{n}) = \frac{\exp \left[-\frac{1}{2} \mathbf{n}^\top \mathbf{C}^{-1} \mathbf{n} \right]}{\sqrt{(2\pi)^N \det \mathbf{C}}}. \quad (1.60)$$

We will see that for this application, knowledge of the numerator will suffice, so we concentrate on this for now.

Assuming that the noise has zero mean and is wide-sense stationary, the elements of the covariance matrix \mathbf{C} are given by the autocorrelation function. The covariance matrix then is a symmetric Toeplitz matrix [111]. A matrix is Toeplitz when its elements $a_{i,j}$ have the property $a_{i,j} = a_{i+1,j+1}$, so the $a_{i,j}$ depend only on the difference $i-j$, and it is symmetric when $a_{i,j} = a_{j,i}$. For large matrices, the inverse of a Toeplitz matrix can be approximated by the inverse of a

circulant matrix. The eigenvectors of circulant matrices are known to be $e^{-2\pi i k p / N} / \sqrt{N}$, with $k = 0, \dots, N-1$ enumerating the eigenvectors and $p = 0, \dots, N-1$ enumerating the components of the vector. Solving for the eigenvalues yields a relation for the components of the inverse of the covariance matrix, \mathbf{C}^{-1} , so that the j, k -th element of the matrix is

$$C^{-1}[j, k] \approx 2\Delta t^2 \mathcal{F}^{-1}(S_n^{-1})[k - j]. \quad (1.61)$$

So the j, k -th element is the discrete inverse Fourier transform of $1/S_n$ at time step $k - j$, with S_n being the PSD of the detector noise. We insert this into the equation for the noise likelihood to find

$$p(\mathbf{n}) \propto \exp \left[-\frac{1}{2} \langle \mathbf{n} | \mathbf{n} \rangle \right], \quad (1.62)$$

again with the inner product $\langle \cdot | \cdot \rangle$ defined as

$$\langle \mathbf{u} | \mathbf{v} \rangle \equiv 4\Re \left\{ \frac{1}{T} \sum_{p=1}^{\lfloor (N-1)/2 \rfloor} \frac{\tilde{u}^*[p] \tilde{v}[p]}{S_n[p]} \right\}, \quad (1.63)$$

where $\tilde{u}[p]$ is the p -th element of the discrete Fourier transform of \mathbf{u}^2 .

From the noise likelihood we also find the likelihood for the case that the data consist of noise and a signal h with parameters $\boldsymbol{\theta}$, $p(\mathbf{s} | \boldsymbol{\theta}, h)$. The residual data after subtracting the signal will consist only of the noise, $\mathbf{s} - \mathbf{h} = \mathbf{n}$. The probability for the data under the signal hypothesis thus is the same as the probability for the residual under the noise hypothesis, and is described by the likelihood found above,

$$p(\mathbf{s} | \boldsymbol{\theta}, h) = p(\mathbf{s} - \mathbf{h} | \mathbf{n}) \propto \exp \left[-\frac{1}{2} \langle \mathbf{s} - \mathbf{h} | \mathbf{s} - \mathbf{h} \rangle \right]. \quad (1.64)$$

We can think of this as the signal shifting the expected value for each sample from zero to the signal value. Thus, if a signal is present in the data, the residual after subtraction of the signal will just be Gaussian noise, and will have the same probability to occur as that realisation of the noise with no signal present.

We notice that the denominator containing the determinant of the covariance matrix will be the same in the cases of the data being only noise or noise with a signal. So when calculating the likelihood ratio Λ comparing the two cases, the denominators cancel, leading to

$$\Lambda = \frac{p(\mathbf{s} | \boldsymbol{\theta}, h)}{p(\mathbf{s} | \mathbf{n})} \quad (1.65)$$

$$= \exp \left[-\frac{1}{2} \langle \mathbf{s} - \mathbf{h} | \mathbf{s} - \mathbf{h} \rangle + \frac{1}{2} \langle \mathbf{s} | \mathbf{s} \rangle \right] \quad (1.66)$$

$$= \exp \left[\langle \mathbf{s} | \mathbf{h} \rangle - \frac{1}{2} \langle \mathbf{h} | \mathbf{h} \rangle \right]. \quad (1.67)$$

The likelihood ratio can be used to determine whether the noise or signal hypothesis is more likely to describe the data, and it can be maximised over different signal parameters. We often use the natural logarithm of these quantities, such as the log-likelihood and log-likelihood ratio.

We now consider a network of K detectors d , each producing a series of samples \mathbf{s}_d , writing the collection of all noise samples \mathbf{n}_{net} . Assuming the noise in the different detectors to be

²This definition excludes the DC-component $p = 0$, and for even N also the Nyquist-frequency component $p = N/2$. These can be treated separately or may be negligible [106].

1 Introduction

uncorrelated, the probability density function then consists of a product of K N -dimensional multivariate Gaussian distributions,

$$p(\mathbf{n}_{\text{net}}) = \frac{\exp\left[-\frac{1}{2} \sum_{d=1}^K \mathbf{n}_d^\top \mathbf{C}_d^{-1} \mathbf{n}_d\right]}{\sqrt{(2\pi)^{NK} \prod_{d=1}^K \det \mathbf{C}_d}}. \quad (1.68)$$

Written for a network of detectors, the likelihood is

$$\mathcal{L}(\mathbf{s}_{\text{net}}|\boldsymbol{\theta}) = p(\mathbf{s}|\boldsymbol{\theta}, h) \propto \exp\left[-\frac{1}{2} \sum_{d=1}^K \langle \mathbf{s}_d - \mathbf{h}(\boldsymbol{\theta}) | \mathbf{s}_d - \mathbf{h}_d(\boldsymbol{\theta}) \rangle\right]. \quad (1.69)$$

Here, \mathbf{s}_d is the data from detector d and \mathbf{h}_d is the signal as it is observed in this detector, taking into account the individual detector's orientation, antenna pattern and relative time delay of the signal due to the detector location.

We can now use the constructed likelihood to calculate the posterior probability for a given signal model and set of parameters. Again we only need to consider the expression proportional to the likelihood, as given in Equation (1.69). Multiplying this with the prior probability for the chosen parameters yields a value proportional to the posterior probability. As both the denominator in the likelihood and the evidence term in Bayes' theorem are constants independent of the choice of parameters, we merely need to normalise our results to find the posterior probability distribution. In practice, we will simply calculate the product of the prior and the expression proportional to the likelihood, which we can now find for a given signal model, PSD estimate, and prior distribution.

We estimate the PSD from the off-source detector data, at times where we expect no signal to be present. Welch's method is typically used in our analyses to produce this estimate [112]. It is based on calculating a periodogram for a given segment of data, i.e. applying the discrete Fourier transform and calculating the squared absolute value for the resulting frequency series. This is an approximation of the power spectral density for the segment as defined in Equation (1.56), where we consider the given realisation of data and thus the mean is just the identity map. This power spectrum is calculated for multiple segments of data and then averaged. The data in each segment is usually windowed before calculating the power spectrum, down-weighting the contributions from the edges of the segment. Therefore, Welch's method uses overlapping data segments to also account for the influence of the data being down-weighted in one segment by including the same data in the unwindowed region of another.

1.5.2 Sampling parameter space

While we can calculate the posterior probability for individual sets of parameters, we now want to find the posterior distribution over the entire parameter space. This is achieved efficiently through stochastic samplers, often using Markov-Chain Monte Carlo (MCMC) methods. In these methods, a set of samples $\boldsymbol{\theta}_n$ from the parameter space is constructed which approximates a set of random draws from the posterior distribution [99]. A Markov chain is an ordered set of samples such that each sample only depends on its immediate predecessor. Each subsequent element of the chain is constructed from the previous element through a two-step process of proposal and acceptance of a new point. Starting from the current element of the chain, $\boldsymbol{\theta}_i$, a new point in parameter space $\boldsymbol{\theta}_{i'}$ is proposed by random draw from a proposal distribution Q . The probability for drawing $\boldsymbol{\theta}_{i'}$ depends only on the current and newly proposed point, $Q(\boldsymbol{\theta}_i, \boldsymbol{\theta}_{i'})$. Based on a suitably constructed acceptance probability, the proposed point is either accepted, $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_{i'}$, or rejected and the new element of the chain remains the same as before,

$\theta_{l+1} = \theta_l$. After sufficiently many iterations of this procedure, the last sample in the chain will approximate a random draw from the posterior distribution. Using an ensemble of chains with different starting points can then produce the desired number of independent samples.

A simple example for the construction of an acceptance probability is found in the Metropolis-Hastings algorithm [113]. For each considered point in parameter space, the product of the likelihood and the prior is calculated. A new candidate point is proposed by one of various schemes used in available sampling tools. If the ratio of the posterior values for the proposed and current points is larger than 1, the new point is accepted. Otherwise, the new point is accepted or rejected with a probability proportional to this ratio. As we only use the ratio of the posterior values, the normalisation factors again cancel and calculating the product of the likelihood and prior suffices.

Care must be taken to ensure that the resulting samples are independent of the starting points of the chains, as samples in the chain separated by only a few steps will still show correlation [99]. So sufficiently many iterations of each chain have to be performed to both remove influences from the initial values and to reach a state where the samples represent the posterior distribution. This also allows using multiple samples from the same chain for the final distribution if they are separated by sufficiently many iterations to remove correlations, effectively producing multiple shorter chains from one longer chain.

The samplers `emcee_pt`, `cpnest` and `dynesty` have primarily been used in this work [114–116]. Parallel-tempering is used in `emcee_pt`, meaning that multiple chains are evolved in parallel at different “temperatures” [117]. For a temperature T , the likelihood for the sampling is raised to the power $1/T$, such that for large temperatures, the posterior being sampled approaches the prior. This allows the sampler to cross low-likelihood regions of the posterior, which might otherwise lead it to miss relevant parts of the distribution separated by such regions from the starting point. The chains at different temperatures periodically exchange elements, allowing each to sample both high- and low-likelihood regions.

Both `cpnest` and `dynesty` use the different approach of nested sampling [118]. While this is intended primarily to calculate the evidence, it can also be used to produce samples from the posterior distribution. Here, samples (live points) are initially drawn randomly from the prior distribution, and the likelihood is calculated for each. In each iteration of the algorithm, the point with the lowest likelihood is recorded and replaced with a new live point drawn from the prior, under the condition that the new point must have a larger likelihood. In addition, a weight is calculated for each replaced point, given by the mass $\int p(\theta|H)d\theta$ of the prior region with likelihood values between those of the point just recorded and that recorded in the previous iteration. The recorded points will represent random draws from the posterior distribution when properly weighted using the calculated weights. Knowledge of these weights then also allows to construct additional samples if desired. Dynamic nested sampling is used in `dynesty`, where the number of live points is dynamically adjusted to better cover the regions of parameter space that will have the greatest influence on the accuracy of the results. As the nested sampling algorithms also provide an estimate of the evidence, they are useful in model selection as described below.

1.5.3 Model selection

The second application of Bayes’ theorem that we are interested in is model selection. In this, we consider different models to explain the data and want to quantify the preference of these models with respect to each other as implied by the data. Here, the evidence for a model H given data D , $p(D|H)$, is relevant, which we have treated as a normalisation constant for parameter estimation. The evidence is the probability of finding the data D given the hypothesis H , marginalised over all parameters θ of the signal model. We can then compare two models

1 Introduction

H_1, H_2 by calculating the Bayes factor as the ratio of their evidences [99],

$$\mathcal{B} = \frac{p(D|H_1)}{p(D|H_2)}. \quad (1.70)$$

Model H_1 is favoured as compatible with the data if $\mathcal{B} > 1$, and vice versa model H_2 is favoured for $\mathcal{B} < 1$. Often, the logarithm of the Bayes factor is quoted instead, where $\log_{10} \mathcal{B} > 0$ means model H_1 being favoured. Suggested interpretations for different values of the log Bayes factor are given in Reference [119], and are shown in Table 1.1.

$\log_{10} \mathcal{B}$	\mathcal{B}	Evidence
[0, 0.5]	[1, 3.2]	Not worth more than a bare mention
[0.5, 1]	[3.2, 10]	Substantial
[1, 2]	[10, 100]	Strong
> 2	> 100	Decisive

Table 1.1: The interpretation of the evidence in favour of one model compared to another represented by different values of the Bayes factor as suggested in Reference [119].

A special case of this model selection procedure is comparing the model of a signal being present in the data versus the data consisting of only the noise. However, determining the preference for a signal being present relies on the accurate description of the noise. While the description of the detector noise as Gaussian is a useful and often sufficient approximation, non-Gaussian features are known to be present. As their effect on the likelihood estimation and interaction with the signal model can be difficult to predict, we can make use of an alternative approach.

A frequentist approach to hypothesis testing is using the p-value under the null-hypothesis that the data are entirely due to noise [98]. Calculating the p-value is based on using a given test statistic, a mapping of the data to the real numbers. We can then find the probability distribution for the values of this mapping based on the probability distribution of the noise. The p-value of a given sample value of the test statistic is the probability of finding a value at least as extreme, i.e. the integral over the probability distribution of the test statistic over the range of more extreme values. The meaning of a value being extreme depends on the considered problem, for example being larger than a threshold value or deviating further from a distinguished value. For the question of whether a signal is present, the estimated signal amplitude or the matched-filter signal-to-noise ratio might be used, where larger values would be considered more extreme. Note that the p-value thus just makes a statement about the probability of finding a result at least as extreme as a given sample if the null hypothesis was true. This is not a statement about the probability of either this or an alternative hypothesis, which might not even be explicitly stated, being correct. Instead, it serves as a measure of whether a result should be considered extraordinary under the null hypothesis.

We notice, however, that this procedure again relies on a description of the noise, or at least knowledge of the test statistic's probability distribution under the null hypothesis. Nevertheless, if many samples of this distribution are available, in our case through large amounts of detector noise free of signals, a simple estimate of the distribution can instead be found empirically. We simply count the number of samples falling in a given range of values and divide by the total number of samples to calculate an approximate probability of finding the value of a random sample to lie in this range.

1.6 Quasi-normal modes

After the non-linear merger phase, the final black hole enters a phase where it can be described by the Kerr solution with additional linear perturbations. For gravitational perturbations, the metric tensor can then be written as $g_{\mu\nu} = g_{\mu\nu}^b + h_{\mu\nu}$, where the background metric $g_{\mu\nu}^b$ is the Kerr metric and $h_{\mu\nu}$ is the perturbation [54]. A master perturbation equation was found by Teukolsky, describing different types of perturbations ψ and in particular gravitational ones, where ψ contains the Weyl scalar Ψ_0 or Ψ_4 [53, 120]. The equation allows for a separation of variables, which reads in Boyer-Lindquist coordinates, (t, r, θ, ϕ) :

$$\psi = e^{-i\omega t} e^{im\phi} S(\theta) R(r). \quad (1.71)$$

This yields linear ordinary differential equations and an eigenvalue problem for $S(\theta)$ and $R(r)$. The eigenvalue problem for $S(\theta)$ has the angular separation constant $A(a\omega)$ as the eigenvalue. Assuming regular boundary conditions at $\theta = 0$ and $\theta = \pi$, the solutions for the eigenfunction are the spin-weighted spheroidal harmonics ${}_sS_{lm}(a\omega, \cos\theta)$, which depend on the complex angular frequency ω . The equation for $R(r)$ can be written in terms of an effective potential, using a change of variables and the tortoise coordinate r_* , $dr_*/dr = (r^2 + a^2)/\Delta$ (confer Equation (1.11)) [121, 122].

For gravitational perturbations, $|s| = 2$. The differential equations for $S(\theta)$ and $R(r)$ then pose an eigenvalue problem with ω as the eigenvalue. Given boundary conditions for $R(r)$, we can find the allowed perturbations using the solutions for ω , A , $S(\theta)$, and $R(r)$. The boundary conditions are physically motivated and given for the event horizon, and for spatial infinity, where the Kerr metric is asymptotically flat. At the horizon, the gravitational wave is purely ingoing as seen by a physical observer, as no radiation should leave the horizon. At spatial infinity, the wave is purely outgoing, so that no gravitational waves are entering the system from infinity. The solution for $s = -2$ is found to be associated with the outgoing radiation, so we restrict to this case as only the outgoing waves can be detected, while the infalling radiation is absorbed by the black hole.

The solutions have the form of exponentially damped sinusoids and are called quasi-normal modes (QNMs). In contrast to normal modes of a non-dissipative system, these are generally neither orthogonal nor do they form a complete set [54, 124]. At large distances, $r \rightarrow \infty$, the gravitational wave strain is then found from the Weyl scalar Ψ_4 via [125]

$$\Psi_4 = \ddot{h}_+ - i\ddot{h}_\times, \quad (1.72)$$

where the dots mean the second time derivative.

The ringdown signal of a perturbed black hole in the linear regime is a superposition of these quasi-normal modes. While there are subtleties involved regarding at what time after the merger this description becomes accurate [126–129], at sufficiently late times $t \geq t_0$, the signal can then be written as [54, 130]

$$h_+ + ih_\times = \sum_{\ell, m, n} {}_{-2}S_{\ell mn}(\iota, \varphi) C_{\ell mn} e^{i(\Omega_{\ell mn}(t-t_0))} \quad (1.73)$$

$$= \sum_{\ell, m, n} {}_{-2}S_{\ell mn}(\iota, \varphi) A_{\ell mn} e^{-t/\tau_{\ell mn}} e^{i(\omega_{\ell mn}t + \phi_{\ell mn})}, \quad (1.74)$$

where we chose $t_0 = 0$ in the second line, only considering the description valid at positive t , and

$${}_{-2}S_{\ell mn}(\iota, \varphi) = e^{im\varphi} {}_{-2}S_{\ell m}(M\chi\Omega_{\ell mn}, \cos\iota). \quad (1.75)$$

1 Introduction




Figure 1.6: The solid lines show the real and imaginary part of the spheroidal harmonics ${}_{-2}S_{\ell mn}$ for a black hole with spin $\chi = 0.8$, for $\ell = 2$, $n = 0$, and all allowed values of m . The dashed lines show the same for the spin-weighted spherical harmonics ${}_{-2}Y_{\ell m}$, which can be used to approximate the spheroidal harmonics. Both are shown for varying polar angle θ , with fixed azimuthal angle $\phi = 0$. The `pykerr` software package was used in generating this and all further figures in this section [123].




Figure 1.7: The same as figure 1.6 for $\ell = 3$.




Figure 1.8: The real and imaginary part of the spheroidal harmonics is shown in the left and right panel, respectively, with $(\ell, m, n) = (3, 3, 0)$ and spin $\chi = 0.2$. For each direction, the surface's distance from the origin shows the absolute value of the real or imaginary part of the spheroidal harmonic. The colour of the surface shows the value of the spheroidal harmonic, with dark colours corresponding to negative values and light colours indicating positive values. The positive z -direction corresponds to the spherical polar coordinates with $\theta = 0$, while the positive x -direction corresponds to $\phi = 0$.

1 Introduction

Here, ${}_{-2}S_{\ell mn}(\theta, \phi)$ are the spin-weighted spheroidal harmonics with spin-weight -2 [123, 131], and ${}_{-2}S_{\ell m}$ corresponds to $S(\theta)$ in Equation (1.71). Examples for the angular dependencies of the spheroidal harmonics are shown in Figures 1.6, 1.7, and 1.8. The integer indices ℓ, m, n enumerate the individual modes of the resulting infinite discrete spectrum. We have $\ell \geq 2$, the azimuthal number m , with $-\ell \leq m \leq \ell$, and n enumerating the overtones, which are ordered by decreasing damping time, starting with $n = 0$. The angles appearing in the spheroidal harmonics are the inclination ι , the angle between the black hole's angular momentum vector and the line-of-sight vector from the black hole to the observer, and φ , the azimuthal angle of the black hole with respect to this line-of-sight vector. The complex frequencies $\Omega_{\ell mn} = \omega_{\ell mn} + \frac{i}{\tau_{\ell mn}}$ are determined by the Teukolsky equation³. Their real part $\omega_{\ell mn}$ describes the frequency of the damped sinusoid, while $\tau_{\ell mn}$ in the imaginary part describes the damping time of the exponential decay. Finally, the complex amplitudes $C_{\ell mn}$ describe the amplitude $A_{\ell mn}$ and initial phase $\phi_{\ell mn}$ of the damped sinusoid through $C_{\ell mn} = A_{\ell mn} \exp[i\phi_{\ell mn}]$, i.e. $A_{\ell mn} = |C_{\ell mn}|$. The amplitudes and phases depend on the initial perturbation of the black hole, and thus on the binary's properties such as its mass ratio [126, 132]. The frequencies and damping times, however, are determined entirely by the mass and spin of the final black hole. These parameters of the spectrum are thus characteristic of the final black hole [131, 133, 134]. Examples for the values of the frequencies and damping times for a selection of modes are shown in Figure 1.9. In analogy to studying matter through its characteristic electromagnetic emission spectrum, the approach of studying the properties of the black hole through measurement of its quasi-normal mode spectrum is known as black hole spectroscopy.

For any choice of ℓ, m, n , there are two solutions, which we refer to as co- and counter-rotating, as they have been shown to correspond to perturbations that are co- or counter-rotating with respect to the angular momentum of the black hole for at least those modes with $\ell = m$ [54]. Instead of introducing an additional index for these two solutions, we use a convention making use of the geometry of the system to label them by the sign of the spin parameter. If we allow negative values for the dimensionless spin χ , we find two redundant descriptions of the same physical system and observer. The orientation of the black hole is given by a vector \vec{s} parallel to its spin axis. A positive spin χ then means that the spin vector is aligned with \vec{s} , while a negative spin means it is anti-aligned. The inclination angle ι is measured between the line-of-sight vector pointing from the black hole to the observer and \vec{s} , thus for positive spin this is the same description as when allowing only $\chi \geq 0$. Consider a black hole with positive spin χ along \vec{s} , with the vector pointing towards the observer in the direction given by ι and the azimuthal angle φ . Then the same black hole and observer are described by replacing $\chi \rightarrow -\chi$, $\iota \rightarrow \pi - \iota$, $\varphi \rightarrow \varphi + \pi$. This redundant description allows us to use the sign of the spin to determine whether we describe the co- or counter-rotating case of the perturbation. Our convention is to use positive χ for co-rotating perturbations and negative χ for counter-rotating ones. The complex frequency and the frequency-dependent spheroidal harmonics will differ between the positive and negative spin values, as these are mapped to either the co- or counter-rotating solution.

When summing over all possible modes, we thus also must sum over the two possible signs of the spin, which we can indicate through the additional index $p = \text{sgn}(\chi)$ added to each ℓmn in Equation (1.74). When considering the perturbed black hole resulting from a binary collision, we can use the binary's orientation to determine the reference direction \vec{s} . We will choose it such that $\vec{s} \cdot \vec{L} \geq 0$ for the orbital angular momentum of the binary \vec{L} . Using the value of χ to determine co- or counter-rotation has the advantage of allowing for a smooth transition between the cases when we analyse probability distributions for the parameters.

In the Kerr case, we find a symmetry for a given $\ell, |m|, n$ between the complex frequencies

³Comparing with Equation (1.71), we have $\Re(\Omega_{\ell mn}) = \Re(\omega)$ and $\Im(\Omega_{\ell mn}) = -\Im(\omega)$.

of the m and $-m$ modes, $\Omega_{\ell-mn} = -\Omega_{\ell mn}^*$ [54, 130]. So the frequencies and damping times are related by $f_{\ell mn} = -f_{\ell-mn}$ and $\tau_{\ell mn} = \tau_{\ell-mn}$, reducing the number of free parameters. We also find a symmetry between the spheroidal harmonics, ${}_{-2}S_{\ell-mn}(\iota, \varphi) = (-1)^\ell {}_{-2}S_{\ell mn}^*(\pi - \iota, \varphi)$ [135, 136].

Describing only a single mode, i.e. the waveform for one specific choice of ℓ, m, n and sign of χ , restricts to the assumption that the wave is circularly polarised, as then $h_+ + ih_\times \propto \exp[-i\Omega_{\ell mn}t + \phi_{\ell mn}]$. To describe generic polarisations, we will typically consider the sum of the m and $-m$ modes for a given ℓ, n ,

$$h_{\ell|m|n} = C_{\ell mn} e^{i\Omega_{\ell mn}t} {}_{-2}S_{\ell mn}(\iota, \varphi) + C_{\ell-mn} e^{i\Omega_{\ell-mn}t} {}_{-2}S_{\ell-mn}(\iota, \varphi), \quad (1.76)$$

with complex amplitudes $C_{\ell mn} = A_{\ell mn} \exp[i\phi_{\ell mn}]$. Using the above symmetries for the complex frequencies and the spheroidal harmonics, we find

$$h_{\ell|m|n} = C_{\ell mn} e^{i\Omega_{\ell mn}t} {}_{-2}S_{\ell mn}(\iota, \varphi) + (-1)^\ell C_{\ell-mn} e^{-i\Omega_{\ell mn}^*t} {}_{-2}S_{\ell mn}^*(\pi - \iota, \varphi). \quad (1.77)$$

If we assume that the initial perturbations are symmetric under reflections with respect to the equatorial plane, as may be expected for a non-precessing system [41, 130, 137], we find for the complex amplitudes

$$C_{\ell-mn} = (-1)^\ell C_{\ell mn}^* \quad (1.78)$$

$$\Rightarrow A_{\ell-mn} = (-1)^\ell A_{\ell mn}, \quad \phi_{\ell-mn} = -\phi_{\ell mn}. \quad (1.79)$$

With ${}_{-2}S_{\ell mn}(\iota, \varphi) = e^{im\varphi} {}_{-2}S_{\ell mn}(\iota, 0)$, and defining ${}_{-2}S_{\ell mn}(\iota) := S_{\ell mn}(\iota, 0)$, this leads to

$$h_{\ell|m|n} = A_{\ell mn} e^{-\frac{t}{\tau_{\ell mn}}} \left[e^{i(\omega_{\ell mn}t + \phi_{\ell mn} + m\varphi)} {}_{-2}S_{\ell mn}(\iota) + e^{-i(\omega_{\ell mn}t + \phi_{\ell mn} + m\varphi)} {}_{-2}S_{\ell mn}^*(\pi - \iota) \right] \quad (1.80)$$

$$= A_{\ell mn} e^{-\frac{t}{\tau_{\ell mn}}} \left[\cos \Phi_{\ell mn} ({}_{-2}S_{\ell mn}(\iota) + {}_{-2}S_{\ell mn}^*(\pi - \iota)) \right. \\ \left. + i \sin \Phi_{\ell mn} ({}_{-2}S_{\ell mn}(\iota) - {}_{-2}S_{\ell mn}^*(\pi - \iota)) \right] \quad (1.81)$$

$$= A_{\ell mn} e^{-\frac{t}{\tau_{\ell mn}}} \left[\cos \Phi_{\ell mn} S_{\ell mn}^+(\iota) + i \sin \Phi_{\ell mn} S_{\ell mn}^\times(\iota) \right], \quad (1.82)$$

with the definitions

$$\Phi_{\ell mn} = \omega_{\ell mn}t + \phi_{\ell mn} + m\varphi, \quad (1.83)$$

$$S_{\ell mn}^+(\iota) = {}_{-2}S_{\ell mn}(\iota, 0) + {}_{-2}S_{\ell mn}^*(\pi - \iota, 0) \quad (1.84)$$

$$= {}_{-2}S_{\ell mn}(\iota, 0) + (-1)^\ell {}_{-2}S_{\ell-mn}(\iota, 0), \quad (1.85)$$

$$S_{\ell mn}^\times(\iota) = {}_{-2}S_{\ell mn}(\iota, 0) - {}_{-2}S_{\ell mn}^*(\pi - \iota, 0) \quad (1.86)$$

$$= {}_{-2}S_{\ell mn}(\iota, 0) - (-1)^\ell {}_{-2}S_{\ell-mn}(\iota, 0). \quad (1.87)$$

We then find for the two polarisations

$$h_{\ell|m|n} = h_{\ell|m|n}^+ + ih_{\ell|m|n}^\times, \quad (1.88)$$

$$h_{\ell|m|n}^+ = A_{\ell mn} e^{-t/\tau_{\ell mn}} \left[\cos \Phi_{\ell mn} \Re \{ S_{\ell mn}^+(\iota) \} - \sin \Phi_{\ell mn} \Im \{ S_{\ell mn}^\times(\iota) \} \right], \quad (1.89)$$

$$h_{\ell|m|n}^\times = A_{\ell mn} e^{-t/\tau_{\ell mn}} \left[\cos \Phi_{\ell mn} \Im \{ S_{\ell mn}^+(\iota) \} + \sin \Phi_{\ell mn} \Re \{ S_{\ell mn}^\times(\iota) \} \right]. \quad (1.90)$$

We can approximate the spin-weighted spheroidal harmonics through spin-weighted spherical harmonics [138, 139], ${}_{-2}S_{\ell mn}(\iota, \varphi = 0) \approx {}_{-2}Y_{\ell m}(\iota, \varphi = 0)$, where ${}_{-2}Y_{\ell m}(\iota, \varphi = 0)$ is real. Inserting this into $h_{\ell mn}$, we find

$$h_{\ell|m|n} = A_{\ell mn} e^{-t/\tau_{\ell mn}} \left[Y_{\ell m}^+ \cos \Phi_{\ell mn} + i Y_{\ell m}^\times \sin \Phi_{\ell mn} \right], \quad (1.91)$$




Figure 1.9: The real and imaginary part ω_R and ω_I of the complex quasi-normal mode frequency $\omega_{\ell mn}$ is shown for $\ell = 2$ in the top panel and for $\ell = 3$ in the bottom panel. For each case, we include all allowed m and the overtones up to $n = 7$. Solid upwards pointing triangles denote co-rotating modes (i.e. with positive spin χ), and hollow downward-pointing triangles refer to counter-rotating modes (i.e. negative spin). The absolute value of the spin is $|\chi| = 0.5$ in this example. The bottom and left axes show the mass-independent $M\omega_R$ and $M\omega_I$, while the top and right axes show $f = \omega_R/2\pi$ and $\tau = 1/\omega_I$ for the example of a black hole with mass $M = 10M_{\odot}$. For each angular mode, the value of m is indicated in the plot. The overtone number is found simply by counting the marked points on each curve, starting with $n = 0$ at the lowest $M\omega_I$, as n is chosen such that the overtones are ordered by decreasing damping time.

with definitions equivalent to those above, with ${}_{-2}Y_{\ell m}^*(\iota, 0) = -{}_{-2}Y_{\ell m}(\iota, 0)$, and

$$Y_{\ell m}^+(\iota) = {}_{-2}Y_{\ell m}(\iota, 0) + {}_{-2}Y_{\ell m}(\pi - \iota, 0) \quad (1.92)$$

$$= {}_{-2}Y_{\ell m}(\iota, 0) + (-1)^\ell {}_{-2}Y_{\ell -m}(\iota, 0), \quad (1.93)$$

$$Y_{\ell m}^\times(\iota) = {}_{-2}Y_{\ell m}(\iota, 0) - {}_{-2}Y_{\ell m}(\pi - \iota, 0) \quad (1.94)$$

$$= {}_{-2}Y_{\ell m}(\iota, 0) - (-1)^\ell {}_{-2}Y_{\ell -m}(\iota, 0). \quad (1.95)$$

The polarisations thus read

$$h_{\ell|m|n}^+ \approx A_{\ell mn} e^{-t/\tau_{\ell mn}} [\cos \Phi_{\ell mn} Y_{\ell mn}^+(\iota)], \quad (1.96)$$

$$h_{\ell|m|n}^\times \approx A_{\ell mn} e^{-t/\tau_{\ell mn}} [\sin \Phi_{\ell mn} Y_{\ell mn}^\times(\iota)]. \quad (1.97)$$

We use the model including the spheroidal harmonics in the analyses presented in Chapter 4, and the model using the approximation through spherical harmonics in Chapters 2, 3, and 7. So far, we have considered $A_{\ell mn}$ to represent the amplitude of the gravitational waves at the location of interest, e.g. at a detector. It is conventional to redefine the amplitude by extracting the mass M and luminosity distance D_L , $A_{\ell mn} \rightarrow (M/D_L)A_{\ell mn}$ [140, 141], with

$$D_L = \sqrt{\frac{L}{4\pi F}}, \quad (1.98)$$

where L is the luminosity of the source and F is the energy flux received by the observer [7]. The energy flux of the gravitational wave is proportional to the square of its amplitude, and so the amplitude falls off with the inverse of the luminosity distance, $F \propto A^2 \Rightarrow A \propto 1/D_L$. The redefined $A_{\ell mn}$ is then a quantity depending only on the source and not on its distance.

This leads to the following waveform model for the quasi-normal mode signal starting at $t = 0$. The QNM-modelled portion of the signal has $h(t < 0) = 0$, not meaning that a physical signal should vanish before $t = 0$, but rather be described by a different appropriate model. We sum over all $\ell \geq 2$, $|m| \geq \ell$, $n \geq 0$, and, if both co- and counter-rotating modes are to be included, over the sign of the spin $p = \text{sgn}(\chi)$.

$$h_+(t) = \frac{M}{D_L} \sum_{\ell, |m|, n, p} A_{\ell mn} e^{-t/\tau_{\ell mn}} [\cos(\omega_{\ell mn} t + \phi_{\ell mn}) \Re \{S_{\ell mn}^+(\iota)\} - \sin(\omega_{\ell mn} t + \phi_{\ell mn}) \Im \{S_{\ell mn}^\times(\iota)\}] \quad (1.99)$$

$$\approx \frac{M}{D_L} \sum_{\ell, |m|, n, p} {}_{-2}Y_{\ell m}^+(\iota) A_{\ell mn} e^{-t/\tau_{\ell mn}} \cos(\omega_{\ell mn} t + \phi_{\ell mn}),$$

$$h_\times(t) = \frac{M}{D_L} \sum_{\ell, |m|, n, p} A_{\ell mn} e^{-t/\tau_{\ell mn}} [\cos(\omega_{\ell mn} t + \phi_{\ell mn}) \Im \{S_{\ell mn}^+(\iota)\} + \sin(\omega_{\ell mn} t + \phi_{\ell mn}) \Re \{S_{\ell mn}^\times(\iota)\}] \quad (1.100)$$

$$\approx \frac{M}{D_L} \sum_{\ell, |m|, n, p} {}_{-2}Y_{\ell m}^\times(\iota) A_{\ell mn} e^{-t/\tau_{\ell mn}} \sin(\omega_{\ell mn} t + \phi_{\ell mn}). \quad (1.101)$$

Here, we have also absorbed the term with the azimuthal angle $m\varphi$ into $\phi_{\ell mn}$, as the initial phases are unknown in our analyses and would be degenerate with the measurement of φ . We will occasionally use the meaning of $A_{\ell mn}$ as the wave's amplitude at the detector. For example, in Chapter 2, we assume a more general waveform model of damped sinusoids but do not relate them to a black hole at a specific mass and distance.

1.6.1 No-hair theorem

We expect that astrophysical black holes are characterised entirely by two intrinsic parameters, their mass M and angular momentum χ . Here, we consider the position, linear momentum, and orientation of the angular momentum to be extrinsic parameters which we control by choice of the reference coordinate system. The third intrinsic parameter, the electric charge Q , is expected to be negligible for astrophysical black holes [8, 133]. Several theorems show that under certain conditions, solutions of the Einstein field equations that contain black holes are restricted to a few unique cases. A static vacuum solution must be the Reissner-Nordström solution and a stationary and axisymmetric vacuum solution must be the Kerr-Newman solution if a black hole is present. In the uncharged case, the Schwarzschild and Kerr solutions are then the only ones possible [8, 142–148].

The Kerr solution can then be considered to approximate an isolated astrophysical black hole. While the above restrictions are not precisely satisfied, the astrophysical case is expected to be sufficiently similar to allow this description. A black hole that is far from other objects for a long time will approach the required symmetries, and sparse matter in the vicinity will have negligible effects [133]. Thus, we use the Kerr solution to approximate an isolated black hole of negligible charge, which is then described by two intrinsic parameters. This has led to the commonly used name for the no-hair theorems in that further characteristics for black holes, such as metaphorical hair, are absent.

The frequencies and damping times of the quasi-normal mode spectrum are functions of only these two parameters. Furthermore, the mapping of the mass and spin of the black hole to the QNM frequencies and damping times is invertible and unique [149, 150]. Measuring these two parameters for one mode with known numbers ℓmn therefore allows to determine the mass and spin of the black hole, and in turn calculation of the entire remaining QNM spectrum.

Measuring at least one additional third parameter then allows to test for consistency between the measured values. We can use two QNM parameters to calculate M , χ , and from this the value expected for the third parameter, to compare with its measured value. By quantifying the agreement between the measured mode parameters, we can test the prediction of the no-hair theorem. As it states that only two independent parameters describe the black hole, all measured mode parameters must thus be compatible with the same M and χ . Equivalently, all measured QNM parameters must agree with the prediction from any pair of such parameters. While this is often referred to as a test of the no-hair theorem, it is the validity of its assumptions that is truly tested. This includes the approximation of the spacetime surrounding the astrophysical black hole by the Kerr metric, the applicability of the perturbative approach, non-degeneracy of the event horizon and more technical requirements of the theorems, as detailed in [142–148], and even the fundamental description through general relativity. Indeed, generalising the theorems to less restrictive conditions is an active area of theoretical research [151, 152].

Knowledge of the numbers ℓmn of a measured mode is necessary to perform these tests. The signal for each mode has the form of a damped sinusoid, differing only in the four parameters of frequency, damping time, amplitude and phase. Measurement of a mode’s frequency and damping time alone would then allow it to be mapped to different masses and spins depending on the choice of ℓmn . The assignment of these numbers can be guided by expectations for the different excitation of individual modes and the effect of the factor from the spheroidal harmonic decomposition. These affect the amplitude of the signal in the detector, where typically the late ringdown is dominated by the $\ell mn = 220$ mode [140]. Furthermore, information from an analysis of the complete inspiral-merger-ringdown signal may be used to constrain the expected mass and spin of the final black hole and the mass ratio of the original binary. This in turn constrains the expected QNM frequencies and damping times, and the relative excitation amplitudes of the modes, which depend on the mass ratio [132].

1.7 Echoes

Alternative models for compact objects as gravitational wave sources have been proposed, which can result in detectable modifications of the emitted signals. One class of models leads to an additional gravitational wave emission following the initial burst from the merger. As this signal appears after a time delay and can take the form of repeated pulses, it is referred to as gravitational wave echoes. The relevant models broadly fall into two categories of either exchanging the black hole with an alternative object or modifying the description of the black hole from that of standard general relativity. Motivations for these models include considerations of quantum-gravitational effects and attempts to address the black hole information paradox [153–155].

If the signal for the inspiral, merger, and ringdown in these models is similar to that of GR, existing searches will be able to detect them. We can then focus on the already available detections instead of searching for new events in the data independently. As the signals detected to date agree well with the predictions of GR, modifications to the signal must be sufficiently small or restricted to features not covered in the employed waveform models. We can then directly search for signals obeying these restrictions in the data of detected events.

Here, we consider models that agree with the Kerr geometry up to a small distance above its outer event horizon, with deviations only appearing below this distance. Such alternative objects are then comparable in compactness to black holes and are sometimes called exotic compact objects. Examples of these are thin-shell wormholes and boson stars. Models with modifications from standard GR include those of a material firewall near the horizon, and gravastars with a thin shell of matter, inspired by near-horizon quantum effects. These models have in common that they possess a photon sphere but either no horizon or some additional structure above it. Infalling gravitational waves therefore do not encounter a surface acting as a perfect absorber, as presented by the event horizon, and need not be completely lost to the outside observer. Instead, a part of the radiation may pass through the object or be reflected at its surface.

For our analyses, we assume that the entire inspiral-merger-ringdown (IMR) signal is not significantly modified from that of a binary black hole merger. While finding the inspiral and merger waveform for a binary containing an exotic compact object is more complicated, they are expected to be similar to that of a black hole binary [122, 153]. The ringdown remains unmodified due to the travel time of gravitational waves between the photon sphere and the location of deviations from the Kerr geometry. We can consider the ringdown radiation to originate at the photon sphere of the object, the region where lightlike orbits are possible, with an ingoing and outgoing component. For a black hole, the ingoing part is lost through the horizon, while the outgoing part results in the measurable ringdown signal. We assume the deviations from the Kerr metric to occur close to the horizon, not affecting the far-away photon sphere. The outgoing radiation is therefore left unchanged, while the ingoing part must traverse the distance to the interior region that deviates from the Kerr geometry. Any waves returning from the interior then traverse this distance again, reaching the photon sphere after a time delay corresponding to the round-trip travel time. They then arrive at the photon sphere too late to interfere significantly with the initial outgoing ringdown radiation. Therefore we consider the condition satisfied that the IMR signal is well matched by that of a black hole binary and focus on additional signals directly following detected events.

The proposed models can be described by placing an at least partially reflective surface close above the horizon. Two physical scenarios are captured by this description, actual reflection at a surface and transmission through the interior. In the first case, the modelled surface represents a physical boundary where actual reflection occurs for ingoing radiation. The same model can

1 Introduction

approximate the behaviour of an object with weak interior interaction with gravitational waves, while no physical reflection occurs. Assuming regularity in the interior, an ingoing wave could enter the central object and pass through its centre, re-emerging after travelling through the object. We can then approximate the object's emission by modelling this process using the reflective surface. Here, the centre acts effectively reflective, so an additional initial time-delay or phase-shift is added to account for the travel through the interior. Introducing the reflective surface changes the boundary conditions for gravitational waves from purely ingoing at the horizon of a black hole to (partial) reflection at the proposed surface above the horizon.

We consider two phenomenologically distinct signals that arise from this modification. First, the change in boundary conditions leads to a different quasi-normal mode spectrum of the object than found for a black hole as described in Section 1.6. While the initial ringdown signal is unchanged, it does not represent the QNMs of the modified object. The partially reflecting boundary conditions instead of the perfect absorption for the black hole case lead to a long-lived, sinusoidal signal when solving the Teukolsky equation [156]. It can be seen as the result of the cavity formed by the inner reflective surface and the effective potential barrier, the peak of the effective potential at the photon sphere, which also acts partially reflective. Low-frequency components of an initial wave pulse become trapped in this cavity, but slowly leak out through the potential barrier [122, 153]. Alternative physical interpretations are proposed, describing an excitation of quantum states of the interior object, with the excitation being radiated as the long-lived modes and passing the potential barrier [157]. The resulting signal has the form of a weakly damped, long-lived sinusoid, and thus a very narrow, almost monochromatic frequency spectrum.

A second distinct signal are pulsed gravitational wave echoes. These result from the same cavity formed by the interior reflective surface and the effective potential barrier that leads to the long-lived sinusoidal modes. However, as the ingoing ringdown pulse contains a broader spectrum of frequencies, a part of the signal is not trapped as effectively as the low-frequency component [153]. The original ingoing pulse is reflected outwards at the interior surface, encountering again the potential barrier located at the photon sphere. A significant portion of the high-frequency part of the signal passes the potential barrier, while the rest is reflected back inwards, and the process repeats. This results in a series of pulses following the ringdown, separated by the round-trip travel time through the cavity, as illustrated in Figure 1.10. The pulses' amplitude decreases as energy is lost at each encounter with the potential barrier, and the frequency content changes with each pulse. The damping of the echo pulses is stronger than that of the quasi-normal modes, and at late times the signal consists only of the previously described, long-lived quasi-normal mode [122, 153].

In a simplified model for the resulting signal, the late part of the IMR waveform is repeated with a time delay. A constant damping factor is applied to each echo compared to the previous, and a phase-inversion between subsequent echoes is introduced to account for the reflection at the interior surface [3, 158]. The resulting waveform is illustrated in Figure 1.11.

By evaluating the support for these alternative models in the data, we can help concentrate modelling efforts on the most promising candidates. We achieve this by ruling out or constraining in supported parameter space the signals deviating from the GR expectation, and thus also the corresponding models. Here we focus on the analysis of the data given proposed signals, treating the details of the mechanisms producing them more superficially. Some additional discussion regarding these mechanisms and waveform models is found in Chapters 5 and 6 for the pulsed echoes and in Chapter 7 for the long-lived modes.

