Fundamental physics in the era of gravitational-wave astronomy: the direct measurement of gravitational-wave polarizations and other topics

Thesis by
Maximiliano Isi Bañales

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2019
Defended October 5, 2018

Maximiliano Isi Bañales
ORCID: 0000-0001-8830-8672

ACKNOWLEDGMENTS

I am a member of the LIGO Laboratory. LIGO was constructed by the California Institute of Technology and the Massachusetts Institute of Technology, with funding from the National Science Foundation, and operates under cooperative agreement PHY-0757058. I carried out most of the work presented here as a member of the LIGO Scientific Collaboration and benefited from interaction with many of its members. However, except where otherwise noted, the opinions expressed here are my own and not necessarily those of the Collaboration. This thesis carries LIGO document number LIGO-P1800336.

I would like to thank my thesis advisor, Alan Weinstein, whose mentorship has been invaluable since back when I was an undergraduate summer student: thank you, Alan, for allowing me to explore my own interests, for pointing me in the right direction when needed and for always being supportive. I am also extremely grateful to my undergraduate advisor, Jonas Mureika, who enriched my college experience to an incalculable degree by exposing me to real-world research and providing much-needed guidance.

I owe much to many colleagues and collaborators. Among them, I especially thank Carver Mead for much inspiration and insight—I am fortunate that we had the chance to work together. I am grateful to Matthew Pitkin, Tjonnie Li, Rory Smith and Leo Stein, all of whom have served as my mentors and have taught me a lot. I would also like to acknowledge my peer LIGO graduate students Surabhi Sachdev, Craig Cahillane and Tom Callister, as well as Matthew Giesler and Vijay Varma.

I am indebted to Nelly and Jim Kilroy, without whose friendship and unwavering support none of this would have been possible. Thank you, Nelly, for always believing in me. In a way, I hope this work serves to honor Jim's memory and is seen as a small addition to his outstanding legacy. I am also grateful to the late Bud Wheelon who, with a generous gesture, opened for me invaluable doors of opportunity.

Mi aventura comenzó en Uruguay. Probablemente hubiera terminado allí de no ser por la visión e iniciativa de Julio Fernández Techera, a quien estaré siempre agradecido. Estoy en deuda también con el Colegio Seminario de Montevideo, donde recibí la sólida preparación que me sirve diariamente de base. Doy gracias a Nelly Noboa por estimular mi curiosidad filosófica y alentar mis esfuerzos por comprender el mundo. En particular, agradezco de todo corazón por Pablo García y Alejandro

Villamil, quienes encendieron la llama de mi pasión por la física y sin quienes nunca hubiera emprendido este camino. Lamento no poder compartir la alegría de este logro con Alejandro, a quien siempre recordaré con mucho afecto y gratitud.

Por último, lo más importante: agradezco a mi familia. Gracias a la Flia. Pena, en especial a Victoria, por estar siempre presentes. También gracias por su cariño a mi abuela, Mirtha, y a la memoria de Tatín. Finalmente, no hay palabras para agradecer a mi hermana y a mis padres, a quienes debo quién soy. Gracias por inculcar en mí la importancia de la educación y el pensamiento crítico; gracias por su apoyo incondicional; gracias por empujarme siempre a seguir mis sueños, aunque éstos me llevasen lejos.

⋆ ⋆ ⋆

Nunc magis atque magis

ABSTRACT

In this thesis, I explore several avenues for learning about fundamental physics from gravitational-wave (GW) observations. In particular, I focus on the phenomenological study of basic properties of GWs in ways that require minimal assumptions about the underlying nature of gravity. I place a special emphasis on GW polarizations, but also consider their speed and possible dispersion. To constrain possible modifications to general relativity, I develop data-analysis frameworks to measure these properties with both transient and persistent signals detected by ground-based detectors. This includes, among other results, the analysis of two LIGO and Virgo compact-binary detections, GW170814 and GW170817, to produce the first direct observational statements about the local geometry of GW polarizations. I also present constraints on the potential amplitude of nontensorial monochromatic signals from 200 known pulsars in the Milky Way and describe in detail the methods used to obtain them. Because stochastic signals will be a great resource for studying GW properties, I also carefully review the assumptions that go into standard stochastic analyses and explore their applicability beyond general relativity, concluding that those measurements will have to be interpreted carefully to make meaningful statements about corrections to Einstein's theory. Besides the properties of the waves themselves, I also study the prospect for using GWs as a means to uncover signatures of new ultralight bosons—an exciting possibility that could bring particle physics into the reach of GW astronomy. I explore the potential of current and future detectors to detect these conjectured particles, concluding that third-generation instruments are certain to place theoretically interesting constraints. Because little can be done in the absence of signals, I also propose data-analysis methods to improve LIGO and Virgo's chances of detecting both transient and continuous signals. The latter have now been used to diagnose real detection candidates on several occasions.

PUBLISHED CONTENT AND CONTRIBUTIONS

[1] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), "Tests of General Relativity with GW170817," (2018), M.I. contributed the polarization analysis and provided feedback on the text, arXiv:1811.00364 [gr-qc] .

[2] M. Isi, L. Sun, R. Brito, and A. Melatos, "Directed searches for gravitational waves from ultralight bosons," (2018), M.I. conceived the project and led the writing of the manuscript, arXiv:1810.03812 [gr-qc] .

[3] M. Isi and L. C. Stein, "Measuring stochastic gravitational-wave energy beyond general relativity," (2018), M.I. conceived the project, carried out most calculations and led the writing of the manuscript, arXiv:1807.02123 [gr-qc] .

[4] M. Isi, R. Smith, S. Vitale, T. J. Massinger, J. Kanner, and A. Vajpeyi, Phys. Rev. **D98**, 042007 (2018), M.I. carried out the analysis and led the writing of the manuscript, arXiv:1803.09783 [gr-qc] .

[5] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **120**, 031104 (2018), M.I. carried out the analysis and wrote the manuscript, arXiv:1709.09203 [gr-qc] .

[6] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **119**, 141101 (2017), M.I. carried out the polarization analysis, wrote the section describing it and provided feedback on the manuscript as a whole, arXiv:1709.09660 [gr-qc] .

[7] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. **D96**, 122006 (2017), [Erratum: Phys. Rev.D97,no.12,129903(2018)]. M.I. contributed to the followup of outliers and provided feedback on the manuscript, arXiv:1710.02327 [gr-qc] .

[8] M. Isi and A. J. Weinstein, "Probing gravitational wave polarizations with signals from compact binary coalescences," (2017), M.I. carried out the analysis and wrote the manuscript, arXiv:1710.03794 [gr-qc] .

[9] M. Pitkin, M. Isi, J. Veitch, and G. Woan, (2017), M.I. developed most of the code to deal with nontensorial polarizations and provided feedback on the manuscript, arXiv:1705.08978 [gr-qc] .

[10] T. Callister, A. S. Biscoveanu, N. Christensen, M. Isi, A. Matas, O. Minazzoli, T. Regimbau, M. Sakellariadou, J. Tasson, and E. Thrane, Phys. Rev. **X7**, 041058 (2017), M.I. advised on statistical methods and provided feedback on the manuscript, arXiv:1704.08373 [gr-qc] .

[11] M. Isi, M. Pitkin, and A. J. Weinstein, Phys. Rev. **D96**, 042001 (2017), M.I. carried out the analysis and wrote the manuscript, arXiv:1703.07530 [gr-qc] .

[12] R. Tso, M. Isi, Y. Chen, and L. Stein, in *Proceedings, 7th Meeting on CPT and Lorentz Symmetry (CPT 16): Bloomington, Indiana, USA, June 20-24, 2016* (2017) pp. 205–208, M.I. contributed the simulations and data analysis, and provided feedback on the manuscript, arXiv:1608.01284 [gr-qc] .

[13] M. Isi, A. J. Weinstein, C. Mead, and M. Pitkin, Phys. Rev. **D91**, 082002 (2015), M.I. carried out the analysis and wrote the manuscript, arXiv:1502.00333 [gr-qc] .

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

NOMENCLATURE

aLIGO. Advanced LIGO.

BBH. Binary black hole.

BCI. Bayes factor comparing signal vs incoherent signal hypotheses.

BCR. Bayesian coherence ratio.

BH. Black hole.

BNS. Binary neutron star.

BSN. Bayes factor comparing signal vs Gaussian-noise hypotheses.

CBC. Compact-binary coalescence.

CS. Chern-Simons.

CW. Continuous gravitational wave.

EM. Electromagnetic.

FAR. False-alarm rate.

FP. Fierz-Pauli.

G4v. Four-vector gravity.

GR. General relativity.

GW. Gravitational wave.

H1. LIGO Hanford 4km instrument.

H2. LIGO Hanford 2km instrument.

HMM. Hidden Markov model.

ICRS. International Celestial Reference System.

IFAR. Inverse false-alarm rate.

L1. LIGO Livingston 4km instrument.

LAL. LIGO Algorithm Library.

LHO. LIGO Hanford Observatory.

LIGO. Laser Interferometer Gravitational-wave Observatory.

LLO. LIGO Livingston Observatory.

LVT. LIGO-Virgo trigger.

O1. Advanced LIGO Observation Run 1 (September 12, 2015 through January 19, 2016).

O2. Advanced LIGO Observation Run 2 (November 30, 2016 through August 25, 2017).

PN. Post-Newtonian.

QCD. Quantum chromodynamics.

RMSE. Root-mean-square error.

ROC. Receiver operator characteristic.

ROQ. Reduced-order quadrature.

S5. LIGO Science Run 5 (November 4, 2005 through October 1, 2007).

S6. LIGO Science Run 6 (July 7, 2009 through October 20, 2010).

SNR. Signal-to-noise ratio.

SSB. Solar-system barycenter.

ST. Scalar-tensor.

VSR. Virgo Science Run.

Chapter 1

INTRODUCTION

On September 14, 2015, the LIGO instruments made history by detecting a passing burst of gravitational waves. After almost exactly 100 years from their prediction by Albert Einstein, this was the first time that these elusive spacetime ripples were directly observed at Earth. The achievement crowned decades of effort by hundreds of scientists in the field, and inaugurated a new era of observational physics and astronomy.

In a sense, gravitational waves are to gravity what light is to electromagnetism. Both are direct manifestations of one of the four fundamental forces (the strong and weak nuclear forces being the other two) and, as such, carry in their nature invaluable information about the most fundamental structure of the Universe. The similarities are also observable: both kinds of waves travel at the same speed through empty space and are created by the acceleration of some type of charge (electric charge for light; gravitational charge, also known as "mass", for gravitational waves). Gravity being drastically weaker than electromagnetism, gravitational waves are much harder to detect than light. Both are, nonetheless, ubiquitous in nature.

In another, richer sense, gravitational waves can be understood not as mere ripples of a field *in* spacetime, but as ripples *of* spacetime itself. Space and time, as we learn from Einstein's general relativity, are dynamic entities that can bend and twist in the presence of matter or energy—gravitational waves are perturbations in this curvature. Accordingly, the main observable effect of these waves is to alternatively stretch and squeeze lengths (producing a "strain") as they whiz by at the speed of light.[1] For any conceivable source, however, these strains are unimaginably small, making their detection a herculean task.

LIGO, the Laser Interferometer Gravitational-wave Observatory, was designed to do just that. Its two instruments, one in Hanford, WA, the other in Livingston, LA, are extremely sophisticated machines that can measure the imperceptible stretching and squeezing caused by a passing gravitational wave. To achieve this, a laser beam is split into two by a translucent mirror, which redirects the laser light down two perpendicular arms 4 km in length; at the end of each arm, the laser is reflected

[1] For a technical characterization of the effect of gravitational waves, see Appendix A.1.

back to the beam splitter, where it is recombined and sent out for readout. As a gravitational wave goes by, it effectively changes the relative length of the two arms—this modifies the distance traveled by light in the instrument and thus becomes imprinted in its phase. Thanks to exquisite optical techniques (most importantly, Fabry-Perot cavities) this method can be used to measure displacements orders of magnitude smaller than the wavelength of the laser light. Advanced seismic isolation and vacuum technology make this possible even in spite of the machine's noisy environment (wind, earthquakes, human activity, to name a few offenders).

At its peak, the first gravitational wave LIGO detected (nicknamed GW150914) changed the lengths of the 4 km arms by a distance roughly a thousand times smaller than the radius of a proton ($\sim 10^{-18}$ m). By thoroughly analyzing the 200ms-long signal, we have learned it originated in the collision of two black holes over 1 billion light-years away. As the two objects, each weighing more than 30 times the mass of our Sun, plunged inexorably towards each other, they rattled the spacetime around them and, in their final moments, produced more power than all of the stars in our Universe combined. But the violent collision was invisible: all that energy, roughly equivalent to three pulverized Suns, was released solely in gravitational waves, not light. The echo of this cosmic crash eventually washed over Earth and was observed by LIGO's instruments. This was not only the first direct detection of gravitational waves, but also of binary black holes.

The detection of GW150914 was only the beginning: since then over six more events like it have been observed. Most remarkably, in August of 2017, LIGO and its European counterpart, Virgo, made the first observation of the coalescence of two neutron stars (GW170817). In that case, the gravitational signal was also accompanied by exuberant emission across the electromagnetic spectrum, which was recorded by 70 telescopes on the ground and in space. The historic discovery confirmed the long-suspected relationship between gamma-ray bursts and neutron star collisions, and delivered a wealth of astrophysical data.

Besides astrophysics, gravitational waves are also an exceptional resource for learning about fundamental physics. This is so, first of all, because they allow us to observe the most violent events in the universe, in regimes in which all of the fundamental forces are active at their most extreme. The inside of a neutron star and the surroundings of a black hole are both examples of exotic environments that cannot be replicated for study in the laboratory. In particular, gravitational waves from black hole collisions offer a uniquely clean probe of gravity in its most interesting, nonlinear regime,

which could offer invaluable clues about the nature of space and time.

Additionally, gravitational waves are direct manifestations of the dynamics of spacetime itself and as such, hold in their basic properties important clues about how the Universe is woven together. For instance, the polarization and speed of these waves can reveal the geometric structure and symmetries underlying fundamental physical laws. In the past few years, LIGO and Virgo observations have already begun to inform these measurements. As an example, with the first detection we were able to measure the dispersion of gravitational waves, allowing for the first dynamical constraints on the mass of the graviton. Later, the gravitational and electromagnetic waves from the neutron-star binary allowed us to place the first direct observational bounds on the speed of gravitational waves. They also informed us about the geometry of gravitational-wave polarizations, significantly adding to the information previously gained with the first three-detector observation of a binary black-hole (GW170814).

From a theoretical perspective, gravity remains the most mysterious of the four fundamental forces. Furthermore, in spite of its prevalence in everyday life, gravity's intrinsic weakness makes it paradoxically hard to study. This makes gravitational-wave observations all the more valuable: gravitational waves provide a unique window into the dynamics of space and time that promises to point us to a better understanding of gravity, possibly beyond Einstein's theory. Furthermore, by allowing us to study otherwise inaccessible exotic sources, they may also teach us about cosmology and nuclear and particle physics.

In this thesis, I will explore some of these possibilities, including a discussion of the first results on gravitational-wave polarizations and the prospects of using gravitational waves to detect new particles. This is outlined below.

Chapter 2

OVERVIEW

In this thesis, I explore several paths for learning about fundamental physics using gravitational waves. In particular, I treat the phenomenological study of their basic properties in ways that require minimal assumptions about the underlying nature of gravity. The main results and focus of this work concern gravitational-wave polarizations. However, I also cover adjacent topics, including the possibility of using gravitational waves to discover new particles beyond the standard model.

Little science can be done in the absence of signals. With that in mind, Chapter 3 and Chapter 4 discuss data-analysis methods to improve the way ground-based instruments make detections without requiring updates to the hardware. The first one concerns signals from compact-binary coalescences, like those detected so far: it shows how a Bayesian measure of signal coherence can increase the number of detections made by LIGO, better distinguishing astrophysical signals from noise. The second one concerns targeted searches for continuous gravitational waves from known pulsars: it presents a technique to establish the significance of detection candidates without idealizing detector noise—this is a big problem in the field, so the method has already been used to diagnose outliers in real data.

We then move on to gravitational-wave polarizations in Chapters 5–9. This includes observational constraints on the local geometry of polarizations, which were obtained from LIGO and Virgo's compact-binary detections. Besides the treatment of transient signals, I also present methods to detect and characterize continuous gravitational waves from known pulsars without assuming their polarization geometry conforms to general relativity. Those methods were applied in a search for signals in actual data from Advanced LIGO's first observation run (O1), the results of which are also included here.

When detected, stochastic gravitational waves will also be another important resource to study the fundamental nature of gravity. For this reason, in Chapter 10 I carefully review the assumptions that go into the traditional searches for these signals and evaluate their applicability in the context of specific extensions of Einstein's theory. This study was motivated by the prospect of analyzing the polarization content of the stochastic background and the realization that tacit assumptions could lead to

misinterpretation of the results.

Next, I offer a short treatment of gravitational-wave speed and dispersion in Chapters 11 and 12. This includes a proposal for how to measure the speed of gravitational waves using persistent signals from known pulsars, as well as how to measure anisotropic dispersion relations with several compact-binary coalescences. These discussions are only exploratory.

Finally, in Chapter 13, I depart from the topic of gravitational-wave properties to study the potential of ground-based detectors to discover new particles. In particular, the presence of a yet-undiscovered ultralight boson could be revealed in continuous gravitational-wave signals originated in the immediate vicinity of fast-spinning black holes. I explore the prospect for detecting these signals by targeting known black holes with second and third generation instruments.

A summary of each chapter is provided below:

Chapter 3: Enhancing confidence in the detection of gravitational waves from compact binaries using signal coherence

We show that gravitational-wave signals from compact binary mergers may be better distinguished from instrumental noise transients by using Bayesian models that look for signal coherence across a detector network. This can be achieved even when the signal power is below the usual threshold for detection. This method could reject the vast majority of noise transients, therefore increasing sensitivity to weak gravitational waves. We demonstrate this using simulated signals, as well as data for GW150914 and LVT151012. Finally, we explore ways of incorporating our method into existing Advanced LIGO and Virgo searches to make them significantly more powerful.

Chapter 4: Establishing the significance of continuous gravitational-wave detections from known pulsars

We provide a much-needed efficient and robust way to empirically estimate the background of searches for continuous gravitational waves targeted at known pulsars, enabling estimates of detection significance that are valid in actual (non-Gaussian) instrumental noise. The technique is based on "off-sourcing" the search, namely blinding the data to actual astrophysical signals by assuming an incorrect sky location for the target. We introduce the rationale behind this strategy, explore its theoretical applicability and study its performance in real and simulated data. We find that off-

sourcing is an excellent tool, outperforming another standard method for estimating search backgrounds. This method has already been put into practice for diagnosing outliers in actual LIGO and Virgo searches.