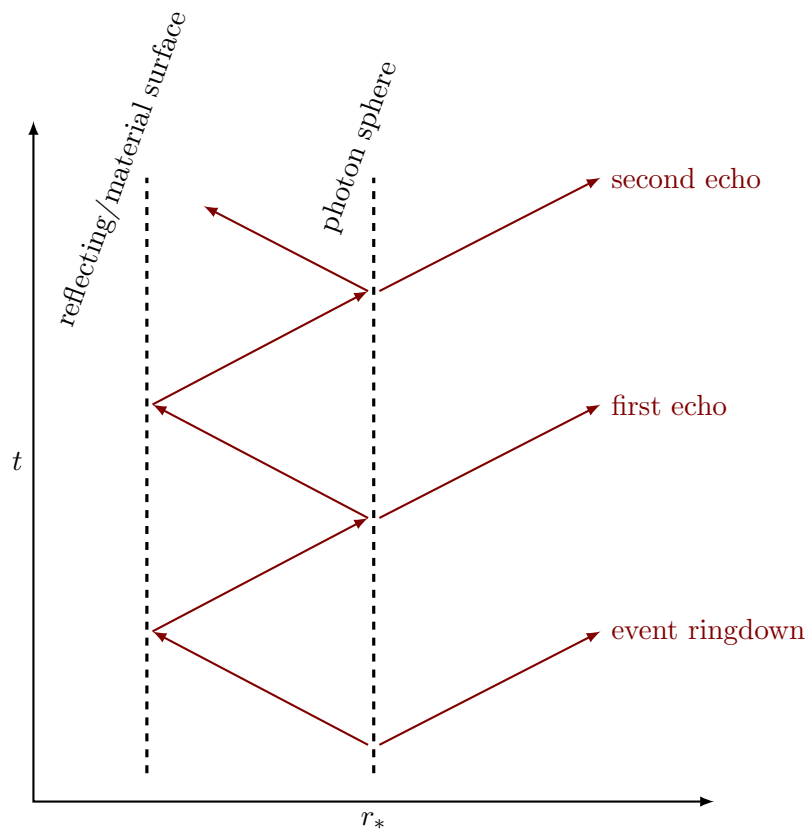


Figure 1.10: Schematic of the occurrence of repeated pulses or echoes when a reflective surface exists outside of the horizon. The axes show the tortoise coordinate r_* and time coordinate t . The initial ringdown emission originates at the photon sphere with an outgoing and ingoing part. While the outgoing part can be detected as the ringdown signal as expected for a black hole described by general relativity, the ingoing part encounters a surface where it is reflected back outwards. Encountering the potential barrier at the photon sphere, the pulse of radiation is partially transmitted, leading to a signal following the ringdown after a time delay. The remainder of the pulse is reflected inwards again, repeating the process and leading to a signal consisting of a sequence of damped pulses following the ringdown.




Figure 1.11: A time-domain inspiral-merger-ringdown waveform with appended echoes as resulting from a simplified model. Dimensionless strain with arbitrary scaling is shown as a function of time. The late part of the inspiral-merger-ringdown waveform is repeated with a delay based on the travel time of the pulse through the cavity. The amplitude of each echo pulse is reduced by the same factor compared to the previous echo, and subsequent echoes are phase-inverted to account for the reflection at the interior surface.

1.8 Results in context

This section provides a short overview of the presented results in the context of existing research.

General relativity is the current standard theory of gravity, and Einstein presented his field equations in 1915 [159]. Gravitational waves were found as an approximate solution in 1916 [160], and the exact solution describing a non-rotating black hole was published in the same year [15]. Although initially debated, both existence of gravitational waves and transmission of energy through these were established over the following decades. The physically more realistic scenario of a rotating black hole was described in 1963 [16]. The first observations recognised as compact objects occurred soon thereafter, with the discovery of pulsars in 1967 and of the first x-ray source thought to be a black hole in 1964 [161–163]. Observations of the first binary pulsar in 1974 provided indirect evidence of gravitational waves through measurement of the orbital period, which changed in agreement with the energy loss of gravitational wave emission [164, 165]. Over the following decades, interferometric detectors for direct measurement of gravitational waves were developed and constructed, undergoing repeated upgrades to improve their sensitivity.

These instruments made the first direct observation of a gravitational wave in 2015 when they detected the signal GW150914 from the merger of a black hole binary [33]. Many further detections have since been made, including the signal GW170817 from a binary neutron star collision, providing multi-messenger data through coincident electromagnetic observations [34, 166–170]. The gravitational-wave signals are extracted from the noise-dominated detector data through matched-filtering searches on large computer clusters, using banks of pre-calculated signal templates [55]. Accurate templates are constructed through a variety of approaches, including post-Newtonian calculations and the effective one-body formalism [43, 44]. Advancements in numerical relativity have made simulations of binary black holes possible, allowing in particular to describe the merger phase, and to approximate the numerical results in efficient phenomenological or interpolated waveform models [46, 50]. Perturbation theory is used to describe the emission from the single perturbed remnant object of the merger [53]. Waveforms for the entire emission may be constructed incorporating these different descriptions for individual stages in the binary evolution.

The perturbed remnant of the merger presents a particularly clean case to study the compact object. Its ringdown emission can probe directly the remnant’s properties, and is accessible to theoretical modelling. The late ringdown of black holes consists of a superposition of damped sinusoids, as first found in numerical studies of gravitational waves scattering off black holes [52]. Later, equations describing perturbations in the Kerr geometry were studied analytically and reduced to an eigenvalue problem, which in turn is solved to yield the characteristic frequencies of the ringdown spectrum [53, 121, 131]. The frequency and damping time of each mode in the ringdown spectrum depend only on the mass and angular momentum of the black hole, allowing to directly measure these properties. This knowledge can be used to evaluate the consistency between the binary inspiral phase and the ringdown of the single remnant, and to test predictions of general relativity, such as the black hole area-increase law [70, 146, 171]. Furthermore, measuring multiple modes of the ringdown allows to test consistency also between the individual modes. Astrophysical black holes of general relativity are described by only two parameters, their mass and angular momentum, according to the no-hair theorem. This prediction can thus be tested by measuring and comparing multiple modes, as all must agree with the same mass and angular momentum. However, detection of all but the dominant quadrupolar mode is difficult, as the subdominant modes’ amplitudes are typically suppressed [132].

We therefore study the prospects of measuring multiple modes in Chapters 2 and 3. Bayesian parameter estimation is performed on a large number of simulated ringdown signals. Both the analysis and the population of signals emulate the study of real detections, as the full param-

1 Introduction

eter estimation is carried out, and the signals are constructed based on current observational population models. We find that the confident detection of subdominant ringdown modes is unlikely using current instruments, but can be achieved with near-future upgrades to the detectors. Similarly, events suitable to place stringent bounds on deviations from the no-hair theorem can be expected to occur during a few years of collected data at these improved sensitivities. By combining data from several detections, we find improved prospects for no-hair tests, while also uncovering systematic biases in this type of analysis and finding strategies to prevent these. We find that deviations from the no-hair prediction can likely be constrained to about $\pm 10\%$, using a few years of data at the target sensitivity of the next observation run.

Searches for subdominant ringdown modes have been conducted on the available detections, but little evidence was found so far [71, 72]. Studies of the fundamental mode's overtones suggest that these can be used to extend the quasi-normal mode description up to the merger, allowing to access larger signal amplitudes [128]. Similar hierarchical analyses have been performed comparing the fundamental mode and its first overtone for no-hair tests, as have other ringdown-based studies profiting from the increased signal-to-noise ratio [71, 72, 172–174]. While details of this approach are still being investigated, overtones appear to be a useful tool to describe the waveform near the merger [129].

In Chapter 4, we analyse the ringdown signal of GW190521, where the high mass of the remnant places the ringdown in a sensitive frequency-band of the detector. A new method is applied, focusing the analysis on the ringdown part of the signal. It prevents contamination from the inspiral and merger portions while retaining the signal inside the analysis window without loss. The ringdown signal is modelled through different combinations of modes, also incorporating one overtone. We find strong evidence for the presence of one subdominant mode in addition to the dominant one. A no-hair test with these modes produces unexpectedly tight bounds on deviations from the prediction.

Compact objects as alternatives to the black holes of vacuum general relativity are a focus of recent research. A number of modifications to black holes and alternative objects have been proposed, often inspired by considerations of quantum effects [122]. Such effects may lead to material structure near the horizon of a black hole, or a horizonless object may take its place. Several such models have in common that they can be described by introducing a surface which reflects infalling radiation, in particular gravitational waves. Instead of being fully absorbed by the black hole horizon, radiation may then re-appear, making the objects observationally distinguishable. Changing the interior boundary condition to allow reflection has been found to lead to a different quasi-normal mode spectrum of the object, consisting of a long-lived sinusoidal signal [156, 175]. More recently, it has been argued that the initial signal from a collision of such objects would be practically indistinguishable from that of black holes [153, 176]. This initial signal would be followed by emission characteristic for the object, which therefore may be found in the data following already detected signals attributed to binary black holes. The modification to the signal takes the form of repeated damped pulses after the merger, gravitational wave echoes, and at late times consists of the long-lived quasi-normal modes of the object.

An initial search for the echo signal, using a simplified waveform model, was reported to find evidence of echoes in the data of the first gravitational wave detections [158]. In Chapter 5, we reproduce the results of this first analysis and characterise in detail the applied methods. We find that these are limited to the detection of very loud echo signals and contain systematic biases, while the estimation of noise statistics suffers from correlated samples. Improving and extending the statistical methods, we find the statistical significance for echoes to be low, while at the same time the presence of echo signals can not be ruled out through these methods. Therefore, we perform a Bayesian parameter estimation analysis using the same waveform model in Chapter 6. This shows the data for the individual events to be either consistent with the noise hypothesis

or to prefer the absence of echoes. We produce upper bounds on the amplitudes of this type of echo signals ruled out by the data.

Further searches with this and generalised waveform models have been conducted, reporting no evidence for echoes [71, 72, 177, 178]. As the matched-filtering approach is dependent on sufficiently accurate modelling of the signal, only those signals very similar to the specific waveform used in the search can be ruled out confidently. Theoretical models of alternative compact objects are being improved, and more accurate waveform models are under development [179].

Chapter 7 presents a search focusing on the long-lived quasi-normal modes of these alternative objects in the data of GW150914. Using new methods adapted to the long duration of the signal, we can place bounds on the position of the reflective surface and thus on the deviations from the Kerr geometry. For a broad class of models, we show such deviations to occur at most extremely close to the horizon.

2 Prospects for black hole spectroscopy in the next decade

Gravitational wave observations of the ringdown of the remnant black hole in a binary black hole coalescence provide a unique opportunity of confronting the black hole no-hair theorem in general relativity with observational data. The most robust tests are possible if multiple ringdown modes can be observed. In this paper, using state-of-the-art Bayesian inference methods and the most up-to-date knowledge of binary black hole population parameters and ringdown mode amplitudes, we evaluate the prospects for black hole spectroscopy with current and future ground based gravitational wave detectors over the next 10 years. For different population models, we estimate the likely number of events for which the subdominant mode can be detected and distinguished from the dominant mode. We show that black hole spectroscopy could significantly test general relativity for events seen by the proposed LIGO Voyager detectors.¹

2.1 Introduction

The remnant black hole (BH) formed after the coalescence of two compact objects emits gravitational radiation while settling down to a Kerr BH. This stage is known as the ringdown. Perturbation theory predicts that, at late enough times, the ringdown consists of a superposition of exponentially damped sinusoids called quasinormal modes (QNM) [52, 180] (see also [54, 181]). The QNMs are characterised by a set of complex frequencies $\Omega_{\ell mn}$ labeled by three integers; ℓ, m are angular quantum numbers while $n = 0, 1, 2 \dots$ is the overtone index. According to the no-hair theorem in standard general relativity (GR), $\Omega_{\ell mn}$ is uniquely defined by the BH mass and spin. The measurement of multiple QNMs in a BH ringdown, known as BH spectroscopy, is crucial for robust observational tests of the no-hair theorem with gravitational waves (GW) based only on the ringdown signal [133, 141].

The excitation of different QNMs depends on the nature of the perturbation, i.e. on the properties of the binary progenitor [126, 132, 182–184]. Thus, for aligned spin systems, the amplitudes of the different modes are determined by the spins of the initial compact objects and the mass ratio $q = m_1/m_2 \geq 1$, with $m_{1,2}$ the mass of each object. The ringdown signature is dominated by the fundamental $(\ell, m) = (2, 2)$ mode [140]. For non-spinning binaries with equal masses ($q = 1$), odd ℓ modes vanish and the loudest subdominant mode is the $(\ell, m) = (4, 4)$ mode. As the mass ratio increases, the $(\ell, m) = (3, 3)$ mode becomes the loudest subdominant mode, with amplitudes larger than 30% of the dominant amplitude ($A_{330}/A_{220} > 0.3$) [126, 132]. Hence, coalescences of two unequal-mass BHs or neutron-star black-hole binaries (NSBH) are the most promising sources for measurability of subdominant modes in the ringdown. For still higher mass ratios, the relative amplitude of the modes can also tell us about the alignment of the orbit relative to the BH spin during the inspiral phase [182–184].

Two main conditions are necessary to test the no-hair theorem: (i) the detectability of at least two modes, and (ii) the resolvability of the frequencies and/or damping times of each mode. Theoretical estimates of the necessary ringdown signal-to-noise ratio (SNR) for each of these

¹This chapter is an adaptation of the work published as [1] as described in Section 1.2, with the copyright for the published article [1] held by the American Physical Society (2020).

conditions can be found in the literature [139, 141]. These studies have predicted that Advanced LIGO should observe several ringdown events at design sensitivity, but will not be able to detect subdominant modes from the coalescence of stellar-mass BBH for BH spectroscopy [185, 186]. In this paper we revisit the prospects for accurate BH spectroscopy with the next decade of LIGO detectors. In general, asymmetric binaries are more likely to produce higher amplitudes for the subdominant ringdown modes. However, based on the gravitational-wave observations to date, more asymmetric systems are also likely to be much fewer in number [63] (although recent public alerts from the third observing run of Advanced LIGO and Virgo suggest possible detections of NSBH [187, 188]). In addition, the orientation of a source relative to the detectors also has an important effect on the observed amplitudes. Systems where the angular momentum is aligned with the line-of-sight to the source are more luminous, but these orientations are not favourable for observing the subdominant modes. Taking all these effects into account, along with the most up-to-date estimates of the ringdown mode amplitudes [132] and state-of-the-art gravitational wave parameter estimation techniques [99, 189], we show that black hole spectroscopy can provide non-trivial limits on general relativity with the LIGO Voyager detector.

At least 10 binary black-hole (BBH) coalescences have been observed in the first two observing runs of Advanced LIGO and Virgo [166, 190–194]. The loudest BBH event is still the first detection, GW150914 [33], with a ringdown signal-to-noise ratio (SNR) $\rho \simeq 8.5$ [69] at 3 ms after merger. This event has not provided significant evidence for the presence of measurable subdominant modes with $\ell \neq 2$ [195]. However, recent work suggests that the inclusion of higher overtones of the dominant $\ell = 2$ mode allows for the modelling of the ringdown immediately after the merger, hence obtaining higher SNR in ringdown signatures [128]. The analysis of the GW150914 ringdown using the fundamental mode and its first overtone provides the first constraints to date of deviations of the no-hair theorem using two QNMs [172]. Here we use the Bayesian inference and model selection frameworks [119] on simulated BBH populations to establish the measurability and accurate resolvability of two ringdown QNMs over the next decade, providing rate estimates for constraining the no-hair theorem to within $\pm 20\%$ at the 90% credible level. We restrict ourselves to the resolvability of subdominant QNMs ($\ell \neq 2$) for two reasons: (1) the excitation amplitudes of overtones on the general parameter space of the binary’s properties are not yet well-understood and we lack predictions to model ringdown signatures that include overtones for a large population of BBH mergers, and (2) the frequencies of the overtones are very similar to each other, hence accurate resolvability of an overtone is more challenging than of a subdominant mode.

This manuscript is organised as follows. Section 2.2 introduces the Bayesian inference and model selection frameworks, as well as the ringdown model used. Section 2.3 describes the details on the BBH population considered. In Section 2.4 we report the rates on measurable subdominant modes and prospects for resolvability of the necessary parameters to perform tests of the no-hair theorem. Finally, we conclude our findings in Sec. 2.5.

2.2 Bayesian framework

We use Bayesian methods to infer the properties of the remnant BH from our data, $d(t)$, and to determine the presence of a measurable subdominant mode in the ringdown signature. Given a model hypothesis of the ringdown signal, H , parameterised by the source properties, $\vec{\vartheta}$, Bayes’ theorem defines the posterior probability distribution:

$$p(\vec{\vartheta}|d, H) = p(\vec{\vartheta}|H) \frac{p(d|\vec{\vartheta}, H)}{p(d|H)}, \quad (2.1)$$

where $p(\vec{\vartheta}|H)$ is the prior knowledge based on astrophysical populations or theoretical models, the likelihood $p(d|\vec{\vartheta}, H)$ is the conditional probability of observing the data $d(t)$ given the model H with parameters $\vec{\vartheta}$, and the evidence $p(d|H)$ is a normalisation constant that only depends on the data and the chosen model. Calculating the evidence requires marginalisation over the entire parameter space, which can become computationally challenging. While this computation can be avoided for Bayesian parameter estimation, model selection between two competing models requires accurate estimates of the evidence.

In Bayesian model selection, the Bayes factor weighs the evidence provided by the data in support of one model versus another [119, 196]:

$$\mathcal{B}_{AB} = \frac{p(d|H_A)}{p(d|H_B)}. \quad (2.2)$$

In this manuscript we follow the nomenclature of [119]: a Bayes factor $\mathcal{B}_{AB} > 3.2$ indicates “substantial” support for H_A over H_B , $\mathcal{B}_{AB} > 10$ indicates “strong” support, while $\mathcal{B}_{AB} > 100$ is “decisive”.

2.2.1 The likelihood function

For a GW detector network with uncorrelated stationary Gaussian noise, the likelihood is given by

$$p(d|\vec{\vartheta}, H) \propto \exp \left[-\frac{1}{2} \sum_{a=1}^N \langle d_a - h_a(\vec{\vartheta}), d_a - h_a(\vec{\vartheta}) \rangle \right], \quad (2.3)$$

where N is the number of detectors, d_a is the data for each detector, and $h_a(\vec{\vartheta})$ is the waveform model evaluated for a set of parameters $\vec{\vartheta}$ as observed by detector a . The noise-weighted inner product is defined as

$$\langle x, y \rangle = 4\Re \int_0^\infty \frac{\tilde{x}^*(f)\tilde{y}(f)}{S_n(f)} df, \quad (2.4)$$

with $S_n(f)$ being the one-sided power spectral density (PSD) of the detector’s noise, $\tilde{x}(f)$ the Fourier transform of $x(t)$, and $*$ indicating the complex conjugate.

In this paper we use the `PyCBC Inference` [99, 189] toolkit to compute the likelihood function and estimate posterior probability distributions. Accurate marginalisation for evidence estimation is achieved using the nested sampling algorithm `cpnest` [115].

2.2.2 The ringdown model

The strain $h(t)$ produced by a gravitational wave at the detector is given by

$$h(t) = F_+(\alpha, \delta, \Psi)h_+(t) + F_\times(\alpha, \delta, \Psi)h_\times(t), \quad (2.5)$$

where $F_{+,\times}$ are the antenna pattern functions determined by the relative orientation between the detector frame and the wave frame [80], i.e. the sky location of the source (right ascension α and declination δ in a geocentric coordinate system) and the polarisation angle Ψ that defines the relative orientation of the wave frame with the geocentric coordinate system. For short transient signals, these orientation angles (and hence $F_{+,\times}$) are assumed to be time independent. For future generations of observatories with improved low frequency sensitivity, it might become necessary to account for the time dependence of $F_{+,\times}$. However, the ringdown itself will be short enough that for our purposes we do not need to consider this effect here.

The ringdown signal of a Kerr BH consists of a sum of exponentially damped sinusoids:

$$h_+ + ih_\times = \frac{M}{D_L} \sum_{\ell, m, n} {}_{-2}S_{\ell mn}(\iota, \varphi) A_{\ell mn} e^{i(\Omega_{\ell mn} t + \phi_{\ell mn})}, \quad (2.6)$$

where M is the mass of the BH in the detector frame and D_L is the luminosity distance to the source. The functions ${}_{-2}S_{\ell mn}(\iota, \varphi)$ are the spin-weighted spheroidal harmonics, which depend on the inclination angle ι between the BH spin and the line-of-sight from the source to the observer, and the azimuth angle φ between the BH and the observer. The complex QNM frequencies $\Omega_{\ell mn}$, determined from the Teukolsky equation [53, 131], define the frequency and damping time of the damped sinusoid, $\Omega_{\ell mn} = \omega_{\ell mn} + i/\tau_{\ell mn}$. The amplitudes $A_{\ell mn}$ and $\phi_{\ell mn}$ depend on the initial perturbation and take different values for different (ℓ, m, n) modes. Henceforth, we restrict ourselves to the $n = 0$ overtone and drop the overtone index n for simplicity.

Assuming that the ringdown begins at $t = 0$, the two gravitational-wave polarisations are given by

$$\begin{aligned} h_+(t) &= \frac{M}{D_L} \sum_{\ell, m} {}_{-2}Y_{\ell m}^+(\iota) A_{\ell m} e^{-t/\tau_{\ell m}} \cos(\omega_{\ell m} t + \phi_{\ell m}), \\ h_\times(t) &= \frac{M}{D_L} \sum_{\ell, m} {}_{-2}Y_{\ell m}^\times(\iota) A_{\ell m} e^{-t/\tau_{\ell m}} \sin(\omega_{\ell m} t + \phi_{\ell m}). \end{aligned} \quad (2.7)$$

Here we have approximated the spheroidal harmonics ${}_{-2}S_{\ell mn}$ by spin-weighted spherical harmonics ${}_{-2}Y_{\ell m}$ [138, 139], which introduces an error of $\sim 1\%$ or less, defining [138]:

$$\begin{aligned} {}_{-2}Y_{\ell m}^+(\iota) &= {}_{-2}Y_{\ell m}(\iota, 0) + (-1)^\ell {}_{-2}Y_{\ell -m}(\iota, 0), \\ {}_{-2}Y_{\ell m}^\times(\iota) &= {}_{-2}Y_{\ell m}(\iota, 0) - (-1)^\ell {}_{-2}Y_{\ell -m}(\iota, 0). \end{aligned} \quad (2.8)$$

The azimuthal angle φ appears in the spherical harmonics as ${}_{-2}Y_{\ell m}(\iota, \varphi) = e^{im\varphi} {}_{-2}Y_{\ell m}(\iota, 0)$. As we will consider the modes' phases $\phi_{\ell m}$ free parameters, they are degenerate with φ , which we absorb into the $\phi_{\ell m}$.

The ringdown analysis in this paper follows the methods developed in [70, 197]. We use two different waveform models, (i) a *Kerr model* where we assume the remnant object to be a Kerr BH, hence the ringdown QNM frequencies $\Omega_{\ell mn}$ are uniquely determined by the mass M and the spin χ of the BH, and (ii) an *agnostic model* where we assume the nature of the remnant object to be unknown, hence the ringdown is parameterised by each individual QNM frequency $\Omega_{\ell mn}$ and we drop the factor M/D_L in Eq. (2.7). The Kerr model (i) is our starting point for determining the measurability of a subdominant mode. Resolvability of the subdominant mode for testing the no-hair theorem is determined using the agnostic model (ii).

2.3 Populations

We construct populations of candidate BBH ringdown signals based on the observational population model B of [63] (we ignore NSBH mergers here because population models including NSBH are largely uncertain). The component-mass and mass-ratio distributions follow power laws with exponents $-\alpha$ and β_q , respectively (see Eq. (2) in [63]). For the component-mass distribution, we use the measured median value $\alpha = 1.6$, with masses in the range $[5.4, 57]M_\odot$ (we use the lowest m_{min} and the largest m_{max} values, to account for uncertainties in the mass bounds of BHs). For the mass-ratio distribution we use two different exponent values: the measured median value $\beta_q = 6.7$, and a uniform distribution $\beta_q = 0$ (which is used in model A of [63]). Mass ratios are restricted to be within the range $[1, 8)$. We assume the individual

BHs to be non-spinning prior to the merger, which is consistent with the population of BBHs observed by LIGO/Virgo thus far. Sources are distributed uniformly in co-moving volume; we choose a maximum luminosity distance, $D_L^{(\max)}$, dependent on the considered detector network. The inclination angle ι is distributed uniformly in $\cos \iota \in [-1, 1)$, and the polarisation angle Ψ uniformly in $\Psi \in [0, 2\pi)$.

The mass and spin of the remnant Kerr BH determine the ringdown frequencies $\Omega_{\ell m}$ [141]. We obtain an estimate of the remnant's source frame mass $M^{(\text{src})}$ and dimensionless spin χ using the fitting formulae to numerical relativity [198, 199] implemented in the `LALSuite` software package [200]. The detector frame mass M is given by $M = (1+z)M^{(\text{src})}$, where z is the redshift calculated from the luminosity distance, D_L , assuming a standard Λ CDM cosmology [201]. The excitation amplitudes $A_{\ell m}$, which depend on the mass ratio q of the binary, are determined using the fitting formulae in [132] at $t = 10M$ after the merger. The phases $\phi_{\ell m}$ of the modes are distributed uniformly in $\phi_{\ell m} \in [0, 2\pi)$, in contrast to previous work in the literature where both phases were fixed for simplicity [134, 139, 202].

The BBH parameters for each candidate are drawn randomly from their respective distributions to generate two-mode ringdown signals with the dominant $(\ell, m) = (2, 2)$ mode and either the $(\ell, m) = (3, 3)$ or the $(\ell, m) = (4, 4)$ subdominant mode. We consider a three-detector LIGO network consisting of the observatories in Hanford (H1), Livingston (L1) and India (I1). We use three different sensitivities for these detectors [203]: Advanced LIGO design sensitivity (Adv. LIGO), A+ and Voyager. We do not consider here the complete third generation detectors, which include the Einstein Telescope [95, 204, 205] and Cosmic Explorer [93], or the space based LISA mission [206], since this would take us beyond the 10-year timeframe.

For each candidate, we calculate the optimal SNR of the subdominant mode in each detector, $\rho_{\text{det}} = \sqrt{\langle h, h \rangle}$, where h is the ringdown signal of the subdominant mode projected into the detector (see Eqs. (2.5) and (2.7)). To avoid a large number of sources with no measurable subdominant mode, we reject candidates with combined optimal SNR $\rho_c = \sqrt{\sum_{\text{det}} \rho_{\text{det}}^2} < 2.5$ in the subdominant mode. For the same reason, the maximum D_L considered is limited to different values for different sensitivities, namely $D_L^{(\max)} = \{1, 3, 5\}$ Gpc for Adv. LIGO, A+ and Voyager, respectively. The number of draws required to find a sample population of 100 signals with $\rho_c \geq 2.5$ in the subdominant mode yields the fraction of interesting candidates out of all BBH signals. Figure 2.1 shows the resulting populations for each detector network considered.

2.4 Analysis and results

2.4.1 Rates of measurable subdominant modes

We add the population of accepted candidate ringdown signals (shown in Fig. 2.1) into different Gaussian noise realisations coloured with the PSD of the desired detector. To determine the measurability of the subdominant mode, we use the Kerr BH ringdown model and perform two separate Bayesian parameter estimation analyses, using: (H_A) templates with the fundamental $(2, 2)$ mode plus the corresponding (ℓ, m) subdominant mode, and (H_B) templates with only the fundamental $(2, 2)$ mode. The Bayes factor \mathcal{B}_{AB} is then calculated as the ratio of the evidences for model H_A versus model H_B . Those sources with $\mathcal{B}_{AB} > 3.2$ are further analysed in the next section to determine the resolvability of the subdominant mode.

The parameters $(M, \chi, A_{\ell m}, \phi_{\ell m}, \iota, \Psi)$ are estimated from the data, which represents a set of 8 parameters in the two-mode ringdown model H_A , and 6 parameters in the single-mode ringdown model H_B . The priors used in the parameter estimation analysis are uniform in all parameters: BH mass $M \in [10, 200]M_\odot$, BH spin $\chi \in [-0.99, 0.99)$, log-amplitude of the fundamental mode




Figure 2.1: Source frame mass, $M^{(\text{src})}$, and spin, χ , of the BHs with optimal SNR $\rho_c \geq 2.5$ in the subdominant mode (either the (3,3) or the (4,4) mode), obtained using the observational population models of [63]. The colours represent the luminosity distance of the source, where the maximum allowed distance was $D_L^{(\text{max})} = \{1, 3, 5\}$ Gpc for Adv. LIGO, A+ and Voyager, respectively.

	$\beta_q = 0$			$\beta_q = 6.7$		
	$\mathcal{B}_{AB} > 3.2$	$\mathcal{B}_{AB} > 10$	$\mathcal{B}_{AB} > 100$	$\mathcal{B}_{AB} > 3.2$	$\mathcal{B}_{AB} > 10$	$\mathcal{B}_{AB} > 100$
Adv. LIGO	$0.036^{+0.039}_{-0.019}$	$0.028^{+0.031}_{-0.015}$	$0.011^{+0.012}_{-0.006}$	$0.008^{+0.009}_{-0.004}$	$0.006^{+0.007}_{-0.003}$	$0.003^{+0.003}_{-0.001}$
A+	$0.46^{+0.51}_{-0.25}$	$0.28^{+0.31}_{-0.15}$	$0.14^{+0.15}_{-0.07}$	$0.08^{+0.09}_{-0.04}$	$0.06^{+0.06}_{-0.03}$	$0.03^{+0.03}_{-0.02}$
Voyager	$2.63^{+2.89}_{-1.42}$	$1.83^{+2.01}_{-0.99}$	$0.89^{+0.97}_{-0.48}$	$0.30^{+0.33}_{-0.16}$	$0.21^{+0.24}_{-0.12}$	$0.11^{+0.12}_{-0.06}$

Table 2.1: Rates of BBH ringdown signals per year (yr^{-1}) with a detectable subdominant (3, 3) or (4, 4) mode for a population with uniform mass-ratio distribution ($\beta_q = 0$) and for a population with a power-law mass-ratio distribution with exponent $\beta_q = 6.7$. The Bayes factors in each column indicate substantial support ($\mathcal{B}_{AB} > 3.2$), strong support ($\mathcal{B}_{AB} > 10$), and decisive support ($\mathcal{B}_{AB} > 100$) for the presence of a second mode.

$\log_{10}(A_{22}) \in [-4, 4)$, relative subdominant mode amplitude $\hat{A}_{\ell m} = A_{\ell m}/A_{22} \in [0, 0.5)$, ringdown phases $\phi_{\ell m} \in [0, 2\pi)$, polarisation angle $\Psi \in [0, 2\pi)$, and inclination angle $\cos \iota \in [-1, 1)$. We fix the start time of the ringdown, the (ℓ, m) of the subdominant mode, the sky location, and the distance to the source to the injected values. While the start time of the ringdown is not uniquely defined in the literature, we do not explore the issue in this paper and assume that this can be determined by other means [70, 195, 207]. Further, we can safely assume that we have some knowledge from the inspiral part of the signal regarding the mass ratio of the binary to determine which is the loudest subdominant mode to look for. Since we are using a network of three detectors, the sky location should be relatively well known from the analysis of the full gravitational-wave signal. Finally, while the distance might not be accurately measured, fixing this parameter to a wrong value will only affect the measurement of the fundamental amplitude A_{22} and not affect our conclusions.

We calculate the rate of ringdown events with detectable subdominant mode in each detector network based on the BBH merger rate density given in [63] ($R = 53.2^{+58.5}_{-28.8} \text{ Gpc}^{-3} \text{ yr}^{-1}$) and the co-moving volume up to $D_L^{(\text{max})}$ for each detector network. Table 2.1 lists the rate of events per year with substantial ($\mathcal{B}_{AB} > 3.2$), strong ($\mathcal{B}_{AB} > 10$), and decisive ($\mathcal{B}_{AB} > 100$) support for the presence of a subdominant mode. These rates are the combination of both the (3, 3) and the (4, 4) modes. While we have made the simplifying assumption that only one subdominant mode will be measurable, some of the considered BBH systems might have two subdominant modes with SNR $\rho_c \geq 2.5$. However, studying the performance of a three-mode ringdown Bayesian analysis is beyond the scope of this paper.

2.4.2 Resolvable subdominant modes for testing GR

In the presence of two measurable ringdown modes, resolvability of the $\Omega_{\ell m}$ frequencies allows for BH spectroscopy tests. However, QNMs of rotating BHs in modified theories of gravity have not been calculated [208], and Kerr-like exotic compact objects can have the same or a similar QNM spectrum as Kerr BHs [122]. While it might be challenging to disprove all BH alternatives, accurate measurements of the QNM spectrum will be crucial to constrain deviations from GR (see however [209] for possible ways of parameterising frequencies and damping times accounting for deviations from GR). It has been shown for non-rotating alternative BH models that GR deviations are more significant in the QNM frequencies than in the damping times [210]. Hence, we focus here on constraining deviations from the subdominant mode's frequency.

We consider those ringdown events with $\mathcal{B}_{AB} > 3.2$ in the previous section and perform

the same parameter estimation analysis, now using the agnostic model defined in Sec. 2.2.2 to estimate the ringdown frequencies $\Omega_{\ell m}$ of the two QNMs. Hence, 10 parameters ($\omega_{\ell m}$, $\tau_{\ell m}$, $A_{\ell m}$, $\phi_{\ell m}$, ι , Ψ) are now estimated from the data. The priors are uniform in the frequencies $f_{\ell m} = \omega_{\ell m}/2\pi \in [50, 1024)$ Hz and damping times $\tau_{\ell m} \in [0.45, 30)$ ms, excluding parameters that yield masses and spins outside of the ranges used in the previous section with the Kerr model. The amplitudes of the (ℓ, m) modes have different orders of magnitude, because of the missing factor M/D_L when dropping the Kerr assumption. Hence, the prior in log-amplitude of the fundamental mode is now $\log_{10}(A_{22}) \in [-25, -17)$. The priors in the remaining parameters are the same as in the previous section. Finally, we apply an additional set of constraints on the subdominant frequency and damping time to be within $\pm 25\%$ of the GR expectation.

Using the fitting formulae in [141], we can compare the mass and spin measurement obtained from the $(2, 2)$ parameters and from the subdominant (ℓ, m) parameters. Furthermore, based on the measurement of the $(2, 2)$ mode, we can infer the measured deviation for the frequency of the subdominant (ℓ, m) mode, $\delta f_{\ell m}$. Table 2.2 lists the rates of BBH ringdown signals per year that constrain deviations from GR to within $\delta f_{\ell m} \pm 20\%$ at the 90% credible level. The results are summarised in Fig. 2.2.

Network	$\delta f_{\ell m} \leq \pm 20\%$
Adv. LIGO	$0.026^{+0.028}_{-0.014}$
A+	$0.27^{+0.30}_{-0.15}$
Voyager	$1.34^{+1.47}_{-0.73}$

Table 2.2: Rates of BBH ringdown signals per year (yr^{-1}) with strong support for the presence of a second mode ($\mathcal{B}_{AB} > 3.2$) where deviations of the frequencies from the GR prediction are constrained to within $\delta f_{\ell m} \leq \pm 20\%$ at the 90% credible level. We only show the rates for the population with uniform mass-ratio distribution ($\beta_q = 0$), since we know from the previous section that rates for a population with $\beta_q = 6.7$ will be lower.

2.5 Conclusions

In this paper we have applied for the first time the full Bayesian inference framework to a population of BH ringdowns derived from the observational population models published by the LIGO Scientific and Virgo Collaborations. Furthermore, we have allowed for completely variable ringdown phases, inclination angles, polarisation angles and sky locations, contrary to previous works that have fixed one or more of these parameters for simplicity [134, 139, 202].

Within the Bayesian model selection framework, future generations of LIGO detectors will likely deliver measurable subdominant QNM modes from BBH mergers over the next decade. However, resolvability of the subdominant frequencies is technically challenging, and accurate tests of the no-hair theorem might only be possible in very few cases. These results are in agreement with previously published works [185, 186], where the ringdown SNR was used to determine the measurability and resolvability of QNMs.

Merger population models from gravitational-wave observations are still largely uncertain. The third observing run of Advanced LIGO and Virgo might be uncovering a new population of NSBH and other previously unobserved types of mergers, which could boost the rates of measurable and resolvable subdominant modes. Hence, the rates obtained in this work might




Figure 2.2: Expected rates of BBH mergers for which two ringdown modes can be observed and resolved. We consider two population models corresponding to $\beta_q = \{0, 6.7\}$. Shown are the rate of events that have “substantial”, “strong”, and “decisive” Bayesian evidence (using the nomenclature of Ref. [119]) for a two-mode Kerr hypothesis relative to a single-mode hypothesis (filled circles, diamonds, and squares, respectively). For the events that have substantial evidence, we perform a followup analysis in which the frequency and damping time of the subdominant mode is allowed to deviate from the expected GR value. The rate of events for which the deviation from GR of the subdominant frequency $|\delta f_{\ell m}|$ is constrained to be $\leq 20\%$ is given by the open circles.

turn out to be pessimistic as more gravitational-wave detections are made available.

Acknowledgements

We are thankful to Swetha Bhagwat, Duncan Brown, Evan Goetz, Scott Hughes, Alexander H. Nitz, Paolo Pani and Frans Pretorius for useful comments and discussions. We are especially grateful to Ssohrab Borhanian and Bangalore Sathyaprakash for providing us with the fits for the ringdown mode amplitudes in [132]. MC acknowledges support from NSF grant PHY-1607449, the Simons Foundation, and the Canadian Institute For Advanced Research (CIFAR). Computations have been performed on the Atlas cluster of the Albert Einstein Institute (Hannover).

2.6 Appendix

Results for the injection with the largest Bayes factor in the (3, 3) population using the Voyager sensitivity are shown in Figure 2.3 for the Kerr model and in Figure 2.4 for the agnostic model. This BBH is located at a distance $D_L \simeq 250$ Mpc.




Figure 2.3: Posterior distributions from the analysis with a Kerr ringdown. The parameters of interest are the BH mass M , BH spin χ , amplitude of the (2, 2) mode, A_{22} , and amplitude ratio of the (3, 3) mode, A_{33}/A_{22} . The red crosses indicate the injected parameters, and the dashed lines in the histograms correspond to the median value and the 90% credible level.




Figure 2.4: Posterior distributions from the analysis with an agnostic ringdown. The parameters of interest are the ringdown frequencies, $f_{\ell m}$, and damping times, $\tau_{\ell m}$, the amplitude of the (2, 2) mode, A_{22} , and the amplitude ratio of the (3, 3) mode, A_{33}/A_{22} . A_{22} has a different order of magnitude compared to the Kerr model because of the missing M/D_L factor in the approximant. The red crosses indicate the injected parameters, and the dashed lines in the histograms correspond to the median value and the 90% credible level.

3 Hierarchical black hole spectroscopy

3.1 Introduction

Gravitational wave detections that allow for precise measurement of multiple ringdown modes are expected to be rare for current and near-future detector sensitivities. By measuring precisely and comparing the parameters of multiple modes in a single event, we can place constraints on potential deviations from the predictions of general relativity. However, in Chapter 2 we have predicted that an event suitable to place appreciable constraints will occur less than once per decade of collected data at the Advanced LIGO design sensitivity, based on recent models of the population of binary black hole systems.

Introducing an additional assumption allows us to combine data from multiple detections and improve the prospects of constraining such deviations in the near future. Assuming that all sources deviate from the prediction of the no-hair theorem in the same manner, we can perform a hierarchical analysis where the results from one event inform the analysis of the next. With each analysed event, our knowledge about the deviations is increased, represented by a change in the probability distribution for the deviation parameters. The posterior distribution of the first event’s analysis is used as the new prior distribution for the analysis of the second, and this process is repeated for all detections.

We construct a population of simulated events by drawing each event’s parameters from their respective distributions, which are based on observational population models. The number of events we draw corresponds to one year of collected data, only including events that are likely detectable according to a minimal signal-to-noise ratio (SNR) threshold for the complete inspiral-merger-ringdown signal. To limit computational cost, we then only perform the full parameter estimation for the 50 candidates that are most promising in terms of the subdominant ringdown mode’s SNR.

To reduce the overall time required to analyse the entire population, it is desirable to analyse the events in parallel instead of sequentially. Instead of using one event’s posterior distribution as the prior for the next analysis, we can equivalently analyse all events individually and combine their individual posterior distributions when ensuring compatibility of the individual priors. If the same prior distributions are used for the parameter estimation of all events, the combined posterior can conveniently be found as a simple product of the individual posterior distributions.

We expect it to be unlikely to find an event in one year of data suitable to individually yield stringent constraints on deviations. However, the combination of many quieter events may suffice to find such constraints. This allows us to lower the threshold for the minimum gain of information imposed on the individual events, compared to the criteria from Chapter 2.

Therefore, this approach contrasts that of Chapter 2. Instead of drawing events until a suitable candidate is found and estimating the required observation time from this, here we fix the observation time and estimate the bounds that may be placed given this amount of data. While the constraints for each individual event are a by-product of this analysis, we change the question from “How long to collect data to constrain deviations to a given magnitude?” to “How tightly can a given duration of collected data constrain the deviations?”.

3.2 Analysis

3.2.1 Combining posteriors

Even with the detection of many gravitational wave signals, we expect events with precisely measurable quasi-normal mode (QNM) spectra to be rare. Each individual event may, however, contain a small amount of information through weak preferences in the posterior distribution of the QNM parameters. A parameter common to different events allows to combine these small changes in the individual posteriors into a single posterior based on the data for multiple events.

As we are interested in measuring deviations from the predicted frequency and damping time of individual modes, we consider as an example a parameter δf . We assume that the true value of δf is the same for two events with the two sets of data d_1 , d_2 , and prior information X . From this, we want to find a single posterior distribution incorporating information from both d_1 and d_2 . Updating the original prior, $p(\delta f | X)$, using all available data means, applying Bayes' theorem,

$$p(\delta f | d_1, d_2, X) = \frac{p(d_1, d_2 | \delta f, X) p(\delta f | X)}{p(d_1, d_2 | X)}. \quad (3.1)$$

Signal parameters that are independent between events result in an increasing total number of parameters in the analysis when many events are included. The high-dimensional parameter space can then impede parameter estimation methods. It is therefore desirable to analyse the events individually and combine the results after the separate analyses. This also prevents us from having to repeat the analysis of the same data when adding data from newly detected events to be included in the combined posterior distribution. Starting from the above expression, we find

$$p(\delta f | d_1, d_2, X) = \frac{p(d_1, d_2 | \delta f, X) p(\delta f | X)}{p(d_1, d_2 | X)} \quad (3.2)$$

$$= \frac{p(d_2 | d_1, \delta f, X) p(d_1 | \delta f, X) p(\delta f | X)}{p(d_2 | d_1, X) p(d_1 | X)} \quad (3.3)$$

$$= \frac{p(d_2 | d_1, \delta f, X) p(\delta f | d_1, X)}{p(d_2 | d_1, X)}, \quad (3.4)$$

where we used the product rule in the first and Bayes' theorem in the second step. We can interpret this expression as simply adding the data d_1 to the prior information X when analysing the data d_2 . We also notice that the last line includes the posterior resulting from the analysis of the first event, $p(\delta f | d_1, X)$. If the likelihood does not depend on the previous data, $p(d_2 | d_1, \delta f, X) = p(d_2 | \delta f, X)$, we can re-interpret the last line as follows. Up to normalisation, it is the expression for the posterior distribution $p(\delta f | d_2, X)$, with the prior replaced by $p(\delta f | X) \rightarrow p(\delta f | d_1, X)$. So we may find the posterior distribution based on multiple sets of data through a hierarchical analysis. For this, we combine the information from multiple events by analysing the first event given our initial prior X , finding the posterior $p(\delta f | d_1, X)$. The second event is then analysed using as prior the posterior distribution of the first analysis, and so forth for subsequent events.

We assume the likelihood of observing some data to be independent of previously observed data. This is encoded in our description of the likelihood as that of stationary Gaussian noise, which remains unchanged when adding data collected at sufficiently separated times to the prior

information. This means that the data d_1 and d_2 are conditionally independent given $\delta f, X$,

$$p(d_1, d_2 | \delta f, X) = p(d_1 | \delta f, X)p(d_2 | \delta f, X) \quad (3.5)$$

$$\Leftrightarrow p(d_2 | d_1, \delta f, X)p(d_1 | \delta f, X) = p(d_1 | \delta f, X)p(d_2 | \delta f, X) \quad (3.6)$$

$$\Leftrightarrow p(d_2 | d_1, \delta f, X) = p(d_2 | \delta f, X). \quad (3.7)$$

We can consider both the first and last line as the definition of conditional independence. The last line gives the expression we use here to combine the results from multiple events.

Using the hierarchical procedure allows us to find the combined posterior from analysing individual events. However, the events must still be analysed sequentially, as one event's posterior is required as the prior for the analysis of the next. We can instead analyse each event using the same prior distribution and then combine the results, allowing for parallel analysis of all events. Assuming conditional independence, we can write

$$p(\delta f | d_1, d_2, X) = \frac{p(d_1, d_2 | \delta f, X) p(\delta f | X)}{p(d_1, d_2 | X)} \quad (3.8)$$

$$= \frac{p(d_1 | \delta f, X) p(d_2 | \delta f, X) p(\delta f | X)}{p(d_1, d_2 | X)} \quad (3.9)$$

$$= \frac{p(\delta f | d_1, X) p(d_1 | X)}{p(\delta f | X)} \frac{p(\delta f | d_2, X) p(d_2 | X)}{p(\delta f | X)} \frac{p(\delta f | X)}{p(d_1, d_2 | X)} \quad (3.10)$$

$$= \underbrace{\frac{p(d_1 | X) p(d_2 | X)}{p(d_1, d_2 | X)}}_c \frac{p(\delta f | d_2, X)}{p(\delta f | X)} \frac{p(\delta f | d_1, X)}{p(\delta f | X)} p(\delta f | X), \quad (3.11)$$

using Bayes' theorem in the third step, and considering c as a normalisation factor depending only on the data. The last line shows a convenient way to calculate the combined posterior distribution from the individual events' posteriors. Ignoring the data-dependent normalisation factor c , the last line is a product of the posterior distributions of the individual events, each divided by the prior $p(\delta f | X)$. This product is then multiplied once by the prior. Our parameter estimation analysis for one event yields as result the posterior distribution, while the prior distribution is known. We can then find the combined posterior by analysing individually N events using the same prior for each, then calculating the product of the individual posteriors, divided by $N - 1$ times the prior, and normalising the result. This is the procedure we use in our analysis, performing the individual events' parameter estimation in parallel.

We note that this represents a special case of the more general approach of reweighting posterior distributions to replace the prior. Assume we are given a posterior distribution for a parameter and the corresponding prior under which the posterior was found, with prior information X . We want to find from this the posterior resulting from a set of prior information X' , which only differs from X in the prior distribution for the parameter, i.e. $p(\delta f | X)$. Then we can write

$$p(\delta f | d, X) \frac{p(\delta f | X')}{p(\delta f | X)} = \frac{p(d | \delta f, X) p(\delta f | X) p(\delta f | X')}{p(d | X) p(\delta f | X)} \quad (3.12)$$

$$= \frac{p(d | \delta f, X) p(\delta f | X')}{p(d | X)} \quad (3.13)$$

$$= p(\delta f | d, X'), \quad (3.14)$$

where the last line holds as we assumed the difference in prior information to lie only in the expected distribution for δf , leaving the likelihood and evidence unchanged. As an additional condition, we require that the new prior distribution covers a subset of the original prior. Otherwise, regions included in the new distribution have not been sampled when calculating the

posterior distribution which we want to reweight with the new prior. The available posterior distribution and the corresponding likelihood samples then contain no information for these regions and cannot be reweighted to represent the posterior for the new prior distribution. We can now for example interpret Eq. (3.11) as reweighting the posterior from d_2 by dividing by its prior and multiplying by the new prior, the posterior based on d_1 .

3.2.2 Populations and analysis framework

The Bayesian analysis framework and the model for the ringdown signal are identical to those presented in Chapter 2. We assume stationary Gaussian noise in the likelihood model, and find the posterior probability distributions for the desired parameters by applying Bayes' theorem. Calculations are performed with `PyCBC Inference`, using the different nested sampling algorithm `dynesty` to sample parameter space and calculate evidences [116]. The ringdown waveform consists of damped sinusoids with frequencies and damping times determined from the mass and spin of the final black hole through the Teukolsky equation for each quasi-normal mode. As before, we approximate spheroidal harmonics through spherical harmonics.

The binary black hole merger population is again constructed based on the observational population model B in [63], where we do not include binary neutron star or neutron star-black hole events. In contrast to Chapter 2, we now only consider the uniform distribution for the mass ratio, with the parameter $\beta_q = 0$ appearing in the exponent of the power-law distribution. This seems reasonable as subsequent high mass-ratio detections have lent increased support to smaller values of β_q (e.g. [211]). However, we defer replacing the population parameter distributions with more recent observational models to future work [64, 212].

The full inspiral-merger-ringdown (IMR) waveform is used to determine the thresholds for inclusion of a simulated event in the analysis. We consider a network of the two Advanced LIGO detectors and the Virgo detector and use their design-sensitivity noise curves [203] to calculate the optimal SNR $\rho_{\text{det}} = \sqrt{\langle h, h \rangle}$ of the IMR signal in each detector. We require an optimal SNR $\rho_{\text{det}} > 4.5$ in each detector, as well as optimal $\rho_{\text{det}} > 6$ in at least one detector to consider the event to be likely detectable. We limit the luminosity distance to $D_L \leq 2 \text{ Gpc}$ to limit the number of drawn candidates, assuming more distant events unlikely to have a detectable quasi-normal mode spectrum. We then draw parameters for as many events as are expected to occur during one year in the corresponding comoving volume according to the optimistic rate estimate from [63]. Restricting the timespan of collected data limits the necessary draws to a manageable number, independent of the signal SNR. This allows us to increase the range compared to Chapter 2, where we drew events until reaching a given number of loud signals and used $D_L \leq 1 \text{ Gpc}$ for this detector sensitivity. There, the larger range would have led to many quiet candidates being drawn and rejected, increasing the computational cost for producing the desired population sample. We find a rate of 1139 events occurring in the given comoving volume per year. About 32% of the events are likely detectable by the network of the two Advanced LIGO and the Virgo detectors at their design sensitivities, according to our chosen SNR threshold.

To produce a more robust statistic of simulated events, we repeat this process and collect candidate events for 4 years of data. We then restrict the parameter estimation to the 200 candidates with the highest SNR in the subdominant mode. When quoting results for one year of collected data, we base these on choosing 50 events at random from the 200. This choice is not equivalent to selecting the 50 events with the loudest subdominant modes from one year, and will over-represent events with higher subdominant mode SNRs. However, most events will have very quiet subdominant modes, with negligible contributions to a combined posterior. Based on the results of Chapter 2, we assume that fewer than 50 events per year will contribute to constraining the deviation parameters of interest. We therefore expect that we can approximate

the contributing events for one year by drawing 50 out of the 200 candidates, without relevant influence from overestimating the SNRs in the non-contributing subset.

We produce simulated gravitational wave data for each of the 200 candidate events in the population, consisting of the sum of noise realisations and signals. The noise is Gaussian, coloured with the power spectral density of the respective detector noise curve. The signals consist of the ringdown QNMs, excluding the inspiral and merger portions of a real signal. We generate only the $\ell = m = 2$ mode and the $\ell = m = 3$ subdominant mode, which is expected to typically be excited more strongly than the $\ell = m = 4$ mode [132]. We do not consider overtones in this work, always setting $n = 0$ and omitting the corresponding index. The parameters of the ringdown are calculated from the BBH parameters of the candidate events. We follow the same procedure as in Chapter 2 to find the remnant black hole’s mass M and spin χ . Assuming the component black holes to be non-spinning, the remnant’s properties are calculated from the component masses and spins through fitting formulae to numerical relativity. The remnant mass and spin in turn directly determine the frequencies and damping times of the QNMs. The modes’ initial phases are drawn randomly from a uniform distribution $\phi_{\ell m} \in [0, 2\pi)$, and the excitation amplitudes are calculated using the fitting formulae from [132] at the time $t = 10M$ after the merger. The injected signal then consists of two quasi-normal modes starting at this time.

We now perform Bayesian parameter estimation on these data. The signal model is the same as for the injection, consisting of the $\ell = m = 2$ and $\ell = m = 3$ QNMs at $t = 10M$ after the merger. However, we allow the subdominant mode’s frequency and damping time to deviate from the values predicted in GR from the dominant mode measurement. In contrast to Chapter 2, we do not restrict the analysis to events where we find evidence for the subdominant mode assuming the Kerr relation between the modes’ parameters. Instead, we directly apply the analysis testing for deviations from this relation to all candidate events. The combination of small preferences in the posteriors of many events may then yield constraints on these deviations, even if the two-mode model is not strongly preferred over a single mode for any individual event. Nevertheless, we still perform the additional analysis using only the dominant $\ell = m = 2$ mode. This allows to quantify the support for the two-mode model relative to the single-mode model, and to correlate this support with the measurability of the parameters.

In our parameter estimation, we adopt a set of 6 parameters for the waveform model with only the dominant mode, with 4 additional parameters when including two modes. The 4 parameters shared between the modes are (M, χ, ι, Ψ) , the same as those in Chapter 2. For each mode, we add the amplitude and phase parameter, $A_{\ell m}, \phi_{\ell m}$, resulting in 6 parameters in total for the single mode model. For the dominant $(\ell, m) = (2, 2)$ mode template, the frequency $f_{\ell m}$ and damping time $\tau_{\ell m}$ are the values calculated for the given M, χ from solving the Teukolsky equation. Adding the second mode introduces the corresponding $A_{\ell m}$ and $\phi_{\ell m}$, while deviations from the no-hair assumption are encoded by two additional parameters, $(\delta f_{33}, \delta \tau_{33})$. These describe the frequency and damping time of the subdominant $(\ell, m) = (3, 3)$ mode through $f_{33} = (1 + \delta f_{33})f_{33}^{\text{GR}}$ and $\tau_{33} = (1 + \delta \tau_{33})\tau_{33}^{\text{GR}}$, where f_{33}^{GR} and τ_{33}^{GR} are the values expected in GR for the given M, χ . All injections obey the GR prediction, so the correct values are $\delta f_{33} = \delta \tau_{33} = 0$. As we only consider deviations for the subdominant mode parameters, we drop the subscript in the following, writing only δf and $\delta \tau$.

Note that no information from the inspiral and merger parts is used to measure the QNMs, other than considering events’ start time, distance, and sky-location known. The values expected for the subdominant mode are determined entirely through the measurement of the dominant mode.

The waveform templates are the same damped sinusoids as in the deviation analysis in Chapter 2, described by the same four parameters $(f_{\ell m}, \tau_{\ell m}, A_{\ell m}, \phi_{\ell m})$. We only change the sampled

3 Hierarchical black hole spectroscopy

parameters from $(f_{\ell m}, \tau_{\ell m})$ to $(M, \chi, \delta f, \delta \tau)$. Mapping the latter to the damped sinusoid parameters results in identical waveforms for appropriate values in both sets of parameters.

The chosen priors for the QNM's frequency and damping time differ between the analyses using each set of parameters. We use uniform priors for the deviation parameters in the range $\delta f_{\ell m}, \delta \tau_{\ell m} \in [-0.6, 0.6]$, meaning a deviation of $\pm 60\%$ from the values expected in GR. For the remaining parameters, we choose the same uniform priors as in Chapter 2: BH final mass $M \in [10, 200]M_{\odot}$, BH final spin $\chi \in [-0.99, 0.99]$, log-amplitude of the fundamental mode $\log_{10}(A_{22}) \in [-4, 4]$, relative subdominant mode amplitude $\hat{A}_{\ell m} = A_{\ell m}/A_{220} \in [0, 0.5)$, ring-down phases $\phi_{\ell m} \in [0, 2\pi)$, polarisation angle $\Psi \in [0, 2\pi)$, and inclination angle $\cos(\iota) \in [-1, 1)$. Figure 3.1 shows a comparison of the resulting priors for $(f_{\ell m}, \tau_{\ell m})$ in this analysis and the one used in Chapter 2. The f_{220}, f_{330} prior is chosen as an example.




Figure 3.1: The prior probability density function used in the current analysis is shown in the left panel, and that from Chapter 2 in the right panel. The axes show the dominant and subdominant mode frequencies. A heat map shows the density function, with light colours representing high values. The 50% and 90% credible regions are marked by dotted and dashed white contours, respectively. The smaller top and right panels show the 1-dimensional marginalised distributions, with dashed lines marking the 90% credible interval and the median value, which are noted next to the corresponding plot.

Due to their construction, the priors on frequency and damping time of the subdominant mode are less restrictive than those in Chapter 2. Through the deviation parameters, f_{33} and τ_{33} may lie outside the ranges allowed by the M, χ -prior under the Kerr assumption. We removed such values in the previous analysis, but allow them here.

Our initial studies have considered both the Advanced LIGO design and the LIGO A+ detector sensitivities. These initial results for the Advanced LIGO design sensitivity yielded few events with appreciable constraints on the deviation parameters. Therefore we focus on the results obtained with the target sensitivities of LIGO A+ and Virgo for the O5 observation run [213]. Our predictions then apply once the detectors reach this target sensitivity when data is next collected. Note that we keep the population of candidate events unchanged, which is based on the Advanced LIGO and Virgo design sensitivities. The data are sampled at 8 kHz, instead of

2 kHz as used in the previous analysis.

We use kernel density estimation to compute the combination of multiple posterior distributions. The parameter estimation analysis produces a set of discrete samples, representing random draws from the posterior distribution. To compute the product of the distributions, we require knowledge of the probability density at a given point in parameter space. An approximate probability density function is calculated with the `gaussian_kde` function provided in the `scipy` software package [214], which uses Gaussian kernels to estimate the probability density function from the set of discrete samples. This estimate can then be evaluated at the desired points in parameter space. To find the product of the given distributions, we evaluate their estimates in a grid in parameter space and multiply the values of the different distributions at each grid point.

Here, we are interested in the combined posterior distribution for δf and $\delta\tau$, so we first marginalise the recovered posterior distribution over the remaining parameters. We then perform the 2-dimensional kernel density estimation and compute the product of the resulting approximate distributions, sampled in a 2-dimensional grid. Finally, we also find the 1-dimensional posterior distributions for δf and $\delta\tau$ by marginalising the 2-dimensional posterior over the opposite of the two parameters, respectively.

We verified the results of the kernel density estimate by a simple 2-dimensional binning of the original posterior samples. For each bin, we count the number of samples found for each distribution and then multiply these numbers. We find that the results from the binning agree with those based on the kernel density estimates.

3.3 Results

Results are quoted for the combination of events occurring during one year of data. For this, we randomly draw 50 events from the 200 candidates produced for 4 years of data. We provide a simple estimate for the frequency of occurrence for these results by averaging over many such random draws. This estimate is inaccurate as the repeated draws are made from the same 200 events, so the exact same posterior distributions will be considered when the same event appears in different draws. We therefore only use it as a simple guideline, which does not require us to extend the computationally expensive parameter estimation to larger numbers of injections.

In Figure 3.2, we show the result of combining all candidate events for one year of collected data. They represent the 50 events with the highest subdominant mode SNR expected to be detected in one year of data. The resulting posterior is centred around negative δf and $\delta\tau$, with the 1-dimensional 90%-credible intervals excluding the injected values of $\delta f = \delta\tau = 0$. Averaging over many draws of 50 events from the total of 200, we find that the injected value is excluded by the 90% credible interval on δf in about 60% of the cases. The posterior in Figure 3.2 is a typical example for the distributions found for different draws of events, and in almost all draws the median value for δf is negative.

We find that combining all events from one year of data leads to a biased result. A systematic bias appears in the recovery of both δf and $\delta\tau$, preferring negative values. A possible explanation for this bias in δf is the overlap of the prior for the subdominant mode with the frequency of the injected dominant mode. The dominant mode signal is present in the data and has a lower frequency than the subdominant mode. If the dominant mode is not matched perfectly by the corresponding part of the template, then a residual of the dominant mode signal remains in the data. This can then be matched by the subdominant mode template, preferring a negative δf to better fit the lower frequency content of the dominant mode signal.

We will study this bias and methods to counteract it through several additional analysis runs. The general setup of the analysis remains unchanged, but in each of the following tests we




Figure 3.2: The centre plot shows the posterior distribution for the combination of all candidate events from one year of simulated data. These are the 50 events which have the highest subdominant mode SNRs, where the data consist of ringdown signal injections into Gaussian noise. The axes are the fractional deviation of the measured subdominant mode parameters from the values expected in GR based on the measurement of the dominant mode. The posterior is shown as a heat map, with light colours representing high values, and the 90% and 50% credible regions are marked by dashed and dotted contours, respectively. As all injections obey GR, the correct value is $\delta f = \delta\tau = 0$. The distribution shows a bias towards negative δf and $\delta\tau$, corresponding to smaller f_{33} and τ_{33} than expected. The top and right panels show the one-dimensional marginalised distributions for δf and $\delta\tau$, respectively. Dashed lines denote the median value and 90% credible interval of the marginalised distribution. The injected value is excluded by the 90% credible interval for both parameters.




Figure 3.3: The same as Figure 3.3, but removing the simulated Gaussian noise, such that the data consist only of the injected signals. The observed bias persists, suggesting it is not due to unfavourable noise realisations. Instead, the bias prefers even lower values for the subdominant mode frequency, deviating further from the injected value.

modify the considered parameter space or the simulated data.

First, we perform the same analysis as before, but remove the simulated Gaussian noise, replacing it by zero strain amplitude. The data then consist only of the signal, while the likelihood calculation still uses the power spectral density of the noise curve for the considered detector. As shown in Figure 3.3, we find the same bias preferring negative δf and small damping times as in the previous analysis, implying it is not due to unfavourable realisations of the noise. Indeed, the bias is stronger in that even smaller frequencies are preferred and the true value is excluded by the 90% credible interval in about 85% of the trials representing one year of data. In Figure 3.4, the same combined posteriors are shown as in Figures 3.2 and 3.3, but we additionally mark the maximum likelihood values and maximum posterior density values found for each event included in the combination. In the presence of Gaussian noise, both the maximum likelihood and maximum posterior density values are distributed over the entire allowed parameter range for δf and $\delta\tau$. When removing the Gaussian noise, the maximum likelihood values found in the parameter estimation converge upon the correct value for δf , without a clear preference in $\delta\tau$. The maximum posterior density values instead cluster around the combined posterior's biased value for δf , while for $\delta\tau$ they again are distributed more broadly over the parameter range. This suggests that the bias appears more pronounced in the absence of noise due to being present in each individual event's posterior.

For all further tests, we again use injections into Gaussian noise. In our next test, we remove the subdominant mode from the injected signal, such that it consists only of the dominant mode. In this case, the bias is present, resulting in posterior distributions very similar to those already shown in 3.3. The bias is observed more frequently, with all one-year trials preferring negative values for δf and the 90% credible interval excluding $\delta f = 0$ in all cases. The median values for δf found in the different trials are clustered around $\delta f = -0.33$. This is close to the value of δf corresponding to the frequency of the dominant mode f_{22} , which is $\delta f \sim -0.37$ for all injections, such that $f_{22} \approx (1 + \delta f)f_{33}$. This relation is based on the values for frequencies and damping times as predicted by GR for the injected mass and spin parameters. We see that the frequency recovered for the subdominant mode approaches that of the dominant mode when only the dominant mode is injected, suggesting that the subdominant template partially matches to the dominant mode signal. This also seems a plausible cause for the bias observed when injecting both modes, as the bias in the posteriors is similar for single- and two-mode injections. The result matches our expectations, as the subdominant mode signal is quiet for many injections, and the signal is therefore similar to injection of only the dominant mode. Injecting no subdominant mode thus approximates many signals well, but emphasises the effect of low subdominant mode content in the signal.

These observations agree with the behaviour of the recovered waveforms compared to the injected signal. We consider again our initial analysis, with both injection and templates consisting of two modes. In Figure 3.5, we show the strain amplitude as a function of frequency for each mode of the injection and for a signal recovered in the analysis. The recovered signal is chosen from the highest density region of the posterior distribution. The dominant mode template is close to the injected dominant mode signal. The subdominant mode template however is shifted towards the dominant mode signal in both frequency and amplitude. This suggests that indeed the subdominant mode template partially matches to the dominant instead of the subdominant mode signal.

We can now test the performance of the analysis when artificially removing the identified cause of the bias. Compared to the initial analysis, we modify the template and injected signal to include only the subdominant mode. In the parameter estimation, we then fix the parameters of the dominant mode to those of the injection. While the dominant mode signal is neither part of the template nor injected into the data, its frequency and damping time are still required as




Figure 3.4: The left panel shows the same as Figure 3.2, the combined result for injections into Gaussian noise. The right is the same as Figure 3.3, simulated signals without noise. Additionally, the maximum likelihood value found in the analysis is marked for each event included in the combined posterior as a white cross. The maximum posterior density values for all included events are marked as empty green circles, and the same value for the combined posterior is shown by a filled green circle. When Gaussian noise is present, both the maximum likelihood and maximum posterior density values are scattered across the entire parameter ranges for δf and $\delta\tau$. When removing the Gaussian noise, the maximum likelihood values are clustered around the injected value for δf , while showing no clear preference for $\delta\tau$. The maximum posterior density values also show no preference in $\delta\tau$, but for δf are clustered around the biased result found in the combined posterior. This bias is more pronounced when no noise is present, both in terms of the position of the combined posterior, and in terms of the number of individual event's posteriors that show a systematically biased peak.




Figure 3.5: The strain amplitude of the injected and recovered signal waveforms are shown as a function of frequency for one example event. The recovered waveform is chosen from the highest density region of the posterior distribution. The two individual modes of each signal are shown, with the dominant mode marked by dark blue and the subdominant mode by light green lines. Dashed lines are used for the injected signal and unbroken lines for the recovered signal. A grey line shows the amplitude spectral density of the detector noise. We see that the dominant mode signal is approached by the dominant mode of the recovered template. The subdominant mode template, however, also approaches the dominant injected mode in frequency and amplitude compared to the injected subdominant mode signal. This suggests that as the dominant mode signal is not perfectly recovered, a residual of this signal remains in the data. This residual is then matched by the subdominant mode template instead of the injected subdominant mode signal.




Figure 3.6: The combined posterior distribution found when injecting only the subdominant mode into Gaussian noise and varying only the corresponding parameters. Again, the combined posterior for the 50 candidate events with the loudest subdominant modes in one year of data is shown, with the one-dimensional marginalised distributions in the top and right panels. While there is a bias towards shorter signals, the frequency deviation is well constrained to less than $\pm 10\%$, with the distribution centred on the correct value. An improved result is not surprising due to the reduced number of free parameters. However, we have also removed the suspected source of the bias, which likely contributes to the improvement.

3 Hierarchical black hole spectroscopy

a reference for the deviation parameters δf and $\delta\tau$ that are being varied for the subdominant mode. We thus fix the mass and spin to the values of the injection, with the free parameters now being δf , $\delta\tau$, A_{33} , ϕ_{33} , ι and Ψ . An example for the combined posterior distributions resulting from this analysis is shown in Figure 3.6. The posterior for δf is now centred on the correct value and well constrained, with the 1-dimensional 90% credible interval encompassing less than $\pm 10\%$ deviation. For $\delta\tau$, we still find a biased result preferring more strongly damped, shorter signals. This preference for short signals may be explained through the late part of the signal being sufficiently quiet to be indistinguishable from the noise background, with the template matching the first cycles of the sinusoid. For the frequency deviation, however, we find that the injected value is within the 90% credible interval in about 93% of the draws for the combined posterior, and the credible interval constrains δf to less than $\pm 20\%$ in about 89% of the draws. The reduced number of free parameters is expected to improve the performance of the analysis, so these results are likely overestimating the bounds that can be placed in a realistic case. Nevertheless, this suggests that the demonstrated bias is a main obstacle for the analysis.

We test a weaker restriction by injecting both modes of the signal, but fixing the dominant mode parameters to the injected values in the parameter estimation analysis. This means that the mass and spin are set to the values of the injected signal, as are the amplitude and phase of the dominant mode. The parameters of the subdominant mode as well as the polarisation and inclination angle are being varied. This restriction proves ineffective in removing the bias. Inaccurate measurement of the inclination angle and polarisation may leave a sufficient residual of the signal in the data to then be matched by the subdominant mode template.

Similarly, the bias is not effectively removed by modifying the prior in δf to exclude the region corresponding to the dominant mode frequency. We repeat the initial analysis with only this modification to the δf -prior. During sampling of the parameter space, we calculate the expected frequency of the dominant mode for the current sample point in mass and spin. We then remove an interval from the δf -prior which corresponds to frequencies of ± 10 Hz around the dominant mode frequency. While the individual events' posteriors show a clear bias less frequently, they now often prefer the values closest to the interval removed from the prior and thus seem to still be matching the dominant mode frequency as closely as allowed by the prior. The correct value for δf is now included by the 90% credible interval in about 63% of the draws of 50 events from the total 200 analysed.

We find an effective strategy to counteract the bias by limiting the analysis to a subset of detected events. We expect the bias to be stronger when the subdominant mode signal is not recovered. A natural choice then is to include only those events where the data actually favour the presence of two modes over that of only the dominant mode according to our analysis. We can calculate the Bayes factor comparing the evidence for the two-mode and single-mode models by performing two separate analyses assuming each of the models. Here, we use no information from the injections, but only the results of the data analysis, as would be the case for real detections with unknown properties. We now repeat the initial analysis, using once the single-mode and once the two-mode model. Only events that favour the presence of two modes over one with a Bayes factor $\mathcal{B} > 1$ are included in the combined posterior, and an example is shown in Figure 3.7. The distribution is centred close to the correct value in δf , with the 90% credible interval covering about $\pm 10\%$ or less compared to the correct value. Through random draws of one-year populations from the analysed events, we find that the correct value is within the 90% credible interval in 92% of the draws. The interval constrains δf to $\pm 20\%$ in 84% of the draws and to $\pm 10\%$ in 53% of the draws. The distribution for $\delta\tau$ is preferring negative values and covers a large portion of the prior interval, however, this is not unexpected as the damping time is typically measured with lower precision than the frequency. Concentrating on the frequency




Figure 3.7: The same as Figure 3.2, but including only those results in the combined posterior where the Bayes factor is larger than 1 in favour of the presence of two modes over only the dominant mode. We find that the distribution in δf is centred close to the correct value, with the one-dimensional 90% credible interval lying within the range of about $\pm 10\%$. The distribution in $\delta\tau$ still shows a bias toward shorter signals and covers a large part of the prior range.

measurements, the prospects to place 10%-level constraints on the deviation seem promising. This is achieved through the combination of posterior distributions from events selected for their evidence of a multimode ringdown spectrum.

3.4 Conclusions

We have found that stringent constraints on deviations from the no-hair theorem can be obtained by combining data from multiple gravitational wave detections, but care must be taken to avoid biased results. We discovered a bias in the recovery of the frequency deviation, which occurs when the data from many events with quiet subdominant modes are included in the combined posterior. This bias can be attributed to the overlap between the sets of waveforms that describe the dominant mode and the subdominant mode with a deviation to lower frequencies. As the frequency deviation prior allows the subdominant mode template to have frequencies close to that of the dominant mode, the subdominant mode template can match to the dominant mode signal. If the dominant mode is not perfectly recovered, a residual remains in the data that can be found by the subdominant mode template. We show that simply combining the data from all available events will then likely lead to a biased result.

However, the performance of the analysis can be significantly improved through a second step, which restricts the combination to suitable events. By performing the same analysis with a signal model containing only the dominant mode, we can calculate the Bayes factor to compare the evidence for the single- and the two-mode models. Including only events where the Bayes factor is larger than 1 in favour of the presence of two modes, we find that the analysis performs well for the LIGO A+ sensitivity. The correct value for δf is included by the 90% credible interval as expected, and a simple estimate suggests that bounds on the deviation can likely be placed to $\pm 20\%$ for a year of collected data, and to $\pm 10\%$ for a few years of data. For the level of precision of our results, our restriction sufficiently mitigates the bias. At higher precision, a bias may still be present, and a method to remove it desirable.

Note that our estimates are optimistic, as they are based on the optimistic merger rate from [63]. New estimates place the merger rate in the optimistic scenario at about half of that considered here, suggesting the required time of collected data to achieve these bounds to be larger by a factor of a few [212]. However, we defer detailed evaluation of the effects of updated population models, as well as inclusion of other types of binary systems, to future work.

We consider several venues for extension of this work promising, given the importance of multimode ringdown measurements for a number of applications. A first step would be extending our analysis to larger numbers of events to improve the robustness of the simple statistics we employed. More recent observational population models can be incorporated to find a more accurate relation between the duration of collected data and the obtainable constraints [64, 212]. Similarly, the evolution of detector sensitivities and the expected amount of data produced for each can be taken into account. In addition to injected signals obeying the relation expected in GR, injections with known deviations can be used to characterise the analysis.

Using methods presented in [2], we can consider a more realistic treatment of the signal before the start of our ringdown template. These methods allow to remove undesired times from the data without introducing artefacts into the likelihood calculation. Through this, we can use injection signals covering the entire ringdown or IMR phases, instead of abruptly starting at the time of the analysis. As our signal and template models are identical, we expect only minor influences from the abrupt start. However, characterising the more realistic case and comparing it to our results will help applying our predictions to analyses on real detector data. Finally, incorporating overtones into the analysis could yield a valuable comparison to its performance using angular modes as studied here. In particular, a similar hierarchical analysis has been

applied to detected events using overtones in [71, 72]. While this offers the advantage of extending the QNM description to times closer to the merger, where higher SNRs can be accessed, details of the applicability are under debate [128, 129]. As different overtones are very close in frequency, a similar bias may not be distinguishable from unbiased results at the sensitivity used in those analyses. Overtones are separated in their damping times, yet these are not well measured and the tests for deviations rely on frequency measurements. It would thus be interesting to characterise the effect of a potential bias also for an overtones-based analysis.

4 Black hole spectroscopy with GW190521

When two black holes merge, the late stage of gravitational wave emission is a superposition of exponentially damped sinusoids. According to the black hole no-hair theorem, this ringdown spectrum depends only on the mass and angular momentum of the final black hole. An observation of more than one ringdown mode can test this fundamental prediction of general relativity. Here we provide strong observational evidence for a multimode black hole ringdown spectrum using the gravitational wave event GW190521, with a Bayes factor of ~ 40 preferring two fundamental modes over one. The dominant mode is the $\ell = m = 2$ harmonic, and the sub-dominant mode corresponds to the $\ell = m = 3$ harmonic. We estimate the redshifted mass and dimensionless spin of the final black hole as $330_{-40}^{+30} M_{\odot}$ and $0.87_{-0.10}^{+0.05}$, respectively. The detection of the two modes disfavors a binary progenitor with equal masses; the mass ratio is constrained to $0.4_{-0.3}^{+0.2}$. We find that the final black hole is consistent with the no hair theorem and constrain the fractional deviation from general relativity of the sub-dominant mode’s frequency to be $-0.01_{-0.11}^{+0.07}$ ¹.

4.1 Introduction

A perturbed black hole approaches equilibrium by emitting a spectrum of damped sinusoidal gravitational-wave signals [69, 180, 215]. Unlike other astrophysical objects, the ringdown spectrum of a black hole is remarkably simple. General relativity predicts that the frequencies and damping times of the entire spectrum of damped sinusoids, or “quasi-normal modes”, are fully determined by just two numbers: the black hole mass M and angular momentum J , as described by the Kerr solution [16]. This prediction, a consequence of the black hole “no-hair theorem”, does not hold in many alternate theories [122]. If astrophysical black holes are observed to violate this property, it indicates new physics beyond standard general relativity.

In order to observationally test this prediction using binary black hole mergers, an important observational challenge must be met: at least two ringdown modes must be observed [133]. The higher the binary mass ratio asymmetry, the more likely it is that sub-dominant ringdown modes are observable. However, more asymmetric binary systems are less likely to be formed, and also lead to weaker signals. Population studies suggested that such multimode ringdown modes were unlikely to be observed until the next generation of gravitational-wave observatories [1, 185], since black-hole population models did not anticipate observations of massive, asymmetric binaries.

Here we confound this expectation with the gravitational-wave event GW190521, detected by the two LIGO detectors and Virgo at 03:02:29 UTC on May 21st, 2019 [216, 217]. This is the heaviest black-hole merger event observed to date [167, 169]. The signal is consistent with the merger of two high mass black holes which merge at a low frequency relative to the detector sensitivity band. As such, it has a barely observable inspiral and the signal is dominated by the merger and ringdown phase.

GW190521 was initially reported as the merger of two comparable mass black holes [216, 217], in which case one would not expect to detect sub-dominant ringdown modes. Subsequent re-analysis of the data suggested the possibility that the progenitor could be an intermediate

¹This chapter is an adaptation of the work in the pre-print [2] as described in Section 1.2.

mass-ratio binary [218], suggesting the possibility of detectable sub-dominant modes. Here we find strong evidence for multimode damped sinusoids in the ringdown phase of the gravitational wave event GW190521.

4.2 Multimode agnostic search

A quasi-normal mode description of the gravitational wave from a binary black hole is not expected to be valid until after the binary has merged to form a perturbed black hole. On the flip side, the damping time of an $O(100 M_\odot)$ black hole is $O(10 \text{ ms})$, leaving a window of only a few tens of milliseconds after merger in which the ringdown is detectable above noise. Accurate identification of the merger time is therefore crucial to extract quasi-normal modes from the data. To account for uncertainty in the merger time of GW190521 due to modelling systematics, we perform a series of analyses in short time increments starting at a geocentric GPS reference time $t_{\text{ref}} = 1242442967.445$. This time is taken from the maximum likelihood merger time obtained via the analysis in Nitz & Capano [218]. We also fix the sky location to the maximum likelihood values from the same analysis.

The ringdown spectrum of a Kerr black hole consists of an infinite set of frequencies $f_{\ell mn}$ and damping times $\tau_{\ell mn}$ labeled by three integers (ℓ, m, n) . Here ℓ and m are the usual angular harmonic numbers. The third index $n \geq 0$ denotes overtones, with $n = 0$ being the fundamental mode. The most agnostic way to search for quasi-normal modes from a perturbed black hole is to search for them individually, without assuming any relation between them. Such a search is complicated by the nature of quasi-normal modes: they are not orthogonal, meaning that modes that overlap in time must be sufficiently separated in frequency or damping time in order to be distinguishable. Simulations of binary black hole mergers have shown that the fundamental $\ell = m = 2$ mode is typically significantly louder than other modes. In order to extract sub-dominant modes from noisy data in an agnostic search it is useful to separate the dominant mode in frequency from the others.

A visual inspection of the time- and frequency-domain data taken at the reference time revealed significant power in the two LIGO detectors between 60–70 Hz (see Supplemental Fig. 4.5). In order to isolate this and search for sub-dominant modes we constructed three frequency ranges: “range A” between 50–80 Hz, “range B” between 80–256 Hz, and “range C” between 15–50 Hz. We search for one quasi-normal mode in each range using Bayesian inference. We use uniform priors on the relative amplitudes of the modes in range B and C between 0 and 0.9 times the mode in range A. No other relation is assumed between the modes.

We repeat this analysis at time steps of $t_{\text{ref}} + 0, 6, 12, 18,$ and 24 ms . As expected from the visual inspection of the data, we find a significant mode in range A at all grid points, which decreases in amplitude at later times. A clear second mode is found in range B. This mode is most visible at $t_{\text{ref}} + 6 \text{ ms}$, the result of which is shown in Fig. 4.1 (results at other times are shown in Supplemental Fig. 4.6). The frequency of the secondary mode at this time is $\sim 98 \text{ Hz}$ with a damping time of $\sim 30 \text{ ms}$, while the primary mode has frequency of $\sim 63 \text{ Hz}$ and damping time $\sim 26 \text{ ms}$. The signal-to-noise ratio of the primary and secondary modes of the maximum likelihood waveform is 12.2 and 4.1, respectively. Results from range C (not shown) are consistent with noise.

The dominant mode found at $\sim 63 \text{ Hz}$ is expected to be the quadrupolar $\ell = m = 2, n = 0$ fundamental mode. Measurement of f_{220} and τ_{220} provides an estimate of the mass and angular momentum of the remnant black hole [141]. This in turn predicts the entire ringdown spectrum of subdominant modes. Figure 4.1 shows that the subdominant mode at $\sim 98 \text{ Hz}$ is consistent with the $\ell = m = 3, n = 0$ mode. This is also in agreement with expectations from numerical simulations of binary black hole mergers [126, 132].




Figure 4.1: Marginal posterior probability distributions on frequency and damping time from an agnostic quasi-normal mode analysis of GW190521 at 6ms after t_{ref} . A single mode is searched for in each of the shown frequency ranges, range from A 50 to 80Hz and range B from 80 to 256Hz. Top panels show the marginal posterior on the mode frequencies, with priors indicated by dotted lines; the gray region in the top left panel shows the highest 90th percentile density interval of the dominant mode. White dotted (dashed) contours in the bottom panels show the 50th (90th) credible regions. Assuming the dominant mode in range A corresponds to the (220) mode of a Kerr black hole, we estimate what the frequency and damping times would be of the (330), (440), and (550) modes (blue, green, and red regions, respectively). The mode in range B is clearly consistent with the expected frequency and damping time of the (330) mode. Here we do not see the (440) and (550) modes, indicating they are weaker than the (330) mode. This is consistent with an asymmetric binary black hole merger.

4.3 Consistency with the Kerr solution

The search for damped sinusoids in section 4.2 assumed no particular relation between different modes, with a corresponding large prior parameter volume. In this section, we assume that the frequency and damping times of the damped sinusoids are related as in the ringdown of a Kerr black hole. This reduces the prior parameter volume and focuses in on particular modes. The amplitudes and phases of the modes are left as free parameters, since they depend on the specific initial state of the remnant black hole immediately after the merger.

For this analysis, we model the ringdown signal based on the final Kerr black hole mass, M_f , and dimensionless spin, $\chi_f = J_f/M_f^2$. We expect only a subset of the entire spectrum of quasi-normal modes to be visible above noise. Including all possible modes in our signal model can lead to overfitting the data. For this reason we perform several analyses which include different combinations of the (330), (440), (210), and (550) modes, in addition to the dominant (220) mode. Numerical simulations of binary black hole mergers have generally shown these modes to be the strongest [132]. Giesler et al. [128] found that overtones of the dominant harmonic are significant close to the merger time, and that including them allows a quasi-normal mode description of the signal to be used at earlier times. We therefore also perform analyses in which we include the first overtone of the dominant harmonic $(\ell mn) = (221)$.

We repeat these analyses in 1 ms intervals between $t_{\text{ref}} + [-9 \text{ ms}, 24 \text{ ms}]$. We use Bayes factors to determine which model is most favoured at each time step. For the model that includes the fundamental dominant harmonic (220) and its overtone (221), the Bayes factor is evaluated against the model with only the (220) mode. For models that include the (330) (or other subdominant modes), the Bayes factor is evaluated against the stronger of the (220) or (220)+(221) models.

The Bayes factors for the various multimode Kerr models are shown in Fig. 4.2. Consistent with the agnostic results, we find strong evidence for the presence of the (330) mode around 6 ms, with the Bayes factor for the (220) + (330) model peaking at 44_{-5}^{+6} one millisecond later, at $t_{\text{ref}} + 7$ ms. The maximum likelihood ringdown waveforms at this time are shown in Supplemental Fig. 4.7. Only moderate evidence is found for other fundamental modes, although at 19 ms after t_{ref} the most favoured model contains the (220), (330), (440) and (210) modes. Extending the analysis to earlier times, we find increasing support for the presence of the (221) overtone (see Supplemental Fig. 4.8), with the Bayes factor reaching a maximum at $t_{\text{ref}} - 5$ ms. This would be consistent with the merger happening prior to t_{ref} , in agreement with a recent reanalysis of GW190521 [219].

Figure 4.3 shows the redshifted mass and the dimensionless spin of the final black hole, measured with the (220) + (330) Kerr model at 7 ms after t_{ref} . We find that the remnant black hole has a redshifted mass $(1+z)M_f = 330_{-40}^{+30} M_{\odot}$ and dimensionless spin $\chi_f = 0.87_{-0.10}^{+0.05}$.

If a quasi-normal model without overtones is used too close to merger, the resulting final mass estimate can be biased toward larger values [69, 128]. Supplemental Fig. 4.8 shows the stability of the final mass estimate between 6 and 12 ms using the (220) + (330) model, and its agreement with the mass estimate at earlier times using the (221) overtone. This indicates that by this time the black hole has reached a regime of constant ringdown frequency, a requirement for the validity of linear-regime, quasi-normal modes.

Given the strong evidence for the presence of the (330) mode at $t_{\text{ref}} + 7$ ms, we can perform the classic no-hair theorem test [133]. We keep the dependence of f_{220} and τ_{220} on (M_f, χ_f) as in the Kerr solution but introduce fractional deviations δf_{330} and $\delta \tau_{330}$ of f_{330} and τ_{330} , respectively. Figure 4.3 shows the Kerr black hole mass M_f and dimensionless spin χ_f associated to the (330) mode frequency $f_{330}(1 + \delta f_{330})$ and damping time $\tau_{330}(1 + \delta \tau_{330})$ measured at 7 ms after t_{ref} . Posterior distributions on the fractional deviations are shown in Supplemental Fig. 4.9. We




Figure 4.2: Bayes factor of models with the indicated modes compared to the stronger of the (220) or the (220) + (221) modes model. The Bayes factor for the (220) + (221) model is calculated against the (220)-only model. The hexagon marks where no-hair tests are performed in section 4.3.




Figure 4.3: Posterior distribution of final redshifted mass $(1+z)M_f$ and dimensionless spin χ_f measured at 7 ms after t_{ref} assuming the identified modes are the (220) and (330) modes of a Kerr black hole. Dashed lines indicate the 90% credible interval. For the Kerr with $\delta(330)$ results, we use fitting formulae [141] to convert the frequency $f_{330}(1 + \delta f_{330})$ and damping time $\tau_{330}(1 + \delta\tau_{330})$ into mass and spin.

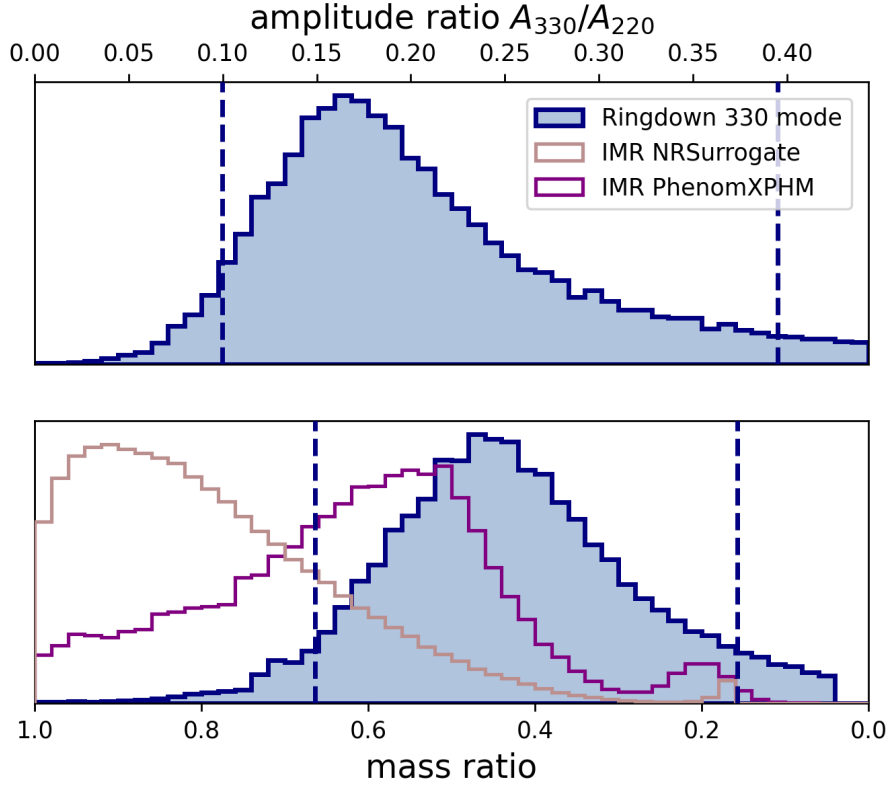


Figure 4.4: Posterior distribution for (*top*) the amplitude ratio of the (330) mode, A_{330}/A_{220} , and (*bottom*) the mass ratio of the binary, $m_2/m_1 < 1$, obtained using numerical fits between the (330) amplitude and mass ratio [132]. Vertical dashed lines indicate the 90% credible interval. In the bottom panel we assume a prior uniform in mass ratio between $1/25 - 1$. For comparison, we also show the mass ratio obtained from the full signal using NR Surrogate [218] and PhenomXPHM models [169] assuming the same prior.

constrain the fractional deviation from Kerr to $\delta f_{330} = -0.01^{+0.07}_{-0.11}$. The damping time is only weakly constrained, with $\delta\tau_{330} = 0.6^{+2.0}_{-1.1}$.

4.4 Discussion and conclusions

The detection of a (330) mode indicates that the progenitor black holes in GW190521 had asymmetric masses, since equal-mass binaries are not expected to excite the (330) mode. Numerical fits [132] provide the relation between the amplitude of the (330) mode and the ratio of the initial black hole masses m_1 and m_2 for quasi-circular, aligned spin binaries. We find $m_2/m_1 = 0.4^{+0.2}_{-0.3}$ from the (330) amplitude measured by the (220) + (330) Kerr model at 7 ms after t_{ref} . Figure 4.4 shows the posterior distribution on the amplitude of the (330) mode and on the corresponding mass ratio.

The redshifted final mass of GW190521 measured by the LVC [167] using a (220) ringdown fit is $(1+z)M_f = 282.2^{+50.0}_{-61.9}M_\odot$, or $259.2^{+36.6}_{-29.0}M_\odot$ when analysing the full signal. The low-mass-ratio part of the posterior of Nitz and Capano [218] found $(1+z)M_f \sim 260M_\odot$ using the full signal [220, 221]. These results are somewhat in tension with the final mass and spin inferred from the ringdown modes found here.

However, the complete waveform models used in the above analyses may not include all relevant physical effects. This, coupled with the fact that GW190521 has a very short inspiral signal, can lead to systematic errors for parameter estimation. For example, the waveform models used in the LVC analysis and Nitz & Capano assume quasi-circular orbits, but several studies have indicated that the binary may have been eccentric at merger [222–224]. These studies have also found slightly larger estimates for the binary’s total mass, making them more consistent with our estimate for the final mass. The ringdown waveforms used in this paper are simpler and more robust than full inspiral-merger-ringdown models for signals like GW190521, provided they are applied sufficiently late in the post-merger regime. This argument would tend to favour the estimates derived in this paper for the total mass. Nevertheless, a full resolution of this tension is beyond the scope of this work.

Evidence for overtones of the (220) mode very close to merger were previously found for the events GW150914 [172] and GW190521.074359 [71] (not to be confused with GW190521). Black hole spectroscopy tests showed consistency with the Kerr hypothesis for these events [71, 172]. However, the resulting constraints were weaker than what we find with the (330) mode here. Furthermore, while there is strong numerical evidence for the presence of ringdown overtones close to the merger [128], a number of theoretical questions remain as to the validity of a quasi-normal description of the black hole close to merger [225–229].