Chapter 5: Probing gravitational wave polarizations with signals from compact binary coalescences

Generic metric theories of gravity may predict up to six polarization modes for metric perturbations: two tensor (helicity ± 2), two vector (helicity ± 1), and two scalar (helicity 0) modes. In contrast, one of the key predictions of general relativity (GR) is that metric perturbations possess only two tensor degrees of freedom. Therefore, a detection of any nontensorial mode would be unambiguous indication of physics beyond GR. In this technical note, we study the possibility of using networks of ground-based detectors to directly measure gravitational-wave polarizations using signals from compact binary coalescences. We present a simple data analysis method to partially achieve this, assuming presence of a strong signal well-captured by a template conforming to general relativity.

Chapter 6: Constraints on gravitational-wave polarizations from compact-binary coalescences

The observation of GW170814 and GW170817 has allowed us to gather the first direct evidence that GWs come in the polarizations predicted by GR, as opposed to the extreme alternatives of purely vector or purely scalar polarizations. This is an important phenomenological check that the weak-field geometry of spacetime does not deviate *drastically* from what Einstein predicted. Here we discuss these results at length, including their limitations and interpretation.

Chapter 7: Detecting beyond-Einstein polarizations of continuous gravitational waves

The direct detection of gravitational waves with the ground-based detectors, like Advanced LIGO, provides the opportunity to measure deviations from the predictions of general relativity. One such departure would be the existence of alternative polarizations. To measure these, we study a single detector measurement of a continuous gravitational wave from a triaxial pulsar source. We develop methods to detect signals of any polarization content and distinguish between them in a model

independent way. We present LIGO S5 sensitivity estimates for 115 pulsars.

Chapter 8: Probing Dynamical Gravity with the Polarization of Continuous Gravitational Waves

The direct detection of gravitational waves provides the opportunity to measure fundamental aspects of gravity which have never been directly probed before, including the polarization of gravitational waves. In the context of searches for continuous waves from known pulsars, we present novel methods to detect signals of any polarization content, measure the modes present and place upper limits on the amplitude of nontensorial components. This will allow us to obtain new model-independent, dynamical constraints on deviations from general relativity. We test this framework on multiple potential sources using simulated data from three advanced-era detectors at design sensitivity. We find that signals of any polarization will become detectable and distinguishable for characteristic strains $h \gtrsim 3 \times 10^{-27}\sqrt{1\ \mathrm{yr}/T}$, for an observation time T. We also find that our ability to detect nontensorial components depends only on the power present in those modes, irrespective of the strength of the tensorial strain.

Chapter 9: First search for nontensorial gravitational waves from known pulsars

We present results from the first directed search for nontensorial gravitational waves. While general relativity allows for tensorial (plus and cross) modes only, a generic metric theory may, in principle, predict waves with up to six different polarizations. This analysis is sensitive to continuous signals of scalar, vector or tensor polarizations, and does not rely on any specific theory of gravity. After searching data from the first observation run of the advanced LIGO detectors for signals at twice the rotational frequency of 200 known pulsars, we find no evidence of gravitational waves of any polarization. We report the first upper limits for scalar and vector strains, finding values comparable in magnitude to previously-published limits for tensor strain. Our results may be translated into constraints on specific alternative theories of gravity.

Chapter 10: Measuring stochastic gravitational-wave energy beyond general relativity

Gravity theories beyond GR can make very different predictions for the properties of gravitational waves: their polarizations, dispersion, speed, and, importantly, energy content are all heavily theory-dependent. All these features can potentially be probed by measuring the stochastic gravitational-wave background. However, most existing treatments of this background beyond GR overlook modifications to the energy carried by gravitational waves, or rely on GR assumptions that are invalid in other theories. This may lead to mistranslation between the observable cross-correlation of detector outputs and gravitational-wave energy density, and thus to errors when deriving observational constraints on theories. In this article, we lay out a generic formalism for stochastic gravitational-wave searches, applicable to a large family of theories beyond GR. We explicitly state the (often tacit) assumptions that go into these searches, evaluating their generic applicability, or lack thereof. Examples of problematic assumptions are: statistical independence of linear polarization amplitudes; which polarizations satisfy equipartition; and which polarizations have well-defined phase velocities. We also show how to correctly infer the value of the stochastic energy density in the context of any given theory. We demonstrate with specific theories in which some of the traditional assumptions break down: Chern-Simons gravity, scalar-tensor theory, and Fierz-Pauli massive gravity. In each theory, we show how to properly include the beyond-GR corrections and how to interpret observational results.

Chapter 11: Measuring the speed of continuous gravitational waves

We explore how to directly measure the speed of gravitational waves by studying continuous monochromatic signals (CWs), for example, from rapidly rotating neutron stars. This technique relies on sensitivity to the Doppler modulation (in this context, also known as "Rømer modulation") of the expected signal caused by Earth's motion with respect to the source. This is similar to the strategy used by Ole Rømer in 1676 to produce the first quantitative estimate of the speed of light itself. We offer a tentative demonstration of the method with a simplistic analysis. Although we do not expect it to be competitive with the bound placed by the joint observation of GWs and light from GW170817, this measurement would offer an independent alternative, not limited by the same systematics.

Chapter 12: Towards constraining generic gravitational-wave dispersion relations

In GR, GWs travel along null geodesics and are, thus, non-dispersive. However, this prediction is a specific consequence of GR's equations of motion and need not hold in generic theories. In fact, modifications to the GW dispersion are equivalent to a change in the wave equation for metric perturbations, which in turn derive from the action of the theory and can take many forms, in principle. Here, we explore the effect of an arbitrary dispersion relation on GW phenomenology. We heuristically derive the observable effect of such a dispersion relation on the GW waveform and apply the procedure to specific forms of the dispersion relation (massive graviton and simplest anisotropic term). We discuss future data analysis implementations of this result.

Chapter 13: Directed searches for gravitational waves from ultralight bosons

Gravitational-wave detectors could be used to search for yet-undiscovered ultralight bosons, including those conjectured to solve problems in particle physics, high-energy physics and cosmology. In particular, ground-based instruments could probe boson masses between 10^{-15} eV to 10^{-12} eV, which are largely inaccessible to other experiments. In this paper, we explore the prospect of searching for the continuous gravitational waves generated by boson clouds around known black holes. We carefully study the predicted waveforms and use the latest-available numerical results to model signals for different black-hole and boson parameters. We then demonstrate the suitability of a specific method (hidden Markov model tracking) to efficiently search for such signals, even when the source parameters are not perfectly known and in the presence of theoretical uncertainty. We empirically study this method's sensitivity and computational cost in the context of boson signals, finding that it will be possible to target remnants from compact-binary mergers localized with at least three instruments. For signals from scalar clouds, we also compute detection horizons for future detectors (Advanced LIGO, LIGO Voyager, Cosmic Explorer and the Einstein Telescope). Among other results, we find that, after one year of observation, an Advanced LIGO detector at design sensitivity could observe these sources up to no more than 100 Mpc for most parameters (at most close to 10^3 Mpc), while Cosmic Explorer could reach up to $\sim 10^4$ Mpc. These projections should be more reliable than previous estimates based on analytic approximations to the signal power or idealized search strategies. Finally, we discuss specific implications

for the followup of compact-binary coalescences and black-holes in x-ray binaries. Along the way, we review the basic physics of the superradiance of bosons around black holes, in the hope of providing a bridge between the theory and data-analysis literatures.

We summarize the results and discuss future work in Chapter 14.

Chapter 3

ENHANCING CONFIDENCE IN THE DETECTION OF GRAVITATIONAL WAVES FROM COMPACT BINARIES USING SIGNAL COHERENCE

M. Isi, R. Smith, S. Vitale, T. J. Massinger, J. Kanner, and A. Vajpeyi, Phys. Rev. **D98**, 042007 (2018), M.I. carried out the analysis and led the writing of the manuscript, arXiv:1803.09783 [gr-qc] .

3.1 Introduction

A pair of neutron stars or black holes merges somewhere in the observable universe roughly every 15–200s, releasing large amounts of energy in the form of gravitational waves (GWs) [1–7]. One of the limiting factors in detecting such GWs with existing detectors, like Advanced LIGO (aLIGO) and Virgo [8, 9], is data contamination by instrumental noise transients (*glitches*) that may mimic astrophysical signals [10]. Glitches can lower the inferred statistical significance of GW signals, making their detection more difficult. In this paper, we show how signal coherence may be used to address this problem by significantly improving our ability to distinguish genuine GW signals from glitches using Bayesian model comparison.

In particular, we demonstrate that Bayesian models—as proposed in [11]—may successfully distinguish real GWs from glitches by using the fact that the former must be *coherent* across detectors, while the latter will generally not be. Here, coherence means that a real GW must produce strain signals in different instruments that: (*i*) are coincident in time (up to a time-of-flight delay); (*ii*) are well-described by a compact-binary-coalescence (CBC) waveform; and (*iii*) share a phase evolution consistent with a single astrophysical source. In contrast, glitches should not be expected to fully satisfy these criteria. Making full use of this information (the expected coherence of signals and incoherence of glitches) may allow us to detect weaker signals than currently possible.

From a subset of glitches and detection candidates (*triggers*) from aLIGO's first observation run (O1), we find that: (*a*) the majority of glitches are markedly more incoherent than coherent across detectors, irrespective of their loudness or the

detection significance assigned by one of the main detection pipelines; (*b*) simulated signals can be identified by their coherence, as long as they are distinguishable from Gaussian noise in at least two detectors; and finally, (*c*) the "gold-plated" detection GW150914 (detection significance $> 5.1\sigma$) [1] and the "silver-plated" *candidate* LVT151012 (detection significance $\sim 2.1\sigma$) [3] are both decidedly more coherent than incoherent. This study of real data thus implies that the Bayesian comparison of coherent and incoherent signal models has the potential to significantly improve the sensitivity of CBC searches, even with currently available computational resources.

3.2 Searches

Templated searches for transient gravitational waves work by constructing a ranking statistic based on matched filtering [12–17]. In principle, to make a rigorous statement about the statistical significance of a pair of time-coincident triggers, it is necessary to know the probability that a given event was produced by instrumental noise, rather than an actual GW. This likelihood may be estimated empirically from the value of the ranking statistic for a large representative set of triggers known with certainty to be spurious. Such a set of signal-free triggers is denoted *background*, in contrast to the *foreground* of candidates that may contain a signal.

Because detectors cannot be physically shielded from gravitational waves, *ad hoc* data analysis techniques must be used to estimate the background. One such strategy is to construct *time slides* by applying relative time offsets (longer than the light-travel time between sites) between the data of different detectors [16, 17]. Detection significance can then be inferred, in a frequentist way, by comparing the value of the ranking statistic for a time-coincident foreground trigger to that of time-slid background triggers. The rate at which background triggers are produced with a given value of the ranking statistic is usually referred to as the *false-alarm rate* (FAR).

Efficient signal detection requires a ranking statistic that extracts the most information from the data, in order to discriminate between noise and weak astrophysical signals. However, existing CBC searches are *not* optimal in this sense: they do not incorporate knowledge of *all* features that may distinguish GWs from noise. Moving towards an optimal statistic is a great challenge, but one large step is to demand that foreground triggers in two or more detectors should be *better* described as coherent gravitational-wave signals, rather than incoherent glitches. Importantly, it is not enough to provide some measure of coherence: one must also prove that an incoherent model is not

more successful at describing the data.

3.3 Coherence vs incoherence

To achieve this, we introduce the *Bayesian coherence ratio* (BCR): the odds between the hypothesis that the data comprise a coherent CBC signal in Gaussian noise ($\mathcal{H}_\mathrm{S}$), and the hypothesis that they instead comprise incoherent instrumental features ($\mathcal{H}_\mathrm{I}$)—meaning each detector has *either* a glitch in Gaussian noise ($\mathcal{H}_\mathrm{G}$), *or* pure Gaussian noise ($\mathcal{H}_\mathrm{N}$). For a network of D detectors:

$$\mathrm{BCR} \equiv \frac{\alpha Z^S}{\prod_{i=1}^{D} \left[\beta Z_i^G + (1-\beta) Z_i^N\right]}, \tag{3.1}$$

where Z^S is the evidence for $\mathcal{H}_\mathrm{S}$, and Z_i^G and Z_i^N are, respectively, the evidences for $\mathcal{H}_{\mathrm{G}i}$ and $\mathcal{H}_{\mathrm{N}i}$ in the i^th detector. The arbitrary weights α and β parametrize our prior belief in each model: $\alpha = P(\mathcal{H}_\mathrm{S})/P(\mathcal{H}_\mathrm{I})$ and $\beta = P(\mathcal{H}_{\mathrm{G}i} \mid \mathcal{H}_\mathrm{I}) = 1 - P(\mathcal{H}_{\mathrm{N}i} \mid \mathcal{H}_\mathrm{I})$ for all i (see, e.g., Eq. (59) in [18]). These priors will be chosen to minimize overlap between the signal and noise trigger populations; their importance is studied in detail in Appendix 3.8.

Evidences (marginalized likelihoods) are the conditional probability (P) of observing some data ($\mathbf{d}_i$, for detector i) given some hypothesis ($\mathcal{H}$). For the coherent-signal hypothesis this is

$$\begin{aligned} Z^S &\equiv P(\{\mathbf{d}_i\}_{i=1}^D \mid \mathcal{H}_\mathrm{S}) \\ &= \int p(\vec{\theta} \mid \mathcal{H}_\mathrm{S})\, p(\{\mathbf{d}_i\}_{i=1}^D \mid \vec{\theta}, \mathcal{H}_\mathrm{S})\, \mathrm{d}\vec{\theta}\,. \end{aligned} \tag{3.2}$$

The vector $\vec{\theta}$ represents a point in the space of parameters that describe the CBC signal, such as the component masses and spins; the terms in the integrand are the prior, $p(\vec{\theta} \mid \mathcal{H}_\mathrm{S})$, and the multi-detector likelihood, $p(\{\mathbf{d}_i\}_{i=1}^D \mid \vec{\theta}, \mathcal{H}_\mathrm{S}) = \prod_{i=1}^D p(\mathbf{d}_i \mid \vec{\theta}, \mathcal{H}_\mathrm{S})$. The specific functional form of the single-detector likelihood, $p(\mathbf{d}_i \mid \vec{\theta})$, is derived from the statistical properties of the noise (e.g. a normal distribution for a Gaussian process). The integral is performed numerically using algorithms like *nested sampling* [19, 20]. In our case, the data $\mathbf{d}_i$ are the calibrated Fourier-domain output of each detector, but could generally be any sufficient statistic produced from it.

Because of their inherently unpredictable nature, it is impossible to produce a template that *a priori* captures all features of a glitch. Therefore, we define a surrogate glitch hypothesis by the presence of simultaneous, but incoherent, CBC-like signals in

different detectors. Thus, for the i^{th} detector, the glitch evidence is

$$\begin{aligned} Z_i^G &\equiv P(\mathbf{d}_i \mid \mathcal{H}_{\rm G}) \\ &= \int p(\vec{\theta}_i \mid \mathcal{H}_{\rm G})\, p(\mathbf{d}_i \mid \vec{\theta}_i, \mathcal{H}_{\rm G})\, \mathrm{d}\vec{\theta}_i\,, \end{aligned} \tag{3.3}$$

where now we allow for a different set of signal parameters $\vec{\theta}_i$ at each detector.[1] We will set $p(\vec{\theta}_i \mid \mathcal{H}_{\rm G}) = p(\vec{\theta}_i \mid \mathcal{H}_{\rm S})$ and $p(\mathbf{d}_i \mid \vec{\theta}_i, \mathcal{H}_{\rm G}) = p(\mathbf{d}_i \mid \vec{\theta}_i, \mathcal{H}_{\rm S})$, but this may be relaxed to better capture specific glitch features, if necessary. The surrogate $\mathcal{H}_{\rm G}$ model captures the portion of glitches that lie within the manifold of CBC signals and, in a sense, corresponds to the worst possible glitch—one that looks exactly like coincident CBC signals. Variations of this strategy have been used before in the analysis of compact binary coalescences [11], minimally-modeled transients [24–26], and continuous waves [27–29]. Other searches also make use of likelihood ratios in the detection process, but they do not rely on signal coherence (e.g. [13, 14]).

Finally, because we assume a perfect measurement of the detector noise power-spectral-density (PSD), the Gaussian-noise evidence is just the usual null likelihood. For our Fourier-domain data, this is just

$$Z_i^N \equiv P(\mathbf{d}_i \mid \mathcal{H}_{\rm N}) = \mathcal{N}(\mathbf{d}_i)\,, \tag{3.4}$$

where $\mathcal{N}(\mathbf{d}_i)$ is a multidimensional normal distribution with zero mean and variance derived from the noise PSD [20]. In principle, this could be easily generalized to marginalize over poorly-known PSD parameters if needed.

3.4 Analysis

During O1, the two aLIGO detectors operated from September 12, 2015 to January 19, 2016. Ideally, we would like to compute the BCR for all triggers produced during this period to show that it can efficiently discriminate between glitches and CBC signals. However, computational limitations prevent this [2]. Instead, we pick a subset of 983 multi-detector background binary-black-hole triggers identified by PyCBC, one of the staple search pipelines [15–17, 30]. We pick the background triggers by sampling from the full trigger-set uniformly in the log of the inverse-FAR

[1] Note that $\mathcal{H}_{\rm S}$ and $\mathcal{H}_{\rm I}$ are disjoint even if we do not explicitly exclude points from the parameter space satisfying $\vec{\theta}_i = \vec{\theta}_j$ for all $i \neq j$, because this condition defines a subspace that offers infinitesimal support to the prior in $\mathcal{H}_{\rm I}$ (see [18, 21], or more general discussions in Ch. 4 in [22] or Ch. 28 in [23]).

[2] There are $O(10^7)$ background triggers in O1. The run time on a single background trigger using the LALInference implementation of nested-sampling is usually between 1 to 5 hours.

(IFAR $\equiv 1/\mathrm{FAR}$) for IFARs in $[\,5\times10^{-5}, 10^{6}\,]$ yr, which is the total range reported by the pipeline. This sampling allows us to analyze common (low IFAR) and rare (high IFAR) background events.

To compute the evidences making up the BCR, Eq. (3.1), we run the nested-sampling algorithm implemented in the LALInference library on 4s-long data segments containing each trigger [20, 31]. Given the large number of triggers involved, this would not be feasible without the reduction in the computational cost of Bayesian inference provided by reduced order quadrature (ROQ) methods (see, e.g., [32]). Using this technique makes no measurable difference for the values of the computed evidences.[3]

Templates are produced using IMRPhenomP, a standard waveform family [32–35]. We restrict the priors on the masses such that we only consider signals that are less than 4s in duration, resulting in a chirp-mass range of $12.3M_\odot \leq \mathcal{M} \leq 44.7M_\odot$. We further restrict the mass ratio to lie within $1 \leq q \leq 8$. The dimensionless spin magnitudes are taken to be within $[0, 0.89]$, and we consider all spin angles. The prior on luminosity distance assigns probability uniformly in volume, with an upper cutoff of 5 Gpc. These priors, as well as the priors for all other parameters, follow the default for standard LALInference analyses with ROQ [20, 32]. The PSD used for matched filtering is calculated using the BayesWave algorithm [25, 36].

The search that originally produced our set of triggers considered a wider range of masses and spins than we do in the BCR computation for the purpose of this demonstration. To accommodate this, we prescreened the background to only allow triggers with masses within our priors. It would be straightforward in principle to broaden our constraints to encompass all triggers produced by the pipelines. However, we refrain from doing so to keep our computational costs manageable. Our preliminary analyses of slightly longer triggers (8s, 16s and 32s) yield results qualitatively similar to those presented below.

We compare the BCRs from our background selection to several foreground triggers. The foreground includes eight *hardware injections*, which were performed by physically actuating the test masses of the detectors to simulate signals similar to GW150914 [37]. We also analyze a set of 648 *software injections*: simulated signals inserted in O1 data, with arbitrary sky location and orientation, and with masses

[3] For example, see Table IV in Appendix B of [3], where Bayes factors computed with and without ROQ can be compared (the values in that example are close, but not identical due to differences in waveform approximants).

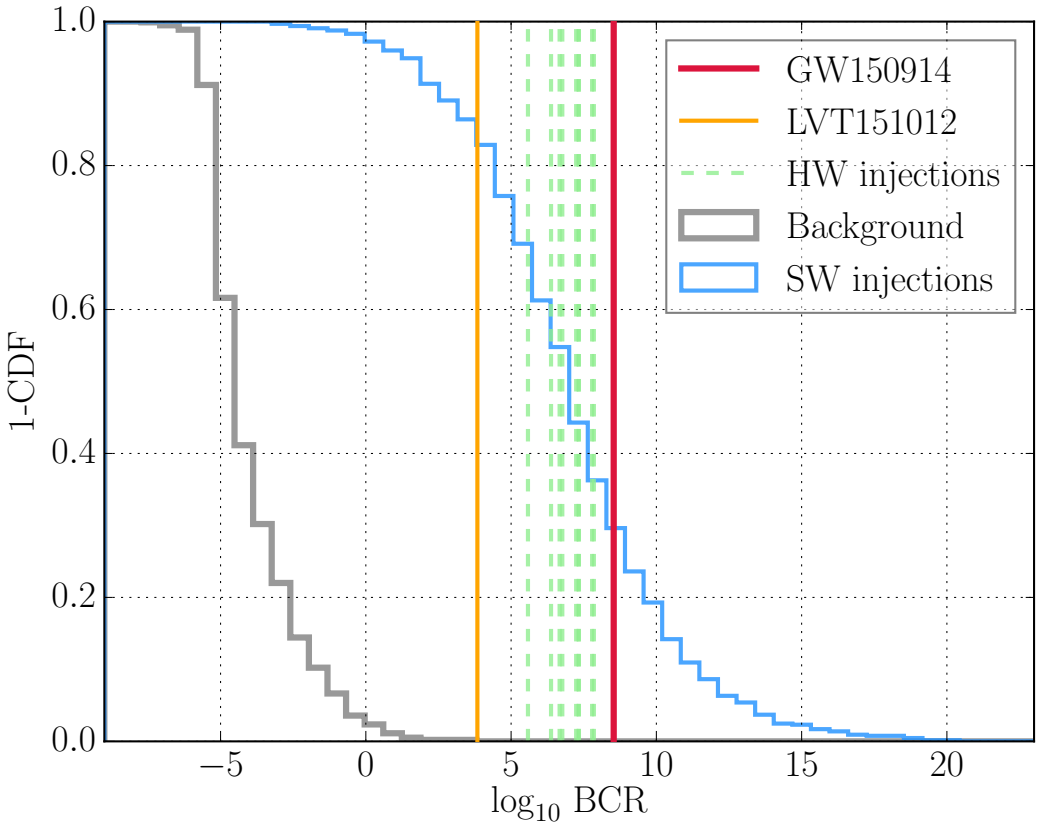

Figure 3.1: *BCR distributions.* Histograms represent the survival function (1-CDF) from our selection of 983 aLIGO O1 background triggers (gray), and 648 simulated signals (blue). Vertical lines mark the BCRs of eight hardware injections (dashed green), LVT151012 (leftmost, orange line), and GW150914 (thick red line). Background triggers were selected to be uniformly distributed in log-IFAR, and 98% yield $\log \mathrm{BCR} < 0$.

and spins that span our priors (in particular, the luminosity distance distribution is uniform in volume with cutoff at 2.5 Gpc). On top of these artificial triggers, we also compute the BCR for GW150914 [1] and LVT151012 [3]. The freedom provided by the α and β parameters in Eq. (3.1) may be used to minimize the overlap between the simulated-signal and background distributions; the results below correspond to values of $\alpha = 10^{-6}$ and $\beta = 10^{-4}$, but may be adjusted in future analyses (see Appendix 3.8).

3.5 Results

Fig. 3.1 shows the BCR distributions obtained for background triggers and software injections. The figure also displays the values obtained for GW150914, LVT151012

and hardware injections, all of which show much stronger evidence for being coherent CBC signals, rather than incoherent glitches (high BCR). We find a clear separation between injections and background events—suggesting that the BCR is good at distinguishing CBC signals from glitches. If we consider the intrinsic probabilistic meaning of the BCR, a value of $\log \mathrm{BCR} < 0$ indicates a preference for the instrumental-artifact hypothesis ($\mathcal{H}_\mathrm{I}$) over the coherent-signal one ($\mathcal{H}_\mathrm{S}$). As expected, the vast majority (98%) of background triggers fall bellow this mark, while the opposite is true for injections. GW150914 and LVT151012 yield $\log_{10} \mathrm{BCR}$ values of 8.5 and 3.8 respectively.

Fig. 3.2 shows the same populations from Fig. 3.1, plotted also as a function of the network signal-to-noise ratio (SNR) recovered by our coherent Bayesian analysis. Fig. 3.2 reveals that the BCR values of the signal population are correlated with SNR, which reflects the fact that we are better able to evaluate the coherence of signals that stand clearly above the noise floor. As a result, the separation between our signal and glitch populations improves with SNR. Because this population of background triggers was purposely selected to be uniform in log-IFAR, the gray contours in Fig. 3.2 should not be taken to be representative of the actual glitch distribution: this would include *vastly* more low-SNR triggers. In any case, BCR is largely independent of SNR for background triggers.

There are three software injections with SNR > 12, but BCR < 1. This is due to two characteristics that make the noise model preferable: (*i*) the ratio of SNRs in two detectors is greater than three, and (*ii*) the signal in at least one detector is too weak to be confidently discernible from Gaussian noise (SNR ~ 5.5). These rare circumstances are caused by source locations and orientations unfavorable to the detector network, and, as such, should be mitigated by adding more instruments.

Irrespective of its Bayesian interpretation, we may treat the BCR as a traditional detection statistic to obtain a frequentist estimate of the significance of any given foreground event based on the measured background (e.g. a p-value, or better, a likelihood ratio). Again, our background triggers were selected to represent common and rare events in equal numbers, so the distribution in Fig. 3.1 need not be the same as that of the entire background, and should not be used for this purpose. However, as shown in Fig. 3.3, we find that there is no evidence for strong correlation between BCR and the IFAR assigned by the detection pipeline. This suggests that the background BCR distribution shown in Fig. 3.1 is likely representative of the whole. Furthermore, Fig. 3.3 implies that the BCR may be used to more easily reject

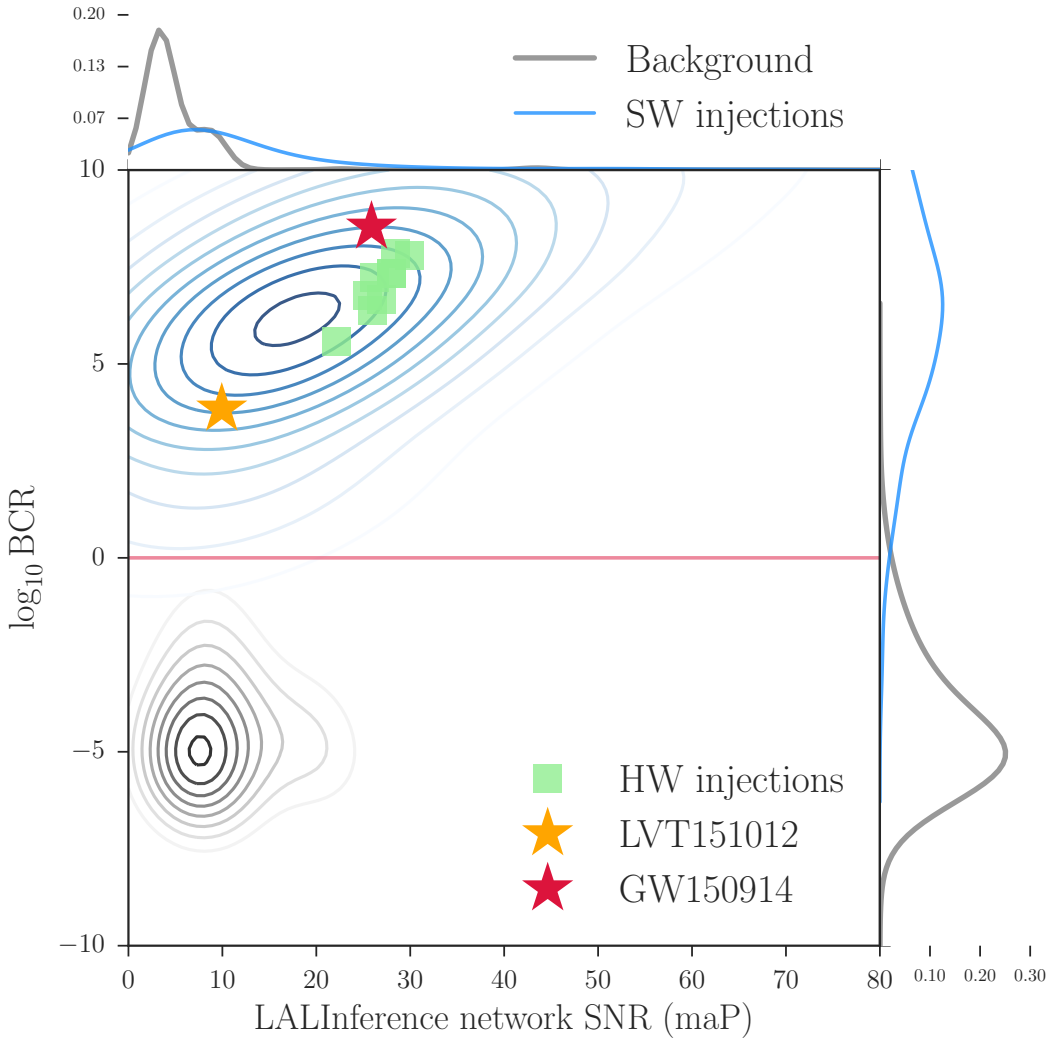

Figure 3.2: *BCR vs SNR distributions.* Contours represent the normalized probability density of selected background triggers (gray) and simulated signals (blue) in log-BCR vs SNR space. The plot also shows eight hardware injections (green squares), LVT151012 (orange star), and GW150914 (red star). The curves shown on the right (top) result from a Gaussian kernel-density estimation of the one-dimensional distribution of log-BCRs (SNRs), obtained after integration over the x-axis (y-axis). Background triggers were selected to be uniformly distributed in log-IFAR, and 98% yield $\log \mathrm{BCR} < 0$ (threshold marked by a horizontal red line for convenience). The SNR on the x-axis is the coherent matched-filter signal-to-noise ratio of the template recovered with maximum *a posteriori* probability (maP) by our inference pipeline (LALInference).

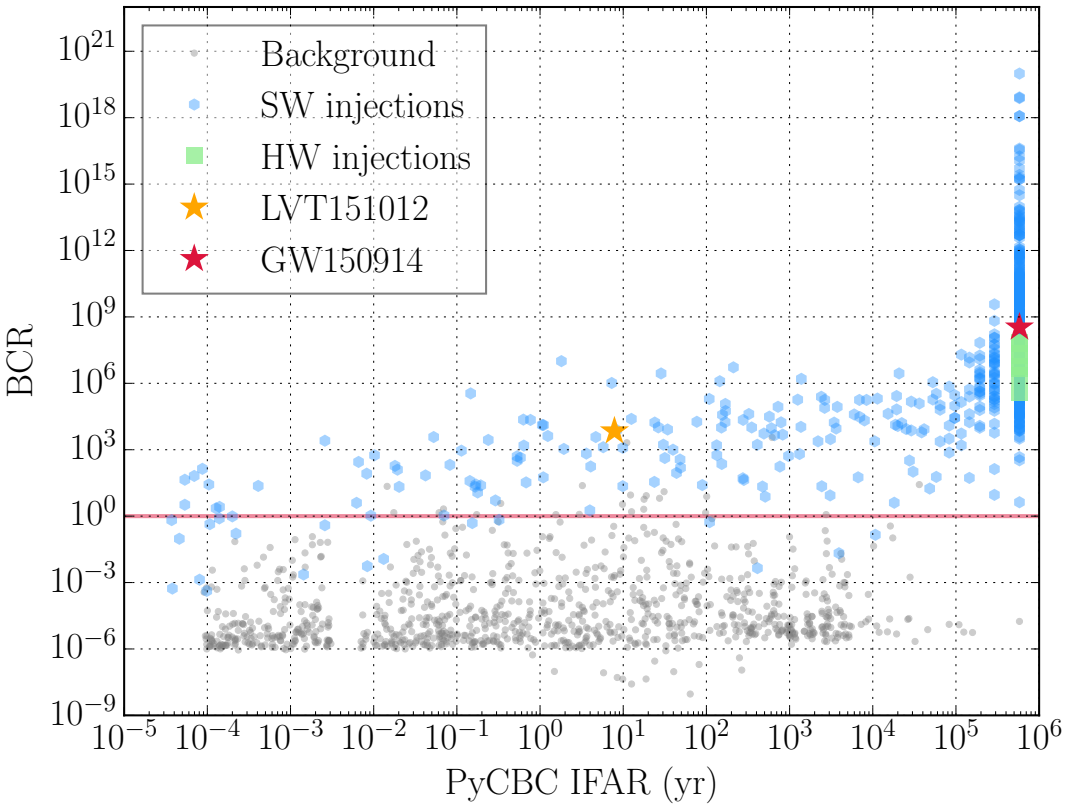

Figure 3.3: *BCR vs IFAR.* BCR for the same data shown in Fig. 3.1, plotted vs the inverse false-alarm-rate (IFAR) assigned to each event by PyCBC, one of the staple aLIGO search pipelines. There are six background triggers with BCR $\ll 10^{-9}$ which fall outside the range of this plot; no foreground triggers are excluded from this plot. High-significance events pile up on the right because their IFAR is a lower limit determined by the most significant trigger in the background. This plot suggests the BCR may be used to more easily reject incoherent glitches.

incoherent glitches, irrespective of IFAR, and thus increase our detection confidence for marginal events like LVT151012.

3.6 Future implementation

Given its ability to separate signals from glitches, the BCR may supplement existing search strategies and help increase their sensitivity, even with existing computational resources. The most straightforward way to achieve this would be to run existing CBC pipelines as usual, with an extra threshold on BCR (e.g. discarding any triggers with, say, BCR < 1). Our results suggest that this would be an efficient way of discarding the vast majority of instrumental artifacts, thereby increasing detection

confidence of real signals [38].

Computational costs would currently preclude obtaining BCRs for *all* triggers (foreground and background) produced during a regular observation run, so this extra step would have to be reserved for the most significant ones, as determined by the main pipeline. However, processing all triggers *would* have the added advantage of potentially enabling the detection of weak GW events that would otherwise be missed (e.g. low-IFAR, but high-BCR, injections in Fig. 3.3). In the future, this would also enable us to move beyond a simple BCR veto, and instead use large numbers of simulated signals and background events to define empirical probability distributions over a space of multiple figures of merit (e.g. BCR and SNR, as in Fig. 3.2). This could be used to obtain likelihood ratios to categorize a trigger as signal or noise—which can be shown to be an optimal strategy for classification problems such as this, and have been used successfully by some existing searches [13, 14, 24]. Future improvements in ROQ methods, like their implementation on graphical processing units, will be vital in making this possible.

The values of the α and β weights in Eq. (3.1) have a strong effect on the shape of the distributions of Fig. 3.2, as discussed in Appendix 3.8. While here we have set them to values that yield a good separation between the signal and background populations, future studies may systematically optimize these parameters using a more comprehensive set of software injections and a large, representative set of background triggers. This may be achieved via any standard optimization scheme that attempts to minimize the overlap between the two populations. The values would, of course, be fixed before analyzing any foreground data.

3.7 Conclusion

We have demonstrated that Bayesian models based on the coherence of GW triggers across detectors may successfully distinguish between real CBC signals and transient instrumental noise (Figs. 3.1 and 3.2). We introduced a specific figure of merit, the BCR, which responds to incoherent glitches in a way that is complementary to that of standard CBC pipelines (Fig. 3.3). Finally, we suggested a few avenues for incorporating this (or similar) measure of coherence into existing GW search strategies, the simplest of which would take the form of a new veto for detection candidates. This could be implemented today to increase the number of gravitational waves confidently detected by LIGO and Virgo, without needing to further improve detector hardware.

Versions of the ranking statistic used by PyCBC in recent analyses have incorporated some measure of coherence [15], and it remains to be seen whether this introduces some correlation between BCR and IFAR in Fig. 3.3. Furthermore, while this study focused on detection candidates produced by the two aLIGO detectors during O1, we are currently investigating how the power of the BCR is affected by the addition of new detectors, like Virgo. Finally, although here we focused on short-duration (4s) triggers from high-mass binary-black-hole mergers, our preliminary results on slightly longer (8s, 16s and 32s) show qualitatively similar behavior.

Acknowledgments

We thank Alan Weinstein, Alex Nitz, Carl-Johan Haster, Stefan Hild, Reed Essick, Ryan Lynch, Colm Talbot, Eric Thrane, John Veitch and Thomas Dent for helpful comments. Rory Smith is supported by the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), through project number CE170100004. The authors thank the LIGO Scientific Collaboration for access to the data and gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO (PHY-0757058), as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, and the Max-Planck-Society (MPS) for support of the construction of Advanced LIGO. Additional support for Advanced LIGO was provided by the Australian Research Council. This manuscript has LIGO Document ID LIGO–P1700414.

3.8 Appendix: Effect of BCR weights

The weights (α, β) that go into the calculation of the BCR in Eq. (3.1) have a critical impact on the degree of separation between the signal and glitch populations. Here we elaborate on this point and show how we improve upon previous work by explicitly taking advantage of the extra freedom afforded by these parameters.