Given the evidence we find for the (221) model at $t_{\text{ref}} - 5$ ms, we also perform a no-hair theorem test on the (221) overtone at this time. The results are shown in Supplemental Fig. 4.10. We find poor constraints, with $\delta f_{221} = -0.11^{+0.33}_{-0.04}$ and $\delta \tau_{221} = 0.45^{+0.31}_{-0.66}$. This highlights the benefits of black hole spectroscopy using the fundamental modes.

The true nature of the gravitational wave event GW190521 has been the subject of much speculation [230–232]. The interpretation of GW190521 as a head-on collision of two highly spinning Proca stars [232] predicts the presence of a (200) mode [233]. We do not find evidence for such a mode. Additionally, the high-mass, multiple-mode ringdown signal observed here does not agree with the scenario of a very massive star collapsing to a black hole of mass $\sim 50M_{\odot}$ and an unstable massive disk [234].

Expectations based on population models were that black hole ringdown signals with multiple modes were unlikely to be observed with the Advanced LIGO and Virgo detectors [1, 185] (although those population models did not include massive binaries). Our results here show that, remarkably, GW190521 displays a distinct subdominant mode and that this mode is consistent with the ringdown of a Kerr black hole.

Acknowledgements

The authors thank Ofek Birnholtz, José Luis Jaramillo, Reinhard Prix, Bruce Allen, Evan Goetz, Saul Teukolsky, Maximiliano Isi, Juan Calderón-Bustillo, Abhay Ashtekar and Bangalore Sathyaprakash for useful discussions. We thank also the Atlas Computational Cluster team at the Albert Einstein Institute in Hanover for assistance. MC acknowledges funding from the Natural Sciences and Engineering Research Council of Canada (NSERC). This research has made use of data obtained from the Gravitational Wave Open Science Center (<https://www.gwopenscience.org/>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) who also gratefully acknowledge the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational

Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain.

4.5 Supplemental

4.6 Materials and Methods

4.6.1 The ringdown signal model

In the quasi-normal mode (QNM) spectrum of a perturbed Kerr black hole, the allowed frequencies $f_{\ell mn} = \omega_{\ell mn}/(2\pi)$ and damping times $\tau_{\ell mn}$ are labeled by three integers $\ell = 2, 3, \dots$, $-\ell \leq m \leq \ell$, and $n = 0, 1, 2, \dots$. These can be combined together in a complex frequency $\Omega_{\ell mn} = \omega_{\ell mn} + i/\tau_{\ell mn}$ such that the ringdown signal of a perturbed Kerr black hole can be expressed as a sum of damped sinusoids:

$$h_+ + ih_\times = \frac{M_f}{D_L} \sum_{\ell mn} {}_{-2}S_{\ell mn}(\iota, \varphi) A_{\ell mn} e^{i(\Omega_{\ell mn} t + \phi_{\ell mn})}, \quad (4.1)$$

where h_+ and h_\times are the plus and cross polarisations of the gravitational wave, M_f is the mass of the black hole in the detector frame and D_L is the luminosity distance to the source. The functions ${}_{-2}S_{\ell mn}(\iota, \varphi)$ are the spin-weighted spheroidal harmonics of spin weight -2 , which depend on the inclination angle ι between the black hole spin and the line-of-sight from the source to the observer, and the azimuth angle φ between the black hole and the observer.

The complex QNM frequencies $\Omega_{\ell mn}$ can be determined from the Teukolsky equation [53, 131]. According to the no-hair theorem, the frequencies and damping times are determined by the mass M_f and spin χ_f of the black hole, with $\chi_f \in (-1, 1)$. Positive (negative) spin means the perturbation is co(counter)-rotating with respect to the black hole. The amplitudes $A_{\ell mn}$ and phases $\phi_{\ell mn}$ depend on the initial perturbation and take different values for different ℓmn modes.

For a given ℓ and n , the $+m$ and $-m$ modes are related to each other by $\omega_{\ell -mn} = -\omega_{\ell mn}$ and $\tau_{\ell -mn} = \tau_{\ell mn}$ [141]. Furthermore, if the initial perturbation is symmetric under reflection at the equatorial plane, the amplitude and phase of the $\pm m$ modes are related to each other by $A_{\ell -mn} e^{i\phi_{\ell -mn}} = (-1)^\ell A_{\ell mn} e^{-i\phi_{\ell mn}}$. Such a symmetry may be expected in the case of non-precessing binaries [130, 137]. To simultaneously sum over the $\pm m$ modes for a given ℓn , we parameterise the waveform as

$$h_{\ell|m|n} = A_{\ell|m|n}^0 e^{-t/\tau_{\ell|m|n}} \times \\ [{}_{-2}S_{\ell mn}(\iota, \varphi) A_{\ell|m|n}^{(+)} e^{i(\omega_{\ell|m|n} t + \phi_{\ell|m|n})} \\ + {}_{-2}S_{\ell -mn}(\iota, \varphi) A_{\ell|m|n}^{(-)} e^{-i(\omega_{\ell|m|n} t + \phi_{\ell|m|n})}],$$

where $A_{\ell|m|n}^0$ is the intrinsic amplitude of the (ℓmn) mode and

$$A_{\ell|m|n}^{(+)} = \sqrt{2} \cos(\pi/4 + \Delta\beta_{\ell|m|n}) \quad (4.2)$$

$$A_{\ell|m|n}^{(-)} = \sqrt{2} \sin(\pi/4 + \Delta\beta_{\ell|m|n}) e^{i(l\pi + \Delta\phi_{\ell|m|n})}. \quad (4.3)$$

If the parameters $\Delta\beta_{\ell mn}$ and $\Delta\phi_{\ell mn}$ are both zero, the waveform reduces to the case of reflection symmetry.

When using the Kerr model, we do two analyses at several times, one in which $\Delta\beta_{\ell mn}$ and $\Delta\phi_{\ell mn}$ are both set to zero, enforcing the symmetry, and one in which a common $\Delta\beta$ and $\Delta\phi$ for all modes are allowed to vary uniformly between $[-\pi/4, \pi/4)$ and $[-\pi, \pi)$, respectively. In all cases we find that the reflection symmetric analysis is favoured over the analysis without the symmetry; we therefore only report results from the former here. In all analyses we fix the azimuthal angle $\varphi = 0$, as it is degenerate with the modes' initial phases. To obtain the frequency and damping times for a given mass and spin we use tabulated values from Berti et al. [141], which we interpolate using a cubic spline. For the spheroidal harmonics we use tabulated values of the angular separation constants (also from Berti et al. [141]) and solve the recursion formula given in Leaver [131]. Our code for doing this is publicly available on GitHub [123].

For the agnostic analysis, we do not assume any mode corresponds to any particular ℓmn . We therefore use arbitrary complex numbers $X_{\ell\pm mn} = e^{i\psi_{\ell\pm mn}}$ in place of the ${}_2S_{\ell\pm mn}$. Here, the $\psi_{\ell\pm mn}$ are allowed to vary uniformly in $[0, 2\pi)$. We vary a common $\Delta\beta$ parameter, but fix the $\Delta\phi$ parameter to zero, as it is degenerate with the $X_{\ell mn}$.

4.6.2 Computational analysis methods

Standard parameter estimation with gravitational waves begins with Bayes' theorem. Given some data \mathbf{s} and a signal model h that depends on some set of parameters λ , we wish to know the posterior probability density function $p(\lambda|\mathbf{s}, h)$. Applying Bayes' theorem we have

$$p(\lambda|\mathbf{s}, h) = \frac{1}{Z} p(\mathbf{s}|\lambda, h) p(\lambda|h),$$

where $p(\mathbf{s}|\lambda, h)$ is the likelihood function, $p(\lambda|h)$ is the prior, and Z is a normalisation constant known as the evidence. Estimates on a single parameter are obtained by marginalising the posterior over all other parameters; marginalising over all parameters yields the evidence. Taking the ratio of evidences Z_A/Z_B for two different signal models yields the ‘‘Bayes factor’’. If our prior belief for the validity of the two models is the same, the Bayes factor gives the odds that model A is favoured over model B. Using a scale by Kass and Raftery [119], a Bayes factor greater than 3.2 is considered ‘‘substantial’’ evidence in favour of model A; greater than 10 is ‘‘strong’’ evidence; greater than 100 is ‘‘decisive’’.

Evaluating the posterior requires a likelihood function $p(\mathbf{s}|\lambda, h)$. Consider a gravitational-wave detector, which we sample every Δt seconds over a time T to obtain $N = \lceil T/\Delta t \rceil$ time-ordered samples $\mathbf{s} = \{s_0, \dots, s_{N-1}\}$. A network of K detectors sampled in this way will produce a set of samples $\mathbf{s}_{\text{net}} = \{\mathbf{s}_1, \dots, \mathbf{s}_K\}$. To obtain a likelihood function we first consider the hypothesis that the set of samples only contains noise $p(\mathbf{s}_{\text{net}}|n) = p(\mathbf{n}_{\text{net}})$.

In gravitational-wave astronomy it is common to assume that, in the absence of a signal, the detectors output stochastic Gaussian noise that has zero mean and is independent across detectors. Under this assumption the probability density function describing the network of time-ordered noise samples \mathbf{n}_{net} is a product of K N -dimensional multivariate normal distributions,

$$p(\mathbf{n}_{\text{net}}) = \frac{\exp\left[-\frac{1}{2} \sum_{d=1}^K \mathbf{n}_d^\top \mathbf{C}_d^{-1} \mathbf{n}_d\right]}{\sqrt{(2\pi)^{NK} \prod_{d=1}^K \det \mathbf{C}_d}}. \quad (4.4)$$

Here, \mathbf{C}_d is the covariance matrix of the noise in detector d . In order to evaluate this function it is necessary to know what the inverse of the covariance matrix is.

If we assume that a detector's noise is wide-sense stationary and ergodic, then its covariance \mathbf{C} is a symmetric Toeplitz matrix with elements given by the autocorrelation function of the data. If the autocorrelation function goes to zero in some finite amount of time that is less than

$T/2$ (for the LIGO and Virgo detectors, this typically happens within a few seconds), then the eigenvectors of the covariance matrix can be well-approximated by that of a circulant matrix. All circulant matrices have the same eigenvectors, $e^{-2\pi i k p/N}/\sqrt{N}$ [111], where $k = 0, \dots, N-1$ enumerates the eigenvectors and $p = 0, \dots, N-1$ the eigenvectors' components. Solving for the eigenvalues yields an analytic expression for \mathbf{C}^{-1} : the j, k -th element is the discrete inverse Fourier transform of $1/S_n$ evaluated at the $k-j$ time step,

$$C^{-1}[j, k] \approx 2\Delta t^2 \mathcal{F}^{-1}(S_n^{-1})[k-j], \quad (4.5)$$

where S_n is the power spectral density of the detector's noise. Substituting this back into Eq. (4.4), yields a canonical likelihood function for the noise [100],

$$p(\mathbf{s}_{\text{net}}|n) \propto \exp \left[-\frac{1}{2} \sum_{d=1}^K \langle \mathbf{n}_d | \mathbf{n}_d \rangle \right]. \quad (4.6)$$

Here, the inner product $\langle \cdot | \cdot \rangle$ is defined as

$$\langle \mathbf{u}_d | \mathbf{v}_d \rangle \equiv 4\Re \left\{ \frac{1}{T} \sum_{p=1}^{\lfloor (N-1)/2 \rfloor} \frac{\tilde{u}_d^*[p] \tilde{v}_d[p]}{S_n^{(d)}[p]} \right\}, \quad (4.7)$$

where \tilde{u} is the discrete Fourier transform of the time series \mathbf{u} and $*$ means complex conjugation².

The signal hypothesis is that the data consists of the signal plus the noise. The likelihood function $p(\mathbf{s}|\lambda, h)$ is therefore Eq. (4.6) with the \mathbf{n}_d replaced by the residuals $\mathbf{s}_d - \mathbf{h}_d$. However, this assumes that \mathbf{h} is an accurate model of the signal across the entire observation time T . As stated above, quasi-normal modes only model the gravitational wave from a binary black hole after the merger, when the two component black holes have formed a single, perturbed black hole. Performing Bayesian inference using quasi-normal modes as the signal model therefore requires excising times from the data when the ringdown prescription is not valid. In other words, instead of considering the full set of time samples $\mathbf{s} = \{s_0, \dots, s_{N-1}\}$, we wish to only evaluate the truncated set $\mathbf{s}_{tr} = \{s_0, \dots, s_a, s_{a+M}, \dots, s_{N-1}\}$, with $M > 1$. The data between the time steps $[a, a+M)$ is said to be ‘‘gated’’.

The probability density function of the truncated noise is still a multivariate normal distribution (excising dimensions from a multivariate normal is equivalent to marginalising over those dimensions), and so Eq. (4.4) still applies. The challenge is that the covariance matrix of the truncated noise \mathbf{C}_{tr} is no longer Toeplitz. Its eigenvectors can no longer be approximated by that of a circulant matrix, and so the expression for the likelihood Eq. (4.6) is no longer valid. The inverse of the covariance matrix needs to be found by other means. One possibility is to numerically invert the covariance matrix. However, this is numerically unstable due to the large dynamic range of the matrix's elements, and computationally impractical for observation times of more than about one second. Instead, we use ‘‘gating and in-painting’’ to find the likelihood of the truncated time series. This was applied to the problem of matched filtering in Zackay et al. [235]. Here we apply it to parameter estimation.

Define $\mathbf{n}' = \mathbf{n}_g + \mathbf{x}$, where \mathbf{n}_g is the noise with the gated times $t \in (a, a+M)\Delta t$ zeroed out, and \mathbf{x} is a vector that is zero everywhere except in the gated times. If the non-zero elements of \mathbf{x} are such that $(\mathbf{C}^{-1}\mathbf{n}')[k] = 0$ for all $k \in (a, a+M)$, then $\mathbf{n}'^T \mathbf{C}^{-1} \mathbf{n}'$ will be the same as the truncated version $\mathbf{n}_{tr}^T \mathbf{C}_{tr}^{-1} \mathbf{n}_{tr}$. Our aim is to solve the equation $\mathbf{C}^{-1}(\mathbf{n}_g + \mathbf{x}) = \mathbf{0}$ in the gated region. Since \mathbf{x} is zero outside of the gated region, $\mathbf{C}^{-1}\mathbf{x}$ only involves the $(a, a+M)$ rows and

²As noted in Section 1.5.1, this definition excludes the DC-component $p=0$ for odd N , and both DC-component and Nyquist-frequency component $p=N/2$ for even N , which can be treated separately or may be negligible [106].

columns of \mathbf{C}^{-1} , which form an $M \times M$ Toeplitz matrix [cf. Eq. (4.5)]. We therefore solve for \mathbf{x} such that

$$\overline{\mathbf{C}^{-1} \mathbf{x}} = -\overline{\mathbf{C}^{-1} \mathbf{n}_g}, \quad (4.8)$$

where the overbar indicates the $(a, a+M)$ rows (and columns) of the given vector (matrix). This can be solved numerically using a Toeplitz solver [214, 236]. Adding \mathbf{x} to the gated noise (“in-painting”) will then yield the same result as if we had truncated the noise and the covariance matrix.

Note that if the gate spans the entire beginning of the data segment, the truncated covariance matrix \mathbf{C}_{tr} is Toeplitz, and so could be inverted numerically using a Toeplitz solver. This is the method used by Isi et al. [172]. The advantage of using in-painting is that it involves solving for an $M \times M$ matrix rather than an $(N - M) \times (N - M)$ matrix. Gating and in-painting also have other applications beyond what we use it for here, such as excising glitches from data when doing parameter estimation.

To evaluate the likelihood for a signal, we use $\mathbf{n}_g = \mathbf{s}_g - \mathbf{h}_g$ (i.e., the residual with the gated region zeroed out) in Eq. (4.8) and solve for \mathbf{x} . We can then use $\mathbf{x} + \mathbf{s}_g - \mathbf{h}_g$ in the standard likelihood, Eq. (4.6). We do not attempt to normalise the likelihood, which would involve finding the determinant of the truncated covariance matrix. For this reason, we calculate and report Bayes factors by comparing models that start at the same time offset from t_{ref} , for which the determinant cancels.

We use the open source PyCBC Inference library for performing Bayesian inference [99, 108], to which we have added the gated likelihood described above. For all analyses we use a gate of two seconds, ending at the start time of the ringdown. For sampling the parameter space we use the dynesty nested sampler [116]. We use data for the event GW190521 made publicly available by the Gravitational Wave Open Science Center [86]. We fix the sky location to the values given by the maximum likelihood result of Nitz & Capano [218], although we have obtained similar results using the LVC’s maximum likelihood sky location [217]. We use a geocentric GPS reference time of $t_{\text{ref}} = 1242442967.445$ [218]. With the sky location used in our analyses, this corresponds to the detector GPS reference times $1242442967 + 0.4259$ at LIGO Hanford, $+0.4243$ at LIGO Livingston and $+0.4361$ at Virgo. Credible intervals in the text are quoted to 90%.

Data availability

Posterior data samples and data necessary to reproduce the figures are available at <https://github.com/gwastro/BH-Spectroscopy-GW190521>. The gravitational-wave data used in this work were obtained from the Gravitational Wave Open Science Center (GWOSC) at <https://www.gw-openscience.org>.

Code availability

All software used in this analysis is open source. Bayesian inference was performed with the PyCBC library, available at <https://github.com/gwastro/pycbc>. Configuration files used to perform all analyses can be found at <https://github.com/gwastro/BH-Spectroscopy-GW190521>. Spheroidal harmonics, Kerr frequencies, and Kerr damping times were generated using pykerr, available at <https://github.com/cdcapano/pykerr>.




Figure 4.5: *Top*: Whitened data in each detector. Dark traces show the data gated at the reference time t_{ref} in each detector, which is indicated by the gray vertical lines. *Bottom*: Frequency domain representation of the Hanford and Livingston data shown in the top panel. Also shown is the frequency domain representation at an off-source time, one second later. The signal is clearly visible in the LIGO time-domain data, and is seen as a spike in the frequency domain data between $\sim 60 - 70$ Hz. The primary frequency bin (“A”) boundaries were set to isolate this spike. Frequencies below (region “C”) and above (region “B”) were searched for additional QNMs in the agnostic analysis.

4 Black hole spectroscopy with GW190521

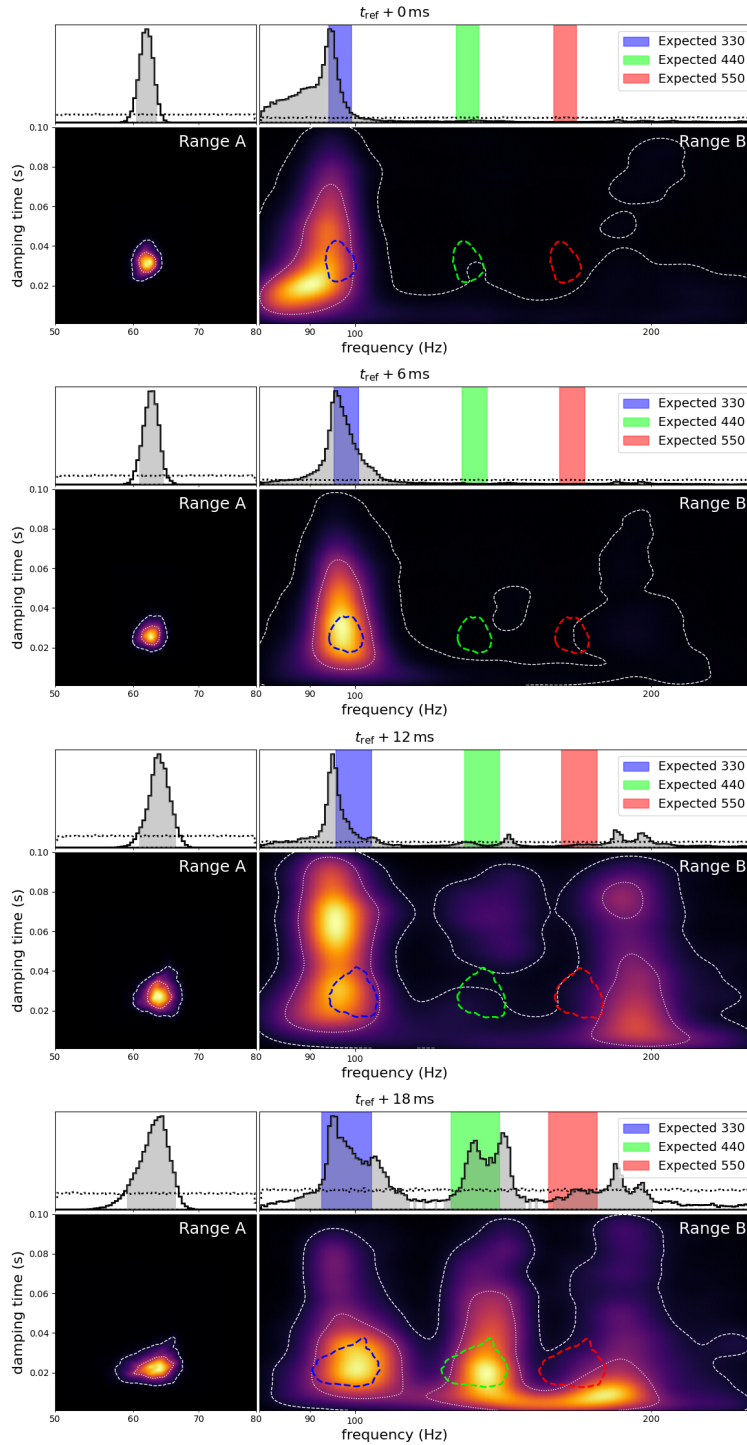


Figure 4.6: Spectra plots at 0, 6, 12, and 18ms showing marginal posterior distributions from frequency range A and B in the agnostic analysis. Also shown are the expected regions for the (330) (blue), (440) (green), and (550) (red) modes, assuming the peak in region A is the (220) mode.




Figure 4.7: Whitened data in each detector, with a gate applied at 7 ms (gray lines) after the detector reference time t_{ref} (gray dotted lines). Semi-transparent traces show the whitened data without the gate. Plotted are the maximum likelihood waveforms using just the (220) mode (black) and the (220) plus (330) mode (orange).




Figure 4.8: Bayes factors as shown in Fig. 4.2 for an extended range of times. (*Top*) Bayes factor of models with the indicated modes compared to the stronger of the (220) or the (220) + (221) modes model. The Bayes factor for the (220) + (221) model is calculated against the (220)-only model. (*Centre*) Median values for the frequency of the (220) mode for the model with the indicated modes. The shading shows the 90% credible interval. (*Bottom*) Median values and 90% credible intervals for the redshifted final mass. Hexagons mark where no-hair tests are performed in section 4.3 for the (220) + (330) model and in section 4.4 for the (220) + (221) model. All values are shown for start times of the analysis relative to the reference time t_{ref} .




Figure 4.9: Posterior on the deviation from Kerr of the (330) frequency δf_{330} and damping time $\delta \tau_{330}$, as well as the resulting modified (330) frequency, using a model in which we include the (220)+(330) modes at $t_{\text{ref}} + 7$ ms. Quoted values are the median and 90% credible interval, and the latter is indicated by the dashed vertical lines. The fractional deviations are expected to be zero for a Kerr black hole (indicated by the green lines). We use a prior (black dotted lines) that is uniform over $\delta f_{330} \in [-0.3, 0.3]$, with the constraint that $f_{330}(1 + \delta f_{330}) > 75$ Hz. This constraint is necessary to avoid label switching with the (220) mode; even with the constraint we clearly measure a lower bound on δf_{330} . For the damping time we use a prior that is uniform over $\delta \tau_{330} \in [-0.9, 3)$, and find that the damping time is only weakly constrained.




Figure 4.10: Posterior on the deviation from Kerr of the (221) frequency δf_{221} and damping time $\delta \tau_{221}$, as well as the resulting modified (221) frequency, using a model in which we include the (220)+(221) modes at $t_{\text{ref}} - 5$ ms. Quoted values are the median and 90% credible interval, and the latter is indicated by the dashed vertical lines. The fractional deviations are expected to be zero for a Kerr black hole (indicated by the green lines). We use a prior (black dotted lines) that is uniform over $\delta f_{221} \in [-0.16, 0.3]$, with the constraint that $f_{221}(1 + \delta f_{221}) > 55$ Hz. This constraint is used to try to exclude additional noise that is present in the Hanford and Virgo detectors at ~ 50 Hz. For the damping time we use a prior that is uniform over $\delta \tau_{221} \in [-0.8, 0.8]$. Despite the tighter prior constraints than that used for the (330) mode, we obtain a larger 90% credible interval on both δf_{221} and $\delta \tau_{221}$. The posterior also peaks toward the prior boundaries. This may be due to the noise at low frequency, or may indicate that the signal is not fully captured by a sum of quasi-normal modes at this time.

5 Statistical significance of evidence for gravitational wave echoes

Recent detections of merging black holes allow observational tests of the nature of these objects. In some proposed models, non-trivial structure at or near the black hole horizon could lead to echo signals in gravitational wave data. Recently, Abedi et al. claimed tentative evidence for repeating damped echo signals following the gravitational-wave signals of the binary black hole merger events recorded in the first observational period of the Advanced LIGO interferometers. We reanalyse the same data, addressing some of the shortcomings of their method using more background data and a modified procedure. We find a reduced statistical significance for the claims of evidence for echoes, calculating increased p-values for the null hypothesis of echo-free noise. The reduced significance is entirely consistent with noise, and so we conclude that the analysis of Abedi et al. does not provide any observational evidence for the existence of Planck-scale structure at black hole horizons.¹

5.1 Introduction

The detections of gravitational wave (GW) signals allow for new tests of the nature of black holes [33, 69, 237–242]. Black holes are characterised by their horizons. In vacuum general relativity these horizons are devoid of material structure. The possibility that additional structure may form at or near the horizon location has been widely discussed in the literature, motivated by a number of different models and theoretical considerations [155, 243]. The Advanced LIGO [73, 244] and Virgo [74] detectors have detected gravitational wave signals from several binary black hole mergers [33, 237–241]. These detections now make those ideas testable in the observational regime.

A generic set of models called Ultra Compact Objects (UCOs) [153, 176, 245–247] can mimic black holes in terms of their gravitational wave emission at early stages of binary inspirals. These models are designed to match the properties of standard black holes at sufficiently large distances, but differ in the near-horizon regime. The gravitational wave signal from the inspiral of two UCOs is expected to be almost identical to that of standard black holes (for possible tidal modifications see [248]). However, the merger and ringdown signals may differ sufficiently to be detectable. Near-horizon material structures motivated by semi-classical and quantum gravity theories could, at least partially, reflect incoming waves which in standard vacuum general relativity would be fully absorbed by the black hole.

Recent works by Abedi, Dykaar and Afshordi (ADA) [158, 249, 250] have claimed to find tentative evidence of near-horizon Planck-scale structure using data [251, 252] from the three Advanced LIGO events GW150914, LVT151012 and GW151226. In the simplified analysis of [158, 249], this near-horizon structure gives rise to so-called echoes [153, 176, 253, 254].

The data used by ADA is from the LIGO Open Science Center (LOSC) [251, 252], which contains a total of 4096 seconds of strain data from both Advanced LIGO detectors around each of the three events. Out of these data ADA used only 32 seconds centred around each event for

¹This chapter is an adaptation of the work published as [3] as described in Section 1.2, with the copyright for the published article [3] held by the American Physical Society (2018).

their analysis. The authors claimed in [249] to find evidence for such echoes in data following the three events with a p-value 3.7×10^{-3} , corresponding to a combined significance of 2.9σ (with the one-sided significance convention used in [33, 237, 238, 255], this value corresponds to 2.7σ). This was subsequently updated to a p-value of $\sim 1\%$ and interpreted as 2.5σ -level tentative evidence in [158]. Nonetheless if such a signal were shown to be present in the data, it would force a major re-evaluation of the standard picture of black holes in vacuum Einstein gravity.

Here we investigate concerns about the methods in [249] and ADA’s updated works [158, 250], and give a different significance estimate for the findings. Our initial caveats concerning [249] appeared as [256]. We do not examine the theoretical motivations for the existence of such near-horizon Planck-scale structure, nor the model templates for which ADA have chosen to search. Rather, we focus on the data analysis methods as reported and on the significance estimates assigned to the results. We identify a number of shortcomings in the analysis and perform an improved analysis, which corrects for several of these problems. We evaluate the echo findings in the gravitational wave data [251, 252], estimate their significance with updated p-values (for a general critique of p-values, see [257]) and conclude that there is as of yet no evidence for the existence of black hole echoes in these data.

5.2 ADA’s Model and search procedure




Figure 5.1: A coalescence template extended to include echoes. The five parameters of the echo waveform model are illustrated, and the phase-inversion between echoes is visible.

The analysis of ADA [158, 249, 250] consists of three parts: a simple waveform model, a search procedure, and a significance estimation method. In this section we briefly review these.

With a partially reflective surface outside the horizon, echo signals may be found as wave solutions in a cavity formed by the near-horizon membrane barrier and the angular momentum barrier (“photon sphere”) that exists further out [153, 176]. In the geometric wave picture, at

each barrier the wave is partially reflected and partially transmitted. Each partially transmitted wave from the outer angular momentum barrier would be detected by distant observers as an echo. The delay time between subsequent echoes results from the travel time between the two barriers. This time may be different for the first echo due to non-linear effects during the merger, as may further parameters of the echo signal such as the damping between successive echoes. For a description of the echoes as poles of the propagators see [246].

An example of such an echo template is shown in Fig. 5.1 and several parameters define its features:

1. Δt_{echo} : The delay time between subsequent echoes, resulting from the travel time between the barriers. $\Delta t_{\text{echo, theory}}$ is the expected value found in [158, 249], based on the inferred final mass and spin parameters for each event [33, 237–239, 258]. In the search, the parameter Δt_{echo} is allowed to vary around the theoretical value $\Delta t_{\text{echo, theory}}$ to account for uncertainties.
2. t_{echo} : The time of the first echo. This is expected to be $t_{\text{merger}} + \Delta t_{\text{echo}}$, where t_{merger} is the time of the merger. It is allowed to deviate from this expectation in the search to account for non-linear effects close to the merger [249].
3. A : The amplitude of the first echo relative to the original signal amplitude.
4. γ : The relative amplitude between subsequent echoes.
5. t_0 : Only the last part of the original waveform is used to produce the echo waveform; this parameter describes how far before t_{merger} the original waveform is tapered down to 0, using the tapering function given in [249].

In addition, the phase is inverted between subsequent echoes. Likewise, the phase-difference between the original signal and the first echo is fixed to $\Delta\phi = \pi$. We use an abbreviated notation for the combination of parameters $x := (t_{\text{echo}} - t_{\text{merger}})/\Delta t_{\text{echo}}$, with an expected value for the first echo of $x = 1$.

The **ADA-search** procedure used in [158, 249] consists of the following steps:

1. Produce a pure echo template for given echo-parameters. The original event is removed from the template.
2. Produce a bank of these templates, with an evenly spaced grid in the parameters listed above.
3. Perform matched filtering with the echo templates. The original event is removed from the data prior to this.
4. Maximise squared signal-to-noise ratio, SNR^2 , over all parameters for each value of x .

The maximisation uses either each single event or combinations of events. The combination assumes some parameters to be different between events, namely A and Δt_{echo} . The parameters x , $t_0/\Delta t_{\text{echo, theory}}$ and γ are kept identical for each event. For combinations of events, the sum of the individual SNR^2 s is maximised.

The **ADA-estimation** uses the following method to estimate the significance of their findings [249]:

1. Find the maximum SNR^2 value in the range $x \in (0.99, 1.01)$ after the event.
2. Calculate and maximise SNR^2 over the time range $9 \leq \frac{t_{\text{echo}} - t_{\text{merger}}}{\Delta t_{\text{echo, theory}}} \leq 38$. The maximisation is slightly adapted for this step.

3. Divide this last range into 1450 segments, each of duration 2% of $\frac{t_{\text{echo}} - t_{\text{merger}}}{\Delta t_{\text{echo, theory}}}$.
4. A p-value is found as the number of segments with higher peak SNR² than after the event in step 1, divided by the total number of segments.

5.3 General remarks

A first immediate problem arises regarding how strong the relative signal should be for the three events. The two binary black hole events GW150914 and GW151226 were detected by the Advanced LIGO detectors with significance levels $> 5.3\sigma$ and signal-to-noise ratios of 23.7 and 13.0 respectively [238]. The other event, LVT151012, had a reported significance of only 1.7σ and a signal-to-noise ratio of 9.7 combined between the two Advanced LIGO detectors. However, in Table II of [158] we see that the signal-to-noise ratio of the claimed echo signal is actually largest for LVT151012.

The higher SNR of LVT151012 cannot be due to the different projected number of echoes between the events. The different Δt_{echo} leads to differing numbers of echoes in a given duration: the 32 seconds of data used would contain ~ 180 for LVT151012 and ~ 110 for GW150914. Although the number of echoes is larger for LVT151012, late echoes are strongly damped. They decrease by a factor of 10 over ~ 22 echoes for the claimed relative amplitude $\gamma \sim 0.9$. Thus in order for the echoes of LVT151012 to have a higher SNR than the echoes of GW150914, their amplitude must be very high. In fact, to account for the reported SNRs, the initial amplitude for the first echo of LVT151012 would have to be about 10% higher than that of GW150914 [256], while the original event's peak is about 2-3 times lower for LVT151012 in comparison to GW150914's. This would require their parameter A to be about 2-3 times larger for LVT151012 than for GW150914. This seems to be confirmed by the best fit search results in Table II of the updated work [158], which gives $A_{\text{GW150914}} = 0.091$ and $A_{\text{LVT151012}} = 0.34$.

We assume that far in the wave zone the gravitational wave signal of the echoes decays similarly to the signal of the event itself, i.e. linearly with the distance from the source. This explicit astrophysical assumption, in addition to those in [158, 249, 250], is the basis for the above concern. The lower significance of LVT151012 is rooted in its distance: its mean estimated distance being more than twice as large as that of GW150914 and GW151226, we expect weaker echo signals. While particular combinations of system parameters and signal morphologies may have significant effects on the generation of echoes and their relative amplitudes, changing their relative significance, there is yet no extensive model to justify abandoning this concern here.

The inferred amplitude parameters suggest that a large amount of gravitational wave energy was emitted in the echoes: a very rough calculation implies that the amount of energy emitted in the echoes was approximately 0.1 solar masses (for GW150914) and 0.2 solar masses (for LVT151012). This should be compared to the total estimated energy emitted by the original signal of 3 solar masses (for GW150914) and 1.5 solar masses (for LVT151012).

We also note an inconsistency in the above procedure, resulting from the use of a fixed waveform for each event as the basis for all echo templates, obtained from the LOSC [251]. The parameters of the echo templates, in particular Δt_{echo} , depend on the mass and spin parameters of the final black hole. Instead of using only one initial waveform and generating all echo templates with this, one should use an initial waveform that corresponds to each set of echo parameters to be varied over. Using the single LOSC waveform is a simplification, restricting to only one choice of final mass and spin parameters for the echoed original event, while simultaneously varying over the final mass and spin values through Δt_{echo} .

5.4 Validation of the matched-filter analysis




Figure 5.2: The matched filtering technique is able to recover signals with a variety of amplitudes. As shown here, the SNR depends on the amplitude of the signal. The amplitude found by ADA ($A = 0.1$) is close to the level that is found in pure Gaussian noise without a signal. An amplitude twice as large as this would be clearly identifiable in the data.

We wrote a separate implementation of the **ADA-search** procedure, that we refer to as **ADA_{AEI}-search**. No changes were made to the algorithm as described before, while the implementation itself is independent. The SNR²-results obtained with our implementation are similar to those shown in [249].

As a first check, we verify that the **ADA_{AEI}-search** procedure can distinguish between pure noise and simulated echo-signals. For this, a known signal is injected into simulated noise by calculating the sum of the noise and signal strain at each sample time. We simulate Gaussian noise with a power spectral density (PSD) similar to that found for the detector data around each event (calculated from the LIGO data). The **ADA_{AEI}-search** is then applied to simulated data consisting of either pure noise or the same noise with added injections of different amplitudes. In this test, we only use echo waveforms with parameters similar to the best-fit results of [158, 249]. Fig. 5.2 shows the dependence of the SNR² peak on the injection amplitude A_{inj} , demonstrating how the effectiveness of the method in finding a signal depends on A_{inj} . We perform this test for different realisations of the simulated noise. The minimum A_{inj} required to find a peak rising above the noise background also depends on the noise instantiation. We find that $A_{\text{inj}} \sim 0.1$

can yield a visible peak. This is the best-fit value of A reported for GW150914 in [249]. In one out of the five trials conducted in this first test, however, a higher amplitude was necessary to distinguish the signal from noise, as shown in Fig. 5.2, where the noise and the quietest injection have almost identical SNR^2 results. This prompted us to perform more detailed statistical and injection-recovery analyses, as described below.

5.5 Prior ranges and template spacing

The estimation of the echo parameters is based on maximising the SNR over a fixed grid in parameter space. The boundaries of the parameter grid are determined by a prior range, where the ranges chosen by ADA are displayed in Table I of [249]. Each template in the bank is produced for a specific value of each parameter. The matched filtering method finds a higher SNR for data similar to the template, but each template can recover signals with a range of parameter values. The recovered signal parameters are defined as the values corresponding to the template in the bank that yields the highest SNR. The SNR maximisation is performed over all templates in the bank and thus over all values in the parameter grid used to create the bank. This determines the parameters γ , Δt_{echo} , t_0 , and t_{echo} . When combining multiple events, the sum of their SNR^2 s is maximised and only Δt_{echo} is varied independently for each event. The same γ is assumed for all events, while t_{echo} and t_0 are related between events by requiring all events to have the same $x = (t_{\text{echo}} - t_{\text{merger}})/\Delta t_{\text{echo}}$ and $t_0/\Delta t_{\text{echo, theory}}$.

The values for γ and t_0 resulting from this maximisation are found to lie very close to the boundary of their prior range, 0.9 and -0.1 respectively [256]. This suggests that there may be support for values of these parameters that lie outside of this range. If these values reflect the priors rather than the data, then they cannot be reliably considered as evidence for a detection claim. Furthermore, a value greater than unity for γ means that each successive echo has an amplitude greater than the previous echo. Such a result would require the echoes to be extracting energy from the black hole spacetime.

We tested whether the preference for these parameter values is an artefact of the method, again using known signals injected into simulated noise. We constructed Gaussian noise with a PSD estimated from the 4096 seconds of LOSC data around GW150914. The injected signals consist of echoes based on the LOSC GW150914-template for various echo parameters. The range of γ is widened to $\gamma \in (0.1, 2.0)$ both in the prior of the search and the injections. The range of t_0 is widened to $t_0 \in (-0.2, 0)\Delta t_{\text{echo, theory}}$ in the search. It is not widened for the injections in this test, as the dependency of the maximised SNR on the wider range in t_0 was found to be much weaker than for γ . The relative amplitude of the injections $A = A_{\text{inj}}$ ranges from ADA's recovered value 0.1 to about 50 times this amplitude. We then compare the best-fit value of γ from the search with the value of the injection. This comparison is shown in Fig. 5.3.

The **ADA-search** method is biased towards finding γ -values close to 1. Ideally, the recovered parameter value would be closest in the grid to the value of the injection. In Fig. 5.3, this would mean lying as close as possible to the plotted diagonal. In this figure, the recovered values are close to the injected ones for higher injection amplitudes. Thus for very high echo amplitudes, the recovery method could in principle be effective. For lower injection amplitudes, there is a preference for recovered values of γ close to 1, independently of the γ value of the injection. Therefore, finding $\gamma \sim 1$ as the best-fit value in the search does not necessarily mean that this is indeed the correct value for an existing signal. The method is biased to find these values for γ in almost all cases. In particular this is also true for relatively low echo amplitudes as found by ADA, and even significantly higher signal amplitudes. We interpret the recovery of $\gamma \sim 1$ as a generic property of the method and finding such a value cannot be considered evidence for the presence of a signal.




Figure 5.3: Injected and recovered values for γ , correctly recovered values would lie on the marked diagonal. The search method's preference for $\gamma = 1$ (dashed line) at lower injection amplitudes is clearly seen.




Figure 5.4: Demonstration of the method’s preference to find $\gamma = 1$ in white Gaussian noise without signals. The histogram shows the number of random noise realisations where the method finds the given value for γ . The line shows the overlap between templates which are neighbours in the γ parameter-grid of the template bank, while all other parameters are the same for both templates.

The bias is introduced through the spacing between templates in the bank, as can be found through an analysis on white Gaussian noise and calculating the overlap between the templates. We use white Gaussian noise with a constant PSD in this test to preclude any influence from the specific PSD of the detector noise. An analysis in this noise using the same parameter range as ADA also shows a strong preference for $\gamma = 0.9$. Further extending the range to $\gamma \in (0.1, 2.0)$ displays preference for $\gamma = 1$ in white noise. The distribution of recovered γ values in this test is shown in Fig. 5.4. The reason for this is revealed by calculating the overlap between neighbouring templates in the parameter grid for different γ , while keeping the other parameters fixed. Here, the overlap between two waveforms is calculated as the SNR when using one waveform as the template and the other as the data. As can be seen in Fig. 5.4, the overlap between neighbouring templates shows the same γ -dependence as the distribution exhibiting the preference for recovering $\gamma = 1$. Templates with γ close to 1 lie further apart in this noise-weighted match-metric than other templates. Each template near $\gamma = 1$ therefore covers a larger region of the signal space than other templates, and thus, more noise realisations are best matched by the (morphology-wise) more scarcely placed templates close to $\gamma = 1$.

We similarly test how the method recovers the echo signal’s amplitude through injections into Gaussian noise². For the results in Fig. 5.5, we chose to show the absolute peak amplitude of the echo signal instead of the parameter A , which gives the echo amplitude relative to the original event amplitude. This allows us to find the minimum echo amplitude to be recovered correctly, independently of that of the event. Fig. 5.5 shows that the recovered values deviate from the injected ones strongly below strain amplitudes of about 10^{-22} . For lower injection amplitudes, values around 10^{-22} are found instead of the injected ones. This suggests that finding such low amplitude values might be expected in pure noise as well. The absolute value is close to those found in [249]: the relative amplitudes of 0.1 for GW150914 and 0.3 for LVT151012, multiplied by the respective events’ peak amplitudes, are shown as horizontal lines in Fig. 5.5, and seem consistent with incorrect recovery of the method for lower injection amplitudes.

Extending the template bank to a wider range in γ and t_0 and performing the same analysis as before leads to a modified SNR structure in x , where additional and higher peaks appear further away from the predicted value for the echoes in GW150914. As we will see below, applying a wider parameter range also for the background estimation results in a further factor ~ 3.5 increase in the p-value for the combination of GW150914, LVT151012, and GW151226 (using the 32-second dataset estimation). The modified p-values for the wider priors of different combinations are shown in Table 5.1, where the widened prior entries refer again to the ranges $\gamma \in (0.1, 2.0)$ and $t_0 \in (-0.2, 0)\Delta t_{\text{echo, theory}}$.

5.6 Extending the background estimation

To calculate a significance for the match found in the templated search, we must assess the noise background. Since an analytical noise model is not known, real data away from the possible signal are used to estimate this background. This relies on an assumption that the data at these times are similar to that during the time of interest. The noise background is calculated by counting how often an equal or larger SNR value is obtained in the off-source data. ADA chose to do this in a short period of time of approximately 16 seconds of data after each event. To obtain sufficient background statistics this period of time was used intensively: they consider 16 second stretches of data as independent when shifted by only 0.1 seconds.

This background estimation is problematic [256] for two reasons: potential contamination of the background samples by existing echo signals, and the lack of independence between background samples. The estimation uses a range of t_{echo} values that is only $\mathcal{O}(10)$ echo periods

²We thank N. Afshordi for suggesting this additional test [259].




Figure 5.5: Injections of echo signals into Gaussian noise are analysed for different injection amplitudes. For injections with amplitudes above 10^{-22} of peak strain, the recovered values are close to the injected ones, indicated by the diagonal black line. For injections with lower amplitudes, the recovered values are around 10^{-22} , independently of the injected value. The amplitudes found in [249] are shown as horizontal lines (“true” values unknown) and are similar to the values incorrectly recovered for low amplitudes. The shown injections are made with $\gamma = 0.8$, with similar results for other values.

away from the merger. If there is indeed an echo signal in the data then this region will not be entirely free of the signal being searched for. At the beginning of the region the amplitude of the echoes would only have dropped by a factor $0.9^9 \sim 0.4$. One therefore expects a contaminated background estimation. Even in the absence of echoes, a random feature mistaken for echoes in one segment may well extend to neighbouring segments, and they cannot be treated as independent (see discussion of template auto-correlations below for the problem of insufficient independence of samples).

Each of the data sets released at the LIGO [251] consists of 4096 seconds of data. Both GW150914 and LVT151012 are located 2048 seconds into these data, equivalent to $\mathcal{O}(10^3)$ echo periods, thus for large stretches of the data, such contamination through the presence of a damped echo signal would be negligibly small.

We have performed a different background estimation as an independent test, which uses the full period of 4096 seconds of LIGO data available around each event. A schematic comparison of the different choices of data used for background estimation is shown in Fig. 5.6.