From a Bayesian perspective, α and β encode our prior beliefs on the relative probabilities of each of the submodels that are compared in the computation of the BCR: α determines by what factor the coherent-signal hypothesis ($\mathcal{H}_{\rm S}$) should be favored over the instrumental-feature hypothesis ($\mathcal{H}_{\rm I}$),

$$\alpha \equiv \frac{P(\mathcal{H}_{\rm S})}{P(\mathcal{H}_{\rm I})}, \tag{3.5}$$

while β gives the probability of the glitch hypothesis ($\mathcal{H}_{\rm G}$) conditional on the

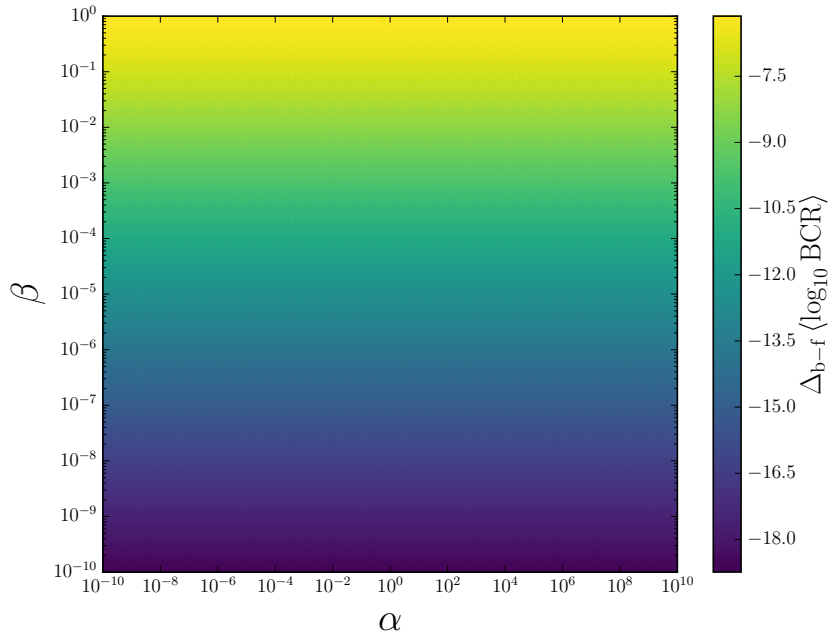

Figure 3.4: *Effect of weight on population separation.* Color represents the difference in mean log BCR between background and foreground, $\Delta_{\mathrm{b-f}} \langle \log \mathrm{BCR} \rangle$ as defined in Eq. (3.8). This is shown as a function of the BCR prior weights, α (x-axis) and β (y-axis), of Eq. (3.1). All values are negative because the foreground always has larger mean, so darker colors correspond to greater distance between the population means.

assumption that there is an instrumental-feature to begin with,

$$\beta \equiv P(\mathcal{H}_{\mathrm{G}i} \mid \mathcal{H}_{\mathrm{I}}) = 1 - P(\mathcal{H}_{\mathrm{N}i} \mid \mathcal{H}_{\mathrm{I}}), \tag{3.6}$$

for any dectector i, as discussed in Sec. 3.3. The last equality in Eq. (3.6) uses the fact that we *define* the instrumental-feature hypothesis as the logical union of the glitch and Gaussian noise ($\mathcal{H}_{\mathrm{N}}$) subhypotheses, i.e. $\mathcal{H}_{\mathrm{I}} \equiv \mathcal{H}_{\mathrm{G}} \vee \mathcal{H}_{\mathrm{N}}$, and that the latter are logically disjoint, i.e. $\mathcal{H}_{\mathrm{G}} \wedge \mathcal{H}_{\mathrm{N}} = \mathrm{False}$, so $P(\mathcal{H}_{\mathrm{N}} \mid \mathcal{H}_{\mathrm{G}}) = P(\mathcal{H}_{\mathrm{G}} \mid \mathcal{H}_{\mathrm{N}}) = 0$.

It follows from the probabilistic interpretation of these parameters that their allowed ranges are $0 < \alpha < \infty$ and $0 \leq \beta \leq 1$. All results presented in the main text were produced using the values

$$\text{Main text:} \quad \left(\alpha = 10^{-6}, \beta = 10^{-4}\right). \tag{3.7}$$

This specific choice was made to yield a good separation between the background and foreground populations, as reflected by Figs. 3.1 and 3.2. These values also result in an overall normalization such that BCR = 1 gives the point at which both

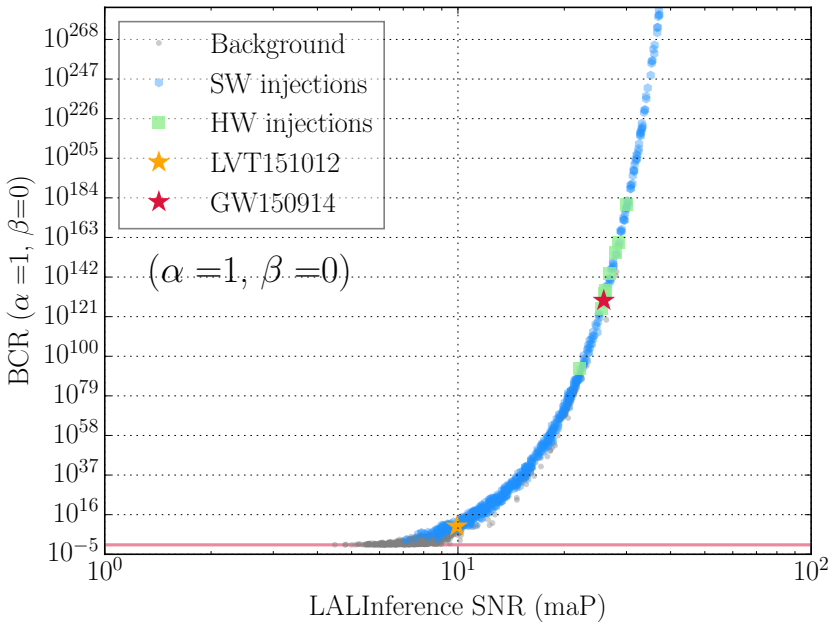

Figure 3.5: *BCR* ($\alpha = 1, \beta = 0$) *vs SNR.* BCR vs SNR for the same data shown in Figs. 3.1–3.3, but analyzed with ($\alpha = 1, \beta = 0$). For this choice of weights, the BCR reduces to the Bayesian odds between signal and Gaussian noise, Eq. (3.9), and scales with SNR according to Eq. (3.10), for both background (gray circles) and foreground (blue hexagons). The SNR on the x-axis is the coherent matched-filter signal-to-noise ratio of the template recovered with maximum *a posteriori* probability (maP) by our inference pipeline (LALInference).

hypotheses are equally likely given *our* trigger set (i.e. the horizontal red line in Fig. 3.2 roughly agrees with the intersection of the blue and gray curves on the right panel).

To see how α and β impact the separation between the background and foreground populations, consider as a proxy the distance between the mean BCRs for the two populations. In particular, define the quantity

$$\Delta_{\mathrm{b-f}} \langle \log \mathrm{BCR} \rangle \equiv \left\langle \log \mathrm{BCR}^{(\mathrm{b})} \right\rangle - \left\langle \log \mathrm{BCR}^{(\mathrm{f})} \right\rangle, \tag{3.8}$$

where the angle brackets on the right denote averaging over triggers, and the superscripts "(b)" and "(f)" refer to "background" and "foreground" respectively. This number then gives a measure of the vertical distance between the centers of the distributions in Fig. 3.2. The effect of α and β on this quantity is shown in Fig. 3.4, where darker colors correspond to greater absolute mean distance. As expected from Eq. (3.1), the separation is a strong function of β, while it is largely independent of α. It can also be seen from Eq. (3.1) that α should merely impact the overall

normalization of the BCR, shifting all values up or down.

By tuning β we may thus control the degree of bias introduced in the computation of the BCR. This can be used to correct for shortcomings in the definitions of the noise submodels themselves, so as to best distinguish foreground and background. The reason this is necessary in the first place is that not all glitches will conform strictly to the "worst-glitch" hypothesis as we have defined it via Eq. (3.3). For instance, the distribution of glitch morphologies and SNRs need not conform to the parameter priors assumed in the computation of Z^G; instead of tuning the parameter priors, one may correct for this effect via β (which is easier to implement).

Looking at Fig. 3.4, one may be tempted to substantially reduce β to maximize the distance between the distribution means. However, the quantity plotted in Fig. 3.4, Eq. (3.8), is insensitive to the fact that the two distributions do not retain their shape when β is varied, and therefore is only useful as a proxy for population overlap when looking at small changes in the weights. In other words, Fig. 3.4 fails to convey the fact that there is a *penalty* in introducing too strong of a bias through β. This is related to the *bias-variance tradeoff*, well known in statistical inference (see, e.g., [39]). Let us explore how this tradeoff is manifested throughout the range of valid values for β.

On one end, setting $\beta = 0$ comes at the price of throwing away all information about the incoherence of the trigger. As can be deduced from Eq. (3.1), in the limit of vanishing β the BCR is nothing but the usual signal vs Gaussian-noise odds (BSN),

$$\mathrm{BCR}(\alpha = 1, \beta = 0) = Z^S/Z^N \equiv \mathrm{BSN}\,, \tag{3.9}$$

and the glitch model is totally ignored. For this choice of β, the BCR will just follow the usual dependence of BSN on SNR (see, e.g., [40]),

$$\log \mathrm{BSN} \propto \mathrm{SNR}^2\,, \tag{3.10}$$

irrespective of whether the trigger is a glitch or a coherent signal, as shown in Fig. 3.5. Although the distance between the means of the two populations in this figure is large (as reflected also by Fig. 3.4, for $\beta \to 0$), this is only because, on average, the background triggers in our set have lower SNR than the foreground.

On the other end, setting $\beta = 1$ is equivalent to ignoring the possibility that the trigger was produced by Gaussian noise. In that case, the BCR reduces to the evidence ratio between the coherent-signal and incoherent-glitch hypotheses, a quantity often called "BCI" by gravitational-wave data analysts (assuming $\alpha = 1$):

$$\mathrm{BCR}(\alpha = 1, \beta = 1) = Z^S/Z^G \equiv \mathrm{BCI}\,. \tag{3.11}$$

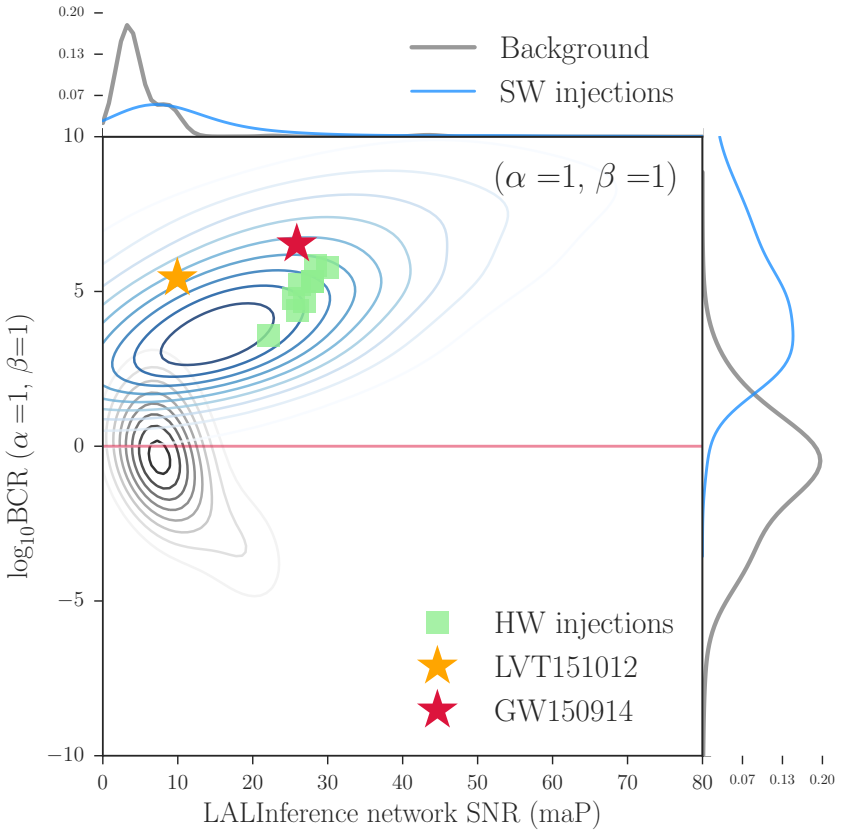

Figure 3.6: *BCR* ($\alpha = 1, \beta = 1$) *vs SNR distributions.* This plot is completely analogous to Fig. 3.2, but with ($\alpha = 1, \beta = 1$) instead of ($\alpha = 10^{-6}, \beta = 10^{-4}$) [cf. Eq. (3.1)]. For this choice of weights, the BCR reduces to the BCI, Eq. (3.11), resulting in greater overlap between the background (gray) and foreground (blue) distributions. For more details about this plot, refer to the caption of Fig. 3.2.

The use of this quantity for glitch-discrimination purposes in CBC searches was proposed in [11]. However, we find that it does not produce a sufficient separation between the background and foreground populations, except for loud triggers. For example, while ($\alpha = 10^{-6}, \beta = 10^{-4}$) yields Fig. 3.2, ($\alpha = 1, \beta = 1$) yields Fig. 3.6. From this plot, it is easy to see that the BCI is good at distinguishing *loud* incoherent glitches from *loud* coherent signals, but is inconclusive for weak triggers.

We can check that changing β indeed affects primarily *weak* glitches by comparing Fig. 3.7 to Fig. 3.3, BCR vs IFAR plots which were produced with $\beta = 1$ and $\beta = 10^{-4}$ respectively. The change in β from Fig. 3.7 to Fig. 3.3 causes low-IFAR (low-SNR) glitches to yield significantly lower BCRs, while high-IFAR (high-SNR) triggers are largely unaffected. Importantly, low-IFAR (low-SNR) signals are also down-ranked after the change, but to a lesser degree on average; hence the separation

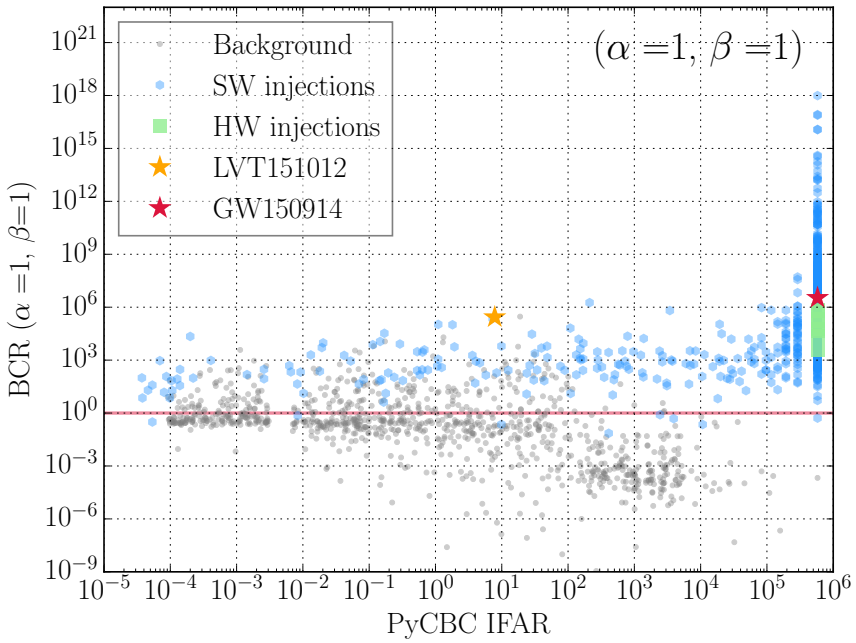

Figure 3.7: *BCR* ($\alpha = 1, \beta = 1$) *vs IFAR*. This plot is completely analogous to Fig. 3.3, but with ($\alpha = 1, \beta = 1$) instead of ($\alpha = 10^{-6}, \beta = 10^{-4}$) [cf. Eq. (3.1)]. For this choice of weights, the BCR reduces to the BCI, Eq. (3.11), resulting in greater overlap between the background (gray) and foreground (blue) distributions. For more details about this plot, refer to the caption of Fig. 3.3.

in BCR improves, as can be seen by comparing the right panels of Fig. 3.6 and Fig. 3.2.

To further quantify the effect of β, we can also look at the fractional change in $\log \mathrm{BCR}$ when going from ($\alpha = 1, \beta = 1$) to ($\alpha = 10^{-6}, \beta = 10^{-4}$),

$$\frac{\Delta\,(\log \mathrm{BCR})}{|\log \mathrm{BCI}|} \equiv \frac{\log \mathrm{BCR}(10^{-6}, 10^{-4}) - \log \mathrm{BCI}}{|\log \mathrm{BCI}|}, \tag{3.12}$$

where vertical bars mark absolute values, and the BCI is defined by Eq. (3.11). This quantity is histogrammed in Fig. 3.8 for the triggers in our set. The fact that the change in β affects weak glitches more significantly than strong ones is reflected in the bimodality of the gray distribution: the left (right) peak corresponds to triggers below (above) an effective threshold of SNR ~ 9. On the other hand, the blue distribution in Fig. 3.8 shows that most (although not all) signals are largely unaffected by the change in β, with a mean increase in BCR but long tails extending mainly to the left. This large variance is due mostly to the weaker signals for which the BCR decreased due to the change in β.

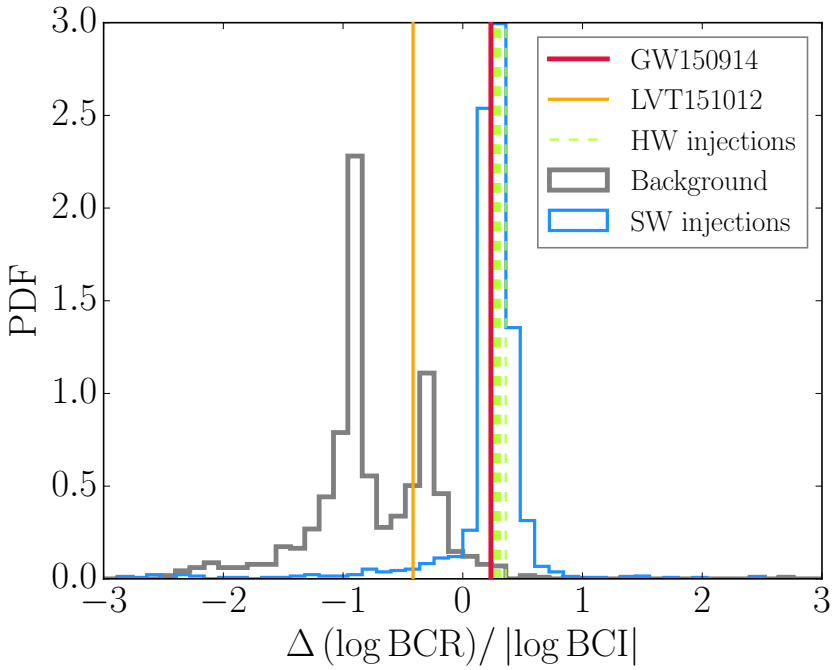

Figure 3.8: *Effect of weights on* log*BCR.* Histogram of the fractional change in log BCR when going from ($\alpha = 1, \beta = 1$) to ($\alpha = 10^{-6}, \beta = 10^{-4}$), Eq. (3.12). This plot summarizes the differences between the BCRs shown in Figs. 3.2 & 3.3 and those in Figs. 3.6 & 3.7.

By tuning the weights, we may attempt to find a sweet spot in which the bias introduced is just enough to separate weak glitches from weak signals, without confounding loud glitches with loud signals. The choice of Eq. (3.7) was found to be close to this ideal, and achieves this by separating the weak glitches in our set from the weak signals to an extent, largely without altering loud triggers (Figs. 3.1–3.3).

Chapter 4

ESTABLISHING THE SIGNIFICANCE OF CONTINUOUS GRAVITATIONAL-WAVE DETECTIONS FROM KNOWN PULSARS

M. Isi, S. Mastrogiovanni, M. Pitkin, and O. J. Piccini, "Establishing the significance of continuous gravitational-wave detections from known pulsars," (2018), M.I. conceived the project and led the writing of the manuscript.

4.1 Introduction

In addition to short-lived gravitational waves (GWs) from compact-binary coalescences like those observed so far [1–3, 41, 42], ground-based detectors like the Advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) [8] and Virgo [9], are also expected to detect persistent, almost-monochromatic signals [29, 43–47]. The primary potential source of such "continuous waves" (CWs) is rapidly spinning neutron stars with an asymmetry in their moment of inertia [48]. This includes galactic pulsars known from electromagnetic observations, which are a main target for searches for continuous signals in LIGO and Virgo data [29, 49]. The detection of gravitational waves from any of these sources would provide a new wealth of astrophysical information, as well as invaluable opportunities to learn about fundamental physics (see, e.g., [50] for a recent review).

There exist a number of efforts to detect gravitational waves from known pulsars [29, 47, 51, 52]. However, an outstanding problem affecting all of these searches is the lack of a well-defined procedure to establish the statistical significance of potential detections without making the assumption that the instrumental noise is Gaussian and (semi-)stationary. Consequently, if evidence for a continuous wave from a known pulsar was found today, we would be unable to establish, with certainty, the probability for this to have arisen from a spurious noise artifact. The need for a systematic and robust way of computing detection significance in the presence of non-Gaussian noise has already become apparent with the appearance of hard-to-diagnose outliers in recent searches in actual aLIGO data [29, 53].

Establishing a robust procedure to assign significance is challenging because the noise artifacts that limit the searches are intrinsically unpredictable and cannot be modeled from first principles. Given this, we may instead attempt to empirically determine the response of the different searches to real detector noise in the *absence* of astrophysical signals. Armed with such knowledge, we would then be able to analyze actual data, or "foreground", and produce empirical likelihood ratios (or weaker measures of detection confidence, like p-values) for the presence of an astrophysical signal vs just instrumental noise, Gaussian or otherwise. This requires several instances of "background"—that is, instrumental data that are known to contain no astrophysical signals, while still retaining all statistical properties representative of real instrumental noise.

Ideally, one would obtain background distributions by physically isolating the instruments from the environment to shield them from actual signals. Because this is impossible in the case of gravitational waves, we must attempt to replicate this shielding digitally after the data have been recorded. Several techniques exist to do this when looking for gravitational-wave transients, the most straightforward of which is probably the use of "time slides": the outputs of different detectors are shifted relative to each other by time offsets longer than the light-travel time between them [16, 17]. This ensures the spuriousness of any signal candidate left in the multi-detector data thus produced, hence allowing us to estimate how likely it is for noise to mimic a signal.

The direct analog of time slides in the context of continuous waves would be "frequency slides": a misalignment of the frequency-domain data of different detectors. However, our ability to effect such frequency shifts is limited by the frequency resolution of the searches (of the order of inverse observation time), and the fact that the properties of actual instrumental noise are heavily dependent on frequency—not only due to a frequency-dependent power spectral density, but also to varying populations of narrow-band noise features. By the same token, time slides themselves would not be feasible in transient analyses if the noise properties of the detectors changed rapidly compared to the sampling time.

In light of this, here we propose a simple method for estimating the background of searches for continuous gravitational waves by analyzing data assuming an incorrect sky location for the targeted source. This "off-sourcing" takes advantage of the expected Doppler modulation of the signal due to the relative motion of detector and source to effectively blind the search to real astrophysical signals. We can use this to

produce a large number of independent noise-only instantiations of data, so as to empirically estimate the background of a search and assign detection significances in the presence of actual detector noise. We demonstrate that this method outperforms another common strategy for estimating the background in realistic situations.

We begin in by providing relevant background about continuous waves and targeted searches in Sec. 4.2. We then introduce the off-sourcing method and explore its applicability in Sec. 4.3. We demonstrate the efficacy of the strategy in Sec. 4.4, where we present several examples in both fabricated and actual detector noise. We conclude in Sec. 4.5.

4.2 Background

In this section, we review the basic morphology of continuous gravitational waves as measured by differential-arm detectors, with an emphasis on the timing corrections on which we will rely for off-sourcing (Sec. 4.2.1). We also make a special point of discussing the relation between the frequency resolution at which a signal is sampled and the ability to localize the source in the sky (Sec. 4.2.1). We next describe the key properties of noise in existing ground-based instruments as they pertain to searches for persistent signals (Sec. 4.2.2). Finally, we provide an overview of the three staple search methods for these signals in LIGO and Virgo data (Sec. 4.2.3): the Bayesian time-domain method, and the frequentist 5-vector and $\mathcal{F}$–statistic methods.

4.2.1 Continuous waves

Morphology

Continuous waves are nearly monochromatic gravitational perturbations with constant intrinsic amplitude that are expected to be sourced by some rapidly spinning bodies, like neutron stars. Within the context of standard physics, there are several ways in which a neutron star could emit CWs, but the most favored is the presence of a nonaxisymmetry in the star's moment of inertia [54]. For this type of *triaxial*, non-precessing source, such a GW will induce a strain in a differential-arm (quadrupolar) detector, like LIGO or Virgo, which can be written as:

$$\begin{aligned} h(t) = h_0 \frac{1}{2}(1 + \cos^2 \iota) F_+(t;\psi) \cos\phi(t) \\ + h_0 \cos\iota F_\times(t;\psi) \sin\phi(t), \end{aligned} \tag{4.1}$$

where the $F_+(t;\psi)$ and $F_\times(t;\psi)$ factors respectively give the instrument's response to the plus (+) and cross (×) GW polarizations, ι is the inclination angle between the spin axis of the source and the observer's line-of-sight, $\phi(t)$ is the phase of the signal, and h_0 is an overall amplitude related to the properties of the source by:

$$h_0 = \frac{16\pi^2 G}{c^4}\frac{\epsilon I_{zz} f_{\rm rot}^2}{r}, \tag{4.2}$$

where r is the source distance, $f_{\rm rot}$ its rotation frequency around the principal axis z, $\mathbf{I}$ the moment-of-inertia tensor and $\epsilon \equiv (I_{xx} - I_{yy})/I_{zz}$ the equatorial ellipticity [48].

The antenna patterns, $F_+(t;\psi)$ and $F_\times(t;\psi)$, encode the amplitude modulation of the signal due to the local geometric effect of a GW acting on a given detector. Thus, they implicitly depend on the relative location and orientation of source and detector by means of the source's right-ascension α, declination δ, and polarization angle ψ. The latter gives the orientation of the frame in which the polarizations are defined, and we set it to be the angle between the line of nodes and the projection of the celestial North onto the plane of the sky. While α and δ are always well known, ψ generally is not, which is why we show this argument explicitly. Importantly, the antenna patterns acquire their time dependence from the rotation of Earth on its axis, and consequently have a characteristic period of a sidereal day ($\sim 10^{-5}$ Hz).

For a simple triaxial source, the GW frequency f is twice the rotational value $f_{\rm rot}$, so we can write:

$$\phi(t) = 2\phi_{\rm rot}(t) + \phi_0, \tag{4.3}$$

where $\phi_{\rm rot}$ is the rotational phase as measured via electromagnetic (EM) observations and ϕ_0 is a fiducial phase offset. The rotational frequency itself is almost constant, with a small spin-down due to energy loss into the environment (via GWs and other mechanisms), which means that the phase evolution can be well described by a simple Taylor expansion on τ, the time measured by a clock inertial with respect to the source:

$$\phi(t) = 2\pi \sum_{j=0}^{N} \frac{\partial_t^{(j)} f_0}{(j+1)!} \left[\tau(t) - T_0\right]^{(j+1)}. \tag{4.4}$$

Here $\partial_t^{(j)} f_0$ is the $j^{\rm th}$ time derivative of the GW frequency measured at the fiducial time T_0, and N is the order of the series expansion (1 or 2 suffices for most sources). Timing solutions are generally obtained through the pulsar timing package TEMPO2 [55, 56]. These solutions present exquisite precision (frequency uncertainty of $10^{-12} \lesssim \delta f_{\rm rot} \lesssim 10^{-8}$ Hz for most pulsars) and are the cornerstone of targeted searches for continuous waves from known pulsars.

The inertial time, τ in Eq. (4.3), is usually taken to be the time measured by a clock at the Solar System barycenter (SSB), which is itself assumed to be inertial with respect to the pulsar. In that case, τ can be written as a function of detector time, t, by taking into account some well-known, time-dependent offsets:

$$\tau(t) = t + \Delta_{\rm E}(t) + \Delta_{\rm S}(t) + \Delta_{\rm binary}(t) + \Delta_{\rm R}(t)\,. \tag{4.5}$$

Here $\Delta_{\rm E}$ is the Solar-System Einstein delay; $\Delta_{\rm S}$ is the Solar-System Shapiro delay; $\Delta_{\rm binary}$ is the delay originating from the motion of the pulsar in its binary (a term that vanishes for isolated sources) [57]; and $\Delta_{\rm R}$ is the kinematic delay due to the relative motion of the detector with respect to the source.

The timing correction of Eq. (4.5) is heavily dependent on the sky-location of the targeted pulsar and will be the key to the off-sourcing method presented in Sec. 4.3. The dependence on sky location is dominated by the last term in Eq. (4.5), $\Delta_{\rm R}$. This is sometimes known as the "Rømer delay" and encodes the Doppler modulation of the signal:

$$\Delta_{\rm R}(t) = -\frac{\hat{\Omega}\cdot\vec{r}(t)}{c}\,, \tag{4.6}$$

where $\vec{r}(t)$ is a vector joining the SSB and the detector at any given time, $\hat{\Omega}$ is a unit vector pointing from the SSB in the direction of the source,[1] and c is the GW speed. For practical purposes, $\vec{r}$ is usually computed by first splitting it into three components:

$$\vec{r} = \vec{r}_{\odot} + \vec{r}_{\oplus} + \vec{R}\,, \tag{4.7}$$

with $\vec{r}_{\odot}$ joining the SSB with the center of Sun, $\vec{r}_{\oplus}$ joining Sun and Earth, and $\vec{R}$ going from the center of Earth to the detector on the surface. One can then use Solar System ephemerides, together with knowledge of the location of the detector on Earth and the source in the sky, to compute the Rømer correction at any given time.

The timing correction of Eq. (4.5) can be understood as inducing extrinsic frequency shifts to the signal, as seen by the detector. This is dominated by the Rømer term, $\Delta_R(t)$, which results in a modulation at the frequency of Earth's orbital rotation, $\omega_{\rm orb} \approx 2\times10^{-7}$ Hz, as well as subdominant daily effects due to its spin, $\omega_{\rm sid} = 2\pi/(\text{sidereal day}) \approx 7\times10^{-5}$ Hz. In the frequency domain, the effect of this correction is to spread the signal power across a narrow band centered on its intrinsic GW frequency, with a characteristic width of $\Delta f \approx f\cdot10^{-4}$ Hz. This frequency modulation will be the key of our approach.

[1] The source-location vector, $\hat{\Omega}$, can be treated as constant over the timescale of our observations.

Frequency and sky resolution

The sky resolution is the minimum angular separation in the sky at which two, otherwise equal, sources could be distinguished. This is a function of the frequency resolution at which the signal is sampled, namely:

$$\delta f = 1/T, \tag{4.8}$$

for an observation time T. This frequency bin is related to the minimum-resolvable sky-bin by the sky-location–dependent frequency modulation of Eq. (4.5). Thus, the angular resolution will be roughly given by the separation in the sky corresponding to a Rømer frequency shift of $\Delta f = \delta f$. In other words, we may define a bin around any point in the sky by the maximum angular distance one can move away from that point before the frequency shift caused by the modulation of Eq. (4.6) reaches a magnitude of $1/T$. Thus, the characteristic size of a bin defined this way will necessarily depend on the integration time.

Proceeding as above, we may cover the sky with a series of such sky bins to obtain a "sky grid" representing the resolvability of points in the sky as a function of angular location. Because the timing correction of Eq. (4.5) is dominated by Δ_R, which is itself mostly due to Earth's *orbital* motion, such a sky grid will be most naturally defined in *ecliptic* coordinates to yield bin sizes given approximately by [58]:

$$\delta\beta = \frac{1}{N_d \sin\beta} \quad , \quad \delta\lambda = \frac{1}{N_d \cos\beta} \tag{4.9}$$

where β and λ are respectively the ecliptic longitude and latitude, and the scale factor is

$$N_d = \frac{f\,\omega_{\mathrm{orb}} R_{\mathrm{orb}} T}{c} \tag{4.10}$$

for R_{orb} Earth's orbital radius. As demonstrated in Fig. 4.1, the sky-grid can be easily computed using ecliptic coordinates (left panel) and then rotated to equatorial coordinates (right panel). Note that this is a conservative sky-grid that implicitly assumes the power of the signal may be split over at most two frequency bins as a result of the timing correction—in practice, the characteristic size of the sky bins may be reduced, but the scaling with f and T will always be as in Eq. (4.9).

4.2.2 Detector noise

The output of ground-based gravitational-wave detectors is vastly dominated by instrumental noise [59, 60]. For this reason, the weak continuous signals discussed

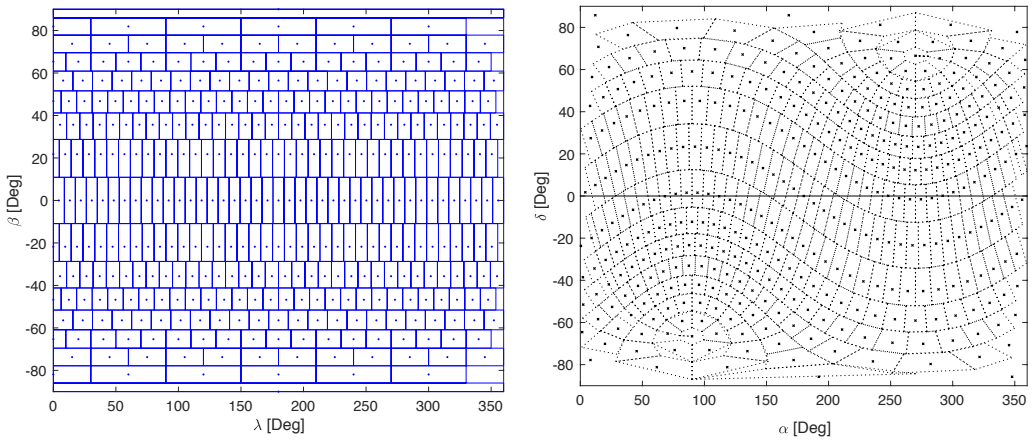

Figure 4.1: *Left:* Sky bins based on ecliptic coordinates for a search at a frequency of 60 Hz and an integration time of 1024 s. *Right:* The same sky grid for the same sky configuration in equatorial coordinates.

in Sec. 4.2.1 are expected to become visible only after long periods of coherently-integrated observation. Thus, understanding the statistical properties of the noise is critical to successfully detecting these signals, irrespective of the specific details underlying the search method of choice (cf. Sec. 4.2.3).

For the most part, the noise in a given detector is well described as a Gaussian random process with a frequency-dependent (colored) power spectral density [59, 60]. Gaussian noise has numerous convenient statistical properties that would drastically simplify many of LIGO and Virgo's analyses. However, this idealization is far from perfect: the data are plagued with uncountable non-Gaussian features with a range of spectral properties and durations. Among these, the most-often discussed are probably the noise transients ("glitches") that haunt searches for compact-binary coalescences [61]. Yet, searches for continuous waves are most affected not by these short-lived glitches, but rather by persistent narrow-band features ("lines") [62]. Many of these spectral lines only become apparent after long-periods of coherent observation, making their identification and eradication especially challenging. Furthermore, their distribution over the sensitive frequency band of the detectors is highly irregular and changes with time with the state of the instruments.

Fully-coherent searches for continuous waves tend to have very high frequency resolution (of order $\delta f \sim 10^{-7}$–10^{-8} Hz), scaling directly with the integration time ($\delta f \sim 1/T$). This fine resolution means that such analyses can fall victim to very narrow (and weak) noise lines. Furthermore, as mentioned at the end of Sec. 4.2.1, a pulsar signal will be spread over a band of width $\Delta f \approx f \cdot 10^{-4}$ Hz around its central

GW frequency f. This means that attempts to find such a signal will be affected by noise over a range of frequencies, wide with respect to the typical resolution of the search. A persistent departure from Gaussianity in that frequency range (e.g. a wandering instrumental line that happens to intermittently cross the targeted band) will confound most searches, potentially yielding false positives ("outliers"). Naturally, the number of outliers due to unmodeled noise found by the pipelines will increase with the searched CW parameter space, as well as with observation time (which increases the frequency resolution).

As in the case with the glitches affecting searches for compact binaries, lines and other non-Gaussian features would not be an issue for continuous-wave searches if there existed a robust way to model them and directly incorporate that knowledge into the statistical analyses (cf. Sec. 4.2.3 below). However, the noise artifacts in the set that interests us are, by definition, impossible to fully model from first principles: any particular noise source that *is* well-understood can usually be physically or digitally removed, so that they are no longer of concern [62–65]; the remaining artifacts are, therefore, those that are intrinsically unpredictable or so-far not understood. Consequently, we are left to try to find ways to *empirically* estimate the true statistical background (i.e. the probability distribution of false-positives) of a search in order to assign significances to potential detection candidates.

4.2.3 Searches

Searches targeted at known pulsars make use of the simple form of the expected signal, Eq. (4.1), to match-filter the data and determine the likelihood that a signal is present. There exist three standard approaches of this kind: the time-domain Bayesian method [51, 57, 66, 67], the time-domain $\mathcal{F}$-statistic method [52, 68], and the frequency-domain 5n-vector method [69–71]. Due to the technical details underlying each implementation, only the Bayesian time-domain method has been broadly applied to a large number of targets [29]. Although the off-sourcing method is applicable to all three of these techniques, in the following sections we will only use the Bayesian and 5-vector searches for concrete examples. Regardless, here we provide a brief overview of the basics of all three approaches for completeness.