The obtained p-value and background estimate are only meaningful if the data in the background are comparable to those at the time of the event. A plot of the noise variations over the full 4096 seconds of data released for each event is shown in Fig. 5.7 and for GW150914 specifically in Fig. 5.8. The variations are seen to be small and we conclude that for the four events considered our background estimate is indeed characteristic of the noise just after the event. For the graph showing the properties of noise in the Hanford detector around the time of LVT151012 in Fig. 5.7, a reduced amount of data was used. This choice is made due to three loud short transient noise features, which we discuss further below. The noise features strongly influence the Rayleigh statistic calculation, while occupying less than 0.1 % of the data. Using data excluding these noise features, the variation as shown in Fig. 5.7 is found. Properties of the data at and around LVT151012 are discussed in [110].

In our case, the 4096 seconds of data for each event are divided into 128 independent, 32-seconds long segments. For each, the echoes analysis is performed as it was on the 32-second segment containing the event. The resulting peak SNR in $x \in (0.99, 1.01)$ is found for each segment. Simply counting the number of segments containing a higher peak SNR in this interval yields an estimate for the p-value. For the combined first three events, GW150914, LVT151012 and GW151226, our resulting p-value of 0.032 is about a factor 3 larger than the value of 0.011 found in [158], where less data and overlapping intervals were used.

An estimate of the p-value significance in this way is susceptible to small number statistics (accounting for the Poisson error as suggested in [259], the p-value can be 0.032 ± 0.016 , still larger than in [158, 249]). The original LIGO templates, before introducing echoes, contain an approximately 16 seconds long waveform followed by 16 seconds of a flat zero-strain template. Echoes were introduced only into this flat region when producing the echo templates. After removing the original event, we are left with a 32-second template with 16 seconds of no signal, followed by the produced echoes. So we can double the number of background samples without losing independence between samples, by dividing the available data into 256 independent segments of 16 seconds length.

The exact number of available data segments varies slightly for each of the events. This is due to the positions of the original signals, and the influence of one of the three mentioned short transient noise features (inconsistent with the echo morphology) in one segment of the LVT151012 data, which was discarded. This short noise feature was found by noticing a very high SNR outlier for one data segment, shown in Figure 5.9. The feature can be seen in the whitened time-domain data, appearing close to the beginning of the data segment. The search procedure always aligns one of the first and thus loudest echoes with the noise feature, yielding the high SNR, shown in Figure 5.10. The effect of not discarding the high SNR noise dataset




Figure 5.6: Schematic comparison of the data segments used to estimate the background between ADA [158, 249] and this work (AEI). Compared to ADA’s 16 seconds, we extend the amount of data used for background estimation to the full 4096 seconds for each event available from the LIGO [109].




Figure 5.7: Rayleigh plot of the noise variation, showing the ratio of the standard deviation to the mean as a function of frequency for PSD samples from different times. The PSD is estimated for 16-second segments of the 4096-second data stretch for both detectors and each event. For each frequency, the corresponding values of the PSD from all segments are collected and their mean and standard deviation are calculated. The ratio of standard deviation and mean is then shown in the plot as a function of frequency, where values smaller/larger than 1 correspond to coherent/incoherent variation, respectively, and the value 1 to Gaussian behaviour. The `gwp` software was used in generating this plot [260, 261].




Figure 5.8: Variation of the noise spectrum during the 4096s around GW150914, calculated using 16-second data segments. The shaded regions cover the 1-st to 99-th percentile variations, the central curves show the mean. In this sense, the data are sufficiently stationary for our background estimation to be reliable during the full 4096-second data stretch. The `gwpv` software was again used in generating this plot [260, 261]

always is an increase of p-value due to the very high SNR, but the effect on the resulting p-values is minimal ($\sim 1/256$). Only the estimation with 16-second-long segments is influenced by this noise feature. The total number of estimates when combining events is thus 125 to 126 for the 32-second segments, and 250 to 251 for the 16-second segments.

The other two short noise features appear late in the respective data segments. For these, the search does not consistently align one of the later and more strongly damped echoes with the noise feature, as the increase in SNR is outweighed by the placement of the first loud echoes in the data. Thus the search is not influenced by these features significantly and we do not exclude the data segments from the estimation.

The results of this alternative approach for the significance estimation, both using 32- and 16-second-long segments, are shown in Table 5.1. Different combinations of the events are considered, denoted chronologically as (GW150914, LVT151012, GW151226) \rightarrow (1, 2, 3). In addition, for comparison and as the first detection after the claims of [249], we also consider the first event in the second observing run, GW170104 [239], denoted (4) in Table 5.1.

For the combination (1, 2, 3), a p-value of 0.011 was found in [249]. Our method finds four out of 125 trials with a maximised combined SNR larger than in the data immediately following the events. So the p-value is $4/125 = 0.032 \sim 3\%$ for the combined SNR value found after the events. Using 250 samples of 16-seconds length each, we find $5/250 \sim 2\%$ for the p-value.

To highlight the role of LVT151012 in obtaining low p-values, we have chosen to make a comparison with a combination of three events excluding LVT151012. When selecting the available events to be combined in the analysis, a reasonable choice seems to be using those of sufficiently high significance. Here, this means GW150914, GW151226 and GW170104, the combination (1, 3, 4), for which we find $9/125 \sim 7\%$ and $50/251 \sim 20\%$ respectively. These values are much




Figure 5.9: 4096 seconds of LOSC data around LVT151012 are divided into ~ 256 sequential segments of 16 seconds, excluding the segment containing the GW event. For each segment, we perform the same analysis as in the segment directly following the event. The maximised SNR^2 is marked by a red circle for each segment, and by a black horizontal line for the data immediately after the GW event, where echoes might be found. A clear outlier in the SNR^2 is visible.




Figure 5.10: The whitened strain data is shown as a function of time for the dataset corresponding to the outlier in maximum SNR visible in Figure 5.9 by the red curve. The black curve shows the whitened best-fit waveform. The lower panel shows the same as the top panel, but restricts the range in time for better visibility of the overlay. A transient noise feature is visible in the data, which is partially matched by one of the echo signals in the template. The morphology of the noise feature differs from that of a single echo pulse, but the similarity is sufficient for the echo template to prefer templates aligning one pulse with it. Only a single such feature is present, thus the data do not match the model of repeated echo pulses.

Event	[158]	Original 16s (32s)	Widened priors 16s (32s)
GW150914	0.11	0.199 (0.238)	0.705 (0.365)
LVT151012	-	0.056 (0.063)	0.124
GW151226	-	0.414 (0.476)	0.837
GW170104	-	0.725	0.757
(1,2)	-	0.004	0.36
(1,3)	-	0.159	0.801
(1,2,3)	0.011	0.020 (0.032)	0.18 (0.144)
(1,3,4)	-	0.199 (0.072)	0.9 (0.32)
(1,2,3,4)	-	0.044 (0.032)	0.368 (0.112)

Table 5.1: p-values obtained by using 4096 seconds of LOSC data divided into segments of 16 or 32 seconds length. Results are given for the priors in t_0 and γ chosen in the original analysis and for widened priors, for the 3 O1 events and the first O2 event individually and for various combinations of events. In combinations, the events are denoted chronologically, with (GW150914, LVT151012, GW151226, GW170104) \rightarrow (1, 2, 3, 4). For the combinations directly comparable to [158], with the original priors, we also record the Poisson errors (as suggested in [259]): for GW150914 our p-values are 0.199 ± 0.028 (0.238 ± 0.043), and for (1,2,3) our p-values are 0.02 ± 0.009 (0.032 ± 0.016). The Poisson errors for the full combination (1,2,3,4) with original priors, are 0.044 ± 0.013 (0.032 ± 0.016). With widened priors, all p-values are significantly larger, and the Poisson error relatively insignificant. Combinations that include LVT151012 have the lowest p-values. The addition of GW170104 to the three O1 events increases the combined p-value and is thus more compatible with pure noise. The lowest p-value out of all 11 combinations using up to four events is found for the combination (1,2). Note however that considering more combinations of events using the same data also leads to a higher effective trials factor to be accounted for.

higher than for combinations including LVT151012 and are fully consistent with the pure noise null-hypothesis.

The combined background is shown in Fig. 5.11 which shows the peak value of SNR^2 found in each segment for both the real detector data and Gaussian noise. For each event, the Gaussian noise was created with the same PSD as estimated from the data of this event. There is no significant difference between the distribution of peaks for detector data and for Gaussian noise. Here we note that there is no obvious trend in the peak SNR over time. By this measure, there is no indication of the noise being unstable and preventing its use for background estimation. These two properties are shared by all single events and the alternate combination (1, 2, 3): all show the similarity of the peak distribution for Gaussian noise and detector data, and lack a trend in time.

A second concern about the background estimation used in [158, 249, 250] arises from the very small shift in time between samples that are considered independent. In this, the quasi-periodic nature of the echo signal has to be considered, leading to potentially long templates with quasi-periodic autocorrelation in time. This property affects the significance estimation as performed in [158, 249]. The estimation is problematic because the template is significantly longer than the shift in time between background sample intervals. If the autocorrelation between templates used in different background samples does not vanish, the results from these samples cannot be considered truly independent: a feature of the data at one point in time then influences the SNR




Figure 5.11: 4096 seconds of LIGO data for GW150914, GW151226, and GW170104 are divided into ~ 256 segments of 16 seconds. The segments containing the GW events are excluded. For comparison, we generated 4096 seconds of Gaussian noise with the same PSD as estimated from the LIGO data for each event and divided it into segments in the same way. For each segment, the same echo search as immediately after the events was performed, combining the data of the selected events and maximising the sum of the SNRs². The resulting maximum SNR² is shown by circles for the detector noise and by triangles for Gaussian noise. The distribution of peak SNRs in the detector data is similar to that in Gaussian noise. A black line marks the combined SNR value found for the data immediately after the GW events. The p-value is calculated from the number of points for detector data lying above this line. There is no obvious trend in the peak SNR over time.




Figure 5.12: Autocorrelation of the pure echo template depending on shift in time. The peaks demonstrate that there is significant correlation between background samples for small shifts.

found for several background samples. The total number of effectively independent samples is thus much lower in this method.

The autocorrelation of the echo templates is shown in Fig. 5.12. We see that the echo signal model leads to a series of peaks in the overlap of the original and the time-shifted wavefunction, depending on the parameters γ , Δt_{echo} and t_0 . A value of overlap to be considered sufficiently independent could be e.g. 1%. To achieve this value, there are two ways to place templates with respect to the original position in time. For $\gamma < 1$, i.e. a damped echo signal, applying a shift in time by a sufficiently large multiple of Δt_{echo} leads to a reduction of the correlation. Using the GW150914 template and $\gamma = 0.5$ shows that at least 7 times Δt_{echo} is necessary. In Fig. 5.12, this corresponds to the very small peak close to 2.1 seconds of time-shift.

Alternatively, the templates can be interlaced, such that the echoes of one template are placed within the time between echoes of the other. This corresponds to the small overlap values between the peaks in Fig. 5.12. Here we again use the GW150914 template and the most favourable values in the prior range, i.e. the shortest echoes ($t_0 = -0.001\Delta t_{\text{echo, theory}}$) and the longest delay ($\Delta t_{\text{echo}} = 0.30166$ s). Then about 7 echoes can be placed between those of the original template. For these parameters we now consider shifts in time of the echo template up to $29\Delta t_{\text{echo}}$, which is the maximum shift used in ADA’s significance estimation. We find about 4 independent samples through sufficient timeshift and a factor 8 through interlacing, giving ~ 32 independent samples.

As the maximisation is performed over a range of parameters, exactly determining the total number of independent samples would require a more detailed calculation. The parameters chosen for this estimate, however, are favourable, as smaller damping or smaller time delay would further lower the total number. For the maximisation combining the different events, Δt_{echo} may vary independently between events, obstructing a clear estimate on the number of samples; the same principles, however, still apply. These considerations suggest the method of

[158] contains only a small number of independent background samples, on the order of a few tens of samples.

The method we employed to estimate the background precludes this effect by only applying the matched filtering procedure to separate sets of data. The template thus is always placed in only one of the background samples and the resulting SNR cannot be influenced by data features in the remaining samples.

The nature of the echo templates leads to a further potential problem: low frequency components may be introduced in the template, resulting from the delay between echoes. Due to the delay times of about 0.1 to 0.3 seconds, these frequencies are expected to be in the range below 20 Hz, down to a few Hz. However, the data as supplied by the LOSC, [251], is not calibrated below 10 Hz, as mentioned in the corresponding notes on data usage. The results of the analysis may thus be influenced by the uncalibrated data. We have repeated the analysis after applying a high-pass filter to the data and the original waveform, removing the data below 30 Hz for the final SNR calculations. The results of this analysis are almost identical to those before applying the high-pass filter in terms of SNR. The resulting p-values similarly show only minor changes compared to the values given in Table 5.1. The combination of introduced low frequency components in the templates and the uncalibrated data thus seems to have no significant effect on the results of the analysis.

The **ADA-search** procedure does not distinguish between inversion and non-inversion of the first echo's phase. The waveform templates used here are based on the simple model presented in [158, 249]. Within this model, the phase-change of the gravitational wave between the original signal and the first echo is described as a simple phase-flip. However, as only the square of the SNR from the matched filter analysis is considered for the maximisation, the result is insensitive to this phase inversion. Repeating the analysis for GW150914 and enforcing the phase inversion as required by this model, we find that the prominent peak in SNR at $x = 1$ vanishes. This is shown in Fig. 5.13. More sophisticated models would be needed to determine whether this phase flip is truly required or not. Nevertheless, it is worth noting that the peak that forms the basis of evidence for echoes in [158, 249] does not contain this phase inversion as required by the simple model, but actually the opposite phase.

5.7 Conclusions

A full analysis of the data at a level necessary to confidently detect echo signals is outside the scope of this work. However, we have analysed the data using a simple templated search similar to [249]. Using an extended background estimated from the full 4096 seconds of data released publicly by the LIGO collaboration for each event in the first observing run, we find a p-value for the null noise-only hypothesis of 0.02 to 0.032, higher than that of 0.011 reported in [249] using the restricted background. We have demonstrated the importance of LVT151012, the weakest LIGO candidate event [110], in obtaining this p-value. A combined analysis of the three events GW150914, GW151226 and GW170104, excluding LVT151012, yields an even larger p-value of 0.199, fully consistent with noise. We have also identified a number of weaknesses in the analysis method of [249] including the role of the prior boundaries and the density of templates. In particular we have examined the role of the γ -parameter and found that the clustering of γ -values near $\gamma = 1$ is entirely consistent with noise. The signal manifold is such that in pure Gaussian noise, one would expect many more triggers with values of $\gamma \sim 1$. This perhaps would not be expected for true echo signals, although a more detailed model of echoes would be needed to make a quantitative prediction. A similar bias in recovered parameters concerns the peak amplitude, which for both GW150914 and LVT151012 was found by [158, 249] just on the boundary of credible signal recovery.




Figure 5.13: The maximised SNR for GW150914 is shown with two different choices for the change in phase $\Delta\phi$ between event and first echo. The original analysis allows for $\Delta\phi = 0 \vee \pi$ for each x and then finds a prominent peak near $x = 1$ (dashed line). However, the model presented in [158] requires $\Delta\phi = \pi$. When we enforce this value in the analysis, the peak vanishes (solid line).

In conclusion, we find that the tentative evidence as presented in [158, 249, 250] is lacking in several key aspects with respect to template placement, background estimation and implementation. Our analysis of these shortcomings shows that the method of Abedi et al. cannot provide observational evidence for or against the existence of near-horizon Planck-scale structure in black holes. This stresses the importance of developing both new theoretical models and analysis methods for gravitational wave echoes from such structures. We hope some of the concerns explored here may be useful to vet other searches for echoes, such as [262], and help in the development of methods which would place black hole near-horizon physics within the realm of gravitational wave observations.

Acknowledgements

The authors thank Andrew Lundgren, Laura Nuttall, Vitor Cardoso, and the authors of [158, 249, 250], for useful discussions, as well as Bruce Allen for helpful comments. Some of the discussions particularly enjoyed the hospitality of meetings at Nikhef and at the Perimeter Institute. This research has made use of data, software and/or web tools obtained from the LIGO Open Science Center (<https://losc.ligo.org>), a service of LIGO Laboratory and the LIGO Scientific Collaboration. LIGO is funded by the U.S. National Science Foundation (NSF). O.B. acknowledges the NSF for financial support from Grant No. PHY-1607520.

6 Bayesian parameter estimation for gravitational wave echoes

Searching for black hole echo signals with gravitational waves provides a means of probing the near-horizon regime of these objects. We demonstrate a pipeline to efficiently search for these signals in gravitational wave data and calculate model selection probabilities between signal and no-signal hypotheses. As an example of its use we calculate Bayes factors for the Abedi-Dykaar-Afshordi (ADA) model on events in LIGO’s first observing run and compare to existing results in the literature. We discuss the benefits of using a full likelihood exploration over existing search methods that used template banks and calculated p-values. We use the waveforms of ADA, although the method is easily extendable to other waveforms. With these waveforms we are able to demonstrate a range of echo amplitudes that is already ruled out by the data.¹

6.1 Introduction

Black holes are defined by their horizons [263]. Although a large amount of astrophysical data is compatible with the existence of black holes [264], a number of theoretical models still predict dark compact objects without horizons or for which the horizon structure is significantly modified from classical vacuum general relativity [153, 176, 245, 248, 254, 265]. These models are typically motivated by quantum effects or attempts to address issues related to black hole information and evaporation [266]. One possible observational signature of such structure is that infalling waves would not be entirely absorbed by the horizon as is generally expected in general relativity, but instead some amount of the infalling wave would be reflected [176]. Similar signals have been studied for stars for a long time [267, 268].

Recent observations of gravitational waves from coalescences of binary black holes [33, 237–241] by the LIGO [73] and Virgo [74] detectors have allowed for a number of new tests of the near horizon structure of black holes [69, 70, 269]. One such test involves searching for echo signals that could potentially be caused by reflective structure forming at or near the location of the black-hole horizon. A number of groups have searched for such signals in gravitational wave data with contrasting conclusions [3, 158, 262], ranging from “tentative evidence” in [158] to “low significance” in [3].

Here we propose a new method to search for these echo signals that provides an explicit probability for the compatibility of the data with the echoes model relative to a noise hypothesis. We demonstrate this method on the binary black hole events detected during the first observing run of the Advanced LIGO detectors; these events are the same events that were the subject of previous studies [3, 158, 262].

The general physical picture of echoes is that infalling radiation is reflected due to some mechanism near the putative horizon location. This radiation is then partially trapped between the near-horizon structure and the angular momentum light-ring barrier [176]. Some of the energy is transmitted away from the system by successive bounces, thereby forming a series of echoes. Generic parameters in the physical models are the amount of wave reflected by the

¹This chapter is an adaptation of the work published as [4] as described in Section 1.2, with the copyright for the published article [4] held by the American Physical Society (2019).

boundary and the effective location where this reflection occurs. These in turn are related to the amplitude of the reflected echo signals and the time separation between the successive echoes. Bounds on the amplitude and time separation of echo signals derived from the data can thus be translated into bounds on the reflectivity and location of the near-horizon structure.

For illustrative purposes here, we focus on the explicit model of Abedi-Dykaar-Afshordi (ADA) [158], which has been the subject of discussion in the literature [3, 250, 256, 259]. However, we note that our methodology can just as well be applied to other, more detailed models with explicit waveforms, including those recently proposed in the literature [246, 270]. Efforts to search for echo templates using Bayesian model selection have been developed with `LALInference` [271] in parallel to our own work, and published concurrently with our own [272]. Other, model-agnostic searches [273], have also been ongoing, along with different techniques to constrain horizonless objects through their impact on the stochastic background [274].

The primary result of [158] is a p-value, calculated as the probability of observing a signal-to-noise ratio (SNR) in noise (assumed to be free of signal) at least as significant as that observed in the on-source data that potentially contains the signal. This by itself does not indicate the probability that the on-source data contains a signal.

A probability that the data contains a signal can however be obtained using Bayes' theorem:

$$P(\text{signal}|\text{data}) = \frac{P(\text{data}|\text{signal})P(\text{signal})}{P(\text{data})}. \quad (6.1)$$

It is most convenient to compare this probability to an alternative hypothesis, for example that the data contains pure noise:

$$\frac{P(\text{signal}|\text{data})}{P(\text{noise}|\text{data})} = \frac{P(\text{data}|\text{signal})}{P(\text{data}|\text{noise})} \frac{P(\text{signal})}{P(\text{noise})}. \quad (6.2)$$

In the above, the first factor on the right-hand side is the likelihood ratio, or Bayes factor, and the second factor is the prior odds. Evaluating the prior odds is difficult without prior data (and in the case of a signal model that violates standard physics, might well be a very small factor). The likelihood ratio, on the other hand, can be calculated by exploring the likelihood function over the model parameters using a stochastic sampling algorithm, such as a Markov chain Monte Carlo (MCMC). In addition to the prior probability of hypotheses, a hypothesis itself may have a number of free parameters, each of which will have its own prior probability distribution. The first term on the right-hand side of Equation (6.2) then is the likelihood ratio for a specific choice of signal parameters. To obtain the Bayes factor from this term, the model parameters must be marginalised over using their respective prior distributions.

The data of gravitational wave detectors is known to be non-Gaussian [110], so a comparison with a Gaussian noise hypothesis does not exhaust the list of possible explanations for the data. However, for sufficiently short periods of time (of the order a few seconds) around the events discussed here, the noise can be well approximated as Gaussian and any signal hypothesis that is disfavoured relative to Gaussian noise is unlikely to be favoured relative to a more accurate noise hypothesis. If the signal hypothesis is favoured over Gaussian noise, then further analysis would be warranted.

To establish that our method can correctly identify echo signals in the data, we test it on simulated echo signals with a variety of different amplitudes. These simulations are added to real detector data, which is made available by the Gravitational Wave Open Science Center (GWOSC) [251, 252]. We choose to inject simulated signals 100 seconds after GW150914. This 100-second delay makes it unlikely that the data at that time is contaminated by a real astrophysical signal [275]. These injections allow us to estimate the sensitivity of our method and establish a signal amplitude that would be unambiguously identified by our method.

Echo param.	Prior Range	GW150914 range	Injected value
Δt_{echo}	inferred	0.2825 to 0.3025 s	0.2925 s
t_{echo}	$\Delta t_{\text{echo}} \pm 1\%$	0.2795 to 0.3055 s	0.2925 s
t_0 trunc.	$(-0.1 \text{ to } 0)\Delta t_{\text{echo}}$	-0.02925 to 0 s	-0.02457 s
γ	0.1 to 0.9	0.1 to 0.9	0.8
A	unconstrained	0.00001 to 0.9	varying

Table 6.1: Table of prior ranges for echoes’ parameters along with the values used here for injection studies. The ranges are adopted from [158]. The injection values are chosen to lie close to the parameter values found in that work, except for γ and t_0 trunc. which are chosen to lie within the prior range rather than at the boundary.

We then apply our method directly to the three binary black hole events in O1: GW150914, GW151012 and GW151226. We find that the data for GW150914 favours Gaussian noise over the echoes hypothesis. The other two events show a marginal preference for echoes, but this is again consistent with noise. Finally, we discuss how our results can be used to place bounds on the reflectivity of structure that has formed a given distance from the location of the would-be horizon. Although our conservative bounds cannot conclusively exclude the signal amplitudes claimed in [158], we show that they are consistent with being noise, having amplitudes below the level that can be reliably estimated and Bayes factors not uncommon in off-source times. Future runs with more sensitive detectors and louder events are likely to lower the level of amplitudes that can be reliably estimated and rule definitively on the generic claims of [158].

6.2 Methodology and analysis pipeline

The example signal hypothesis we consider here is based on that of ADA [158]; we refer the reader to that work for more detail on the model and the meaning of the various model parameters. The most important of these parameters are the overall amplitude of the echoes relative to the original signal’s peak A , the relative amplitude between successive echoes γ , and the time separation between successive echoes Δt_{echo} ; these and the other parameters t_{echo} and t_0 trunc. are explained more fully in [158]. Table 6.2 gives the prior ranges employed here for the relevant parameters. These are adapted for our purposes from the template bank search performed in [158].

In the ADA model the range for Δt_{echo} is inferred from the published parameters of GW150914 [238], using 50% ranges, and assuming Gaussian distributions for the error. The Kerr metric formula is used for the light travel time between the light ring and a perfectly reflecting surface. This surface is assumed to be at a proper distance one Planck length along Boyer-Lindquist time slices from the Kerr metric event horizon. The parameter γ was chosen to reflect the physical expectation that the amplitude of successive echoes should decrease due to energy loss through one or both of the boundaries. We allow the parameter t_{echo} to vary independently from Δt_{echo} within 1% of its maximum values, and choose an explicit prior for the relative amplitude.

Since the value of the amplitude will have a direct influence on the signal strength, and hence the signal likelihood, its prior range is of central importance to our results. In the template bank search of [158] a prior for the amplitude is not explicitly given. Instead, it is maximised over the template bank. To replicate as closely as possible the method of [158] we choose a flat amplitude prior from 10^{-5} to 0.9. This ensures we are sensitive to relatively quiet amplitude signals, although not arbitrarily quiet, and implements the reasonable assumption that the first

echo should not be louder than the main signal.

For simplicity we choose to fix the number of echoes to 30. In principle this could be allowed to vary, but for values of γ less than 0.9, 30 echoes capture the main part of the signal that influences the SNR. In testing, we found that varying this number did not change the results substantially.

The pipeline we use is based on `PyCBC Inference` [99]. It employs a parallel-tempered MCMC algorithm, `emcee_pt` [114, 117], to sample the likelihood function for a hypothesis based on the existence of a signal in the data. The likelihood function is chosen to be compatible with the assumption that the underlying noise is Gaussian with a given power spectral density. Once the likelihood has been mapped, the marginalisation over the model parameters is performed using thermodynamic integration to obtain a probability for the hypothesis given the data. Although it is known that LIGO data is not Gaussian over long periods of time, over shorter periods it is approximately Gaussian [110, 275]. To account for the non-Gaussianities without a model hypothesis for them, it is possible to sample the Gaussian Bayes factor over many realisations of the true detector noise.

In the results presented here we used 100 Markov chains to sample the likelihood. We require that each chain run for at least five auto-correlation lengths (ACL) beyond 1000 iterations of the sampler. The ACL is measured by averaging parameter samples over all chains, then taking the maximum ACL over all parameters. For the thermodynamic integration of the likelihood function, care has to be taken that it is sufficiently sampled both near its peak, but also at lower values of the likelihood. In tests we found that using 16 different temperatures, each placed by inspection, was sufficient to guarantee a consistent value of the Bayes factor. Convergence of this result was checked by running with double the number of temperatures and ensuring that the results were consistent. The posterior distributions are constructed from the coldest temperature chain.

6.3 Injections based on GW150914

To test our method we choose to examine simulated echo signals based on GW150914. This is, to date, the loudest binary black hole signal that has been observed via gravitational waves, and should play a central role in constraints derived from the data. While we focus here on GW150914, we expect similar results will apply to echo signals of other events when suitably scaled in amplitude.

Following ADA for simplicity, we choose to fix the base inspiral-merger-ringdown (IMR) waveform to be echoed for both injections and for the search templates. The parameters for these base IMR waveforms are given in the appendix and are obtained from the maximum likelihood results of [99]. The waveforms are constructed using the phenomenological IMR waveform family `IMRPhenomPv2` [276, 277] which is freely available as part of `LALSuite` [200]. These IMR signals are then used to produce echo signals with echo parameters given in Table 6.2. The simulated echo injections are added linearly at varying amplitudes to real detector noise (chosen to be 100 seconds after GW150914, far enough away to be uncontaminated by echo signals or any pre-merger signal). We then attempt to recover them with our analysis pipeline. Example results are shown in Figs 6.1 and 6.2.

Figure 6.1 shows a very loud injection with a relative amplitude of 0.4 and a maximum likelihood SNR of ~ 17.7 . The log Bayes factor for this injection is 140.57, showing a strong preference for the echoes hypothesis over the pure Gaussian noise hypothesis. In this case the echo parameters are well recovered, with the injected values lying within the 90% credible intervals of the marginalised one-dimensional posterior distributions.

Figure 6.2 shows a much quieter injection with a relative amplitude of 0.0125 and a maximum




Figure 6.1: Posterior on the echo parameters for a loud ($\text{SNR} \sim 18$) simulated signal. The signal has GW150914-like parameters at a fiducial distance of 400 Mpc. An amplitude factor of 0.4 is used for the echoes. Off-diagonal plots show 2D marginal posteriors; the white contours show the 50% and 90% credible regions. Each point represents a random draw from the posterior, coloured by the SNR (ρ) at those parameters. The diagonal plots show the 1D marginal posteriors, with the median and 90% credible intervals indicated by the dashed lines. The reported values are the median of the 1D marginal posterior plus/minus the 5th/95th percentiles. The injected parameter values, shown by the red lines, are all within the 90% credible intervals. The log Bayes factor for this signal is 140.57.




Figure 6.2: Posterior on the echo parameters for a quiet ($\text{SNR} \sim 4$) simulated signal. The signal has GW150914-like parameters at a fiducial distance of 400 Mpc. An amplitude factor of 0.0125 is used for the echoes. Again, the injected values are shown by the red lines, while points are coloured by the SNR at that point in the parameter space. The log Bayes factor for this injection is -1.55, indicating what to expect when the signal is indistinguishable from noise. The prior ranges are largely saturated and lines appear in the 1D marginal posterior for t_{echo} .




Figure 6.3: The 90% credible regions of the 2D marginal posteriors of Δt_{echo} and γ for GW150914-like simulated signals. Shown are a range of echo amplitudes (relative to the peak amplitude of the original signal) A . The injected values are given by the horizontal and vertical red lines. For small values of A , the 90% contour covers most of the prior range, whereas for larger amplitudes the contours narrow down onto the injected values.




Figure 6.4: Values of the maximum likelihood SNR and log Bayes factors for GW150914-based injections with amplitudes from 0.025 to 0.4 at a distance of 400 Mpc. A linear fit is possible through the SNR points down to an amplitude around 0.1. The log Bayes factor is negative for amplitude values below ~ 0.07 (indicating formal preference for Gaussian noise over the echoes hypothesis).

likelihood SNR of only 3.8. The log Bayes factor for this injection is -1.55 showing a preference for the pure Gaussian noise hypothesis. In this case most echo parameters are not well recovered and their posterior distributions are close to the original prior distributions.

Figure 6.3 shows the recovery of γ and Δt_{echo} for a range of different injection amplitudes. As the amplitude is increased, the recovered value is increasingly constrained to the injected value.

The log Bayes factor and likelihood SNR for injected signals with different amplitudes is shown in Fig. 6.4. Here we find that below an injected amplitude of ~ 0.1 the recovered maximum likelihood SNR no longer falls off linearly, and flattens out to an approximately constant value of ~ 4 , independently of the signal injected. At amplitudes below ~ 0.07 the log Bayes factor becomes negative.

A plot of the recovered amplitudes versus injected amplitudes is shown in Fig. 6.5. This figure can be compared with Fig. 4 of [3], which also shows recovered amplitudes relative to injected amplitudes for a template bank search. In that work it was found that below a certain injection strength, the recovered echo amplitude was no longer reliable using the template bank method. Our results here are consistent with that finding. This gives a clear sense of the amplitude of signals which it is possible to reliably recover, relative to the detector noise level. Signals below this level, such as the one claimed in [158], cannot be clearly distinguished from noise.

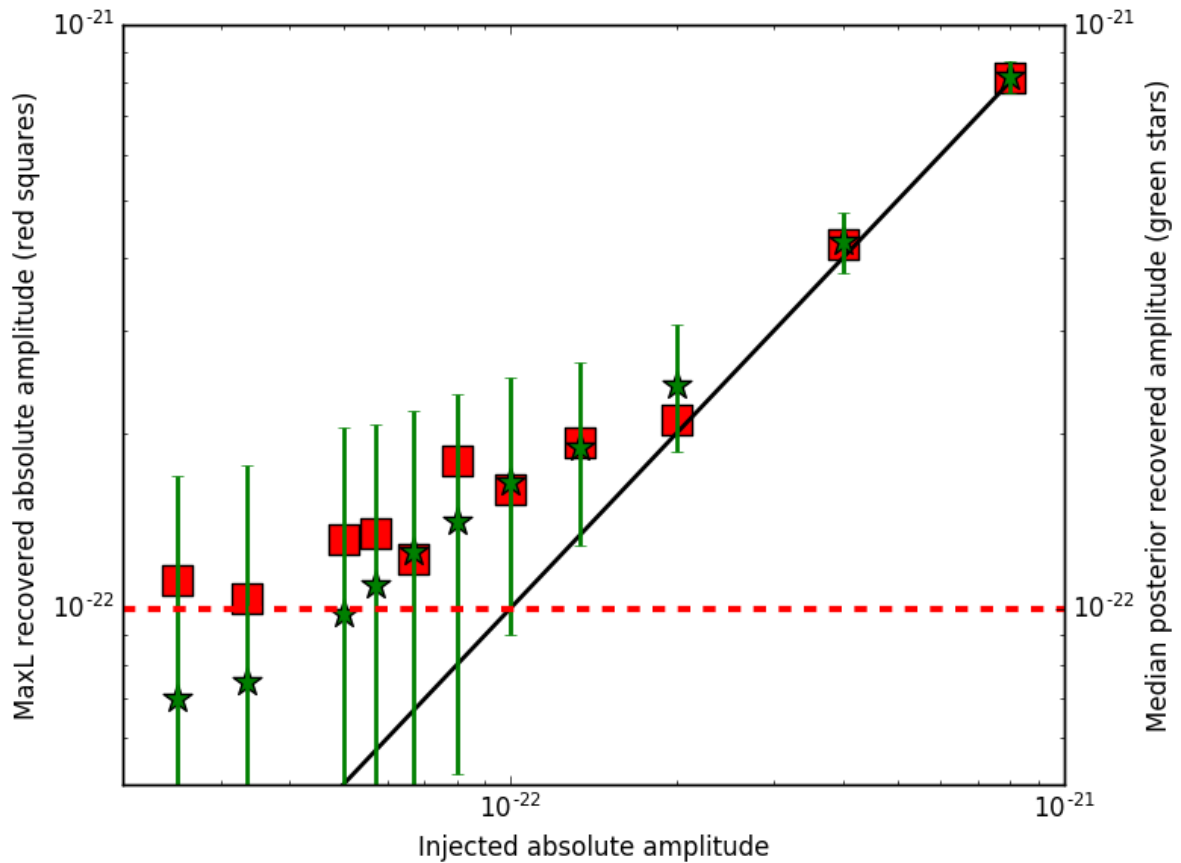


Figure 6.5: Injected versus recovered amplitudes expressed in terms of strain values for GW150914-based injections. At lower injected amplitude values the recovered amplitude and hence the recovered SNR, saturate around an amplitude value of 10^{-22} . Median posterior values are given by the green stars with 90% credible intervals given by the green bars. The amplitudes of the maximum likelihood waveforms are given by the red boxes. The red dashed line is the approximate amplitude of the signal claimed in [158].

6.4 Events in the first observing run

The developed pipeline can be run directly on data immediately after the observed GW events (without injections). We show results for the three events of the first LIGO observing run in Table 6.2. This shows that Gaussian noise is favoured over the echoes hypothesis for GW150914 with a log Bayes factor of ~ -1.81 . GW150914 is the loudest binary black hole merger yet detected. A corner plot of the posterior distributions for the echo parameters for GW150914 is shown in Fig. 6.6. The 90% credible interval for the marginalised posterior of the parameter γ is almost as wide as the prior range. The posterior of the amplitude, A , prefers lower values of the amplitude. The posterior for t_{echo} shows distinct lines at certain values of time. These lines are unlikely to be associated with an astrophysical signal and are also seen in tests on simulated Gaussian noise with the same pipeline.

As seen in Table 6.2, both GW151226 and GW151012 prefer the echoes hypothesis over Gaussian noise, but only marginally. The log Bayes factor for GW151012 is ~ 1.25 , indicating that the echoes hypothesis is ~ 3.5 times more likely than the Gaussian noise hypothesis. As mentioned earlier, the detector noise is known not to be truly Gaussian for the LIGO detectors [110]. To estimate how likely the GW151012 Bayes factor is in true detector noise, we performed 20 background tests on off-source data that lies before or after the time of GW151012 at intervals of 50 seconds. Each of these tests is sufficiently separated in time from the others that it will not be contaminated by a common signal. In these background tests, two examples were found with a Bayes factor larger than the result for GW151012 shown in Table 6.2. A total of four intervals returned Bayes factors that favoured the echo hypothesis over Gaussian noise. Backgrounds for similar (but not identical) echoes hypotheses were also studied in [272] which found evidence for significant tails in the distribution of Bayes factors in real detector noise versus simulated Gaussian noise.

While it is interesting to speculate whether a signal model could be developed that postdicts echo signals for certain events, such as GW151012, but not for others, such as GW150914, we do not pursue that here. The argument that GW151012 should be accepted as a genuine binary black hole merger was given recently in [191], however we do not feel that the echoes data for GW151012 is sufficiently strong to seriously entertain a model where GW151012-like events display echoes, but GW150914-like events do not.

The SNR values found for the maximum likelihood templates in Table 6.2 are comparable, although not identical to those found in [158] and [3]. The finite template spacing in the template banks of [158] and [3] causes a minor difference in recovered SNR values. The main differences are the different base IMR waveform employed and the different power spectral density (PSD) used to calculate the matches. Here we have used a PSD computed using standard `pycbc` [55, 278] routines based on Welch’s method. We estimate the PSD by taking the median value over 64 8 second-long segments (each overlapped by 4 seconds), centred on the main event. A different routine was used in [158] and [3] based on examples provided at [251].

With the simplistic hypothesis that all three binary black hole events should show evidence for echo signals in the range of parameters assumed, we can simply add the log Bayes factor together to obtain an overall log Bayes factor for this model relative to Gaussian noise of $-1.81 + 1.25 + 0.42 = -0.14$. This is negative, indicating a preference for Gaussian noise, but not by much. It is worth noting that this simplistic combination assumes that the values for the echo parameters can lie anywhere in their prior ranges for any of the three events. This is slightly different from the hypothesis of [158] that assumes certain echo parameters should have the same value in all three events. With a hypothesis that fixes the values of certain echo parameters to be the same in all cases, it is possible that the overall Bayes factor would be different from our result. But this issue also raises the question of how these common parameters should be




Figure 6.6: Corner plot for ADA echoes templates in data just after the merger of GW150914. The log Bayes factor for this data is -1.81 , indicating a preference for the Gaussian noise hypothesis over the Echoes hypothesis. Lines are visible in the t_{echo} subplots, but the SNR associated with these is still not high. These lines are also seen in quiet injections into noise Fig. 6.2 and even tests of the pipeline on simulated Gaussian noise (not shown).

Event	Log Bayes factor	Max. SNR
GW150914	-1.8056	2.86
GW151012	1.2499	5.57
GW151226	0.4186	4.07

Table 6.2: Table of Bayes factor results. Negative values indicate that the Gaussian noise hypothesis is preferred. Positive values indicate that the echoes hypothesis is preferred after marginalisation over parameters. For an approximate indication of scale, log Bayes values with magnitude < 1 are “not worth more than a bare mention” in the nomenclature of [119].

fixed; a simple maximisation of the sum of the squares of the template SNRs as in [158], or as a maximisation or marginalisation of the likelihood function introduced here. The two events, GW151012 and GW151226, have lower amplitudes for the main signal than for GW150914 and thus echoes signals with the same relative amplitude would have a lower absolute amplitude relative to the ambient noise [3]. We defer investigation of these subtle issues to future work.

6.5 Discussion and conclusions

With knowledge of how sensitive our pipeline is from the injection test runs of Sec. 6.3 we can determine the amplitude of echoes that would have been detectable had they been present in the data. This allows us to place a bound on the amplitude of echoes emitted from the events considered here. We remind the reader that bounds from our search only relate to the family of echo waveforms considered here. These are based on the waveform model assumed in [158] and adopting the prior ranges of Table 6.2.

As shown in Fig. 6.6, the posterior amplitude recovery for GW150914 has a 90% confidence interval from 0.0051 to 0.1789. For this realisation of the noise, amplitudes above 0.1789 are ruled out at 90% confidence. This is consistent with the injection studies depicted in Fig. 6.4 which show that (for noise at a different time, 100 seconds after the main event) echo signals with amplitudes $\gtrsim 0.15$ would have been unambiguously identified in the data.

Echo signals of amplitude 0.1 relative to GW150914 would correspond to approximately 0.1 solar masses of energy being reflected from near the black hole horizon [256]. Although this value of the amplitude is not conclusively ruled out with the current data, an amplitude as high as 0.2 is conclusively ruled out by our results.

For numerical simulations of systems similar to GW150914 within general relativity, it is conservatively estimated that ~ 2 solar masses of gravitational energy flows across the horizon [127]. Our constraints here on the amplitude of echoes within the model of [158] suggest that at most 20% of this energy is being reflected by near-horizon structure and re-emitted as echoes. In the context of this specific model, it implies that a significant amount of energy is still being lost into the black hole. Our results in this context are qualitatively similar to those in [279], although an exact quantitative comparison is not possible because of the different models assumed.

We have seen that Gaussian noise is preferred over ADA echo-like signals in the data of GW150914. Although there is some evidence of echoes in GW151012 and GW151226, as both show positive log Bayes factors, this evidence is not very strong. Sampling the true detector noise by running over off-source times, shows that the log Bayes factor found for GW151012 is not unusual, and contrarily to [158], cannot be reliably distinguished from pure noise. A number of improved echo waveform models have been proposed; we defer running with these on further

events to future work.

Acknowledgements

O.B. acknowledges the National Science Foundation (NSF) for financial support from Grant No. PHY-1607520. This work was supported by the Max Planck Gesellschaft and we thank the Atlas cluster computing team at AEI Hanover. This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (<https://gw-openscience.org>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

6.6 Appendix

6.6.1 Fiducial IMR waveform parameters

We list here the parameters of the base IMR waveforms used to construct the echo templates both for injections and for the searches. These values are obtained from the maximum likelihood values of [99].

Parameter	GW150914	GW151012	GW151226
mass1	39.03	22.87	18.80
mass2	32.06	18.67	6.92
spin1x	-0.87	0.12	0.44
spin1y	-0.43	0.19	0.59
spin1z	-0.06	-0.20	0.33
spin2x	-0.11	0.018	0.00
spin2y	-0.03	-0.019	-0.017
spin2z	-0.15	0.062	0.0033
distance	477	751	315
ra	1.57	0.65	2.23
dec	-1.27	0.069	0.98
tc	1126259462.42	1128678900.46	1135136350.66
polarisation	5.99	5.64	1.43
inclination	2.91	2.32	0.68
coa_phase	0.69	4.44	1.64
phase_shift	-0.92	-0.91	1.86

7 Constraining the Kerr-geometry of GW150914 with horizonless modes

We obtain stringent constraints on near-horizon deviations of a black hole from the Kerr geometry by performing a long-duration Bayesian analysis of the gravitational-wave data immediately following GW150914. GW150914 was caused by a binary system that merged to form a final compact object. We parameterise deviations of this object from a Kerr black hole by modifying its boundary conditions from full absorption to full reflection, thereby modelling it as a horizonless ultracompact object. Such modifications result in the emission of long-lived monochromatic quasinormal modes after the merger. These modes would extract energy on the order of a few solar masses from the final object, making them observable by LIGO. By putting bounds on the existence of these modes, we show that the Kerr geometry is not modified in this way down to distances as small as 5.8×10^{-19} meters away from the horizon. Our results indicate that the post-merger object of GW150914 is well described as a black hole of the Kerr geometry.¹

7.1 Introduction

General relativity (GR) predicts the existence of black holes which possess a horizon, a surface that acts as a perfect absorber. The exterior vacuum geometry of stationary rotating black holes in GR is that of the Kerr geometry [16].

A binary black hole merger results in a rotating, perturbed black hole which then relaxes to equilibrium by emitting gravitational waves (GWs) at specific frequencies, the frequencies of its ringdown or quasinormal modes (QNMs). In GR, the spectrum of the QNMs is completely determined by the black hole mass and spin. Previous QNM analyses of the GW ringdown from binary black hole mergers have yielded broad consistency with the remnant being a Kerr black hole [2, 69, 71, 72, 172]. The first overtone of the dominant QNM was found in LIGO's GW150914 event by Ref. [172] (and in other events in Ref. [71, 72]). In Ref. [2], a subdominant fundamental mode was found in GW190521. In all cases, the recovered modes were consistent with GR.

Here we present a method for testing the validity of the Kerr geometry down to microscopic distances away from the horizon, in the region where gravity becomes strong, and apply it to the LIGO GW150914 data. In contrast to a black hole, a horizonless object is not a perfect absorber of GWs, and could be distinguished from a black hole by its post-merger GW emission. As the interaction of the GWs with the interior matter of the object is expected to be weak, the infalling waves could propagate into the object and re-emerge after some time delay [280]. Thus, to model a horizonless object, we modify the boundary conditions to allow GW reflection at a surface that is a relative distance $\epsilon \ll 1$ away from the would be horizon (see Eq. (7.3)). This description is equivalent to the wave passing through the interior and leads to the same functional dependence [156].