Bayesian approach

Bayesian statistics provide a complete and straightforward framework for computing the probability that a given set of data contain a signal vs Gaussian noise, and for inferring the parameters that best describe the signal if present. LIGO's implementation [51] takes advantage of the fact that the phase evolution $\phi(t)$ is known from electromagnetic observations to remove the high-frequency components of the signal early in the process—this dramatically simplifies the Bayesian inference step itself [67].

First, the data are digitally heterodyned [57, 66] so that the signal they contain becomes:

$$h'(t) \equiv h(t)e^{-i\phi(t)} = \Lambda(t) + \Lambda^*(t)e^{-i2\phi(t)}\,, \tag{4.11}$$

with $*$ indicating complex conjugation, and

$$\Lambda(t) \equiv \frac{1}{4}F_+(t)h_0(1+\cos^2\iota) - \frac{i}{2}F_\times(t)h_0\cos\iota\,. \tag{4.12}$$

A series of low-pass filters are then applied to remove the second term in Eq. (4.11), which enables the down-sampling of the data by averaging over minute-long time bins. As a result, $\Lambda(t)$ is the only contribution from the original signal left in our binned data, $\mathbf{B}$, which will now look like

$$B_{\text{expected}}(t_k) = \Lambda(t_k) + n(t_k), \tag{4.13}$$

where $n(t_k)$ is the heterodyned, filtered and downsampled noise in bin k, which carries no information about the GW signal.

Eq. (4.13) implies that $B(t_k) - \Lambda(t_k)$ should have the statistical properties of noise, and that Eq. (4.12) should be taken to be the template in our search. This knowledge can be used to compute the marginalized-likelihood ratio (Bayes factor) that the data contain a signal buried in noise ($\mathcal{H}_\text{S}$), vs just Gaussian noise ($\mathcal{H}_\text{N}$):

$$\mathcal{B}^\text{S}_\text{N} = \frac{P(\mathbf{B}\mid\mathcal{H}_\text{S})}{P(\mathbf{B}\mid\mathcal{H}_\text{N})}\,. \tag{4.14}$$

If the detector noise were indeed Gaussian, this single quantity would suffice to define a detection criterion: a value greater than unity would indicate the signal model is favored by that factor (in terms of betting odds), and *vice versa*. However, since actual noise cannot be guaranteed to be Gaussian (and, generally, will not), the probability ratio of Eq. (4.14) does not inform us about the relative likelihoods

of a signal vs *actual* (non-Gaussian) noise. To address this, one may attempt to capture instrumental artifacts by defining a construction similar to Eq. (4.14) but using signal-coherence across detectors to distinguish spurious effects from actual astrophysical signals [18, 27, 29]. Nevertheless, it cannot be shown that any such construction will always capture all the features of real instrumental noise (in the language of formal logic, our hypothesis set is never complete). Therefore, we would benefit from a method to empirically test the efficacy of our Bayesian constructions at actually distinguishing signals from (non-Gaussian) detector noise.

5-vector approach

The frequentist 5-vector method [69] builds a detection statistic using the sidereal modulation given by the interferometer antenna response to the two CW polarizations, encoded by $F_{+/\times}$ in Eq. (4.1). Similar to the procedure outlined in Sec. 4.2.3, the first step is to remove all the possible phase modulations, apart from the sidereal ones caused by the antenna patterns. Depending on the type of search, this may be achieved through different techniques, including subheterodyning, nonuniform resampling, or a combination thereof [72, 73]. After this step, the signal can be modeled via two sidereal responses, $A_{+/\times}(t)$, analogous to $F_{+/\times}(t)$ but which do not depend on the polarization angle ψ (see [69] for more details). It can be shown that: $A_+(t) \equiv F_+(t; \psi = 0)$ and $A_\times(t) \equiv F_\times(t; \psi = \pi/4)$. By doing this, the signal assumes the complex-valued form:

$$h(t) = H_0(\eta)\left[H_+(\psi, \eta)A_+(t) + H_\times(\psi, \eta)A_\times(t)\right], \tag{4.15}$$

where η is related to the ratio of the two polarization amplitudes given in Eq. (4.1),

$$\eta = -\frac{2\cos\iota}{1 + \cos^2\iota^2} \tag{4.16}$$

and with $H_{+/\times}$ defined by

$$H_+ = \frac{\cos(2\psi) - i\eta\sin(2\psi)}{\sqrt{1+\eta^2}},$$
$$H_\times = \frac{\sin(2\psi) - i\eta\cos(2\psi)}{\sqrt{1+\eta^2}}.$$

Just as in the case of $\Lambda(t)$ in Eq. (4.12), the frequency components of a signal described by Eq. (4.15) are simply those corresponding to the sidereal modulations

encoded in $A_{+/\times}(t)$. These frequency components ($f^i_{\text{5-vec}}$) are integer multiples of the sidereal rotation frequency of Earth, namely:

$$f^i_{\text{5-vec}} = f_{\text{gw}} + 2\pi k^i \, \omega_{\text{sid}} \, , \; \vec{k} = [-2, -1, 0, 1, 2] \, . \tag{4.17}$$

Therefore, any signal like Eq. (4.15) may be described as a vector in the space spanned by the five δ-functions corresponding to the frequencies in Eq. (4.17).

To search for signals, the frequency domain GW data can be projected onto the 5-vector space to obtain a set of projections $\vec{X}$. This resulting vector now lives in the same space as the sidereal templates, which can be represented as 5-vectors $\vec{A}_{+/\times}$. We may those obtain the matched-filter between the data and the antenna patterns by taking a simple scalar product between $\vec{X}$ and $\vec{A}_{+/\times}$. By maximizing this matched-filter, one obtains an estimator for the GW polarization amplitudes:

$$\widehat{H}_{+/\times} = \frac{\vec{X} \cdot \vec{A}_{+/\times}}{|\vec{A}_{+/\times}|^2} \, , \tag{4.18}$$

which can be in turn used to define a detection statistic:

$$S_5 \equiv |\vec{A}_+|^4 |H_+|^2 + |\vec{A}_\times|^4 |H_\times|^2 \, . \tag{4.19}$$

After carrying out the above procedure for templates corresponding to different parameters, detection candidates (i.e. values of the parameters that might match a potential signal) are identified by their value of S_5. In particular, the statistic is required to exceed a threshold corresponding to a preset false alarm probability. To do this, one must know or measure the distribution of S_5 over noise. Traditionally this has been computed analytically by assuming purely Gaussian noise with known variance [69].

Alternatively, since real data are not Gaussian, one may try to approximate the background distribution by computing S_5 over frequency bands far from the expected signal ("off-frequency" analysis). The frequency regions should be far enough from a possible CW signal such that only the noise contribution is present in the detection statistic, and close enough to the analyzed band to share its statistical properties. Given that the noise strongly frequency dependent, finding this sweetspot is far from trivial (if at all possible) and one can never guarantee that the conditions required for an unbiased estimation of the background are being satisfied.

$\mathcal{F}$-statistic

The $\mathcal{F}$-statistic was first introduced in [52] for gravitational-waves searches from neutron stars, and was later extended for other astrophysical objects [74, 75]. In the case of Gaussian noise, the $\mathcal{F}$-statistic is defined as the natural logarithm of the maximum-likelihood ratio between the signal and noise hypotheses:

$$\mathcal{F} = \max \left[\ln \frac{P(\mathbf{d} \mid \vec{\theta}, \mathcal{H}_{\mathrm{S}})}{P(\mathbf{d} \mid \mathcal{H}_{\mathrm{N}})} \right]_{\vec{\theta}}, \tag{4.20}$$

where $\mathbf{d}$ is usually the calibrated detector data, and the maximization is over the signal-template parameters, $\vec{\theta}$. It can be shown that the $\mathcal{F}$ statistic can be analytically maximized over the "extrinsic parameters" (h_0, ψ, ι and ϕ_0), thus reducing the dimensionality of the numerical computations to the so-called "intrinsic parameters" (α, δ, f and $\dot{f}$). Since in a targeted search the intrinsic parameters are supposed to be perfectly known, a targeted search based on the $\mathcal{F}$-statistic would reduce to the computation of one value for $\mathcal{F}$, which is later compared to the expected noise-only distribution for Gaussian noise in order to assign p-value.

The analysis proceeds by match-filtering the data against four different templates, each of them corresponding to a particular combination of intrinsic phase evolution and sidereal modulation. The outcome of these four filters is the $\mathcal{F}$-statistic, which, if the data are composed purely of Gaussian noise, can be shown to follow a χ^2-distribution with 4 degrees of freedom [52]. Detection candidates ("outliers") are selected according to their false alarm probability, which can be computed analytically if one assumes Gaussian noise. However, false-positive outliers arise when the noise is not Gaussian, and thus not properly handled by Eq. (4.20). If this is the case and one cannot trust the background distribution of the statistic to be simple χ^2, this distribution must be estimated empirically by producing sets of data known with certainty to contain no astrophysical signals.

4.3 Method

Having reviewed the basics behind targeted searches for continuous waves, including the difficulties inherent to non-idealized instrumental noise, we here introduce off-sourcing as way to empirically assign detection significances. In Sec. 4.3.1 we describe the basic ideas behind this simple procedure and explain how it can be easily applied to the Bayesian and 5-vector searches. In Sec. 4.3.2, we heuristically explore the limits of applicability of this technique, concluding that off-sourcing is

a viable method for estimating the background distribution of detection statistics, as long as a few simple conditions are satisfied. This will be demonstrated in the following section (Sec. 4.4) with concrete examples.

4.3.1 Off-sourcing

Lacking a satisfactory way to model all noise artifacts and their effect on CW searches from first principles, we may instead attempt to empirically determine the distribution of the different search statistics in response to real detector noise and in the absence of astrophysical signals. As discussed in Sec. 4.1, a naive attempt at blinding the data to astrophysical CWs using methods analogous to those used for CBCs is doomed to failure. Therefore, we may instead look for a solution in specific properties exclusive to real gravitational signals, as opposed to noise.

One example of such a feature is the requirement of consistency between the phase evolution observed by EM astronomers and the sky location of the source: while the two properties, as encoded in the signal itself, must necessarily agree in the case of a real GW, there is no special link between them in the case of noise artifacts. Furthermore, as explained in Sec. 4.2.1, the location of the source is independently imprinted in the morphology of the signal twice: in the amplitude modulation due to the antenna patterns, Eq. (4.1), and in the frequency modulation due to Rømer and other timing delays, Eq. (4.5). Since these three properties (frequency, amplitude modulation, and phase modulation) must all agree for an astrophysical signal, we may ask: how likely is it for an instrumental artifact to randomly satisfy this condition and thus mimic a real signal from a given source?

A priori, an instrumental artifact with frequency close to that expected from a given source is no more likely to also show the amplitude and phase modulations corresponding to the true location of the pulsar than those of any other arbitrary sky location. In other words, there is no reason for instrumental noise at the target frequency to "know" what the true sky location of the source is. *By carrying out our analysis assuming incorrect sky locations, we may blind ourselves to astrophysical signals and empirically estimate the probability that instrumental artifacts in the narrow frequency region corresponding to a given source also present the modulation matching its location in the sky.*

The above idea may be rephrased in the language of function spaces. Continuous wave signals can be represented as vectors in the space of square-integrable functions (L^2) or, after discretization, the space of square-summable sequences (ℓ^2). We would

like to estimate the overlap between the subspace of L^2 (or, rather, ℓ^2) occupied by noise features and the much narrower one spanned by the signal template, Eq. (4.1). We attempt to empirically achieve this by computing an inner product (defined by the detection statistic itself) between the data and signal templates (basis elements) corresponding to different sky locations. We expect this to work partly because templates for different sky locations will be morphologically very similar to the true template, while the same is not true for any arbitrary function. This also allows us to explore the statistical properties of the noise in the same region of frequency space occupied by the expected signal.

Implementation examples

Background distributions may be estimated via off-sourcing in the context of any of the searches described in Sec. 4.2.3. This is true regardless of whether the search is carried out in the time or frequency domains, for one or several detectors, maximizing or marginalizing over nuisance parameters. This generality stems from the fact that off-sourcing is largely insensitive to the specific details behind the computation, as long as the sky location and phase evolution are assumed to be known.

Let us first illustrate this by using the time-domain Bayesian search as a concrete example. As outlined in Sec. 4.2.3, this approach is split into two stages: (*i*) heterodyning of the data to put the signal in the shape of Eq. (4.12); and (*ii*) Bayesian inference to compute the relative likelihood of a signal being present, Eq. (4.14). It is important for our purposes that information about the location of the pulsar is only needed in the first step, making it straightforward to apply our suggested strategy. In particular, we may intentionally heterodyne the data assuming an off-source sky location, and then carry out the inference stage as usual, assuming the true ("on-source") sky location. Rather than being indicative of a signal, a large Bayes factor obtained this way would necessarily reveal the presence of a noise artifact. This process may be repeated for different sky locations to obtain an estimate of how likely noise is to mimic a signal from this pulsar.

As another example, consider the frequency-domain 5-vector approach of Sec. 4.2.3. In that case, we may also resample or reheterodyne the data assuming an off-source sky-location during a preprocessing stage. This procedure is expected to spread the power of a possible GW signal over many different frequency bins, making it too weak to be detectable and thus blinding the analysis to it. Next, we compute the S_5

statistic by using the 5-vector sidereal function $A_{+/\times}(t)$ computed for the on-source sky position. We can then repeat these steps for many different sky locations to obtain a collection of background values for the S_5 statistic. This yields a noise-only distribution for the S_5 statistic that quantifies the probability for a noise disturbance to mimic the sidereal antenna patterns corresponding to the true sky location.

4.3.2 Blinding and draw-independence

In order for the method above to work, we need to make sure that: (*i*) the different sky locations used are actually distinct, so that the results can be treated as independent draws of the probability distribution we are trying to estimate; and (*ii*) off-sourcing really does blind the data to foreground signals. The first requirement is easy to satisfy and translates into the need for picking sky locations with angular separations greater than the worst (largest) sky bin resolvable by the search, as explained in Sec. 4.2.1. As we show below, the second requirement can also be satisfied by picking off-source locations far-enough away from the true position of the source.

Signal contribution to off-source statistic

To understand this, let us examine the conditions under which off-sourcing effectively removes contributions from real continuous waves. For simplicity, consider a signal of fixed frequency (f) originating from some known location ($\hat{\Omega}_{\rm on}$). Now imagine heterodyning the data containing it by using a mismatched timing correction corresponding to some off-source sky location ($\hat{\Omega}_{\rm off}$), as proposed in Sec. 4.3.1. In full analogy to Eq. (4.13), we would then obtain binned data like:

$$B(t) = \Lambda(t)e^{2\pi i f \Delta(t)} + n(t) \equiv \Lambda'(t) + n(t), \tag{4.21}$$

where we no longer take the instrumental noise $n(t)$ to necessarily be normally distributed, and where

$$\Delta\tau(t; \Delta\hat{\Omega}) \equiv \tau(t; \hat{\Omega}_{\rm on}) - \tau(t; \hat{\Omega}_{\rm off}) \tag{4.22}$$

represents the timing-correction mismatch between the two sky locations with angular separation $\Delta\hat{\Omega} = \hat{\Omega}_{\rm off} - \hat{\Omega}_{\rm on}$ [cf. Eq. (4.5)]. As a proxy for a generic search statistic, consider the evaluation of a simple inner product between the data and the expected template:

$$\langle B(t) \mid \Lambda(t) \rangle = \langle n(t) \mid \Lambda(t) \rangle + \langle \Lambda'(t) \mid \Lambda(t) \rangle \,. \tag{4.23}$$

Our goal is to estimate the distribution of the overlap between the noise and the template, $\langle n(t) \mid \Lambda(t)\rangle$, by studying our proxy statistic, $\langle B(t) \mid \Lambda(t)\rangle$. Consequently, we would like the contribution of the true signal, $\langle \Lambda'(t) \mid \Lambda(t)\rangle$, to Eq. (4.23) to be sufficiently small to be effectively undetectable.

Explicitly, the contribution of the signal to the inner product of Eq. (4.23) can be written in terms of a time integral over the observation time T,

$$\langle \Lambda'(t) \mid \Lambda(t)\rangle = \left| \int_0^T \Lambda^2(t)\, e^{2\pi i f \Delta\tau(t;\Delta\hat{\Omega})} \mathrm{d}t \right| , \tag{4.24}$$

where $\Lambda^2(t) \equiv \Lambda^*(t)\Lambda(t)$. The first key feature of this result is that a signal with greater signal-to-noise ratio (SNR) will tend to contaminate the off-source statistic more strongly. This is not at all surprising: a strong signal can be detected even if there is a small error in its assumed sky location, because enough coherent power can remain even after the timing correction spreads it over several frequency bins. In fact, the contribution of Eq. (4.24) is bounded from above by the squared-norm of the signal template,

$$\langle \Lambda'(t) \mid \Lambda(t)\rangle \leq \langle \Lambda(t) \mid \Lambda(t)\rangle = \left| \int_0^T \Lambda^2(t)\, \mathrm{d}t \right|, \tag{4.25}$$

which, for a flat power spectrum, is directly proportional to the square of the SNR.

The second relevant feature of Eq. (4.24) is that the dependence of $\langle \Lambda'(t) \mid \Lambda(t)\rangle$ on sky location will be determined solely by the angular structure of $\Delta\tau$, and how well that can be resolved given f and T (see Sec. 4.2.1). The inequality of Eq. (4.25) is, of course, saturated if and only if the off-source location is such that $\Delta\tau(t;\Delta\hat{\Omega}) = 0$ at all times: this takes place, for instance, if the "off-source" location is really just the original location of the source ($\Delta\hat{\Omega} = 0$). On the other hand, for most other values of $\Delta\hat{\Omega}$ and for the range of frequencies we are interested in, the exponential term in Eq. (4.24) is highly oscillatory—this means that we should expect $\langle \Lambda'(t) \mid \Lambda(t)\rangle$ to quickly vanish as we move away from the true location of the source. This is consistent with the sky-bin definition given in Sec. 4.2.1, from which it is possible to see that the sky-bin size decreases with increasing frequency.

Angular pattern and magnitude

The detailed angular structure of our proxy off-source statistic is represented in Fig. 4.2. To produce this plot, we began with a set of binned data, Eq. (4.13), containing

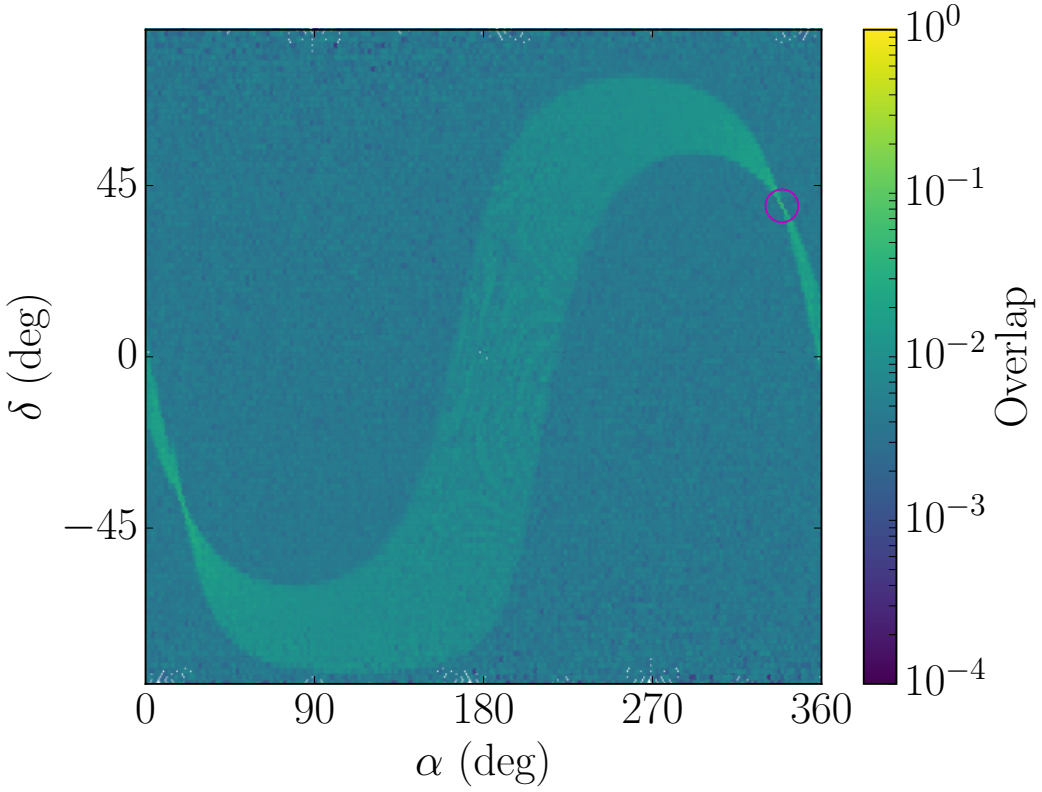

Figure 4.2: Value of the correlation term (colorbar) over the entire sky area for a loud signal injected at ($\alpha = 22^{\mathrm{h}}35^{\mathrm{m}}40.73^{\mathrm{s}}$, $\delta = 39°40'44.76''$).

Gaussian noise and a very strong (SNR = 70) simulated signal from an arbitrary sky location on the ecliptic plane (indicated by a magenta circle). We then reheterodyned these data assuming different (off-source) locations covering the whole sky, and for each instantiation computed the overlap (normalized cross-correlation),

$$\text{Overlap} = \frac{\langle B'(t) \mid B(t)\rangle}{\langle B(t) \mid B(t)\rangle} \approx \frac{\langle \Lambda'(t) \mid \Lambda(t)\rangle}{\langle \Lambda(t) \mid \Lambda(t)\rangle}, \tag{4.26}$$

between the off-source data, $B'(t)$, and the on-source data, $B(t)$. This quantity (shown in color in Fig. 4.2) represents the normalized contribution of the injected signal to the off-source statistic for different sky locations, as desired. This is because $\langle n'(t) \mid n(t)\rangle \approx 0$ for Gaussian noise, yielding the approximate equality in Eq. (4.26).

As expected, the contribution of the signal falls off steeply as we move away from the source location: while the normalized overlap of Fig. 4.2 equals unity if $\Delta\hat{\Omega} = 0$ (center of the magenta circle), it is orders of magnitude smaller for all other choices of $\hat{\Omega}_{\text{off}}$. The rest of the structure in this plot reflects the symmetries of the timing correction Eq. (4.5), which are themselves dominated by the Rømer term in Eq. (4.6): locations across lines of fixed ecliptic latitude remain somewhat correlated to the

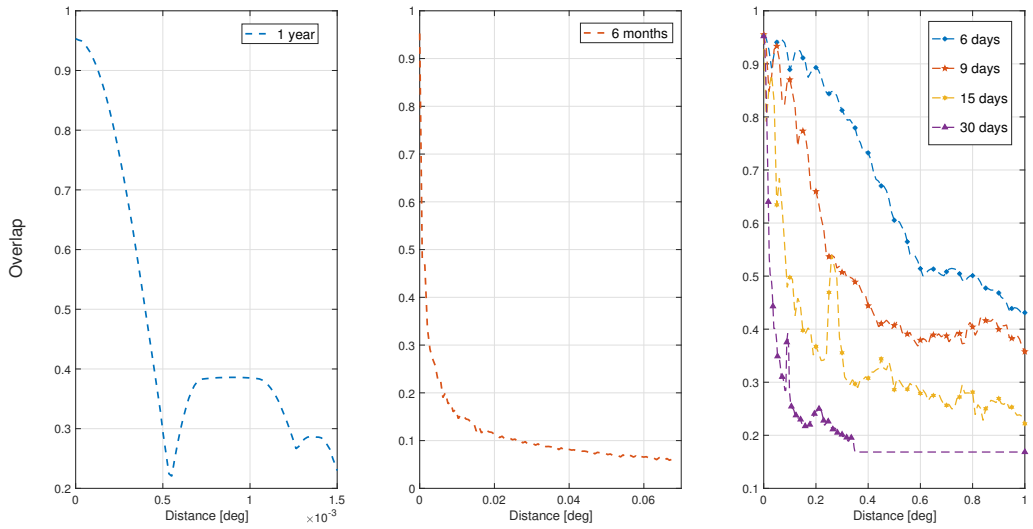

Figure 4.3: *Left panel:* Overlap function computed for a software injection with the same parameters of the one generated for Fig. 4.2 for 1 year of integration and close to the source. The overlap function drops very rapidly within a distance of 0.005 deg from the source. *Central panel:* Overlap function computed for a software injection with the same parameters of the one generated for Fig. 4.2 for 6 months of integration. The overlap function drops within a larger distance with respect to 1 year of integration. *Right panel:* Overlap function computed for a software injection with the same parameters of the one generated for Fig. 4.2 for various integration times (see legend). We see that shorter integration times will correspond to a wider sky area in which the templates will have non-negligible overlap. The overlap function at the source location (null distance) is not 1 dueto the presence of the noise.

on-source location, and the whole pattern is symmetric under reflections through the ecliptic (see Sec. 4.2.1 for more details). This important observation means that, for any given on-source location, we will want to sample our off-source points from a *grid* covering only one of the ecliptic hemispheres. Another possibility is to *randomly* pick points anywhere in the sky, while enforcing a minimum separation determined by the decay of the overlap function in Fig. 4.3. For 6 months of data analyzed coherently, this distance would have to be $\gtrsim 0.01$ deg.

For a sufficiently loud signal, sky bins neighboring the source will yield contaminated off-source data (i.e. data that still contain measurable coherent power due to the on-source signal). Unlike the angular dependence, the overall magnitude of this contamination will be determined by the SNR of the signal and, as such, will depend on the integration time and intrinsic amplitude. For the same example as in Fig. 4.2, Fig. 4.3 shows the rate at which the overlap with the on-source location decreases

as one moves at constant ecliptic latitude away from the source and for different integration times.

Because latitude is held constant in this plot, Fig. 4.3 represents the slowest-possible decrease in the contamination by this source (cf. Fig. 4.2). Furthermore, this example was chosen to have very high SNR and to lie on the ecliptic plane, where the sky resolution is poorer (cf. Sec. 4.2.1)—all of which makes this close to a worst-case scenario. In spite of this, the overlap vanishes quite quickly, plateauing far away from the source at a value of the order 0.01. For realistic CW signals (lower SNR) a 1% overlap is small enough to make the contribution of the signal in Eq. (4.23) negligible with respect to the noise. Hence for 1 month of data integration or more, taking as off-source sky resolution 1 deg is enough to remove any measurable correlation between sky bins in any realistic situation. (Examples of this are given in Sec. 4.4 below.)

Contaminated backgrounds

In any case, an analysis that draws part of the background from a measurably contaminated region may underestimate the significance of the true signal, but never overestimate it. This is because a contaminated background will show artificial tails towards higher values of the detection statistic, due to coherent power left over after off-sourcing in some of the "noise-only" instantiations. Thus, in a sense, such an analysis would, at worst, be conservative. Yet, as we will show in Sec. 4.4, a signal that is sufficiently loud to cause such contamination over a non-negligible region of the sky will itself yield an on-source detection statistic that is significantly higher than any of the contaminated-background tails. Therefore, the significance (e.g. p-value) assigned to such a signal will be the same with or without the tails.

In any case, sky bins in the immediate vicinity of the source may always be removed from the background estimation to prevent contamination. However, the excision of a large area of the sky will have the detrimental consequence of effectively reducing the number of independent sky bins available for background estimation. Furthermore, such procedure is only justified if we (implicitly) assume that the on-source data do contain an astrophysical signal. In a way, this is analogous to how a very loud CBC signal may pollute the time-slid background in searches for transient gravitational waves (e.g. see caption to Fig. 3 in [3], or Fig. 7 in [76]). In that case, the standard procedure has been to first compute significances with the "polluted" background to

determine whether the zero-lag detection candidate is a real signal and only if that is the case remove it from the background.[2] The same can be done here if necessary.

In summary, we conclude that off-sourcing, as described in Sec. 4.3.1, is a viable method for estimating the background distribution of detection statistics in targeted searches for continuous waves, as long as off-source sky-locations are chosen such that: (*i*) they are distributed over only half the sky; (*ii*) the angular distance between them is no shorter than the sky-resolution of the search (cf. Sec. 4.2.1). This will guarantee that the different draws of the background distribution (obtained from different off-source sky locations) are independent and uncontaminated by a true signal, were it to be there.

4.4 Analysis

We study the efficacy of off-sourcing (Sec. 4.3.1) as a viable method to empirically estimate the background distribution of detection statistics in targeted searches for continuous gravitational waves from known pulsars. We do this in the context of both the Bayesian (Sec. 4.2.3) and frequentist 5-vector (Sec. 4.2.3) analyses to demonstrate the generality of the approach. We discuss specific case-studies in Sec. 4.4.1 and systematically compare to different methods by computing false-dismissal and false-alarm rates in Sec. 4.4.2.

The following results make use of both simulated and actual noise from interferometric detectors. In all cases, we begin with a set of data representing the (fabricated or actual) output of a detector after applying the preprocessing required to target some chosen pulsar (e.g. filtering and downsampling)—these are the on-source data. We then proceed as described in Sec. 4.3.1 to generate multiple new sets of off-source data, and then evaluate the distribution of the detection statistic over all such instantiations (excluding the original, on-source one). We can then compare the value of the on-source statistic to the off-source background, as we would in a real analysis.

4.4.1 Case studies

Here we provide several concrete examples of off-sourcing at work in the presence of pure noise, realistic signals and strong signals, as summarized in Table 4.1. Background distributions are estimated from 10^4 off-source locations in the hemisphere of

[2] These two kinds of background are known colloquially as with and without "little dogs", since this distinction first arose during the analysis of an injection in the direction of *Canis Major* [77].

Table 4.1: Parameters for the case-study signals (Sec. 4.4.1).

	f_{GW} (Hz)	α	δ	Figs.	Data	Comment
J0534+2200	59.33	$5^{\mathrm{h}}34^{\mathrm{m}}31.97^{\mathrm{s}}$	$22°00'52.07''$	4.5, 4.6	Gauss. (H,L,V) design	Assumed $\iota = 61.3°$, $\psi = 124.0°$
J1932+17	47.81	$19^{\mathrm{h}}32^{\mathrm{m}}07.17^{\mathrm{s}}$	$17°56'18.70''$	4.7	Real O1 (H,L)	Published in [29]
P03	108.86	$11^{\mathrm{h}}53^{\mathrm{m}}29.42^{\mathrm{s}}$	$-33°26'11.77''$	4.8	Real O1 (H,L)	Hardware injection [78]

the source. The simulations of Gaussian noise (Sec. 4.4.1) were carried out assuming an observation time of 6 months and PSDs corresponding to the aLIGO and Virgo design sensitivities. With the exception of Fig. 4.4, the simulated data for LIGO Hanford ("H"), LIGO Livingston ("L") and Virgo ("V") detectors were then analyzed coherently with the Bayesian method of Sec. 4.2.3, to obtain the signal vs noise Bayes factors of Eq. (4.14) as our detection statistic.

The examples with real instrumental noise correspond to LIGO's first observation run (O1). The data streams start on 2015 Sep 11 at 01:25:03 UTC for Hanford and 18:29:03 UTC for Livingston and finish on 2016 Jan 19 at 17:07:59 UTC at both sites. The first example consists of data prepared for the pulsar PSR J19312+17, for which search results were presented in [29]. All analysis settings are the same as in [29], except for a log-uniform prior in the signal amplitude (same as in [79]). The second example is for a hardware injection presented in [78]. Both these examples are offered merely to demonstrate the performance of off-sourcing under realistic circumstances—we present no new observational results.

Pure Gaussian noise

We first demonstrate that off-sourcing works as expected in pure Gaussian noise and in the absence of signal. In this case, the on-source data are just a set of samples from a Gaussian distribution with zero-mean and standard deviation given by the value of the detector PSD at the GW frequency expected from the targeted pulsar. The off-sourcing process should correspondingly produce multiple instantiations of *independent* Gaussian noise, a fact that should be reflected in the resulting background distribution of the detection statistic. This distribution is shown in Fig. 4.4 for an example using the 5-vector statistic of Eq. (4.19). When computed over Gaussian noise, it can be shown that this statistic *must* follow a Γ distribution with 2 degrees of freedom [69]. Fig. 4.4 shows that this is the case, in agreement with our expectation that off-sourcing should produce independent draws from the background distribution.

Injections in Gaussian noise

Ideally, the background distribution should be unaffected by the presence of a signal: while the value of the on-source statistic should rise to reveal it, the off-source values should be insensitive to it. We demonstrate that this is the case by injecting signals

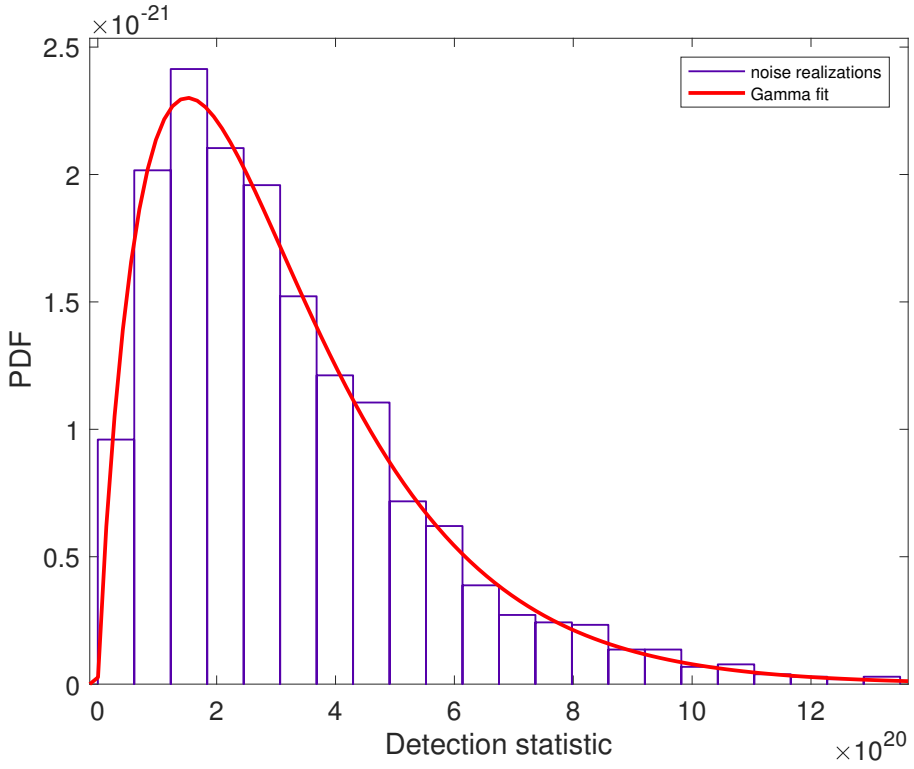

Figure 4.4: Histogram of the noise-only distribution obtained using the off-sourcing method in case of Gaussian noise. Red line: best fit given by a Γ distribution with $a = 1.95$. In this test case the hypothetical source was assumed at a frequency $f \approx 108.85$ Hz with a spin-down $\dot{f} \approx 10^{-17}$ Hz/s . The on-source position was chosen on the ecliptic plane ($\alpha = 22^{\mathrm{h}}35^{\mathrm{m}}40.73^{\mathrm{s}}$, $\delta = 39°40'44.76''$) as in the example for the overlap function reported in Sec. 4.3.1
and the search was performed for a hypothetical GW detector at the position of the LIGO Hanford site.

of different amplitudes in Gaussian noise. We simulate a signal from the Crab pulsar (PSR J0534+2200) as seen by three advanced detectors (H1, L1, V1) at design sensitivity over 6 months and recover it using the Bayesian method of Sec. 4.2.3.