Imposing boundary conditions in the Kerr geometry that allow reflection near the horizon leads to the appearance of additional modes, the QNMs of the horizonless object. The initial ringdown modes are very similar to those of a Kerr black hole, as they result from excitations of

¹This chapter is an adaptation of the work in the pre-print [5] as described in Section 1.2.

the photon sphere. The additional modes are long-lived, nearly monochromatic GWs, expected to appear after a time delay and dominate the emission at times long after the merger (see for example Fig. 3 of [281]). Their frequency is proportional to the rotational frequency of the black hole, while their lifetime is $\tau \sim M |\ln \epsilon|^2$, where M is the black hole mass ($\ln = \log_e$) [156, 175, 282]. For GW150914, the frequency of such modes would be ~ 210 Hz — well within LIGO’s sensitive band — with lifetimes in the range $30 \text{ s} \lesssim \tau \lesssim 8000 \text{ s}$ (assuming mass and spin estimates from [190], and allowing $\epsilon \in [10^{-45}, 10^{-5}]$).

The amplitude of the additional modes is determined by the total energy falling in through the initially formed trapped surface [127, 246]. Since about the same amount of energy falls into the trapped surface as is emitted during the merger, [127, 246], we expect that the same amount will be channeled to the additional QNMs. In GW150914, we estimate that the total amount of extracted energy is $\sim 3 M_\odot$ [see Eq. (7.10)] and therefore should be detectable with high signal-to-noise ratio (SNR). The additional signal is weak but extremely long lived. By using a long integration time we can place stringent constraints on ϵ .

In this work, we directly constrain ϵ by performing a long-duration Bayesian analysis of the GW150914 post-merger data. We develop new parameter estimation methods to overcome the challenges posed by the long duration of the signal and analysed data. Through these, we can probe the near-horizon region of a rotating black hole with unprecedented accuracy, and constrain its geometry down to microscopic distances away from the horizon.

The additional ringdown modes are related to the so-called black hole echoes [153, 176], as both are associated with reflection from the black hole and produce a long-duration post-merger GW signal. Indeed, the additional modes can be seen as the superposition of many echoes’ low-frequency components at very late times. However, the additional modes differ in some significant aspects from models for echoes previously considered. In these echoes models, the initial merger signal repeats itself at regular intervals, with a decay rate that is treated as a free parameter. The model has five free parameters in total. In our model, the resulting GW signal is a damped sinusoid which resembles in form the standard black hole ringdown modes. The frequency, decay time, and amplitude are all determined by the modified boundary conditions at the reflecting surface, and the mass and angular momentum of the black hole.

Several echoes searches were performed in [3, 4, 71, 158, 177, 256, 262]. While some of the searches reported evidence for near-horizon structure [158, 262], others [3, 4, 177, 256] found low statistical evidence for echoes. An extended search that uses the model proposed in [71] was done using the LIGO-Virgo gravitational-waves transient-catalog-2 (GWTC-2) for 31 black hole events. That search reported no statistically significant evidence for echoes in the data. Some implicit constraints on ϵ can be deduced from the null results of these searches [3, 4, 122, 256]. However, these constraints depend on several uncertain modelling assumptions.

Previous efforts to constrain ϵ using electromagnetic emission from black holes were based on the idea that if the horizon of a black hole is replaced by a hard surface at a fractional distance ϵ away from the horizon, the electromagnetic emission from such a surface can be observed and could be used to place limits on the luminosity of black holes [283]. Several analyses [284–286] eventually led to impressive nominal results $\epsilon \lesssim 10^{-16}$ [283, 287] (corresponding to a distance of $\sim 10^{-6}$ m). However, obtaining concrete limits using this method requires making many assumptions [283], including about the surrounding matter. For additional discussions of the caveats and limitations of this method, see [283, 288].

Fortunately, assuming that the Einstein equivalence principle holds, the dynamics of GWs are only sensitive to the geometry, and the interaction between GWs and matter is extremely weak, and therefore independent of specific environmental models. This allows us to obtain extremely strong constraints: we find $\epsilon < 3.3 \times 10^{-24}$ (90%-credible interval), which corresponds to a distance between the reflective surface and the Kerr horizon of no more than 5.8×10^{-19} m

in the Boyer-Lindquist coordinate distance.

7.2 Theoretical framework

The invariant line-element of a Kerr black hole in Boyer-Lindquist coordinates is

$$ds^2 = - \left(1 - \frac{2Mr}{\Sigma} \right) dt^2 - \frac{4Mr}{\Sigma} a \sin^2 \theta d\phi dt + \frac{\Sigma}{\Delta} dr^2 + \Sigma d\theta^2 + \left((r^2 + a^2) \sin^2 \theta + \frac{2Mr}{\Sigma} a^2 \sin^4 \theta \right) d\phi^2 . \quad (7.1)$$

Here a is the spin parameter, $\Sigma = r^2 + a^2 \cos^2 \theta$, and $\Delta = r^2 + a^2 - 2Mr = (r - r_+)(r - r_-)$, with $r_{\pm} = M \pm \sqrt{M^2 - a^2}$. The angular velocity of the horizon, Ω , is related to a through $\Omega = (a/M)/2r_+ = \chi/2r_+$, with the dimensionless spin parameter $\chi = a/M$.

Gravitational perturbations in the exterior vacuum Kerr geometry obey the Teukolsky equations [53, 120], which reduce to an eigenvalue problem when regularity of the solution is imposed. The resulting radial equation can be simplified by changing variables [121] and using tortoise coordinates $dr_*/dr = (r^2 + a^2)/\Delta$, taking the final form

$$\frac{d^2 {}_s\Psi_{\ell m}}{dr_*^2} - V(r, \omega) {}_s\Psi_{\ell m} = 0 . \quad (7.2)$$

For gravitational perturbations, the spin is $s = \pm 2$. In tortoise coordinates, the spatial coordinates are Euclidean and hence Eq. (7.2) describes potential scattering in flat space. The expression for the effective potential $V(r, \omega)$ can be found in [121].

We find the spectrum of the additional QNMs by imposing boundary conditions at infinity and at the near-horizon surface r_{NH} , which is at a relative distance ϵ above r_+ ,

$$\epsilon = \frac{r_{NH} - r_+}{r_+} . \quad (7.3)$$

The solutions of Eq. (7.2) behave approximately as follows,

$$\Psi \sim e^{i\omega r_*}, \quad r_* \rightarrow \infty, \quad (7.4)$$

$$\Psi \sim e^{-i\omega r_*} + \mathcal{R} e^{i\omega r_*}, \quad r_* \rightarrow r_*(r_{NH}), \quad (7.5)$$

where \mathcal{R} is the reflection coefficient of the surface, and the complex frequency $\omega = \omega_R + i\omega_I$ has to satisfy Eq. (7.2). The real and imaginary part of ω are related to the frequency f and damping time τ of the QNM by $\omega_R = 2\pi f$ and $\omega_I^{-1} = \tau$. An additional unknown phase accounts for the propagation through the interior and is absorbed into the phase ϕ in the waveform of Eq. (7.8), while we marginalise over the phase of the signal in the numerical analysis.

For a Kerr black hole, the reflection coefficient is zero at the horizon. We modify the boundary conditions at $r = r_{NH}$ such that \mathcal{R} is non-vanishing. In general, \mathcal{R} may depend on the frequency. However, for the range in ϵ that we consider here, the possible signal frequencies are limited to a small range near $m\Omega$, with $M|\omega_R - m\Omega| \ll 1$. Therefore, we take \mathcal{R} to be a constant.

We choose a perfectly reflecting boundary condition, $\mathcal{R} = 1$. This choice is justified on grounds that if the Einstein equivalence principle holds for the interaction of GWs with the black hole, then the object's surface can only either be fully absorbing ($\mathcal{R} \ll 1$), or fully reflecting ($1 - \mathcal{R} \ll 1$). Partial absorption ($0 < \mathcal{R} < 1$) would require the object to contain a membrane or other viscous fluid capable of dissipating GWs [243, 289]. However, such models only yield non-negligible absorption when unknown exotic matter is considered [289, 290]. Heuristically, if

the matter is not exotic, then the absorption through the object's surface scales as $1/\tau$. This means that the deviation from total reflection should scale as $r_+/\tau \ll 1$, which means that $1 - \mathcal{R} \ll 1$. Conversely, firewall and fuzzball models yield almost full absorption due to the large density of black hole microstates and the small energy gaps between them [291, 292]. This makes them functionally indistinguishable from classical GR black holes. We therefore focus on the pure reflection case and fix $\mathcal{R} = 1$. A more detailed argument is found in the Appendix.

For perfect reflection and $s = -2$, the solution for the dominant contribution $\ell = 2$ can be found analytically [156, 175] (also see [282]), yielding

$$\omega_R \simeq m\Omega \pm \frac{\pi}{2|r_*^0|} (\nu + 1), \quad (7.6)$$

$$\omega_I \simeq \frac{2M(\omega_R - m\Omega)r_+}{225|r_*^0|(r_+ - r_-)} [\omega_R(r_+ - r_-)]^5. \quad (7.7)$$

Here, $|r_*^0| \sim \int dr \sqrt{g_{rr}} \sim M(1 + (1 - \chi^2)^{-1/2}) |\ln \epsilon|$; we choose the dominant overtone number $\nu = 1$ (not to be confused with the BH QNM-overtone number n). The amount of energy stored in the higher overtones $\nu \geq 2$ is expected to be much lower than that stored in the dominant mode. For very small ϵ , the remaining modes approach the $\nu = 1$ mode in frequency.

The solutions contain two types of signals, damped or superradiant for a positive or negative sign of $\omega_R(\omega_R - m\Omega)$, respectively [156, 175]. Only two absolute values of ω_R appear for each value of $|m|$, as changing the sign of both m and the second term in Eq.(7.6) in turn only changes the sign of ω_R .

Alternatively, the damping properties of the modes can be explained from an interior perspective where, similar to [157], the scattering cross-section of the outgoing waves is positive and leads to a damped rather than amplified waveform, see [175] and Appendix for further details. As noted in [157, 280], a heuristic description is that the would-be BH is effectively in an excited state and it decays to equilibrium with a lifetime τ .

We focus on the case $\epsilon \ll 1$ such that $|\ln \epsilon| \gg 1$. Then Eq. (7.6) is mostly governed by the angular frequency of the object $\omega_R \approx \chi/r_+$ and Eq. (7.7) corresponds to a large damping time $\tau \sim r_+ |\ln \epsilon|^2$. The large damping time allows us to constrain ϵ by analysing a long duration of post-merger data.

7.3 Signal model

Our signal model reflects the damped oscillatory properties of the modes, and relies on the knowledge of the initial merger phases from which we can extract all other parameters of the black hole. We then assume a smooth transition between the early to late time phases [293, 294].

We use a quasi-normal mode to model the late-time post merger signal,

$$(h_+ + ih_\times)(t) = {}_{-2}S_{\ell m}(\iota, \varphi) A e^{-t/\tau} e^{i(2\pi f t + \phi)} \Theta(t - t_0), \quad (7.8)$$

which is parameterised by five intrinsic parameters. These are the amplitude A , frequency $f = \omega_R/2\pi$, damping time $\tau = \omega_I^{-1}$ and initial phase ϕ of the damped sinusoid, and a start time t_0 of the signal. If the prompt ringdown emission occurs at $t = 0$, then t_0 describes the time delay between this and the start of the additional QNM signal. The spin-weighted spheroidal harmonics ${}_{-2}S_{\ell m}$ depend on the inclination ι and azimuthal angle φ . Here, we consider the dominant $l = m = 2$ mode and approximate the spheroidal harmonics by spin-weighted spherical harmonics [138, 139]. For $\epsilon \ll 1$, the frequency ω_R in Eq. (7.6) is governed




Figure 7.1: An example of the proposed waveform with parameters corresponding to the maximum likelihood values for the inspiral-merger-ringdown analysis of GW150914, with $\epsilon = 10^{-25}$. The upper panels show the waveform in the frequency domain, while the lower panels show the same waveform in the time domain. In the right panels, the same waveform is shown as on the left, with a restricted range in frequency and time, respectively, to improve visibility of the waveform features.

by the object's angular velocity,

$$M\omega_R = \frac{\chi}{(1 + \sqrt{1 - \chi^2})} + \frac{\pi\sqrt{1 - \chi^2}}{|\ln \epsilon| (1 + \sqrt{1 - \chi^2})}. \quad (7.9)$$

For a set of example parameters compatible with GW150914, $M \approx 62M_\odot$, $\chi \approx 0.67$, and for $\epsilon = 10^{-25}$, we would find $M\omega_R \approx 0.4$ and $f \approx 211$ Hz. An example of the waveform is shown in Figure 7.1. This range of parameters guarantees the validity of Eq. (7.7), since as pointed out in [175], the derivation relies on the assumption that $M\omega_R < 1$, $a\omega_R < 1$ and $M(\omega_R - m\Omega) \sim \frac{1}{|\ln \epsilon|} \ll 1$.

The amplitude A is determined by the total energy (and angular momentum) that is carried away by the GWs to infinity (see [280]). The total emitted energy is determined at the merger [127], we label it by $\Delta E = E_{init} + E_{rot}$. Here, E_{init} is the energy of the infalling gravitational radiation from the merger, while E_{rot} is the rotational energy extracted as this infalling radiation is forced to rotate through frame-dragging. Then, by using the non-relativistic approximation, such that $E_{rot} = \frac{1}{2}E_{init}\Omega^2 r_+^2$, we find

$$\Delta E = E_{init} \left(1 + \frac{\chi^2}{8} \right). \quad (7.10)$$

For the same example parameters, rotational effects lead to a correction of the emitted energy by an increase of $\sim 5\%$ compared to the non-spinning case, yielding $\Delta E \approx 3.2M_\odot$. In the

superradiant case in contrast, most of the rotational energy is extracted by the emitted GW, such that $E_{rot} \sim M\Omega^2 r_+^2 \sim 5M_\odot$.

To calculate the amplitude, we evaluate the emitted energy ΔE by using the leading order GW flux formula,

$$\dot{E}_{GW} = \frac{D_L^2}{32\pi} \int \langle \dot{h}_{\mu\nu} \dot{h}^{\mu\nu} \rangle d\Omega. \quad (7.11)$$

Here the dot denotes a time derivative, D_L is the luminosity distance, $h_{\mu\nu}$ is the waveform in the transverse–traceless gauge, $d\Omega$ is an element of solid angle, and angular brackets denote averaging over short wavelengths. We approximate the integral in Eq. (7.11) by noticing that the emitted GWs are approximately monochromatic with $\omega_R \simeq 2\Omega$, yielding $\dot{E}_{GW} \approx \frac{1}{4} D_L^2 \langle |\dot{h}|^2 \rangle$. Then, by taking $h(t)$ from Eq. (7.8) and for $\epsilon \ll 1$ such that $\omega_R \tau \gg 1$, the final expression for the amplitude becomes

$$A = \frac{4}{\omega_R D_L} \left(\frac{\Delta E}{\tau} \right)^{1/2}. \quad (7.12)$$

In Eq. (7.12) the parameters ω_R , τ and ΔE are given in Eqs. (7.6), (7.7), and (7.10), respectively. The explicit form of ω_R is given in Eq. (7.9), while for τ it is

$$\tau = \frac{225M}{32\pi} \left(\frac{1 + \sqrt{1 - \chi^2}}{\sqrt{1 - \chi^2}} \right)^6 \frac{|\ln \epsilon|^7}{\left(\chi |\ln \epsilon| + \pi \sqrt{1 - \chi^2} \right)^5} \quad (7.13)$$

We fix the parameter t_0 to an arbitrary value some time after the merger. To prevent contamination of the analysis from the standard ringdown modes, we choose a time that is large compared to the lifetime of these modes, but short compared to the lifetime of the additional signal, $t_0 = 32$ s. Because the amount of energy emitted during this relatively short time is small and because the SNR is determined by the total collected energy, we do not lose much diagnostic power by this choice. As the damping time increases for smaller ϵ , this approximation is more accurate for the expected small values of ϵ .

In addition to ϵ , the parameters varied in the analysis are right ascension α , declination δ , polarisation Ψ , inclination ι , luminosity distance D_L , final mass M , final spin χ , and energy radiated in the primary GW emission, ΔE . Equations (7.9), (7.10), (7.12) and (7.13) then determine the parameters of the damped sinusoid template. The phase ϕ of the signal is marginalised over analytically. We use as priors for the source parameters the posteriors found in [190], calculating M , χ and ΔE from the component parameters via fitting formulae to numerical relativity [198–200]. For the only additional parameter of our model, ϵ , we use a log-uniform prior in the interval $-45 \leq \log_{10} \epsilon \leq -5$.

We use Bayesian methods to estimate the signal parameters from the data. The toolkit `PyCBC Inference` [99, 108] is used to compute the likelihood and estimate the posterior probability distributions. The parameter space is sampled using the parallel-tempered Markov-chain Monte Carlo sampler `emcee_pt` [114, 117].

We modify the standard parameter estimation analysis to prevent influences from boundary effects. The expected signal persists for a longer time than the currently manageable duration of the analysis. We therefore need to restrict the time series data to a shorter time window, which introduces a discontinuity from the sharp cut-off at the window edges. This leads to artefacts in the time-domain response function of the whitening filter. To avoid this, we remove the times containing these artefacts, and we employ a heterodyning procedure to reduce the computational cost of generating long template waveforms (see Appendix).

7.4 Results

Applying our analysis to 128 s of data starting 32 s after GW150914 yields the posterior on ϵ as shown in Fig. 7.2. The results are consistent with the absence of the searched signal, as the posterior peaks toward the lower boundary. Our upper bound on the 90%-credible interval is $\log_{10} \epsilon = -23.5$. For the post-merger black hole of GW150914, this bound corresponds to a distance between the reflective surface and the Kerr event horizon of no more than 5.8×10^{-19} m in the Boyer-Lindquist coordinate distance. The corresponding proper distance is 2.8×10^{-7} m.

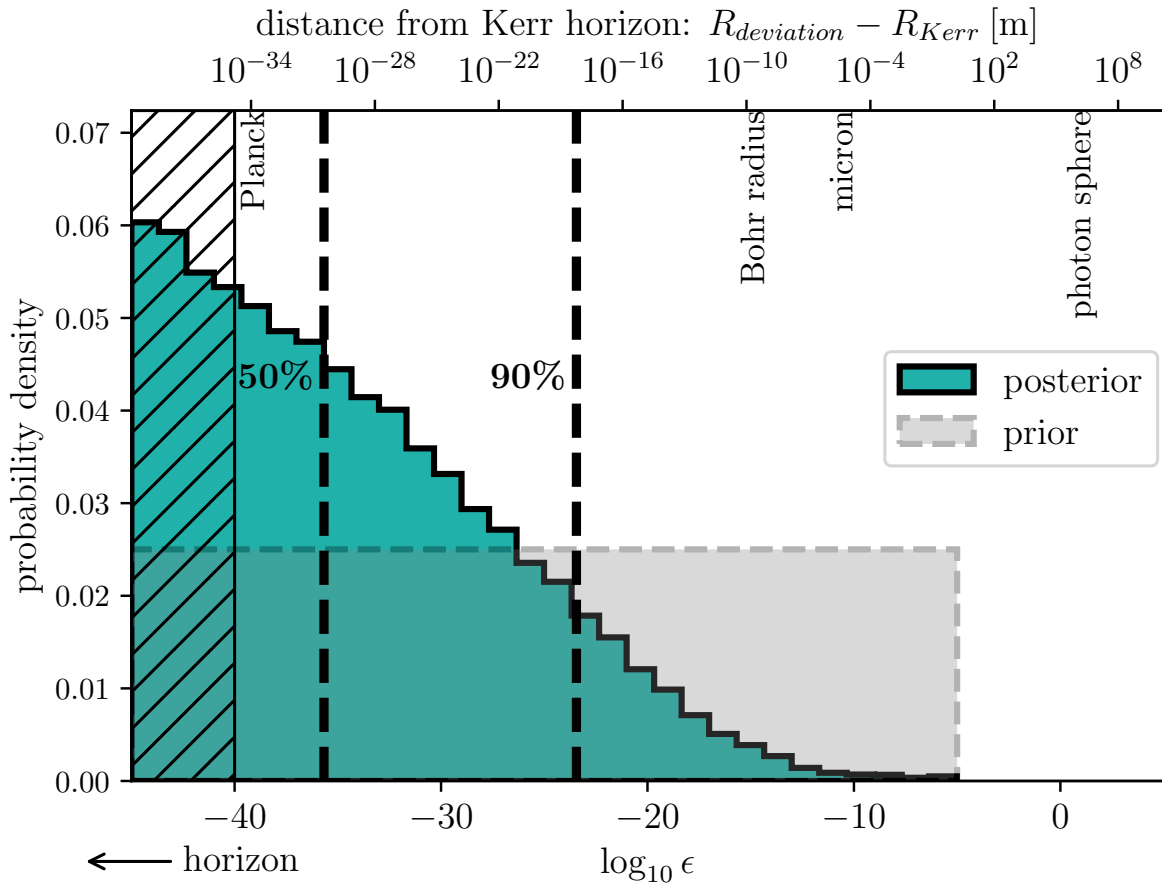


Figure 7.2: The histogram shows the marginal posterior for the fractional deviation from the Kerr geometry, ϵ , measured for the proposed signal for GW150914. The prior for $\log_{10} \epsilon$ is flat, as shown in the shaded region. The dashed lines mark the one-sided 50th and 90th percentile upper bound. On the top axis the coordinate distance between reflective surface and horizon corresponding to ϵ is shown for the post-merger black hole in GW150914, and hatching indicates distances below the Planck length. As the distance posterior is virtually identical to the posterior for $\log_{10} \epsilon$, we only show the latter and use the maximum likelihood values for mass and spin from [190] to convert from $\log_{10} \epsilon$ to the distance scale.

To validate this result, we repeat the analysis on data before GW150914, when no signal is expected, as well as on Gaussian noise. We also inject a simulated signal with $\log_{10} \epsilon = -21$ into detector noise to verify that the analysis can detect a louder signal when present. We find that the posterior on ϵ does peak toward the injected value in the latter case, whereas in noise the posterior and limits are similar to what we obtain for the GW150914 post-merger data (see




Figure 7.3: The 90% bounds placed on ϵ for the analysis of different durations of data are shown as blue dots. A curve of the form $T(|\ln \epsilon|) \sim c|\ln(\epsilon)|^2 \ln\left(1 - (a - b/|\ln(\epsilon)|)^2\right)$, with constants a, b, c , is fitted to the retrieved bounds (see Appendix). The fit asymptotically approaches $\log_{10} \epsilon \approx -23.7$. For longer segments of data, the bound increases again (empty circle), as the posterior begins to be dominated by lines in the power spectral density of the noise.

Appendix).

To investigate the effect of the amount of time analysed on the bound on ϵ , we repeat the analysis using time segments 16 s, 32 s, 64 s, and 256 s. The results are shown in Fig. 7.3. As expected, the upper bound on ϵ increases as we analyse shorter time segments than the 128 s we use above. This suggests that analysing longer times would yield even better limits. However, in the 256 s analysis the bound is worse than what we obtain with 128 s. This is because the posterior on ϵ begins to be dominated by lines in the power spectral density of the noise as the analysis time increases, leading to weaker constraints. Overcoming this would require removing lines from the data, which is outside the scope of this work.

Using the 16 s to 128 s results (for which lines are not an issue) we estimate the best limit that could theoretically be obtained with GW150914. Fitting the expected relationship between observation time and ϵ , we find that the best 90%-constraint using arbitrary lengths of data would be $\log_{10} \epsilon \approx -23.7$. This limit arises due to a combination of the SNR of GW150914, the energy available in this system that could be converted to the long duration QNMs, and the transfer function of the analysis. Since the potential signal is a damped sinusoid, the recoverable SNR asymptotes to a fixed value for infinite observation time. This in turn puts a limitation on the smallest ϵ that can be measured. While the transfer function is unknown in detail, we can use the relation of SNR and analysis duration as the basis for a fit to the measured bounds (see Appendix). As can be seen in Fig. 7.3, we are close to the resulting empirical limit with the 128 s analysis time.

7.5 Conclusion and outlook

We performed the first long-duration QNM analysis of the post-merger data of GW150914, and ruled out the existence of long-lived additional QNMs for a broad class of alternative compact objects. For this class of models, we put a bound on the validity of the Kerr geometry down to fractional distances from the horizon as small as $\epsilon < 3.3 \times 10^{-24}$, which is equivalent to a coordinate distance $< 5.8 \times 10^{-19}$ m. Our result improves existing bounds by many orders of magnitude and indicates that the GW150914 post-merger object is a black hole that is well described by the Kerr geometry.

Based on the fit in Fig. 7.3, we conclude that to significantly improve our bounds will require a black hole merger with larger SNR than GW150914.

By combining results over multiple events, and with improving sensitivity of future detectors, it should be possible to eventually probe spacetime geometry down to Planck scales above the horizon. This could provide confirmation of the Kerr nature of astrophysical black holes all the way to their horizons.

Acknowledgements

The authors thank Ofek Birnholtz, Alex B. Nielsen, and Reinhard Prix for valuable discussions, and Paolo Pani and Enrico Barausse for helpful comments. This work benefited from discussions at the ‘‘Gravitational wave searches and parameter estimation in the era of detections’’ workshop held January 12-18, 2020 at Ringberg castle, Tegernsee lake. Calculations were performed on the Atlas computer cluster of the Albert Einstein Institute Hannover. The research of RB and YS was supported by the Israel Science Foundation grant no. 1294/16. The research of YS was supported by the Negev scholarship.

This research has made use of data obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org/>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO Laboratory and Advanced LIGO are funded

by the United States National Science Foundation (NSF) who also gratefully acknowledge the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain.

7.6 Appendix

7.6.1 Justifying the assumption of full reflection

Here we elaborate on the arguments given in the main text and provide further explanations for justifying full reflection.

One can understand, heuristically, the scaling of ω_R and τ , the frequency and decay time of the additional modes. Here, we consider the interpretation of gravitational waves entering the compact object and re-emerging after passing through its centre, equivalent to reflection at the surface. In tortoise coordinates, the near horizon geometry looks flat and Eq. (7.2) can be viewed in terms of a wave propagating in a cavity of length $r_+|\ln \epsilon|$. The scaling of the decay time τ can be understood in terms of ideas that were introduced in [157] and elaborated on in [280]. We briefly review them here and refer the reader to the original articles for further details.

First recall from Eq. (7.6) that the “proper” angular frequency of the additional modes is $\omega_R \sim \frac{1}{r_+|\ln \epsilon|}$. This means that a co-rotating GR external observer would view them as having a wavelength $\lambda \sim r_+|\ln \epsilon|$. The source of the GWs is the ultracompact object which has an area of about $A \sim Mr_+$. The transmission cross-section for such long wavelength modes for an area A is proportional to the ratio A/λ^2 , which scales as $Mr_+/\lambda^2 \sim \frac{1}{|\ln \epsilon|^2}$. The decay time is inversely proportional to the transmission rate, so scales as $\tau \sim |\ln \epsilon|^2$. In the GR-limit, $\epsilon \rightarrow 0$, and so the transmission through the object vanishes. The scaling A/λ^2 results from the assumption that gravity acts equally on all forms of matter according to the Einstein equivalence principle.

The heuristic argument that we have just reviewed can also be applied to the case of imperfect reflection at the surface $r_{NH} = r_+(1 + \epsilon)$. Such scenarios require exotic matter which in some cases may violate fundamental principles [290] and are therefore disfavoured. If the reflection is not parametrically small, which corresponds to nearly full absorption and so, effectively, to a horizon, the mode’s decay time would scale as it does for the case of total reflection. The key point is that partial absorption occurs at the surface r_{NH} . Then, the absorption through this surface would scale as $A/\lambda^2 \sim \frac{1}{|\ln \epsilon|^2}$. When the angular momentum of the GW is taken into account, one finds that the absorption through the object’s outer surface scales precisely as $1/\tau$. This means that the deviation from total reflection should scale similarly. Consequently, $1 - \mathcal{R} \ll 1$ since $r_+/\tau \ll 1$.

In the majority of echo models, the reflection coefficient is an arbitrary constant that is put by hand; see [295] and reference therein. None of the reviewed models elaborate on the underlying mechanism that provides the absorption properties of the would-be black hole. Many of them refer to the fundamental papers that motivate horizon scale corrections, such as the firewall and fuzzballs proposals. However, a closer look reveals that a partial absorption of GWs that is comparable to black hole absorption is an unrealistic situation that is not compatible with fundamental physical properties.

For example, in the firewall-inspired models and the fuzzball proposal [291, 292], due to the large entropy and density of states and the small energy gap between the black-hole microstates, an infalling quantum is almost fully absorbed. Fuzzball absorption is therefore almost identical to the black hole absorption (see [296] for specific examples). In [297] it was argued that (Eq. 3), for $\omega \ll T_H$, with the Hawking temperature T_H ,

$$\mathcal{R} = \exp(-\omega/T_H)(\gamma\omega)^{-\omega/T_H}. \quad (7.14)$$

This means that $\mathcal{R} = 1$ to exponential accuracy, or

$$\mathcal{R} = 1 - \omega/T_H. \quad (7.15)$$

The Hawking temperature T_H in natural units is $1/r_+$, so the intrinsic frequencies that we discuss obey this condition. Similarly in [298], they argue that \mathcal{R} is close to one, except for special frequencies that correspond to the intrinsic frequencies of the quantum black hole, which are of order $1/R_S$, where $R_S = 2M$ is the Schwarzschild radius.

If one wishes to model the object's absorption by an alternative dissipation mechanism as in the membrane paradigm, one needs to assume the existence of an exotic matter. To show this, it is possible to model the object's intrinsic dissipation in terms of its effective viscosity as in the membrane paradigm [243]. In [290] it is shown that the absorption coefficient γ_{abs} scales as $\gamma_{abs} \sim \eta/\eta_{BH}$, where η_{BH} is the BH viscosity. The absorption is negligible for all known matter forms. For example, a highly viscous cold neutron star has $\gamma_{abs} \sim 10^{-8}$, while non-rotating strongly magnetised neutron stars and fictitious highly viscous bosonic matter have $\gamma_{abs} \sim 10^{-4}$. Obviously, for these extreme examples the reflection coefficients $\mathcal{R}^2 = 1 - \gamma_{abs} \simeq 1$. The conclusion is that physical matter cannot mimic the effect of full absorption as the BH membrane does, and is almost completely transparent to GWs.

The orthogonal case is represented by models with approximately full absorption, which are indistinguishable from GR BHs. Since the latter is irrelevant for the post-merger measurements we will focus on the former case, where no absorption is present, and therefore fix the reflection coefficient to one.

We stress that ultracompact objects without a horizon and that obey the equivalence principle are plausible. Examples include anisotropic stars, gravastars, and possibly other compact objects [122, 299–301]. These objects, under some unique circumstances, allow for such reflection properties.

To summarise, the above arguments indicate that having a partially absorbing surface is not a realistic scenario. Therefore, the absorption properties are binary: either full reflection, or complete absorption.

7.6.2 Lower bound for ϵ

To derive the lower bound on epsilon shown in Fig. 7.3 we first recall the formula for the optimal SNR of the signal,

$$\rho^2 = 4 \int_0^\infty \frac{|\tilde{h}(f)|^2}{S_n(f)} df, \quad (7.16)$$

where $\tilde{h}(f)$ is the Fourier transform of Eq. (7.8) and $S_n(f)$ is LIGO's strain sensitivity. Since the signal is approximately monochromatic, Eq. (7.9), the strain sensitivity is constant, $S_n(f) = S_n(f_R)$, where f_R is the signal's frequency. This allows us to use the Plancherel (Parseval's) theorem $\int |\tilde{h}(f)|^2 df = \int |h(t)|^2 dt$ such that the SNR becomes

$$\rho^2 = \frac{4}{S_n(f_R)} \int_0^\infty |h(t)|^2 dt. \quad (7.17)$$

Next, we take the time domain waveform Eq. (7.8) and replace the integral upper bound by some arbitrary time T , which corresponds to the analysis integration time. Integration over time leads to

$$\rho^2 \approx \frac{\tau A^2}{2S_n(f_R)} \left(1 - e^{-2T/\tau}\right). \quad (7.18)$$

We use the amplitude from Eq. (7.12) and assume $\tau^2 \omega_R^2 \gg 1$, finding

$$\rho^2 \approx \frac{8\Delta E}{\omega_R^2 D_L^2 S_n(f_R)} \left(1 - e^{-2T/\tau}\right). \quad (7.19)$$

Finally, we extract the analysis time T ,

$$T(|\ln \epsilon|) \sim c |\ln(\epsilon)|^2 \ln \left(1 - \left(a + \frac{b}{|\ln(\epsilon)|}\right)^2\right), \quad (7.20)$$

where the constants a, b, c are to be determined by the numerical fit to the data points of the 90% credible interval of $\log_{10} \epsilon$, see Fig. 7.3. In general, these constants are functions of the mass, spin, strain, SNR and additional unknown systematic errors. We quantify our lack of knowledge regarding the additional errors by the constants that are determined by the fit. Providing an exact analytical expression for the constant in terms of the physical parameters requires a transfer function that includes the additional errors, nevertheless the fit to data is mostly governed by the logarithmic asymptotic behaviour which is insensitive to these changes. Further details regarding the external effects are provided in the main text. Eventually, the numerical fit for the data is found to be bounded from below by $\log_{10} \epsilon = -23.7$. The interpretation is that, given sufficiently long analysis time, the lowest possible bound that can be measured is $\epsilon = 10^{-23.7}$.

7.6.3 Data analysis details

To analyse data spanning times $[t_0, t_1]$, we first consider a slightly longer stretch of data corresponding to $[t_0 - \Delta t, t_1 + \Delta t]$. The template is generated with duration $(t_1 - t_0) + 2\Delta t$, starting at $t_0 - \Delta t$. Both data and template are Fourier-transformed to the frequency domain and the whitening filter is applied to both. We then transform both back to the time domain and remove the times previously added, $[t_0 - \Delta t, t_0]$ and $(t_1, t_1 + \Delta t]$, from each timeseries. We choose Δt such that the effects of the discontinuity at the boundaries are restricted to the times we remove. The resulting timeseries' are Fourier-transformed back to the frequency domain to calculate the likelihood from the inner product of the whitened data and template. For the damped sinusoid signal, the earlier start time is compensated in the template by increasing the initial amplitude by a factor $\exp[\Delta t/\tau]$.

We use heterodyning to minimise the computational cost of generating signal templates. The frequency-domain representation of the signal is restricted to a very narrow range around its central frequency. This allows us to generate the time-domain signal at a low sampling frequency with lower computational cost, and then shift the frequency-representation of this signal to the desired frequency, equivalent to generating the signal directly at a higher sampling frequency. We first generate a time-domain damped sinusoid signal, with the desired damping time τ , but at frequency $f = 8$ Hz. The sampling rate is chosen to be 32 Hz to accommodate signal components up to Nyquist-frequency of 16 Hz, which encompasses the narrow frequency band of relevant signal content. This signal is then Fourier-transformed to the frequency domain, using the natural frequency sampling-rate for the full duration of the signal, $(t_1 - t_0) + 2\Delta t$. Finally, we shift the signal to the desired frequency f , by placing the content of the frequency series from

range $[0 \text{ Hz}, 16 \text{ Hz}]$ into the range $[f - 8 \text{ Hz}, f + 8 \text{ Hz}]$. The resulting frequency domain waveform is then used for the likelihood calculation.

While a damped sinusoid can be Fourier-transformed analytically, we can not use this to generate the template less expensively directly in the frequency domain. For our analysis, the required time-domain template is a damped sinusoid that ends at the analysis window boundaries. The time-domain representation of the template generated in the frequency domain instead extends to times outside the window. Due to the periodicity assumption of the discrete Fourier transform, the late-time signal from earlier repetitions of the window would appear in the analysis window and overlap with the desired signal. Therefore, we generate the signal in the time domain following the described procedure.

For long analysis durations, the Doppler shift due to the orbital motion of the Earth becomes time-dependent. However, we find this to be negligible for the durations of less than ~ 1000 s used in this analysis, and consider only a static Doppler shift.

7.6.4 Validation with noise and simulated signals

To validate our results we repeat our analysis on off-source detector noise (before GW150914) and on simulated Gaussian noise. These serve to determine the analysis' diagnostic power when no signal is present in the noise. We also add simulated signals to the off-source data to verify the effectiveness of the analysis in detecting known signals.

In each case we analyse 128 seconds of data for the presence of a signal and use 512 seconds of data before the analysis window to estimate the power spectral density (PSD). For the Gaussian noise case, the noise is coloured to agree with the PSD estimated from off-source data at times before GW150914.

For the off-source analysis on real detector noise, we find that the source-parameter posteriors are unchanged from their priors. The posteriors for $\log_{10} \epsilon$, τ and A are consistent with the expectation for noise without a signal. Smaller ϵ corresponds to smaller signal amplitudes and longer damping times, as the same total energy is radiated away over increasingly long times. We find that the posteriors prefer large τ and small A and ϵ , with the latter peaking at the lower prior boundary. The one-sided 90%-credible-interval bound for $\log_{10} \epsilon$ is -23.7 . This is compatible with the 90%-result of $\log_{10} \epsilon \leq -23.5$ for the data following GW150914 within the observed fluctuations from different realisations of noise.

The frequency posterior shows narrow peaks for specific frequencies, often associated with increased SNRs. These peaks appear only for long analysis durations and become more dominant with increasing duration. We can attribute the most prominent peaks to lines in the power spectral density of the noise, such as the 180 Hz harmonic of the 60 Hz line resulting from the AC power grid frequency. Figure 7.5 shows an example of the off-source PSD and marks several frequencies where posterior peaks have been found in the analyses for some starting times and window lengths. The simulated Gaussian noise analysis yields similar results as the off-source detector noise case, with the source-parameter posteriors unchanged from their priors. Large τ and small A and ϵ are preferred, with ϵ peaking at the lower prior boundary, and the 90% bound being $\log_{10} \epsilon = -22.9$.

Both cases show the narrow peaks in the frequency posterior described before. The peaks are more pronounced for real detector noise than for simulated Gaussian noise. Prominent peaks often coincide with visible lines of excess power in the PSD for the detector noise, but not for simulated Gaussian noise coloured with the same PSD. This suggests the presence of non-Gaussian noise features in the real noise that are partially matched by the sinusoidal templates.

Slow variations of the PSD in the detector noise may amplify this effect. For the analysis, the PSD has to be estimated from off-source data, such that slow variations in the line parameters cannot be corrected for in long-duration analyses.

We perform several analyses with simulated signals added to the off-source noise, for example with $\log_{10} \epsilon = -18$ or $\log_{10} \epsilon = -21$. For each simulation, the injected value lies within the 90% credible interval of the ϵ -posterior, and the posterior peaks away from the lower prior boundary and near the correct value. Figure 7.4 shows this for the $\log_{10} \epsilon = -21$ injection. For all injections, the frequency posterior is concentrated in a narrow peak around the injected signal frequency, limited by the frequency-resolution of the data. The one-sided 90% bounds for these injections into detector noise are $\log_{10} \epsilon = -13.3$ and $\log_{10} \epsilon = -14.3$, respectively, larger than those found for noise without a signal. As we are expecting a signal in the injection case, we also use the two-sided credible interval as shown in Figure 7.4. The range recovered then is $\log_{10} \epsilon = -18.3^{+6.3}_{-10.0}$ for the injection with $\log_{10} \epsilon = -18$ and $\log_{10} \epsilon = 19.5^{+6.4}_{-10.4}$ for the injection with $\log_{10} \epsilon = -21$.




Figure 7.4: Same as Fig. 7.2 with a simulated signal injected into detector noise. The histogram shows the marginalised posterior for ϵ , the shaded region is the prior. The red line marks the value for $\log_{10} \epsilon$ of the simulated signal. The dashed lines indicate the (two-sided) 90% credible interval and the median value, respectively. The posterior clearly prefers non-zero values of ϵ in the presence of the simulated signal, and the recovered value for $\log_{10} \epsilon$ is within the 90% credible interval.




Figure 7.5: The blue curve shows the power spectral density of the LIGO Hanford detector data as estimated in the analysis from 512 seconds of data before GW150914. Several line-like features are visible at specific frequencies, indicating excess power. Dashed lines mark several of the frequencies where the posterior for f , the central frequency of the damped sinusoidal waveform template, peaks narrowly for some analyses, depending on start time and window duration. These coincide with the lines in the PSD, suggesting that noise features corresponding to the lines are matched by templates in the analysis. These noise features are thus likely non-Gaussian in nature, or vary sufficiently over the duration of the analysis that they cannot be removed entirely through the estimate of the PSD from the data before the analysis window.

8 Conclusions and outlook

Gravitational-wave observations have quickly proven to be a valuable tool in diverse areas, ranging from astronomy and astrophysics to nuclear physics [34, 64, 212, 302–306]. In particular, they provide a new window to study fundamental physics in the strong-gravity regime. The work presented in this thesis contributes several such studies, and extensions are now being pursued.

The ringdown modes of a perturbed compact object are closely tied to its nature, and offer a direct probe of its properties. Compared to the full evolution of a binary system and the merger phase, the individual perturbed object provides a case more accessible to theoretical study.

For black holes, measurement of the ringdown spectrum allows to directly infer the mass and angular momentum. Predictions such as the black hole area-increase law can be tested using the measurement of at least a single mode from this spectrum. Measuring multiple ringdown modes allows to test their consistency with the no-hair theorem. The measurement of the ringdown modes is therefore desirable, but challenging due to their short-lived nature and typically quiet signal. In addition, all modes but the dominant quadrupolar one are suppressed in amplitude, making multi-mode studies even more difficult to achieve.

We study the prospects for multi-mode tests through Bayesian parameter estimation on large populations of simulated events based on observational population models. We estimate the rate of detections where multiple modes are found, and the rate of those suitable to perform no-hair tests. For current detector sensitivities, these are unlikely to occur for reasonable amounts of observation time. Planned near-future sensitivities, however, are estimated to lead to a limited number of multi-mode detections in a few years of collected data, some constraining deviations from general relativity to within $\pm 20\%$ at the 90%-credible level.

We show that combining information from multiple detections significantly improves the prospects of no-hair tests. Assuming that a possible deviation is the same for all events, we perform a hierarchical analysis using the target sensitivity for the next observation run of current instruments. Again, realistic populations of simulated events are analysed, showing that constraints on deviations of about $\pm 10\%$ can likely be placed using a few years of data with this sensitivity. However, this requires to account for systematic biases we discovered in this type of analysis.

Similar hierarchical analyses are now being performed using available detections, but little evidence for subdominant modes has been found [71, 72]. In addition, studies focusing instead on the overtones of the fundamental mode are conducted, allowing to extend the waveform templates to the merger time and access increased signal amplitudes [128]. These can be used for the same type of no-hair and general ringdown-based tests [172–174]. While details in the interpretation of these findings are still being investigated, including overtones seems a promising approach to improve the ringdown modelling near the merger [129].

Our analysis of the ringdown spectrum of GW190521 shows strong evidence for a subdominant angular mode. While this seems surprising in light of our injection studies, the population models used there did not cover this binary’s component masses. We introduce a new method to restrict the parameter estimation to the ringdown data. It prevents contamination from data outside the desired range, introducing no windowing artefacts and no loss of signal. We include one overtone to model the signal near the merger, but focus on the angular mode supported by the data. With this mode, we perform a no-hair test, yielding unexpectedly tight bounds on the

deviation of the subdominant mode.

A large injection study with simulated signals similar to GW190521 has been conducted to solidify the statistical evidence, and the corresponding publication is in preparation. Additionally, our methods are now being applied to the population of detections, further characterising the behaviour of both angular mode- and overtone-based analyses in this population.

The study of modified or alternative compact objects is a very active area of research. Here, again, the individual object under certain perturbations is more accessible to theoretical modelling than the full merger scenario. The ringdown signal of such a perturbed object allows to test these models observationally, constraining both the models and their parameters through their support in the data.

The most direct approach is the search for the specific signal resulting from a given model. This approach was followed for a simplified waveform model in a previous first search for gravitational wave echo signals from alternative compact objects. Our analysis of this first study demonstrates the methodological challenges met in the search for echoes. Only very loud echo signals could possibly be detected through these methods, while the recovery of signal parameters suffers from systematic biases. Extending and improving the estimation methods, we find the statistical significance of the results to be low. However, the presence of the proposed signal is not ruled out by our results.

We apply a Bayesian parameter estimation analysis to the same data with the same signal model. This allows us to measure the support for echoes confidently, finding the data to be consistent with or preferring the noise hypothesis. However, we also place upper bounds on the amplitudes of echoes that are compatible with our results. Further searches have been carried out using this model and methodology, reporting no evidence for these specific echo signals [71, 177].

In these approaches, our results apply only to the specific echo waveform model used in the analysis, and to a limited degree to very similar waveforms. Sufficiently accurate theoretical modelling of the waveform is therefore necessary to confidently constrain the proposed models, and such improved models are under development [179]. Nevertheless, uncertainties remain about many details regarding the nature of alternative objects, and a number of different models have been proposed. The correspondingly large space of waveform morphologies and parameters thus currently makes the targeted searches difficult. Relaxing the restrictions on the waveforms and using more generic echo templates allows to search for more diverse classes of signals, albeit at the cost of reduced sensitivity compared to an accurate waveform [72, 178, 273].

Alternatively, we can focus on signals that are common to a broad class of models for horizonless objects, and can thus place constraints on this class. We follow this approach for their long-lived quasi-normal modes, with new methods developed to accommodate the duration of the signal. Applying these to GW150914 allows us to place a tight bound on deviations from the Kerr geometry for this class of models. Immediate extensions are the application of these methods to further detections and to similar signals proposed for a second type of emission mechanism [307]. For a common deviation parameter, the application to multiple events would allow us to construct a combined posterior as in the study on black hole ringdown populations.

The existing gravitational wave detectors are currently undergoing work to improve their sensitivity for the next observation run. Construction and development of the next generation of instruments and of new space-based observatories are underway. With a growing number of detections and increased sensitivity, the presented approaches are likely to yield improving bounds for deviations from general relativity. Similarly, growing opportunities arise to confidently assess the presence of signatures from alternative compact objects, stressing the need for improved theoretical models. As such, there is a promising future for gravitational-wave observations of black holes to illuminate fundamental physics.

Bibliography

- [1] Miriam Cabero et al. ‘Black hole spectroscopy in the next decade’. In: *Phys. Rev. D* 101.6 (2020), p. 064044. DOI: 10.1103/PhysRevD.101.064044. arXiv: 1911.01361 [gr-qc].
- [2] Collin D. Capano et al. *Observation of a multimode quasi-normal spectrum from a perturbed black hole*. May 2021. arXiv: 2105.05238 [gr-qc].
- [3] Julian Westerweck et al. ‘Low significance of evidence for black hole echoes in gravitational wave data’. In: *Phys. Rev. D* 97.12 (2018), p. 124037. DOI: 10.1103/PhysRevD.97.124037. arXiv: 1712.09966 [gr-qc].
- [4] Alex B. Nielsen, Collin D. Capano, Ofek Birnholtz and Julian Westerweck. ‘Parameter estimation and statistical significance of echoes following black hole signals in the first Advanced LIGO observing run’. In: *Phys. Rev. D* 99.10 (2019), p. 104012. DOI: 10.1103/PhysRevD.99.104012. arXiv: 1811.04904 [gr-qc].
- [5] Julian Westerweck, Yotam Sherf, Collin D. Capano and Ram Brustein. *Sub-atomic constraints on the Kerr geometry of GW150914*. Aug. 2021. arXiv: 2108.08823 [gr-qc].
- [6] Robert M. Wald. *General Relativity*. Chicago, USA: Chicago Univ. Pr., 1984. DOI: 10.7208/chicago/9780226870373.001.0001.
- [7] Bernard F. Schutz. *A First Course in General Relativity*. 2nd ed. Cambridge, UK: Cambridge Univ. Pr., 2009.
- [8] Charles W. Misner, K. S. Thorne and J. A. Wheeler. *Gravitation*. San Francisco: W. H. Freeman, 1973. ISBN: 978-0-7167-0344-0, 978-0-691-17779-3.
- [9] Sean M. Carroll. *Spacetime and Geometry*. Cambridge University Press, July 2019. ISBN: 978-0-8053-8732-2, 978-1-108-48839-6, 978-1-108-77555-7.
- [10] Michele Maggiore. *Gravitational Waves. Vol. 1: Theory and Experiments*. Oxford Master Series in Physics. Oxford University Press, 2007. ISBN: 978-0-19-857074-5, 978-0-19-852074-0.
- [11] Abhay Ashtekar and Vesselin Petkov, eds. *Springer Handbook of Spacetime*. Berlin: Springer, 2014. ISBN: 978-3-642-41991-1, 978-3-642-41992-8. DOI: 10.1007/978-3-642-41992-8.
- [12] Hermann Minkowski. ‘Die Grundgleichungen für die elektromagnetischen Vorgänge in bewegten Körpern’. In: *Mathematische Annalen* 68 (1910), pp. 472–525. DOI: 10.1007/BF01455871.
- [13] Hermann Minkowski. ‘Raum und Zeit’. In: *Physikalische Zeitschrift* 10 (1909), pp. 104–115.
- [14] Albert Einstein, Hendrik Antoon Lorentz, Hermann Minkowski and Hermann Weyl. *The Principle of Relativity*. Dover Books on Physics. Dover Publications, 2013. ISBN: 0-486-60081-5.
- [15] Karl Schwarzschild. ‘Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie’. In: *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1916 (1916), pp. 189–196. arXiv: physics/9905030.

Bibliography

- [16] Roy P. Kerr. ‘Gravitational field of a spinning mass as an example of algebraically special metrics’. In: *Phys. Rev. Lett.* 11 (1963), pp. 237–238. DOI: 10.1103/PhysRevLett.11.237.
- [17] Alexander Friedmann. ‘Über die Krümmung des Raumes’. In: *Zeitschrift für Physik* 10 (1922), pp. 377–386. DOI: 10.1007/BF01332580.
- [18] Alexander Friedmann. ‘Über die Möglichkeit einer Welt mit konstanter negativer Krümmung des Raumes’. In: *Zeitschrift für Physik* 21 (1924), pp. 326–332. DOI: 10.1007/BF01328280.
- [19] Georges Lemaître. ‘A Homogeneous Universe of Constant Mass and Growing Radius Accounting for the Radial Velocity of Extragalactic Nebulae’. In: *Annales Soc. Sci. Bruxelles A* 47 (1927), pp. 49–59. DOI: 10.1007/s10714-013-1548-3.
- [20] H. P. Robertson. ‘Kinematics and World-Structure’. In: *Astrophys. J.* 82 (1935), pp. 284–301. DOI: 10.1086/143681.
- [21] H. P. Robertson. ‘Kinematics and World-Structure. 2’. In: *Astrophys. J.* 83 (1935), pp. 187–201. DOI: 10.1086/143716.
- [22] H. P. Robertson. ‘Kinematics and World-Structure. 3’. In: *Astrophys. J.* 83 (1936), pp. 257–271. DOI: 10.1086/143726.
- [23] A. G. Walker. ‘On Milne’s Theory of World-Structure’. In: *Proceedings of the London Mathematical Society* s2-42.1 (1937), pp. 90–127. DOI: 10.1112/plms/s2-42.1.90.
- [24] H. Reissner. ‘Über die Eigengravitation des elektrischen Feldes nach der Einsteinschen Theorie’. In: *Annalen der Physik* 50 (1916), pp. 106–120.
- [25] Robert H. Boyer and Richard W. Lindquist. ‘Maximal analytic extension of the Kerr metric’. In: *J. Math. Phys.* 8 (1967), p. 265. DOI: 10.1063/1.1705193.
- [26] E. T. Newman and A. I. Janis. ‘Note on the Kerr spinning particle metric’. In: *J. Math. Phys.* 6 (1965), pp. 915–917. DOI: 10.1063/1.1704350.
- [27] E T. Newman et al. ‘Metric of a Rotating, Charged Mass’. In: *J. Math. Phys.* 6 (1965), pp. 918–919. DOI: 10.1063/1.1704351.
- [28] K. S. Thorne. ‘Multipole Expansions of Gravitational Radiation’. In: *Rev. Mod. Phys.* 52 (1980), pp. 299–339. DOI: 10.1103/RevModPhys.52.299.
- [29] Albert Einstein. ‘Über Gravitationswellen’. In: *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1918 (1918), pp. 154–167.
- [30] G.D. Birkhoff and R.E. Langer. *Relativity and Modern Physics*. Cambridge, Massachusetts, USA: Harvard University Press, 1923.
- [31] J. T. Jebsen. ‘Über die allgemeinen kugelsymmetrischen Lösungen der Einsteinschen Gravitationsgleichungen im Vakuum’. In: *Ark. Mat. Ast. Fys. (Stockholm)* 15.18 (1921), pp. 1–9.
- [32] Nils Andersson et al. ‘The Transient Gravitational-Wave Sky’. In: *Class. Quant. Grav.* 30 (2013), p. 193002. DOI: 10.1088/0264-9381/30/19/193002. arXiv: 1305.0816 [gr-qc].
- [33] B. P. Abbott et al. ‘Observation of Gravitational Waves from a Binary Black Hole Merger’. In: *Phys. Rev. Lett.* 116.6 (2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102. arXiv: 1602.03837 [gr-qc].
- [34] B. P. Abbott et al. ‘Multi-messenger Observations of a Binary Neutron Star Merger’. In: *Astrophys. J. Lett.* 848.2 (2017), p. L12. DOI: 10.3847/2041-8213/aa91c9. arXiv: 1710.05833 [astro-ph.HE].

- [35] R. Abbott et al. ‘Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences’. In: *Astrophys. J. Lett.* 915.1 (2021), p. L5. DOI: 10.3847/2041-8213/ac082e. arXiv: 2106.15163 [astro-ph.HE].
- [36] B. P. Abbott et al. ‘A First Targeted Search for Gravitational-Wave Bursts from Core-Collapse Supernovae in Data of First-Generation Laser Interferometer Detectors’. In: *Phys. Rev. D* 94.10 (2016), p. 102001. DOI: 10.1103/PhysRevD.94.102001. arXiv: 1605.01785 [gr-qc].
- [37] B. Steltner et al. ‘Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data’. In: *Astrophys. J.* 909.1 (2021), p. 79. DOI: 10.3847/1538-4357/abc7c9. arXiv: 2009.12260 [astro-ph.HE].
- [38] R. Abbott et al. ‘All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data’. In: *Phys. Rev. D* 104.8 (2021), p. 082004. DOI: 10.1103/PhysRevD.104.082004. arXiv: 2107.00600 [gr-qc].
- [39] C. J. Moore, R. H. Cole and C. P. L. Berry. ‘Gravitational-wave sensitivity curves’. In: *Class. Quant. Grav.* 32.1 (2015), p. 015014. DOI: 10.1088/0264-9381/32/1/015014. arXiv: 1408.0740 [gr-qc].
- [40] Benjamin P. Abbott et al. ‘The basic physics of the binary black hole merger GW150914’. In: *Annalen Phys.* 529.1-2 (2017), p. 1600209. DOI: 10.1002/andp.201600209. arXiv: 1608.01940 [gr-qc].
- [41] Luc Blanchet. ‘Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries’. In: *Living Rev. Rel.* 17 (2014), p. 2. DOI: 10.12942/lrr-2014-2. arXiv: 1310.1528 [gr-qc].
- [42] Luc Blanchet. ‘Analyzing Gravitational Waves with General Relativity’. In: *Comptes Rendus Physique* 20 (2019), pp. 507–520. DOI: 10.1016/j.crhy.2019.02.004. arXiv: 1902.09801 [gr-qc].
- [43] Luc Blanchet, Bala R. Iyer, Clifford M. Will and Alan G. Wiseman. ‘Gravitational wave forms from inspiralling compact binaries to second postNewtonian order’. In: *Class. Quant. Grav.* 13 (1996), pp. 575–584. DOI: 10.1088/0264-9381/13/4/002. arXiv: gr-qc/9602024.
- [44] A. Buonanno and T. Damour. ‘Effective one-body approach to general relativistic two-body dynamics’. In: *Phys. Rev. D* 59 (1999), p. 084006. DOI: 10.1103/PhysRevD.59.084006. arXiv: gr-qc/9811091.
- [45] Clifford M. Will. ‘On the unreasonable effectiveness of the post-Newtonian approximation in gravitational physics’. In: *Proc. Nat. Acad. Sci.* 108 (2011), p. 5938. DOI: 10.1073/pnas.1103127108. arXiv: 1102.5192 [gr-qc].
- [46] Frans Pretorius. ‘Evolution of binary black hole spacetimes’. In: *Phys. Rev. Lett.* 95 (2005), p. 121101. DOI: 10.1103/PhysRevLett.95.121101. arXiv: gr-qc/0507014.
- [47] Luca Baiotti and Luciano Rezzolla. ‘Binary neutron star mergers: a review of Einstein’s richest laboratory’. In: *Rept. Prog. Phys.* 80.9 (2017), p. 096901. DOI: 10.1088/1361-6633/aa67bb. arXiv: 1607.03540 [gr-qc].
- [48] Miguel Alcubierre. *Introduction to 3+1 numerical relativity*. International series of monographs on physics. Oxford: Oxford Univ. Press, 2008. DOI: 10.1093/acprof:oso/9780199205677.001.0001.
- [49] Nigel T. Bishop and Luciano Rezzolla. ‘Extraction of Gravitational Waves in Numerical Relativity’. In: *Living Rev. Rel.* 19 (2016), p. 2. DOI: 10.1007/s41114-016-0001-9. arXiv: 1606.02532 [gr-qc].

Bibliography

- [50] Sebastian Khan, Katerina Chatziioannou, Mark Hannam and Frank Ohme. ‘Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects’. In: *Phys. Rev. D* 100.2 (2019), p. 024059. DOI: 10.1103/PhysRevD.100.024059. arXiv: 1809.10113 [gr-qc].
- [51] Anthony L. Piro, Bruno Giacomazzo and Rosalba Perna. ‘The Fate of Neutron Star Binary Mergers’. In: *Astrophys. J. Lett.* 844.2 (2017), p. L19. DOI: 10.3847/2041-8213/aa7f2f. arXiv: 1704.08697 [astro-ph.HE].
- [52] C. V. Vishveshwara. ‘Scattering of Gravitational Radiation by a Schwarzschild Black-hole’. In: *Nature* 227 (1970), pp. 936–938. DOI: 10.1038/227936a0.
- [53] S. A. Teukolsky. ‘Rotating black holes - separable wave equations for gravitational and electromagnetic perturbations’. In: *Phys. Rev. Lett.* 29 (1972), pp. 1114–1118. DOI: 10.1103/PhysRevLett.29.1114.
- [54] Emanuele Berti, Vitor Cardoso and Andrei O. Starinets. ‘Quasinormal modes of black holes and black branes’. In: *Class. Quant. Grav.* 26 (2009), p. 163001. DOI: 10.1088/0264-9381/26/16/163001. arXiv: 0905.2975 [gr-qc].
- [55] Samantha A. Usman et al. ‘The PyCBC search for gravitational waves from compact binary coalescence’. In: *Class. Quant. Grav.* 33.21 (2016), p. 215004. DOI: 10.1088/0264-9381/33/21/215004. arXiv: 1508.02357 [gr-qc].
- [56] Reinhard Prix. ‘Gravitational Waves from Spinning Neutron Stars’. In: *Astrophys. Space Sci. Lib.* 357 (2009), pp. 651–685. DOI: 10.1007/978-3-540-76965-1_24.
- [57] Paul D. Lasky. ‘Gravitational Waves from Neutron Stars: A Review’. In: *Publ. Astron. Soc. Austral.* 32 (2015), e034. DOI: 10.1017/pasa.2015.35. arXiv: 1508.06643 [astro-ph.HE].
- [58] Santiago Caride, Ra Inta, Benjamin J. Owen and Binod Rajbhandari. ‘How to search for gravitational waves from r -modes of known pulsars’. In: *Phys. Rev. D* 100.6 (2019), p. 064013. DOI: 10.1103/PhysRevD.100.064013. arXiv: 1907.04946 [gr-qc].
- [59] B. P. Abbott et al. ‘Narrow-band search for gravitational waves from known pulsars using the second LIGO observing run’. In: *Phys. Rev. D* 99.12 (2019), p. 122002. DOI: 10.1103/PhysRevD.99.122002. arXiv: 1902.08442 [gr-qc].
- [60] R. Abbott et al. *Search for continuous gravitational waves from 20 accreting millisecond X-ray pulsars in O3 LIGO data*. Sept. 2021. arXiv: 2109.09255 [astro-ph.HE].
- [61] Anjana Ashok et al. ‘New Searches for Continuous Gravitational Waves from Seven Fast Pulsars’. In: *Astrophys. J.* 923.1 (2021), p. 85. DOI: 10.3847/1538-4357/ac2582. arXiv: 2107.09727 [astro-ph.HE].
- [62] B. P. Abbott et al. ‘All-Sky Search for Short Gravitational-Wave Bursts in the Second Advanced LIGO and Advanced Virgo Run’. In: *Phys. Rev. D* 100.2 (2019), p. 024017. DOI: 10.1103/PhysRevD.100.024017. arXiv: 1905.03457 [gr-qc].
- [63] B. P. Abbott et al. *Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo*. 2018. arXiv: 1811.12940 [astro-ph.HE].
- [64] R. Abbott et al. ‘Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog’. In: *Astrophys. J. Lett.* 913.1 (2021), p. L7. DOI: 10.3847/2041-8213/abe949. arXiv: 2010.14533 [astro-ph.HE].

- [65] B. P. Abbott et al. ‘Tests of General Relativity with GW170817’. In: *Phys. Rev. Lett.* 123.1 (2019), p. 011102. DOI: 10.1103/PhysRevLett.123.011102. arXiv: 1811.00364 [gr-qc].
- [66] Keith Riles. ‘Recent searches for continuous gravitational waves’. In: *Mod. Phys. Lett. A* 32.39 (2017), p. 1730035. DOI: 10.1142/S021773231730035X. arXiv: 1712.05897 [gr-qc].
- [67] Maximiliano Isi, Matthew Pitkin and Alan J. Weinstein. ‘Probing Dynamical Gravity with the Polarization of Continuous Gravitational Waves’. In: *Phys. Rev. D* 96.4 (2017), p. 042001. DOI: 10.1103/PhysRevD.96.042001. arXiv: 1703.07530 [gr-qc].
- [68] Benjamin P. Abbott et al. ‘First search for nontensorial gravitational waves from known pulsars’. In: *Phys. Rev. Lett.* 120.3 (2018), p. 031104. DOI: 10.1103/PhysRevLett.120.031104. arXiv: 1709.09203 [gr-qc].
- [69] The LIGO Scientific Collaboration and the Virgo Collaboration. ‘Tests of general relativity with GW150914’. In: *Phys. Rev. Lett.* 116.22 (2016). [Erratum: *Phys. Rev. Lett.* 121, no.12, 129902(2018)], p. 221101. DOI: 10.1103/PhysRevLett.116.221101, 10.1103/PhysRevLett.121.129902. arXiv: 1602.03841 [gr-qc].
- [70] Miriam Cabero et al. ‘Observational tests of the black hole area increase law’. In: *Phys. Rev. D* 97.12 (2018), p. 124069. DOI: 10.1103/PhysRevD.97.124069. arXiv: 1711.09073 [gr-qc].
- [71] R. Abbott et al. ‘Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog’. In: *Phys. Rev. D* 103.12 (2021), p. 122002. DOI: 10.1103/PhysRevD.103.122002. arXiv: 2010.14529 [gr-qc].
- [72] R. Abbott et al. *Tests of General Relativity with GWTC-3*. Dec. 2021. arXiv: 2112.06861 [gr-qc].
- [73] J. Aasi et al. ‘Advanced LIGO’. In: *Class. Quant. Grav.* 32 (2015), p. 074001. DOI: 10.1088/0264-9381/32/7/074001. arXiv: 1411.4547 [gr-qc].
- [74] F. Acernese et al. ‘Advanced Virgo: a second-generation interferometric gravitational wave detector’. In: *Class. Quant. Grav.* 32.2 (2015), p. 024001. DOI: 10.1088/0264-9381/32/2/024001. arXiv: 1408.3978 [gr-qc].
- [75] T. Akutsu et al. ‘KAGRA: 2.5 Generation Interferometric Gravitational Wave Detector’. In: *Nature Astron.* 3.1 (2019), pp. 35–40. DOI: 10.1038/s41550-018-0658-y. arXiv: 1811.08079 [gr-qc].
- [76] Bernard F. Schutz and Massimo Tinto. ‘Antenna patterns of interferometric detectors of gravitational waves – I. Linearly polarized waves’. In: *Monthly Notices of the Royal Astronomical Society* 224.1 (Jan. 1987), pp. 131–154. ISSN: 0035-8711. DOI: 10.1093/mnras/224.1.131. eprint: <https://academic.oup.com/mnras/article-pdf/224/1/131/4099535/mnras224-0131.pdf>. URL: <https://doi.org/10.1093/mnras/224.1.131>.
- [77] Alex Abramovici et al. ‘LIGO: The Laser interferometer gravitational wave observatory’. In: *Science* 256 (1992), pp. 325–333. DOI: 10.1126/science.256.5055.325.
- [78] Kostas D. Kokkotas. ‘Gravitational Wave Physics’. In: *Encyclopedia of Physical Science and Technology (Third Edition)*. Ed. by Robert A. Meyers. Third Edition. New York: Academic Press, 2003, pp. 67–85. ISBN: 978-0-12-227410-7. DOI: <https://doi.org/10.1016/B0-12-227410-5/00300-8>.

Bibliography

- [79] Piotr Jaranowski and Andrzej Krolak. *Analysis of gravitational-wave data*. Cambridge: Cambridge Univ. Press, 2009. DOI: 10.1017/CB09780511605482.
- [80] Kip S. Thorne. ‘300 years of gravitation’. In: *Gravitational radiation*. Ed. by Stephen Hawking and Werner Israel. Cambridge University Press, 1987.
- [81] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman and Kip S. Thorne. ‘Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries’. In: *Phys. Rev. D* 49 (1994), pp. 6274–6297. DOI: 10.1103/PhysRevD.49.6274.
- [82] Michael Zeilik and Stephen A. Gregory. *Introductory Astronomy and Astrophysics*. Fort Worth: Saunders College Pub., 1998.
- [83] Great Britain and United States Naval Observatory. *Explanatory supplement to the Astronomical ephemeris and the American ephemeris and nautical almanac*. London: Her Majesty’s Stationery Office, 1961.
- [84] Lee Samuel Finn and David F. Chernoff. ‘Observing binary inspiral in gravitational radiation: One interferometer’. In: *Phys. Rev. D* 47 (1993), pp. 2198–2219. DOI: 10.1103/PhysRevD.47.2198. arXiv: gr-qc/9301003.
- [85] A. Krolak and Bernard F. Schutz. ‘Coalescing binaries — Probe of the universe’. In: *Gen. Rel. Grav.* 19 (1987), pp. 1163–1171. DOI: 10.1007/BF00759095.
- [86] Rich Abbott et al. ‘Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo’. In: *SoftwareX* 13 (2021), p. 100658. DOI: 10.1016/j.softx.2021.100658. arXiv: 1912.11716 [gr-qc].
- [87] P. B. Covas et al. ‘Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO’. In: *Phys. Rev. D* 97.8 (2018), p. 082002. DOI: 10.1103/PhysRevD.97.082002. arXiv: 1801.07204 [astro-ph.IM].
- [88] LIGO Scientific Collaboration and Virgo Collaboration. “*Gravitational Wave Open Science Center*”. 2022. URL: <https://www.gw-openscience.org/o1speclines/> (visited on 18/01/2022).
- [89] LIGO Scientific Collaboration and Virgo Collaboration. “*Gravitational Wave Open Science Center*”. 2022. URL: <https://www.gw-openscience.org/o2speclines/> (visited on 18/01/2022).
- [90] H. Grote. ‘The GEO 600 status’. In: *Class. Quant. Grav.* 27 (2010). Ed. by Zsuzsa Marka and Szabolcs Marka, p. 084003. DOI: 10.1088/0264-9381/27/8/084003.
- [91] James Lough et al. ‘First Demonstration of 6 dB Quantum Noise Reduction in a Kilometer Scale Gravitational Wave Observatory’. In: *Phys. Rev. Lett.* 126.4 (2021), p. 041102. DOI: 10.1103/PhysRevLett.126.041102. arXiv: 2005.10292 [physics.ins-det].
- [92] LIGO India Scientific Collaboration. “*LIGO India*”. 2022. URL: <https://www.ligo-india.in/about/site/> (visited on 18/01/2022).
- [93] David Reitze et al. ‘Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO’. In: *Bull. Am. Astron. Soc.* 51.7 (2019), p. 035. arXiv: 1907.04833 [astro-ph.IM].
- [94] Benjamin P Abbott et al. ‘Exploring the Sensitivity of Next Generation Gravitational Wave Detectors’. In: *Class. Quant. Grav.* 34.4 (2017), p. 044001. DOI: 10.1088/1361-6382/aa51f4. arXiv: 1607.08697 [astro-ph.IM].

- [95] M. Punturo et al. ‘The Einstein Telescope: A third-generation gravitational wave observatory’. In: *Class. Quant. Grav.* 27 (2010). Ed. by Fulvio Ricci, p. 194002. DOI: 10.1088/0264-9381/27/19/194002.
- [96] Pau Amaro-Seoane et al. *Laser Interferometer Space Antenna*. Feb. 2017. arXiv: 1702.00786 [astro-ph.IM].
- [97] Shane L. Larson, William A. Hiscock and Ronald W. Hellings. ‘Sensitivity curves for spaceborne gravitational wave interferometers’. In: *Phys. Rev. D* 62 (2000), p. 062001. DOI: 10.1103/PhysRevD.62.062001. arXiv: gr-qc/9909080.
- [98] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003. ISBN: 978-0-521-59271-0.
- [99] C. M. Biwer et al. ‘PyCBC Inference: A Python-based parameter estimation toolkit for compact binary coalescence signals’. In: *Publ. Astron. Soc. Pac.* 131.996 (2019), p. 024503. DOI: 10.1088/1538-3873/aaef0b. arXiv: 1807.10312 [astro-ph.IM].
- [100] Lee S. Finn. ‘Detection, measurement and gravitational radiation’. In: *Phys. Rev. D* 46 (1992), pp. 5236–5249. DOI: 10.1103/PhysRevD.46.5236. arXiv: gr-qc/9209010.
- [101] John A. Gubner. *Probability and Random Processes for Electrical and Computer Engineers*. Cambridge University Press, 2006. ISBN: 978-0-511-22023-4.
- [102] Ionuț Florescu. *Probability and Stochastic Processes*. New Jersey, USA: Wiley, 2006. ISBN: 978-0-470-62455-5.
- [103] Chris Chatfield and Haipeng Xing. *The analysis of time series: An introduction*. Boca Raton, FL, USA: Chapman and Hall, CRC Press, 2019. ISBN: 978-1-4987-9563-0.
- [104] E. J. Hannan. ‘Stationary time series’. In: *Time Series and Statistics*. Ed. by John Eatwell, Murray Milgate and Peter Newman. The Maxmillan Press Limited, 1990. ISBN: 978-1-349-20865-4.
- [105] L. A. Wainstein and V. D. Zubakov. *Extraction of Signals from Noise*. Dover books on physics and mathematical physics. Englewood Cliffs, NJ: Prentice-Hall, 1962. ISBN: 978-0-486-62625-3.
- [106] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown and Jolien D. E. Creighton. ‘FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries’. In: *Phys. Rev. D* 85 (2012), p. 122006. DOI: 10.1103/PhysRevD.85.122006. arXiv: gr-qc/0509116.
- [107] Benjamin P Abbott et al. ‘A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals’. In: *Class. Quant. Grav.* 37.5 (2020), p. 055002. DOI: 10.1088/1361-6382/ab685e. arXiv: 1908.11170 [gr-qc].
- [108] A. H. Nitz et al. *PyCBC Software*. 2021. URL: <https://github.com/gwastro/pycbc,%20GitHub>.
- [109] LIGO Scientific Collaboration. “*LOSC Event tutorial v1.4*”. 2016. URL: https://losc.ligo.org/s/events/GW150914/LOSC_Event_tutorial_GW150914.html.
- [110] B. P. Abbott et al. ‘Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914’. In: *Class. Quant. Grav.* 33.13 (2016), p. 134001. DOI: 10.1088/0264-9381/33/13/134001. arXiv: 1602.03844 [gr-qc].
- [111] Robert M. Gray. ‘Toeplitz and Circulant Matrices: A Review’. In: *Foundations and Trends® in Communications and Information Theory* 2.3 (2006), pp. 155–239. ISSN: 1567-2190. DOI: 10.1561/0100000006. URL: <http://dx.doi.org/10.1561/0100000006>.

Bibliography

- [112] Peter D. Welch. ‘The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms’. In: *IEEE Transactions on Audio and Electroacoustics* 15.2 (1967), pp. 70–73. DOI: 10.1109/TAU.1967.1161901.
- [113] W. R. Chatfield, S. Richardson and D. J. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Boca Raton, FL, USA: Chapman and Hall, CRC Press, 1996. ISBN: 0-412-05551-1.
- [114] Daniel Foreman-Mackey, David W. Hogg, Dustin Lang and Jonathan Goodman. ‘emcee: The MCMC Hammer’. In: *Publ. Astron. Soc. Pac.* 125 (2013), pp. 306–312. DOI: 10.1086/670067. arXiv: 1202.3665 [astro-ph.IM].
- [115] John Veitch et al. *cpnest*. <https://doi.org/10.5281/zenodo.835874>. Version 0.1.4. July 2017. DOI: 10.5281/zenodo.835874.
- [116] Joshua S. Speagle. ‘dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences’. In: *Mon. Not. Roy. Astron. Soc.* 493.3 (2020), pp. 3132–3158. DOI: 10.1093/mnras/staa278. arXiv: 1904.02180 [astro-ph.IM].
- [117] W. D. Vousden, W. M. Farr and I. Mandel. ‘Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations’. In: *Monthly Notices of the Royal Astronomical Society* 455.2 (Nov. 2015), pp. 1919–1937. ISSN: 0035-8711. DOI: 10.1093/mnras/stv2422. eprint: <https://academic.oup.com/mnras/article-pdf/455/2/1919/18514064/stv2422.pdf>. URL: <https://doi.org/10.1093/mnras/stv2422>.
- [118] John Skilling. ‘Nested sampling for general Bayesian computation’. In: *Bayesian Analysis* 1.4 (2006), pp. 833–859. DOI: 10.1214/06-BA127.
- [119] Robert E. Kass and Adrian E. Raftery. ‘Bayes Factors’. In: *J. Am. Statist. Assoc.* 90.430 (1995), pp. 773–795. DOI: 10.1080/01621459.1995.10476572.
- [120] Saul A. Teukolsky. ‘Perturbations of a rotating black hole. 1. Fundamental equations for gravitational electromagnetic and neutrino field perturbations’. In: *Astrophys. J.* 185 (1973), pp. 635–647. DOI: 10.1086/152444.
- [121] Steven L. Detweiler. ‘Resonant oscillations of a rapidly rotating black hole’. In: *Proc. Roy. Soc. Lond. A* 352 (1977), pp. 381–395. DOI: 10.1098/rspa.1977.0005.
- [122] Vitor Cardoso and Paolo Pani. ‘Testing the nature of dark compact objects: a status report’. In: *Living Rev. Rel.* 22.1 (2019), p. 4. DOI: 10.1007/s41114-019-0020-4. arXiv: 1904.05363 [gr-qc].
- [123] C. D. Capano. *pykerr*. 2021. URL: <https://github.com/cdcapano/pykerr>.
- [124] Hans-Peter Nollert. ‘TOPICAL REVIEW: Quasinormal modes: the characteristic ‘sound’ of black holes and neutron stars’. In: *Class. Quant. Grav.* 16 (1999), R159–R216. DOI: 10.1088/0264-9381/16/12/201.
- [125] Jolien D. E. Creighton and Warren G. Anderson. *Gravitational-wave physics and astronomy: An introduction to theory, experiment and data analysis*. Weinheim, Germany: Wiley-VCH, 2011.
- [126] Ioannis Kamaretsos, Mark Hannam, Sascha Husa and B. S. Sathyaprakash. ‘Black-hole hair loss: learning about binary progenitors from ringdown signals’. In: *Phys. Rev. D* 85 (2012), p. 024018. DOI: 10.1103/PhysRevD.85.024018. arXiv: 1107.0854 [gr-qc].

- [127] Anshu Gupta, Badri Krishnan, Alex Nielsen and Erik Schnetter. ‘Dynamics of marginally trapped surfaces in a binary black hole merger: Growth and approach to equilibrium’. In: *Phys. Rev. D* 97.8 (2018), p. 084028. DOI: 10.1103/PhysRevD.97.084028. arXiv: 1801.07048 [gr-qc].
- [128] Matthew Giesler, Maximiliano Isi, Mark Scheel and Saul Teukolsky. *Black hole ringdown: the importance of overtones*. 2019. arXiv: 1903.08284 [gr-qc].
- [129] Swetha Bhagwat, Xisco Jimenez Forteza, Paolo Pani and Valeria Ferrari. ‘Ringdown overtones, black hole spectroscopy, and no-hair theorem tests’. In: *Phys. Rev. D* 101.4 (2020), p. 044033. DOI: 10.1103/PhysRevD.101.044033. arXiv: 1910.08708 [gr-qc].
- [130] Maximiliano Isi and Will M. Farr. *Analyzing black-hole ringdowns*. July 2021. arXiv: 2107.05609 [gr-qc].
- [131] E. W. Leaver. ‘An Analytic representation for the quasi normal modes of Kerr black holes’. In: *Proc. Roy. Soc. Lond.* A402 (1985), pp. 285–298. DOI: 10.1098/rspa.1985.0119.
- [132] Ssohrab Borhanian, K. G. Arun, Harald P. Pfeiffer and B. S. Sathyaprakash. ‘Comparison of post-Newtonian mode amplitudes with numerical relativity simulations of binary black holes’. In: *Class. Quant. Grav.* 37.6 (2020), p. 065006. DOI: 10.1088/1361-6382/ab6a21. arXiv: 1901.08516 [gr-qc].
- [133] Olaf Dreyer et al. ‘Black hole spectroscopy: Testing general relativity through gravitational wave observations’. In: *Class. Quant. Grav.* 21 (2004), pp. 787–804. DOI: 10.1088/0264-9381/21/4/003. arXiv: gr-qc/0309007 [gr-qc].
- [134] S. Gossan, J. Veitch and B. S. Sathyaprakash. ‘Bayesian model selection for testing the no-hair theorem with black hole ringdowns’. In: *Phys. Rev.* D85 (2012), p. 124056. DOI: 10.1103/PhysRevD.85.124056. arXiv: 1111.5819 [gr-qc].
- [135] Gregory B. Cook and Maxim Zalutskiy. ‘Gravitational perturbations of the Kerr geometry: High-accuracy study’. In: *Phys. Rev. D* 90.12 (2014), p. 124021. DOI: 10.1103/PhysRevD.90.124021. arXiv: 1410.7698 [gr-qc].
- [136] William H. Press and Saul A. Teukolsky. ‘Perturbations of a Rotating Black Hole. II. Dynamical Stability of the Kerr Metric’. In: *Astrophys. J.* 185 (1973), pp. 649–674. DOI: 10.1086/152445.
- [137] Juan Calderón Bustillo et al. *Comparison of subdominant gravitational wave harmonics between post-Newtonian and numerical relativity calculations and construction of multi-mode hybrids*. Jan. 2015. arXiv: 1501.00918 [gr-qc].
- [138] Emanuele Berti, Vitor Cardoso and Marc Casals. ‘Eigenvalues and eigenfunctions of spin-weighted spheroidal harmonics in four and higher dimensions’. In: *Phys. Rev. D* 73 (2006). [Erratum: Phys.Rev.D 73, 109902 (2006)], p. 024013. DOI: 10.1103/PhysRevD.73.109902. arXiv: gr-qc/0511111.
- [139] Emanuele Berti, Jaime Cardoso, Vitor Cardoso and Marco Cavaglia. ‘Matched-filtering and parameter estimation of ringdown waveforms’. In: *Phys. Rev. D* 76 (2007), p. 104044. DOI: 10.1103/PhysRevD.76.104044. arXiv: 0707.1202 [gr-qc].
- [140] Eanna E. Flanagan and Scott A. Hughes. ‘Measuring gravitational waves from binary black hole coalescences: 1. Signal-to-noise for inspiral, merger, and ringdown’. In: *Phys. Rev.* D57 (1998), pp. 4535–4565. DOI: 10.1103/PhysRevD.57.4535. arXiv: gr-qc/9701039 [gr-qc].
- [141] Emanuele Berti, Vitor Cardoso and Clifford M. Will. ‘Gravitational-wave spectroscopy of massive black holes with the space interferometer LISA’. In: *Phys. Rev.* D73 (2006), p. 064030. DOI: 10.1103/PhysRevD.73.064030. arXiv: gr-qc/0512160 [gr-qc].

Bibliography

- [142] Werner Israel. ‘Event horizons in static vacuum space-times’. In: *Phys. Rev.* 164 (1967), pp. 1776–1779. DOI: 10.1103/PhysRev.164.1776.
- [143] Werner Israel. ‘Event horizons in static electrovac space-times’. In: *Commun. Math. Phys.* 8 (1968), pp. 245–260. DOI: 10.1007/BF01645859.
- [144] B. Carter. ‘Axisymmetric Black Hole Has Only Two Degrees of Freedom’. In: *Phys. Rev. Lett.* 26 (1971), pp. 331–333. DOI: 10.1103/PhysRevLett.26.331.
- [145] B. Carter. ‘Black holes equilibrium states’. In: *Les Houches Summer School of Theoretical Physics: Black Holes.* 1973, pp. 57–214.
- [146] S. W. Hawking. ‘Black holes in general relativity’. In: *Commun. Math. Phys.* 25 (1972), pp. 152–166. DOI: 10.1007/BF01877517.
- [147] D. C. Robinson. ‘Uniqueness of the Kerr black hole’. In: *Phys. Rev. Lett.* 34 (1975), pp. 905–906. DOI: 10.1103/PhysRevLett.34.905.
- [148] P. O. Mazur. ‘PROOF OF UNIQUENESS OF THE KERR-NEWMAN BLACK HOLE SOLUTION’. In: *J. Phys. A* 15 (1982), pp. 3173–3180. DOI: 10.1088/0305-4470/15/10/021.
- [149] Steven L. Detweiler. ‘Black Holes and Gravitational Waves. III. The Resonant Frequencies of Rotating Holes’. In: *Astrophys. J.* 239 (1980), pp. 292–295. DOI: 10.1086/158109.
- [150] F. Echeverria. ‘Gravitational Wave Measurements of the Mass and Angular Momentum of a Black Hole’. In: *Phys. Rev. D* 40 (1989), pp. 3194–3203. DOI: 10.1103/PhysRevD.40.3194.
- [151] Alexandru D. Ionescu and Sergiu Klainerman. ‘Rigidity Results in General Relativity: a Review’. In: *Surveys in Differential Geometry.* Ed. by Lydia Bieri and Shing-Tung Yau. Vol. 20. International Press, July 2015. DOI: 10.4310/SDG.2015.v20.n1.a6. arXiv: 1501.01587 [gr-qc].
- [152] Piotr T. Chrusciel, Joao Lopes Costa and Markus Heusler. ‘Stationary Black Holes: Uniqueness and Beyond’. In: *Living Rev. Rel.* 15 (2012), p. 7. DOI: 10.12942/lrr-2012-7. arXiv: 1205.6112 [gr-qc].
- [153] Vitor Cardoso, Seth Hopper, Caio F. B. Macedo, Carlos Palenzuela and Paolo Pani. ‘Gravitational-wave signatures of exotic compact objects and of quantum corrections at the horizon scale’. In: *Phys. Rev. D* 94.8 (2016), p. 084031. DOI: 10.1103/PhysRevD.94.084031. arXiv: 1608.08637 [gr-qc].
- [154] Steven B. Giddings. ‘Gravitational wave tests of quantum modifications to black hole structure – with post-GW150914 update’. In: *Class. Quant. Grav.* 33.23 (2016), p. 235010. DOI: 10.1088/0264-9381/33/23/235010. arXiv: 1602.03622 [gr-qc].
- [155] Ahmed Almheiri, Donald Marolf, Joseph Polchinski and James Sully. ‘Black Holes: Complementarity or Firewalls?’ In: *JHEP* 02 (2013), p. 062. DOI: 10.1007/JHEP02(2013)062. arXiv: 1207.3123 [hep-th].
- [156] A. Vilenkin. ‘Exponential Amplification of Waves in the Gravitational Field of Ultrarelativistic Rotating Body’. In: *Phys. Lett. B* 78 (1978), pp. 301–303. DOI: 10.1016/0370-2693(78)90027-8.
- [157] Ram Brustein and A. J. M. Medved. ‘Quantum hair of black holes out of equilibrium’. In: *Phys. Rev. D* 97.4 (2018), p. 044035. DOI: 10.1103/PhysRevD.97.044035. arXiv: 1709.03566 [hep-th].

- [158] Jahed Abedi, Hannah Dykaar and Niayesh Afshordi. ‘Echoes from the Abyss: Tentative evidence for Planck-scale structure at black hole horizons’. In: *Phys. Rev. D* 96.8 (2017), p. 082004. DOI: 10.1103/PhysRevD.96.082004. arXiv: 1612.00266 [gr-qc].
- [159] Albert Einstein. ‘Die Feldgleichungen der Gravitation’. In: *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1915 (1915), pp. 844–847.
- [160] Albert Einstein. ‘Näherungsweise Integration der Feldgleichungen der Gravitation’. In: *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1916 (1916), pp. 688–696.
- [161] A. Hewish, S. J. Bell, J. D. H Pilkington, P. F. Scott and R. A. Collins. ‘Observation of a rapidly pulsating radio source’. In: *Nature* 217 (1968), pp. 709–713. DOI: 10.1038/217709a0.
- [162] S. Bowyer, E. T. Byram, T. A. Chubb and H. Friedman. ‘Cosmic X-ray Sources’. In: *Science* 147.3656 (1965), pp. 394–398. DOI: 10.1126/science.147.3656.394.
- [163] H. L. Shipman. ‘The Implausible History of Triple Star Models for Cygnus X-1: Evidence for a Black Hole’. In: *Astrophys. Lett.* 16 (Feb. 1975), p. 9.
- [164] R. A. Hulse and J. H. Taylor. ‘Discovery of a pulsar in a binary system’. In: *Astrophys. J. Lett.* 195 (1975), pp. L51–L53. DOI: 10.1086/181708.
- [165] J. H. Taylor, L. A. Fowler and P. M. McCulloch. ‘Measurements of general relativistic effects in the binary pulsar PSR 1913+16’. In: *Nature* 277 (1979), pp. 437–440. DOI: 10.1038/277437a0.
- [166] 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’. In: *Phys. Rev. X* 9.3 (2019), p. 031040. DOI: 10.1103/PhysRevX.9.031040. arXiv: 1811.12907 [astro-ph.HE].
- [167] R. Abbott et al. ‘GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run’. In: *Phys. Rev. X* 11 (2021), p. 021053. DOI: 10.1103/PhysRevX.11.021053. arXiv: 2010.14527 [gr-qc].
- [168] 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 [gr-qc].
- [169] Alexander H. Nitz et al. *3-OGC: Catalog of gravitational waves from compact-binary mergers*. May 2021. arXiv: 2105.09151 [astro-ph.HE].
- [170] Alexander H. Nitz et al. *4-OGC: Catalog of gravitational waves from compact-binary mergers*. Dec. 2021. arXiv: 2112.06878 [astro-ph.HE].
- [171] S. W. Hawking. ‘Gravitational radiation from colliding black holes’. In: *Phys. Rev. Lett.* 26 (1971), pp. 1344–1346. DOI: 10.1103/PhysRevLett.26.1344.
- [172] Maximiliano Isi, Matthew Giesler, Will M. Farr, Mark A. Scheel and Saul A. Teukolsky. ‘Testing the no-hair theorem with GW150914’. In: *Phys. Rev. Lett.* 123.11 (2019), p. 111102. DOI: 10.1103/PhysRevLett.123.111102, 10.1103/physrevlett.123.111102. arXiv: 1905.00869 [gr-qc].
- [173] Maximiliano Isi, Will M. Farr, Matthew Giesler, Mark A. Scheel and Saul A. Teukolsky. ‘Testing the Black-Hole Area Law with GW150914’. In: *Phys. Rev. Lett.* 127.1 (2021), p. 011103. DOI: 10.1103/PhysRevLett.127.011103. arXiv: 2012.04486 [gr-qc].
- [174] Shilpa Kastha et al. ‘Model systematics in time domain tests of binary black hole evolution’. In: *Phys. Rev. D* 105.6 (2022), p. 064042. DOI: 10.1103/PhysRevD.105.064042. arXiv: 2111.13664 [gr-qc].

Bibliography

- [175] A. A. Starobinsky. ‘Amplification of waves reflected from a rotating “black hole”’. In: *Sov. Phys. JETP* 37.1 (1973), pp. 28–32.
- [176] Vitor Cardoso, Edgardo Franzin and Paolo Pani. ‘Is the gravitational-wave ringdown a probe of the event horizon?’ In: *Phys. Rev. Lett.* 116.17 (2016). [Erratum: *Phys.Rev.Lett.* 117, 089902 (2016)], p. 171101. DOI: 10.1103/PhysRevLett.116.171101. arXiv: 1602.07309 [gr-qc].
- [177] Nami Uchikata et al. ‘Searching for black hole echoes from the LIGO-Virgo Catalog GWTC-1’. In: *Phys. Rev. D* 100.6 (2019), p. 062006. DOI: 10.1103/PhysRevD.100.062006. arXiv: 1906.00838 [gr-qc].
- [178] Ka Wa Tsang et al. ‘A morphology-independent search for gravitational wave echoes in data from the first and second observing runs of Advanced LIGO and Advanced Virgo’. In: *Phys. Rev. D* 101.6 (2020), p. 064012. DOI: 10.1103/PhysRevD.101.064012. arXiv: 1906.11168 [gr-qc].
- [179] Elisa Maggio, Adriano Testa, Swetha Bhagwat and Paolo Pani. ‘Analytical model for gravitational-wave echoes from spinning remnants’. In: *Phys. Rev. D* 100.6 (2019), p. 064056. DOI: 10.1103/PhysRevD.100.064056. arXiv: 1907.03091 [gr-qc].
- [180] S. Chandrasekhar and Steven L. Detweiler. ‘The quasi-normal modes of the Schwarzschild black hole’. In: *Proc. Roy. Soc. Lond.* A344 (1975), pp. 441–452. DOI: 10.1098/rspa.1975.0112.
- [181] Kostas D. Kokkotas and Bernd G. Schmidt. ‘Quasinormal modes of stars and black holes’. In: *Living Rev. Rel.* 2 (1999), p. 2. DOI: 10.12942/lrr-1999-2. arXiv: gr-qc/9909058 [gr-qc].
- [182] Scott A. Hughes, Anuj Apte, Gaurav Khanna and Halston Lim. ‘Learning about black hole binaries from their ringdown spectra’. In: *Phys. Rev. Lett.* 123.16 (2019), p. 161101. DOI: 10.1103/PhysRevLett.123.161101. arXiv: 1901.05900 [gr-qc].
- [183] Anuj Apte and Scott A. Hughes. ‘Exciting black hole modes via misaligned coalescences: I. Inspiral, transition, and plunge trajectories using a generalized Ori-Thorne procedure’. In: *Phys. Rev. D* 100.8 (2019), p. 084031. DOI: 10.1103/PhysRevD.100.084031. arXiv: 1901.05901 [gr-qc].
- [184] Halston Lim, Gaurav Khanna, Anuj Apte and Scott A. Hughes. ‘Exciting black hole modes via misaligned coalescences: II. The mode content of late-time coalescence waveforms’. In: *Phys. Rev. D* 100.8 (2019), p. 084032. DOI: 10.1103/PhysRevD.100.084032. arXiv: 1901.05902 [gr-qc].
- [185] Emanuele Berti, Alberto Sesana, Enrico Barausse, Vitor Cardoso and Krzysztof Belczynski. ‘Spectroscopy of Kerr black holes with Earth- and space-based interferometers’. In: *Phys. Rev. Lett.* 117.10 (2016), p. 101102. DOI: 10.1103/PhysRevLett.117.101102. arXiv: 1605.09286 [gr-qc].
- [186] Vishal Baibhav and Emanuele Berti. ‘Multimode black hole spectroscopy’. In: *Phys. Rev. D* 99.2 (2019), p. 024005. DOI: 10.1103/PhysRevD.99.024005. arXiv: 1809.03500 [gr-qc].
- [187] LIGO Scientific Collaboration and Virgo Collaboration. *GraceDB: S190814bv*. <https://gracedb.ligo.org/superevents/S190814bv/>.
- [188] LIGO Scientific Collaboration and Virgo Collaboration. *GraceDB: S190910d*. <https://gracedb.ligo.org/superevents/S190910d/>.

- [189] Alex Nitz et al. *gwastro/pycbc*. <https://doi.org/10.5281/zenodo.3265452>. July 2019. DOI: 10.5281/zenodo.3265452.
- [190] Alexander H. Nitz et al. *2-OGC: Open Gravitational-wave Catalog of binary mergers from analysis of public Advanced LIGO and Virgo data*. 2019. arXiv: 1910.05331 [astro-ph.HE].
- [191] Alexander H. Nitz et al. ‘1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO data’. In: *Astrophys. J.* 872.2 (2019), p. 195. DOI: 10.3847/1538-4357/ab0108. arXiv: 1811.01921 [gr-qc].
- [192] Javier M. Antelis and Claudia Moreno. ‘An independent search of gravitational waves in the first observation run of advanced LIGO using cross-correlation’. In: *Gen. Rel. Grav.* 51.5 (2019), p. 61. DOI: 10.1007/s10714-019-2546-x. arXiv: 1807.07660 [gr-qc].
- [193] Barak Zackay, Liang Dai, Tejaswi Venumadhav, Javier Roulet and Matias Zaldarriaga. *Detecting Gravitational Waves With Disparate Detector Responses: Two New Binary Black Hole Mergers*. 2019. arXiv: 1910.09528 [astro-ph.HE].
- [194] Tejaswi Venumadhav, Barak Zackay, Javier Roulet, Liang Dai and Matias Zaldarriaga. *New Binary Black Hole Mergers in the Second Observing Run of Advanced LIGO and Advanced Virgo*. 2019. arXiv: 1904.07214 [astro-ph.HE].
- [195] Gregorio Carullo, Walter Del Pozzo and John Veitch. *Observational Black Hole Spectroscopy: A time-domain multimode analysis of GW150914*. 2019. arXiv: 1902.07527 [gr-qc].
- [196] Jeffrey Annis, Nathan J Evans, Brent J Miller and Thomas J Palmeri. ‘Thermodynamic integration and steppingstone sampling methods for estimating Bayes factors: A tutorial’. In: *Journal of mathematical psychology* 89 (2019), pp. 67–86.
- [197] Miriam Anabel Cabero Müller. ‘Gravitational-wave astronomy with compact binary coalescences: from blip glitches to the black hole area increase law’. PhD thesis. Leibniz Universität Hannover, 2018. DOI: 10.15488/3422. URL: <https://www.repo.uni-hannover.de/handle/123456789/3452>.
- [198] Wolfgang Tichy and Pedro Marronetti. ‘The Final mass and spin of black hole mergers’. In: *Phys. Rev. D* 78 (2008), p. 081501. DOI: 10.1103/PhysRevD.78.081501. arXiv: 0807.2985 [gr-qc].
- [199] Fabian Hofmann, Enrico Barausse and Luciano Rezzolla. ‘The final spin from binary black holes in quasi-circular orbits’. In: *Astrophys. J.* 825.2 (2016), p. L19. DOI: 10.3847/2041-8205/825/2/L19. arXiv: 1605.01938 [gr-qc].
- [200] LIGO Scientific Collaboration. *LIGO Algorithm Library - LALSuite*. free software (GPL). 2018. DOI: 10.7935/GT1W-FZ16.
- [201] P. A. R. Ade et al. ‘Planck 2015 results. XIII. Cosmological parameters’. In: *Astron. Astrophys.* 594 (2016), A13. DOI: 10.1051/0004-6361/201525830. arXiv: 1502.01589 [astro-ph.CO].
- [202] Swetha Bhagwat, Miriam Cabero, Collin D. Capano, Badri Krishnan and Duncan A. Brown. ‘Detectability of the subdominant mode in a binary black hole ringdown’. In: (2019). arXiv: 1910.13203 [gr-qc].
- [203] The LIGO Scientific Collaboration. *Unofficial sensitivity curves (ASD) for aLIGO, Kagra, Virgo, Voyager, Cosmic Explorer and ET*. <https://dcc.ligo.org/LIGO-T1500293/public>. 2019.

Bibliography

- [204] B. Sathyaprakash et al. ‘Scientific Objectives of Einstein Telescope’. In: *Class. Quant. Grav.* 29 (2012). [Erratum: *Class. Quant. Grav.*30,079501(2013)], p. 124013. DOI: 10.1088/0264-9381/29/12/124013, 10.1088/0264-9381/30/7/079501. arXiv: 1206.0331 [gr-qc].
- [205] B. Sathyaprakash et al. ‘Scientific Potential of Einstein Telescope’. In: *Proceedings, 46th Rencontres de Moriond on Gravitational Waves and Experimental Gravity: La Thuile, Italy, March 20-27, 2011*. 2011, pp. 127–136. arXiv: 1108.1423 [gr-qc].
- [206] K. Danzmann and A. Rudiger. ‘LISA technology - Concept, status, prospects’. In: *Class. Quant. Grav.* 20 (2003), S1–S9. DOI: 10.1088/0264-9381/20/10/301.
- [207] Swetha Bhagwat et al. ‘On choosing the start time of binary black hole ringdowns’. In: *Phys. Rev. D* 97.10 (2018), p. 104065. DOI: 10.1103/PhysRevD.97.104065. arXiv: 1711.00926 [gr-qc].
- [208] Emanuele Berti, Kent Yagi, Huan Yang and Nicolás Yunes. ‘Extreme Gravity Tests with Gravitational Waves from Compact Binary Coalescences: (II) Ringdown’. In: *Gen. Rel. Grav.* 50.5 (2018), p. 49. DOI: 10.1007/s10714-018-2372-6. arXiv: 1801.03587 [gr-qc].
- [209] Andrea Maselli, Paolo Pani, Leonardo Gualtieri and Emanuele Berti. *Parametrized ringdown spin expansion coefficients: a data-analysis framework for black-hole spectroscopy with multiple events*. 2019. arXiv: 1910.12893 [gr-qc].
- [210] Flora Moulin, Aurélien Barrau and Killian Martineau. ‘An overview of quasinormal modes in modified and extended gravity’. In: *Universe* 5.9 (2019), p. 202. DOI: 10.3390/universe5090202. arXiv: 1908.06311 [gr-qc].
- [211] R. Abbott et al. ‘GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses’. In: *Phys. Rev. D* 102.4 (2020), p. 043015. DOI: 10.1103/PhysRevD.102.043015. arXiv: 2004.08342 [astro-ph.HE].
- [212] 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].
- [213] The LIGO Scientific and Virgo Collaborations. *Noise curves used for Simulations in the update of the Observing Scenarios Paper*. <https://dcc.ligo.org/LIGO-T2000012-v1/public>. 2021.
- [214] Pauli Virtanen et al. ‘SciPy 1.0—Fundamental Algorithms for Scientific Computing in Python’. In: *Nature Meth.* 17 (2020), p. 261. DOI: 10.1038/s41592-019-0686-2. arXiv: 1907.10121 [cs.MS].
- [215] C. V. Vishveshwara. ‘Stability of the schwarzschild metric’. In: *Phys. Rev. D* 1 (1970), pp. 2870–2879. DOI: 10.1103/PhysRevD.1.2870.
- [216] R. Abbott et al. ‘GW190521: A Binary Black Hole Merger with a Total Mass of $150M_{\odot}$ ’. In: *Phys. Rev. Lett.* 125.10 (2020), p. 101102. DOI: 10.1103/PhysRevLett.125.101102. arXiv: 2009.01075 [gr-qc].
- [217] R. Abbott et al. ‘Properties and Astrophysical Implications of the $150 M_{\odot}$ Binary Black Hole Merger GW190521’. In: *Astrophys. J. Lett.* 900.1 (2020), p. L13. DOI: 10.3847/2041-8213/aba493. arXiv: 2009.01190 [astro-ph.HE].
- [218] Alexander H. Nitz and Collin D. Capano. ‘GW190521 may be an intermediate mass ratio inspiral’. In: *Astrophys. J. Lett.* 907.1 (2021), p. L9. DOI: 10.3847/2041-8213/abcccc5. arXiv: 2010.12558 [astro-ph.HE].

- [219] Héctor Estellés et al. *A detailed analysis of GW190521 with phenomenological waveform models*. May 2021. arXiv: 2105.06360 [gr-qc].
- [220] Vijay Varma et al. ‘Surrogate model of hybridized numerical relativity binary black hole waveforms’. In: *Phys. Rev. D* 99.6 (2019), p. 064045. DOI: 10.1103/PhysRevD.99.064045. arXiv: 1812.07865 [gr-qc].
- [221] Geraint Pratten et al. ‘Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes’. In: *Phys. Rev. D* 103.10 (2021), p. 104056. DOI: 10.1103/PhysRevD.103.104056. arXiv: 2004.06503 [gr-qc].
- [222] Juan Calderón Bustillo, Nicolas Sanchis-Gual, Alejandro Torres-Forné and José A. Font. ‘Confusing Head-On Collisions with Precessing Intermediate-Mass Binary Black Hole Mergers’. In: *Phys. Rev. Lett.* 126.20 (2021), p. 201101. DOI: 10.1103/PhysRevLett.126.201101. arXiv: 2009.01066 [gr-qc].
- [223] Isobel M. Romero-Shaw, Paul D. Lasky, Eric Thrane and Juan Calderon Bustillo. ‘GW190521: orbital eccentricity and signatures of dynamical formation in a binary black hole merger signal’. In: *Astrophys. J. Lett.* 903.1 (2020), p. L5. DOI: 10.3847/2041-8213/abbe26. arXiv: 2009.04771 [astro-ph.HE].
- [224] V. Gayathri et al. *GW190521 as a Highly Eccentric Black Hole Merger*. Sept. 2020. arXiv: 2009.05461 [astro-ph.HE].
- [225] Hans-Peter Nollert. ‘About the significance of quasinormal modes of black holes’. In: *Phys. Rev. D* 53 (1996), pp. 4397–4402. DOI: 10.1103/PhysRevD.53.4397. arXiv: gr-qc/9602032.
- [226] Hans-Peter Nollert and Richard H. Price. ‘Quantifying excitations of quasinormal mode systems’. In: *J. Math. Phys.* 40 (1999), pp. 980–1010. DOI: 10.1063/1.532698. arXiv: gr-qc/9810074.
- [227] Ramin G. Daghighi, Michael D. Green and Jodin C. Morey. ‘Significance of Black Hole Quasinormal Modes: A Closer Look’. In: *Phys. Rev. D* 101.10 (2020), p. 104009. DOI: 10.1103/PhysRevD.101.104009. arXiv: 2002.07251 [gr-qc].
- [228] Wei-Liang Qian, Kai Lin, Cai-Ying Shao, Bin Wang and Rui-Hong Yue. ‘Asymptotical quasinormal mode spectrum for piecewise approximate effective potential’. In: *Phys. Rev. D* 103.2 (2021), p. 024019. DOI: 10.1103/PhysRevD.103.024019. arXiv: 2009.11627 [gr-qc].
- [229] José Luis Jaramillo, Rodrigo Panosso Macedo and Lamis Al Sheikh. ‘Pseudospectrum and Black Hole Quasinormal Mode Instability’. In: *Phys. Rev. X* 11.3 (2021), p. 031003. DOI: 10.1103/PhysRevX.11.031003. arXiv: 2004.06434 [gr-qc].
- [230] Jeremy Sakstein, Djuna Croon, Samuel D. McDermott, Maria C. Straight and Eric J. Baxter. ‘Beyond the Standard Model Explanations of GW190521’. In: *Phys. Rev. Lett.* 125.26 (2020), p. 261105. DOI: 10.1103/PhysRevLett.125.261105. arXiv: 2009.01213 [gr-qc].
- [231] Mohammadtaher Safarzadeh and Zoltán Haiman. ‘Formation of GW190521 via gas accretion onto Population III stellar black hole remnants born in high-redshift minihalos’. In: *Astrophys. J. Lett.* 903.1 (2020), p. L21. DOI: 10.3847/2041-8213/abc253. arXiv: 2009.09320 [astro-ph.HE].
- [232] Juan Calderón Bustillo et al. ‘GW190521 as a Merger of Proca Stars: A Potential New Vector Boson of 8.7×10^{-13} eV’. In: *Phys. Rev. Lett.* 126.8 (2021), p. 081101. DOI: 10.1103/PhysRevLett.126.081101. arXiv: 2009.05376 [gr-qc].

Bibliography

- [233] C. Palenzuela, I. Olabarrieta, L. Lehner and Steven L. Liebling. ‘Head-on collisions of boson stars’. In: *Phys. Rev. D* 75 (2007), p. 064005. DOI: 10.1103/PhysRevD.75.064005. arXiv: gr-qc/0612067.
- [234] Masaru Shibata, Kenta Kiuchi, Sho Fujibayashi and Yuichiro Sekiguchi. ‘Alternative possibility of GW190521: Gravitational waves from high-mass black hole-disk systems’. In: *Phys. Rev. D* 103.6 (2021), p. 063037. DOI: 10.1103/PhysRevD.103.063037. arXiv: 2101.05440 [astro-ph.HE].
- [235] Barak Zackay, Tejaswi Venumadhav, Javier Roulet, Liang Dai and Matias Zaldarriaga. *Detecting Gravitational Waves in Data with Non-Gaussian Noise*. Aug. 2019. arXiv: 1908.05644 [astro-ph.IM].
- [236] Charles R. Harris et al. ‘Array programming with NumPy’. In: *Nature* 585.7825 (2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2. arXiv: 2006.10256 [cs.MS].
- [237] B. P. Abbott et al. ‘GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence’. In: *Phys. Rev. Lett.* 116.24 (2016), p. 241103. DOI: 10.1103/PhysRevLett.116.241103. arXiv: 1606.04855 [gr-qc].
- [238] B. P. Abbott et al. ‘Binary Black Hole Mergers in the first Advanced LIGO Observing Run’. In: *Phys. Rev. X* 6.4 (2016). [Erratum: Phys.Rev.X 8, 039903 (2018)], p. 041015. DOI: 10.1103/PhysRevX.6.041015. arXiv: 1606.04856 [gr-qc].
- [239] B. P. Abbott et al. ‘GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2’. In: *Phys. Rev. Lett.* 118.22 (2017). [Erratum: Phys.Rev.Lett. 121, 129901 (2018)], p. 221101. DOI: 10.1103/PhysRevLett.118.221101. arXiv: 1706.01812 [gr-qc].
- [240] B. P. Abbott et al. ‘GW170608: Observation of a 19-solar-mass Binary Black Hole Coalescence’. In: *Astrophys. J. Lett.* 851 (2017), p. L35. DOI: 10.3847/2041-8213/aa9f0c. arXiv: 1711.05578 [astro-ph.HE].
- [241] B. P. Abbott et al. ‘GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence’. In: *Phys. Rev. Lett.* 119.14 (2017), p. 141101. DOI: 10.1103/PhysRevLett.119.141101. arXiv: 1709.09660 [gr-qc].
- [242] B. P. Abbott et al. ‘Properties of the Binary Black Hole Merger GW150914’. In: *Phys. Rev. Lett.* 116.24 (2016), p. 241102. DOI: 10.1103/PhysRevLett.116.241102. arXiv: 1602.03840 [gr-qc].
- [243] Kip S. Thorne, R. H. Price and D. A. Macdonald, eds. *Black Holes: The Membrane Paradigm*. 1986. ISBN: 978-0-300-03770-8.
- [244] Gregory M. Harry. ‘Advanced LIGO: The next generation of gravitational wave detectors’. In: *Class. Quant. Grav.* 27 (2010). Ed. by Zsuzsa Marka and Szabolcs Marka, p. 084006. DOI: 10.1088/0264-9381/27/8/084006.
- [245] Vitor Cardoso and Paolo Pani. *The observational evidence for horizons: from echoes to precision gravitational-wave physics*. July 2017. arXiv: 1707.03021 [gr-qc].
- [246] Zachary Mark, Aaron Zimmerman, Song Ming Du and Yanbei Chen. ‘A recipe for echoes from exotic compact objects’. In: *Phys. Rev. D* 96.8 (2017), p. 084002. DOI: 10.1103/PhysRevD.96.084002. arXiv: 1706.06155 [gr-qc].
- [247] Sebastian H. Völkel and Kostas D. Kokkotas. ‘Ultra Compact Stars: Reconstructing the Perturbation Potential’. In: *Class. Quant. Grav.* 34.17 (2017), p. 175015. DOI: 10.1088/1361-6382/aa82de. arXiv: 1704.07517 [gr-qc].

- [248] Vitor Cardoso, Edgardo Franzin, Andrea Maselli, Paolo Pani and Guilherme Raposo. ‘Testing strong-field gravity with tidal Love numbers’. In: *Phys. Rev. D* 95.8 (2017). [Addendum: *Phys.Rev.D* 95, 089901 (2017)], p. 084014. DOI: 10.1103/PhysRevD.95.084014. arXiv: 1701.01116 [gr-qc].
- [249] Jahed Abedi, Hannah Dykaar and Niayesh Afshordi. *Echoes from the Abyss: Tentative evidence for Planck-scale structure at black hole horizons*. 2016. arXiv: 1612.00266v1 [gr-qc].
- [250] Jahed Abedi, Hannah Dykaar and Niayesh Afshordi. *Echoes from the Abyss: The Holiday Edition!* Jan. 2017. arXiv: 1701.03485 [gr-qc].
- [251] LIGO Scientific Collaboration. “*LIGO Open Science Center - Data Releases for Observed Transients*”. 2017. DOI: 10.7935/K5MW2F23, 10.7935/K5CC0XMZ, 10.7935/K5H41PBP, 10.7935/K53X84K2. URL: <https://losc.ligo.org/events>.
- [252] Michele Vallisneri, Jonah Kanner, Roy Williams, Alan Weinstein and Branson Stephens. ‘The LIGO Open Science Center’. In: *J. Phys. Conf. Ser.* 610.1 (2015). Ed. by Giacomo Ciani, John W. Conklin and Guido Mueller, p. 012021. DOI: 10.1088/1742-6596/610/1/012021. arXiv: 1410.4839 [gr-qc].
- [253] Takashi Nakamura, Hiroyuki Nakano and Takahiro Tanaka. ‘Detecting quasinormal modes of binary black hole mergers with second-generation gravitational-wave detectors’. In: *Phys. Rev. D* 93.4 (2016), p. 044048. DOI: 10.1103/PhysRevD.93.044048. arXiv: 1601.00356 [astro-ph.HE].
- [254] Bob Holdom and Jing Ren. ‘Not quite a black hole’. In: *Phys. Rev. D* 95.8 (2017), p. 084034. DOI: 10.1103/PhysRevD.95.084034. arXiv: 1612.04889 [gr-qc].
- [255] V. Connaughton et al. ‘Fermi GBM Observations of LIGO Gravitational Wave event GW150914’. In: *Astrophys. J. Lett.* 826.1 (2016), p. L6. DOI: 10.3847/2041-8205/826/1/L6. arXiv: 1602.03920 [astro-ph.HE].
- [256] Gregory Ashton et al. *Comments on: "Echoes from the abyss: Evidence for Planck-scale structure at black hole horizons"*. Dec. 2016. arXiv: 1612.05625 [gr-qc].
- [257] Ronald L. Wasserstein and Nicole A. Lazar. ‘The ASA Statement on p-Values: Context, Process, and Purpose’. In: *The American Statistician* 70.2 (2016), pp. 129–133. DOI: 10.1080/00031305.2016.1154108. eprint: <https://doi.org/10.1080/00031305.2016.1154108>. URL: <https://doi.org/10.1080/00031305.2016.1154108>.
- [258] J. Abedi. Private communication, applying the method of [250] to the posteriors of [239], yielding $\Delta t_{\text{echo}} = 0.2412 \pm 0.018449$ s.
- [259] Jahed Abedi, Hannah Dykaar and Niayesh Afshordi. *Comment on: "Low significance of evidence for black hole echoes in gravitational wave data"*. Mar. 2018. arXiv: 1803.08565 [gr-qc].
- [260] Duncan Macleod et al. *gwpv/gwpv: 0.6.2*. Version v0.6.2. Oct. 2017. DOI: 10.5281/zenodo.1009536. URL: <https://doi.org/10.5281/zenodo.1009536>.
- [261] Duncan Macleod et al. *gwpv/gwpv: GWpy 2.1.3*. Version v2.1.3. Dec. 2021. DOI: 10.5281/zenodo.5788702. URL: <https://doi.org/10.5281/zenodo.5788702>.
- [262] Randy S. Conklin, Bob Holdom and Jing Ren. ‘Gravitational wave echoes through new windows’. In: *Phys. Rev. D* 98.4 (2018), p. 044021. DOI: 10.1103/PhysRevD.98.044021. arXiv: 1712.06517 [gr-qc].

Bibliography

- [263] S. W. Hawking and G. F. R. Ellis. *The Large Scale Structure of Space-Time*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, Feb. 2011. ISBN: 978-0-521-20016-5, 978-0-521-09906-6, 978-0-511-82630-6, 978-0-521-09906-6. DOI: 10.1017/CB09780511524646.
- [264] Ramesh Narayan and Jeffrey E. McClintock. *Observational Evidence for Black Holes*. Dec. 2013. arXiv: 1312.6698 [astro-ph.HE].
- [265] Matt Visser, Carlos Barcelo, Stefano Liberati and Sebastiano Sonego. ‘Small, dark, and heavy: But is it a black hole?’ In: *PoS BHGRS* (2008), p. 010. DOI: 10.22323/1.075.0010. arXiv: 0902.0346 [gr-qc].
- [266] Joseph Polchinski. ‘The Black Hole Information Problem’. In: *Theoretical Advanced Study Institute in Elementary Particle Physics: New Frontiers in Fields and Strings*. 2017, pp. 353–397. DOI: 10.1142/9789813149441_0006. arXiv: 1609.04036 [hep-th].
- [267] Subrahmanyan Chandrasekhar, Valeria Ferrari and John Edwin Enderby. ‘On the non-radial oscillations of a star. III. A reconsideration of the axial modes’. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 434.1891 (1991), pp. 449–457. DOI: 10.1098/rspa.1991.0104. eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.1991.0104>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1991.0104>.
- [268] V. Ferrari and K. D. Kokkotas. ‘Scattering of particles by neutron stars: Time evolutions for axial perturbations’. In: *Phys. Rev. D* 62 (2000), p. 107504. DOI: 10.1103/PhysRevD.62.107504. arXiv: gr-qc/0008057.
- [269] Alex B. Nielsen and Ofek Birnholtz. ‘Testing pseudo-complex general relativity with gravitational waves’. In: *Astron. Nachr.* 339.4 (2018), pp. 298–305. DOI: 10.1002/asna.201813473. arXiv: 1708.03334 [gr-qc].
- [270] Hiroyuki Nakano, Norichika Sago, Hideyuki Tagoshi and Takahiro Tanaka. ‘Black hole ringdown echoes and howls’. In: *PTEP* 2017.7 (2017), 071E01. DOI: 10.1093/ptep/ptx093. arXiv: 1704.07175 [gr-qc].
- [271] J. Veitch et al. ‘Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library’. In: *Phys. Rev. D* 91.4 (2015), p. 042003. DOI: 10.1103/PhysRevD.91.042003. arXiv: 1409.7215 [gr-qc].
- [272] R. K. L. Lo, T. G. F. Li and A. J. Weinstein. ‘Template-based Gravitational-Wave Echoes Search Using Bayesian Model Selection’. In: *Phys. Rev. D* 99.8 (2019), p. 084052. DOI: 10.1103/PhysRevD.99.084052. arXiv: 1811.07431 [gr-qc].
- [273] Ka Wa Tsang et al. ‘A morphology-independent data analysis method for detecting and characterizing gravitational wave echoes’. In: *Phys. Rev. D* 98.2 (2018), p. 024023. DOI: 10.1103/PhysRevD.98.024023. arXiv: 1804.04877 [gr-qc].
- [274] Enrico Barausse, Richard Brito, Vitor Cardoso, Irina Dvorkin and Paolo Pani. ‘The stochastic gravitational-wave background in the absence of horizons’. In: *Class. Quant. Grav.* 35.20 (2018), 20LT01. DOI: 10.1088/1361-6382/aae1de. arXiv: 1805.08229 [gr-qc].
- [275] Alex B. Nielsen, Alexander H. Nitz, Collin D. Capano and Duncan A. Brown. ‘Investigating the noise residuals around the gravitational wave event GW150914’. In: *JCAP* 02 (2019), p. 019. DOI: 10.1088/1475-7516/2019/02/019. arXiv: 1811.04071 [astro-ph.HE].

- [276] Sebastian Khan et al. ‘Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era’. In: *Phys. Rev. D* 93.4 (2016), p. 044007. DOI: 10.1103/PhysRevD.93.044007. arXiv: 1508.07253 [gr-qc].
- [277] Mark Hannam et al. ‘Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms’. In: *Phys. Rev. Lett.* 113.15 (2014), p. 151101. DOI: 10.1103/PhysRevLett.113.151101. arXiv: 1308.3271 [gr-qc].
- [278] Tito Dal Canton et al. ‘Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors’. In: *Phys. Rev. D* 90.8 (2014), p. 082004. DOI: 10.1103/PhysRevD.90.082004. arXiv: 1405.6731 [gr-qc].
- [279] Adriano Testa and Paolo Pani. ‘Analytical template for gravitational-wave echoes: signal characterization and prospects of detection with current and future interferometers’. In: *Phys. Rev. D* 98.4 (2018), p. 044018. DOI: 10.1103/PhysRevD.98.044018. arXiv: 1806.04253 [gr-qc].
- [280] Ram Brustein, A. J. M. Medved and K. Yagi. ‘Lower limit on the entropy of black holes as inferred from gravitational wave observations’. In: *Phys. Rev. D* 100.10 (2019), p. 104009. DOI: 10.1103/PhysRevD.100.104009. arXiv: 1811.12283 [gr-qc].
- [281] Vitor Cardoso, Luís C. B. Crispino, Caio F. B. Macedo, Hirotada Okawa and Paolo Pani. ‘Light rings as observational evidence for event horizons: long-lived modes, ergoregions and nonlinear instabilities of ultracompact objects’. In: *Phys. Rev. D* 90.4 (2014), p. 044069. DOI: 10.1103/PhysRevD.90.044069. arXiv: 1406.5510 [gr-qc].
- [282] Elisa Maggio, Vitor Cardoso, Sam R. Dolan and Paolo Pani. ‘Ergoregion instability of exotic compact objects: electromagnetic and gravitational perturbations and the role of absorption’. In: *Phys. Rev. D* 99.6 (2019), p. 064007. DOI: 10.1103/PhysRevD.99.064007. arXiv: 1807.08840 [gr-qc].
- [283] Wenbin Lu, Pawan Kumar and Ramesh Narayan. ‘Stellar disruption events support the existence of the black hole event horizon’. In: *Mon. Not. Roy. Astron. Soc.* 468.1 (2017), pp. 910–919. DOI: 10.1093/mnras/stx542. arXiv: 1703.00023 [astro-ph.HE].
- [284] Sheperd S. Doeleman et al. ‘Jet Launching Structure Resolved Near the Supermassive Black Hole in M87’. In: *Science* 338 (2012), p. 355. DOI: 10.1126/science.1224768. arXiv: 1210.6132 [astro-ph.HE].
- [285] Kazunori Akiyama et al. ‘First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole’. In: *Astrophys. J. Lett.* 875 (2019), p. L1. DOI: 10.3847/2041-8213/ab0ec7. arXiv: 1906.11238 [astro-ph.GA].
- [286] Kazunori Akiyama et al. ‘First M87 Event Horizon Telescope Results. VI. The Shadow and Mass of the Central Black Hole’. In: *Astrophys. J. Lett.* 875.1 (2019), p. L6. DOI: 10.3847/2041-8213/ab1141. arXiv: 1906.11243 [astro-ph.GA].
- [287] Anna Zulianello, Raúl Carballo-Rubio, Stefano Liberati and Stefano Ansoldi. ‘Electromagnetic tests of horizonless rotating black hole mimickers’. In: *Phys. Rev. D* 103.6 (2021), p. 064071. DOI: 10.1103/PhysRevD.103.064071. arXiv: 2005.01837 [gr-qc].
- [288] Raúl Carballo-Rubio, Francesco Di Filippo, Stefano Liberati and Matt Visser. ‘Phenomenological aspects of black holes beyond general relativity’. In: *Phys. Rev. D* 98.12 (2018), p. 124009. DOI: 10.1103/PhysRevD.98.124009. arXiv: 1809.08238 [gr-qc].
- [289] Nicolas Yunes, Kent Yagi and Frans Pretorius. ‘Theoretical Physics Implications of the Binary Black-Hole Mergers GW150914 and GW151226’. In: *Phys. Rev. D* 94.8 (2016), p. 084002. DOI: 10.1103/PhysRevD.94.084002. arXiv: 1603.08955 [gr-qc].

Bibliography

- [290] Yotam Sherf. ‘Tidal-heating and viscous dissipation correspondence in black holes and viscous compact objects’. In: *Phys. Rev. D* 103.10 (2021), p. 104003. DOI: 10.1103/PhysRevD.103.104003. arXiv: 2104.03766 [gr-qc].
- [291] Samir D. Mathur. ‘The Fuzzball proposal for black holes: An Elementary review’. In: *Fortsch. Phys.* 53 (2005). Ed. by E. Kiritsis, pp. 793–827. DOI: 10.1002/prop.200410203. arXiv: hep-th/0502050.
- [292] Bin Guo, Shaun Hampton and Samir D. Mathur. ‘Can we observe fuzzballs or firewalls?’ In: *JHEP* 07 (2018), p. 162. DOI: 10.1007/JHEP07(2018)162. arXiv: 1711.01617 [hep-th].
- [293] Richard H. Price and Jorge Pullin. ‘Colliding black holes: The Close limit’. In: *Phys. Rev. Lett.* 72 (1994), pp. 3297–3300. DOI: 10.1103/PhysRevLett.72.3297. arXiv: gr-qc/9402039.
- [294] Alessandra Buonanno, Gregory B. Cook and Frans Pretorius. ‘Inspiral, merger and ring-down of equal-mass black-hole binaries’. In: *Phys. Rev. D* 75 (2007), p. 124018. DOI: 10.1103/PhysRevD.75.124018. arXiv: gr-qc/0610122.
- [295] Baoyi Chen, Qingwen Wang and Yanbei Chen. ‘Tidal response and near-horizon boundary conditions for spinning exotic compact objects’. In: *Phys. Rev. D* 103.10 (2021), p. 104054. DOI: 10.1103/PhysRevD.103.104054. arXiv: 2012.10842 [gr-qc].
- [296] Alexander Tyukov, Robert Walker and Nicholas P. Warner. ‘Tidal Stresses and Energy Gaps in Microstate Geometries’. In: *JHEP* 02 (2018), p. 122. DOI: 10.1007/JHEP02(2018)122. arXiv: 1710.09006 [hep-th].
- [297] Qingwen Wang, Naritaka Oshita and Niayesh Afshordi. ‘Echoes from Quantum Black Holes’. In: *Phys. Rev. D* 101.2 (2020), p. 024031. DOI: 10.1103/PhysRevD.101.024031. arXiv: 1905.00446 [gr-qc].
- [298] Vitor Cardoso, Valentino F. Foit and Matthew Kleban. ‘Gravitational wave echoes from black hole area quantization’. In: *JCAP* 08 (2019), p. 006. DOI: 10.1088/1475-7516/2019/08/006. arXiv: 1902.10164 [hep-th].
- [299] Pawel O. Mazur and Emil Mottola. ‘Gravitational vacuum condensate stars’. In: *Proc. Nat. Acad. Sci.* 101 (2004), pp. 9545–9550. DOI: 10.1073/pnas.0402717101. arXiv: gr-qc/0407075.
- [300] Paolo Pani. ‘I-Love-Q relations for gravastars and the approach to the black-hole limit’. In: *Phys. Rev. D* 92.12 (2015). [Erratum: Phys.Rev.D 95, 049902 (2017)], p. 124030. DOI: 10.1103/PhysRevD.95.049902. arXiv: 1506.06050 [gr-qc].
- [301] Guilherme Raposo, Paolo Pani, Miguel Bezares, Carlos Palenzuela and Vitor Cardoso. ‘Anisotropic stars as ultracompact objects in General Relativity’. In: *Phys. Rev. D* 99.10 (2019), p. 104072. DOI: 10.1103/PhysRevD.99.104072. arXiv: 1811.07917 [gr-qc].
- [302] B. P. Abbott et al. ‘GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral’. In: *Phys. Rev. Lett.* 119.16 (2017), p. 161101. DOI: 10.1103/PhysRevLett.119.161101. arXiv: 1710.05832 [gr-qc].
- [303] B. P. Abbott et al. ‘Properties of the binary neutron star merger GW170817’. In: *Phys. Rev. X* 9.1 (2019), p. 011001. DOI: 10.1103/PhysRevX.9.011001. arXiv: 1805.11579 [gr-qc].
- [304] B. P. Abbott et al. ‘GW170817: Measurements of neutron star radii and equation of state’. In: *Phys. Rev. Lett.* 121.16 (2018), p. 161101. DOI: 10.1103/PhysRevLett.121.161101. arXiv: 1805.11581 [gr-qc].

- [305] Collin D. Capano et al. ‘Stringent constraints on neutron-star radii from multimessenger observations and nuclear theory’. In: *Nature Astron.* 4.6 (2020), pp. 625–632. DOI: 10.1038/s41550-020-1014-6. arXiv: 1908.10352 [astro-ph.HE].
- [306] Salvatore Vitale. ‘The first 5 years of gravitational-wave astrophysics’. In: *Science* 372.6546 (2021), abc7397. DOI: 10.1126/science.abc7397. arXiv: 2011.03563 [gr-qc].
- [307] Ram Brustein, A. J. M. Medved and K. Yagi. ‘Discovering the interior of black holes’. In: *Phys. Rev. D* 96.12 (2017), p. 124021. DOI: 10.1103/PhysRevD.96.124021. arXiv: 1701.07444 [gr-qc].

9 Acknowledgements

Many have contributed to the scientific work presented in this thesis, but also to the journey that led to it, and for this I am deeply grateful. They helped me become a better scientist and made the experience a rich and joyful one. I am looking forward to continue research at their side, and they have already contributed much also to the shared further work not included in this thesis.

I thank my thesis advisors, Badri Krishnan, Bruce Allen, and Domenico Giulini for their advice, enthusiasm, and encouragement, teaching me about physics and being a physicist alike, and giving me the opportunity to do research at the Albert Einstein Institute. Badri Krishnan, for many enlightening discussions, for advice on projects to pursue and papers to write, balancing progress on the thesis and following what fascinated me, and for his insights into the path of a scientist. Bruce Allen, for insightful questions and discussions, for instructive anecdotes, and for welcoming me at the institute. Domenico Giulini for deepening my appreciation and understanding of theoretical physics and its conceptual aspects.

Collin Capano, who became an advisor for me as well, for his insight, ideas, explanations and discussions, and for getting into the details of the code. Ofek Birnholtz and Alex Nielsen, for their guidance in working as a scientist.

My co-authors, from whom I have learned much, for many discussions, ideas, the shared research and excitement about doing science.

My colleagues working at or visiting the institute, giving me the opportunity to learn about such diverse topics, and making the institute such a fascinating and welcoming place. And also to the colleagues met elsewhere, on conferences and meetings, for exploring ideas and surroundings alike.

Badri Krishnan and Karsten Danzmann for the opportunity to teach, which I found equally enjoyable and instructive.

The administrative team at the institute, including Gabriele Richardson, Oksana Khiari, Luisa Gärner, and Sandra Bruns, for their always friendly help in making travel plans and processes effortless.

The Atlas computer cluster team, including Henning Fehrmann and Carsten Aulbert, for maintaining and improving the machine that ran our analyses and promptly assisting whenever needed.

The thesis referees, Frank Ohme, Badri Krishnan, and Alex Nielsen, for their attentive reading and providing helpful comments and corrections.

The defence committee, Gunther Seckmeyer, Vitor Cardoso, Bruce Allen, Badri Krishnan, and Alex Nielsen, for the interesting questions and discussion, and for making the defence an experience that I enjoyed.

My office neighbours, for sharing stories both scientific and humorous, and for making the office aesthetically pleasing.

For proof-reading, Lars Nieder and Leonie Vieweg, and also the latter for her appreciation of the no-hair theorem.

My parents and friends, for their support, kindness, the shared adventures, and their humour; or what I choose to interpret as such, for the latter.

For her affection and for making it so, my glorious partner in shining armour, Leonie.

To all of you, my thanks, and may you live long and prosper.

10 Curriculum Vitae

JULIAN WESTERWECK

Personal Data

PLACE AND DATE OF BIRTH: Germany — 24 February 1991
NATIONALITY: German
ADDRESS: Steigertahlstr. 45, 31582 Nienburg, Germany
PHONE: +49 511 76217090
EMAIL: julian.westerweck@aei.mpg.de

Education

- 06/2022 **Ph.D. in Physics at Albert Einstein Institut Hannover, Leibniz Universität Hannover**
“Observational tests of fundamental physics from gravitational wave detections”
Advisors: Badri Krishnan, Bruce Allen, Domenico Giulini
- 07/2016 **M.Sc. in Physics at Leibniz Universität Hannover**
“Interpretation and Ambiguity of Bowen-York Type Black Hole Initial Data”
Advisors: Domenico Giulini, Badri Krishnan
- 10/2013 **B.Sc. in Physics at Leibniz Universität Hannover**
“Quantum Dynamics in Complex Time for the Description of High Harmonics”
Advisor: Manfred Lein

Teaching Experience

- | | |
|-------------------|---|
| 10/2017 - 03/2020 | Teaching Assistant at Leibniz Universität Hannover
Organising exercise classes and exams for courses :
Gravitational Physics, Mechanics and Relativity, Electricity,
Data Analysis Lab |
| 04/2012 - 10/2016 | Student Assistant at Leibniz Universität Hannover
Tutoring undergraduate students in various courses :
Mathematical Methods in Physics, Theoretical Electrodynamics,
Analytical Mechanics, Advanced Quantum Theory,
Experimental Physics |

Presentations & Workshops

- 14/12 - 17/12 2021 **Gravitational Wave Physics and Astronomy Workshop**, Albert Einstein Institute Hannover, Germany
Poster presentation: “Sub-atomic constraints on the Kerr geometry of GW150914”
- 05/07 - 10/07 2021 **16th Marcel Grossmann Meeting**, Online
Talk: “Observation of a multimode quasi-normal spectrum from a perturbed black hole”
- 12/01 - 18/01 2020 **Gravitational wave searches and parameter estimation in the era of detections**, Schloss Ringberg, Kreuth, Germany
- 14/10 - 17/10 2019 **Gravitational Wave Physics and Astronomy Workshop**, University of Tokyo, Japan
Talk: “Population of measurable black hole ringdown modes”
- 07/10 - 11/10 2019 **Yukawa International Seminar: Black Holes and Neutron Stars with Gravitational Waves**, Kyoto University, Japan
Poster presentation: “Population of observable black hole ringdown modes”
- 07/07 - 12/07 2019 **22nd International Conference on General Relativity and Gravitation**, Valencia, Spain
Talk: “Measuring higher order black hole ringdown modes”
- 14/05 - 16/05 2019 **Gravitational-wave parameter estimation with PyCBC Inference workshop**, University of Portsmouth, United Kingdom
- 01/12 - 04/12 2018 **Gravitational Wave Physics and Astronomy Workshop**, University of Maryland, Maryland, USA
Talk: “Upper limits on gravitational wave echoes”
- 04/06 - 08/06 2018 **Summer School in Astroinformatics**, Penn State University, Pennsylvania, USA
- 29/05 - 02/06 2018 **Summer School in Statistics for Astronomers**, Penn State University, Pennsylvania, USA
- 08/11 - 10/11 2017 **Quantum Black Holes in the Sky?**, Perimeter Institute, Canada
Invited talk: “Improvements on the methods for searching echoes”
- 14/08 - 17/08 2017 **Physics and Astrophysics at the Extreme**, Nikhef, The Netherlands
Poster presentation: “Templated search for black hole GW echoes in LIGO data”
- 30/05 - 02/06 2017 **Gravitational Wave Physics and Astronomy Workshop**, Annecy, France
Poster presentation: “Templated search for black hole GW echoes in LIGO data”

Publications

Shilpa Kastha, Collin D. Capano, Julian Westerweck, Miriam Cabero, Badri Krishnan and Alex B. Nielsen. ‘Model systematics in time domain tests of binary black hole evolution’. In: *Phys. Rev. D* 105.6 (2022), p. 064042. DOI: 10.1103/PhysRevD.105.064042. arXiv: 2111.13664 [gr-qc]

Julian Westerweck, Yotam Sherf, Collin D. Capano and Ram Brustein. *Sub-atomic constraints on the Kerr geometry of GW150914*. Aug. 2021. arXiv: 2108.08823 [gr-qc]

Collin D. Capano, Miriam Cabero, Julian Westerweck, Jahed Abedi, Shilpa Kastha, Alexander H. Nitz, Alex B. Nielsen and Badri Krishnan. *Observation of a multimode quasi-normal spectrum from a perturbed black hole*. May 2021. arXiv: 2105.05238 [gr-qc]

Miriam Cabero, Julian Westerweck, Collin D. Capano, Sumit Kumar, Alex B. Nielsen and Badri Krishnan. ‘Black hole spectroscopy in the next decade’. In: *Phys. Rev. D* 101.6 (2020), p. 064044. DOI: 10.1103/PhysRevD.101.064044. arXiv: 1911.01361 [gr-qc]

Alex B. Nielsen, Collin D. Capano, Ofek Birnholtz and Julian Westerweck. ‘Parameter estimation and statistical significance of echoes following black hole signals in the first Advanced LIGO observing run’. In: *Phys. Rev. D* 99.10 (2019), p. 104012. DOI: 10.1103/PhysRevD.99.104012. arXiv: 1811.04904 [gr-qc]

Julian Westerweck, Alex Nielsen, Ofek Fischer-Birnholtz, Miriam Cabero, Collin Capano, Thomas Dent, Badri Krishnan, Grant Meadors and Alexander H. Nitz. ‘Low significance of evidence for black hole echoes in gravitational wave data’. In: *Phys. Rev. D* 97.12 (2018), p. 124037. DOI: 10.1103/PhysRevD.97.124037. arXiv: 1712.09966 [gr-qc]

Gregory Ashton, Ofek Birnholtz, Miriam Cabero, Collin Capano, Thomas Dent, Badri Krishnan, Grant David Meadors, Alex B. Nielsen, Alex Nitz and Julian Westerweck. *Comments on: "Echoes from the abyss: Evidence for Planck-scale structure at black hole horizons"*. Dec. 2016. arXiv: 1612.05625 [gr-qc]

Publications as member of the LIGO Scientific Collaboration are not listed.