We first choose a realistic signal amplitude of $h_0 = 10^{-26}$, which is weak enough to be consistent with the latest upper limits for this source [29], but strong enough to yield a nonnegligible network-SNR = 5 for the chosen PSDs and observation time. The data containing the injected signal are then reheterodyned for 10^4 off-source sky-locations to yield the survival functions (1-CDF) in Fig. 4.5. Each colored trace in this figure represents the distribution of off-source background computed from the Northern celestial hemisphere, *excluding* any points closer to the source than the indicated angular distance, i.e. excluding points with (α, δ) such that $|\alpha - \alpha_\star| < \Delta$

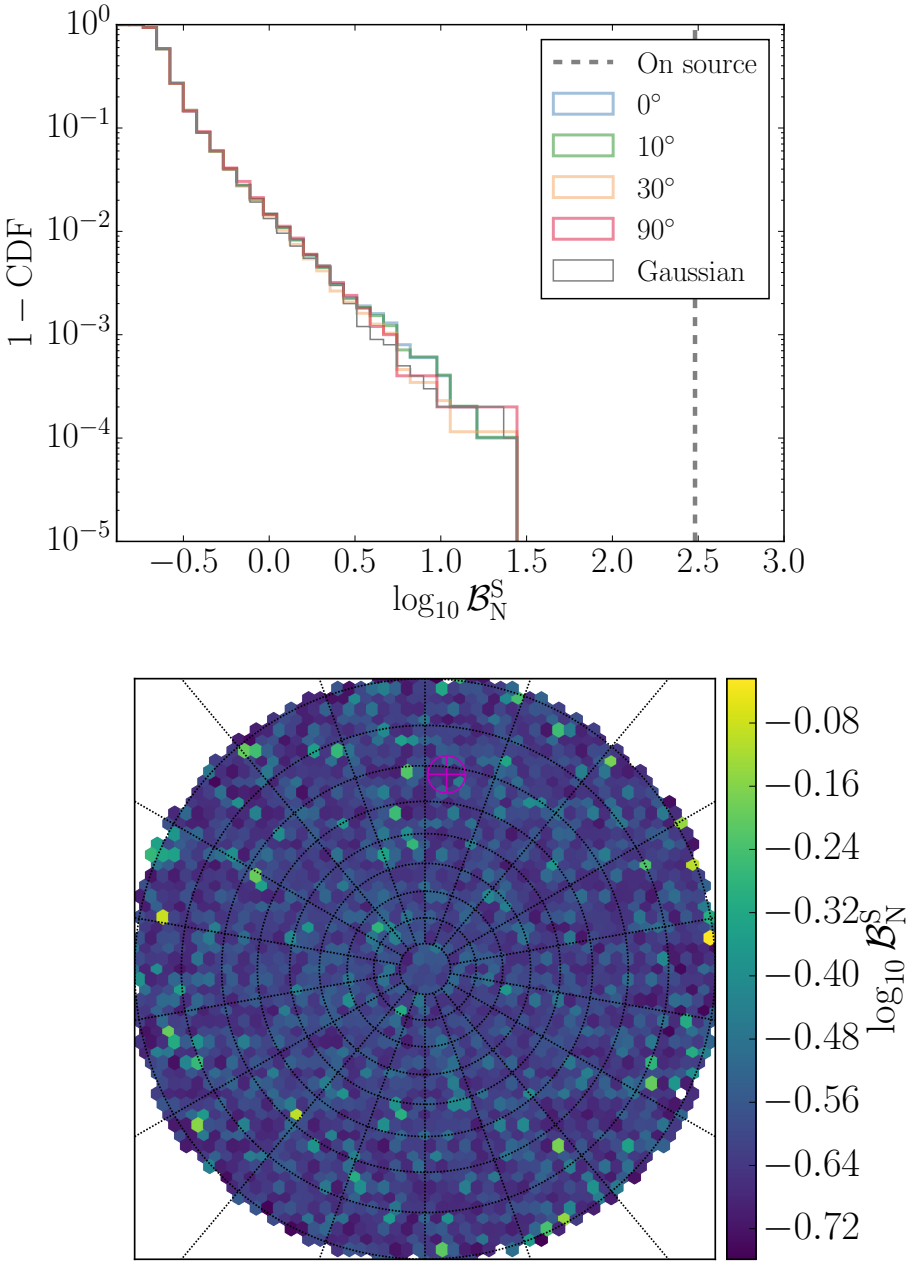

Figure 4.5: *Realistic Crab signal in Gaussian noise*. We simulate a 6-month-long signal with $h_0 = 10^{-26}$ (network SNR 5) and parameters consistent with the Crab pulsar (PSR J0534+2200), inject it in Gaussian noise for aLIGO and Virgo design PSDs and recover it using the Bayesian analysis of Sec. 4.2.3 (see Table 4.1). The *left panel* shows the survival function (1-CDF) of $\log_{10}\mathcal{B}^{S}_{N}$, Eq. (4.14), for the off-sourced background produced from the injected data for different excision areas around the source (different colors, blue and green overlap almost perfectly), as well as from pure Gaussian noise (gray, thin histogram); the on-source statistic for the injection is $\log_{10}\mathcal{B}^{S}_{N} = 2.5$ (thick dashed line), higher than any of the 10^4 off-source instantiations. The *right panel* shows the distribution of the off-source statistic over the sky in a North-polar stereographic projection, with the Crab's location marked by the magenta crosshairs; the color of each hexagon gives the average of $\log_{10}\mathcal{B}^{S}_{N}$ over several sky bins.

and $|\delta - \delta_\star| < \Delta$ with a "$\star$" indicating the true location of the Crab and Δ one of the values given in the legend of Fig. 4.5: 0° (blue), 10° (green), 30° (yellow) or 90° (red). In particular, the blue curve corresponds to background from points sampled over the whole hemisphere, while the red curve corresponds to points sampled over the half-hemisphere not containing the source.

In this case, the choice of sky-region does not have a strong effect on the background: we may take advantage of the whole hemisphere, getting quite close to the source (as allowed by the frequency resolution of this search). In fact, note that the blue and green curves in Fig. 4.5 are essentially identical. For reference, the distribution of the off-source statistic over the whole Northern sky (0° curve on the left) is represented on the right panel of Fig. 4.5 via a stereographic map, with the true location of the source indicated by magenta crosshairs.

As expected, the background produced via off-sourcing is practically indistinguishable from results in pure Gaussian noise (gray, thin histogram). Indeed, these two samples yield a Kolmogorov-Smirnov (KS) p-value of 0.77, favoring the hypothesis that they were both drawn from a common distribution. This agreement is in spite of the fact that the on-source statistic (red line) takes a significantly-increased value, revealing the presence of the injection. Completely ignoring the intrinsic probabilistic meaning of $\mathcal{B}^{\mathrm{S}}_{\mathrm{N}}$, a background like the blue curve in the left panel of Fig. 4.5 would allow us to place a p-value of at most 10^{-4} on the null hypothesis that the on-source data are noise.

As anticipated in Sec. 4.3.2, there is a limit to how loud the injection can be without noticeably contaminating nearby sky-bins and, therefore, biasing the background distribution obtained through off-sourcing. However, this threshold is quite high: for the same detector configuration as above, we find that the injection must reach $h_0 \sim O\left(10^{-24}\right)$, or a network-SNR ~700 at design sensitivity, before off-sourcing is unable to effectively remove it. We show an example of this in Fig. 4.6 for a signal from the Crab pulsar with $h_0 = 1.4 \times 10^{-24}$, which roughly corresponds to the spin-down limit for this source [29].[3] This time, as seen from the panel on the left, the full-hemisphere off-sourcing distribution (blue curve) is visibly inconsistent with a pure-Gaussian background (gray, thin curve), and a KS p-value of 10^{-77} strongly disfavors a shared distribution between the two sample sets. From the right panel, it is clear that the culprits are noticeably-contaminated sky locations in the neighborhood

[3]For an isolated pulsar (no accretion), the spin-down limit is the maximum power that could possibly be emitted in gravitational waves given the observed decay in the pulsar's angular momentum.

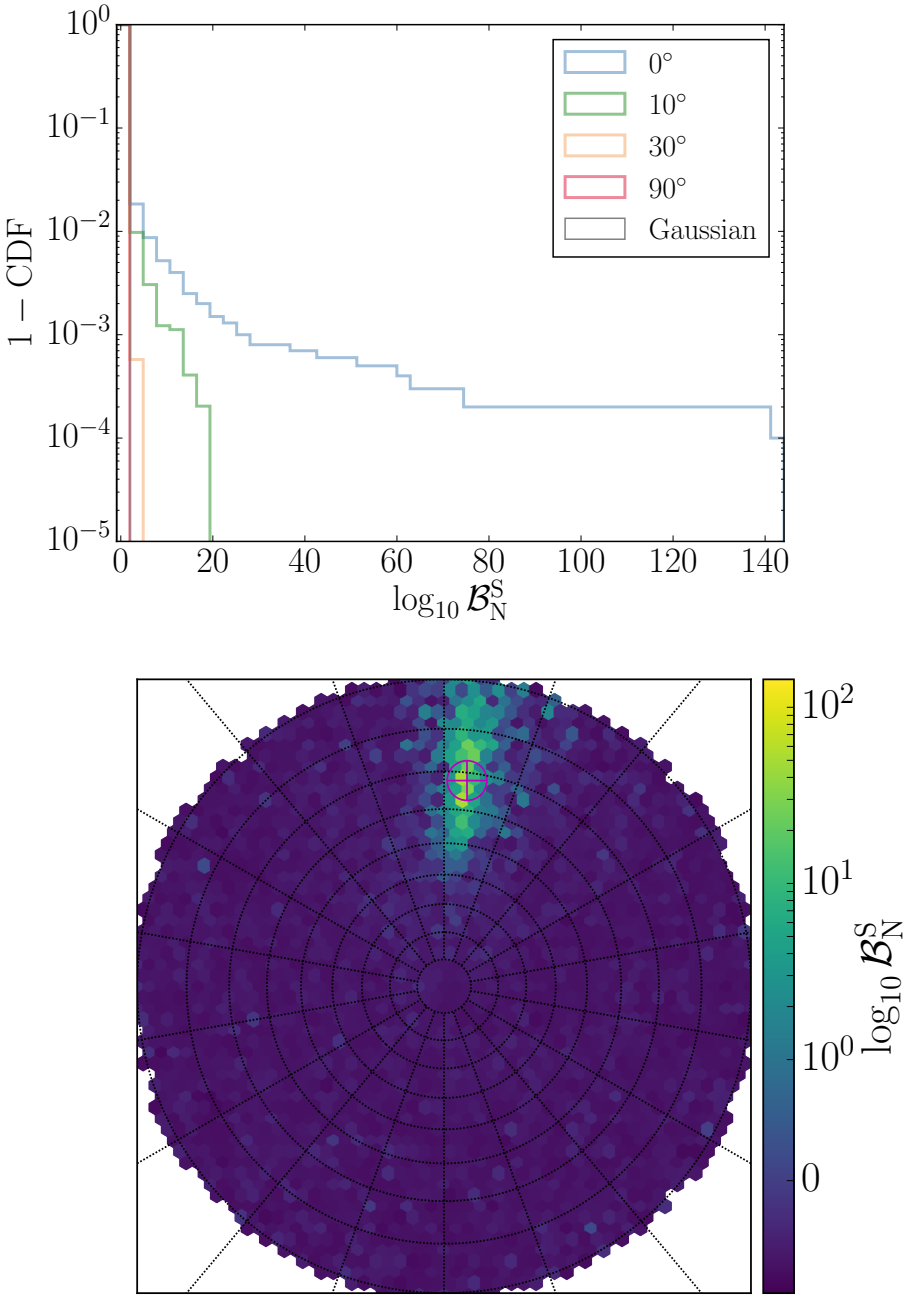

Figure 4.6: *Very loud Crab signal in Gaussian noise*. We simulate a 6-month-long signal with $h_0 = 1.4 \times 10^{-24}$ (network SNR 700) and parameters consistent with the Crab pulsar (PSR J0534+2200), inject it in Gaussian noise for aLIGO and Virgo design PSDs and recover it using the Bayesian analysis of Sec. 4.2.3 (see Table 4.1). The *left panel* shows the survival function (1-CDF) of $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}}$, Eq. (4.14), for the off-sourced background produced from the injected data for different excision areas around the source (different colors), as well as from pure Gaussian noise (gray, thin histogram); the on-source statistic for the injection is $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}} = 9 \times 10^4$, which is vastly higher than any of the 10^4 off-source instantiations (off the scale). The *right panel* shows the distribution of the off-source statistic over the sky in a North-polar stereographic projection, and with the Crab's location marked by the magenta crosshairs; the color of each hexagon gives the average of $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}}$ over several sky bins, in semi-log scale linearly interpolated between $(-1, 1)$.

of the source (magenta crosshairs). These polluted sky bins are arranged in the same pattern predicted in Fig. 4.2, although under a different guise due to the logarithmic color scale.

As discussed in Sec. 4.3.2, background contamination can at worst cause us to *underestimate*, never overestimate, the significance of a detection. However, this is not the case in the example of Fig. 4.6 because the signal is too loud ($\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}} = 9 \times 10^4$, off the scale of the histogram in Fig. 4.6). This is a general feature: in Gaussian noise, if a signal is loud enough to contaminate a large region of the sky, it will also be louder than the loudest background produced from it.

Real noise

The above behavior is replicated in the presence of actual noise from LIGO and Virgo, with the difference that the background naturally shows tails due to the non-Gaussianities in the data. An example of this is shown in Fig. 4.7, which was produced using actual data from aLIGO's first observation run, prepared for the pulsar PSR J1932+17 and with both detectors analyzed coherently using the Bayesian method of Sec. 4.2.3. As before, the left panel show the off-sourced background distributions for different excision areas around the source (different colors). Note that the excision process does not have any significant impact on the distribution, which is what one would expect in the absence of a very loud signal at the on-source location. The presence of artifacts in the data becomes apparent in the slower drop of the survival function with respect to, e.g., Fig. 4.5. The on-source value of the signal vs noise Bayes factor for this source was published in [29], and is marked here by a vertical dashed line—clearly, there is no evidence for a signal in the data. As before, the distribution of the off-source statistic over the Southern sky is represented on the right-hand side of Fig. 4.7.

To study the effectiveness of off-sourcing in detecting a signal in real noise, we analyze data for the hardware injection referred to as "P03" in [78]. Hardware injections are produced by physically actuating on the test masses to mimic the effect of a true gravitational wave, providing a valuable end-to-end test of the instrumental calibration and analysis pipelines. In the case of P03, the signal was injected at 108.86 Hz with a loud amplitude of $h_0 = 8.2 \times 10^{-25}$ (network-SNR = 50). As shown on the left panel of Fig. 4.8, this signal seems to be sufficiently strong to slightly contaminate the sky bins in its immediate vicinity, but this pollution is easily

removed via a narrow excision (compare the blue trace to the rest in Fig. 4.8). In any case, the value of the on-source Bayes factor for this signal is $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}} = 504$, which is significantly louder than the loudest background. Given that 10^4 off-source noise instantiations were used to estimate the background, this implies that Fig. 4.8 would allow us to claim a detection of P03 with $p \leq 10^{-4}$ (ignoring the probabilistic meaning of $\mathcal{B}^{\mathrm{S}}_{\mathrm{N}}$).

4.4.2 Comparison to other methods

In order to determine whether off-sourcing offers an improvement over other strategies, we must go beyond specific examples and study false-alarm and false-dismissal rates. That is, respectively, how likely is off-sourcing to conclude that a noise artifact is a signal (false alarm), and how likely is it to conclude that a signal is a noise artifact (false dismissal), as a function of confidence level? We estimate those rates from a large number simulations in simulated and actual noise and use them to directly compare with the standard background-estimation method for the 5-vector search (Sec. 4.2.3) We find that off-sourcing outperforms the usual methods in real LIGO data.

False-dismissal rate

First, in order to study the false dismissal rate, for a selection of SNRs, we simulate 250 signals over the sky, with extrinsic parameters (ψ, η, ϕ_0) picked randomly over their allowed ranges. We then inject these in idealized (Gaussian) and actual O1 noise for the LIGO Hanford and Livingston detectors (4 months observation time). In the case of Gaussian noise, the frequency and spin-down parameters of the injections were decided to be equal to the ones of P03, this choice will reflect only on size of the sky-patches that we should use in the off-sourcing as pointed out in Eq. (4.9). On the other hand, when using detectors real data, the frequency of the injections was decided to be around 54.5 Hz, since we knew *a priori* that O1 data in this frequency band was polluted by noise artifacts, especially in the Livingston detector [80], making this a good frequency region to test the ability of off-sourcing to distinguish between astrophysical GW signals and noise.

In each case, we analyze the resulting data using the method of Sec. 4.2.3 to obtain the on-source value of the detection statistic, as well as 2×10^4 off-source values to produce a background drawn with a minimum distance of 0.01 deg between each

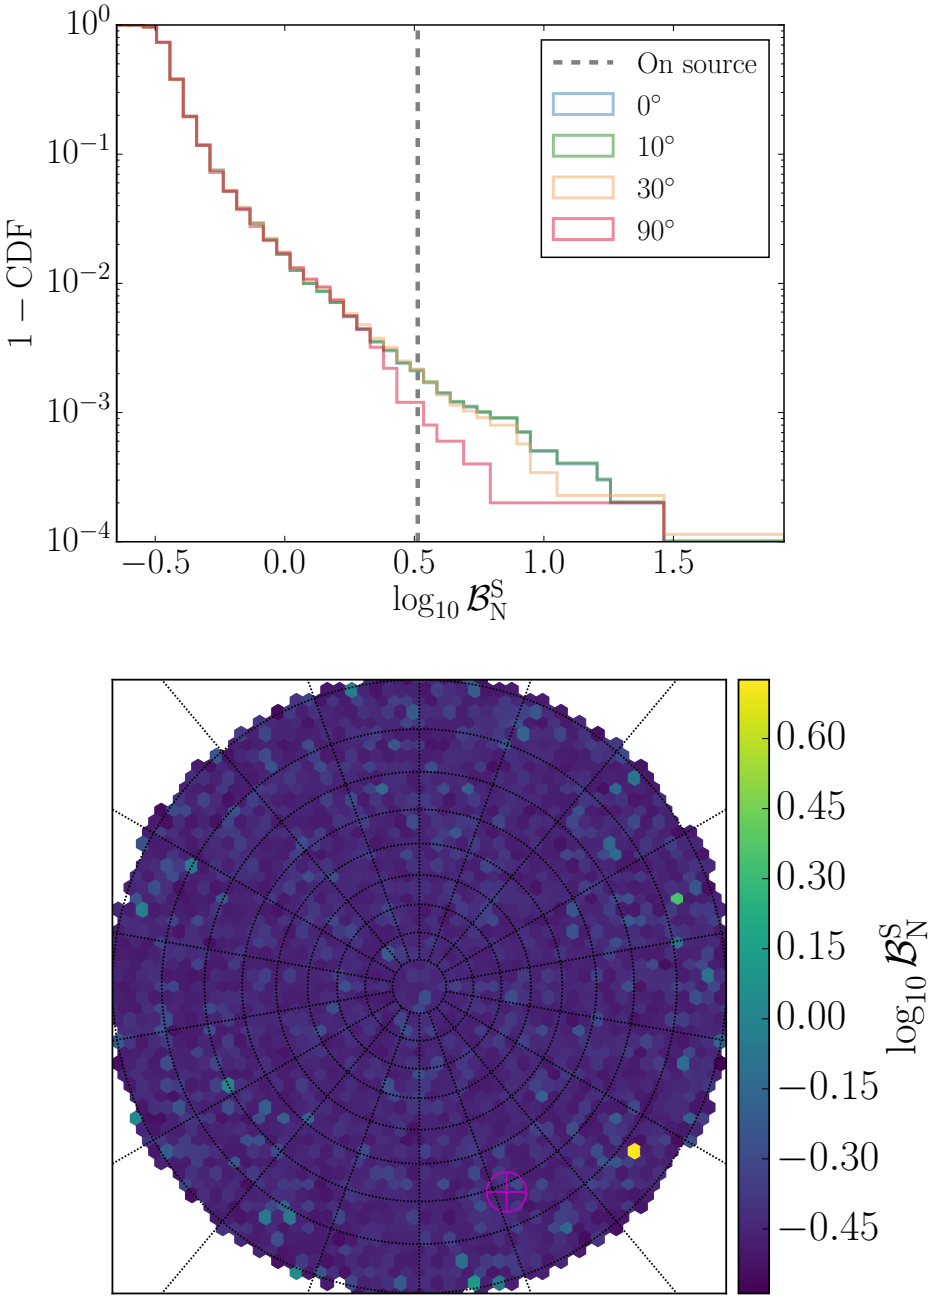

Figure 4.7: *Real O1 noise for J1932+17.* Off-sourced background produced from real O1 LIGO data prepared for PSR J1932+17, analyzed coherently with the Bayesian method of Sec. 4.2.3 (see Table 4.1). The *left panel* shows the survival function (1-CDF) of $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}}$, Eq. (4.14), for the off-sourced background for different excision areas around the source (different colors); the on-source statistic for this pulsar is $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}} = 0.5$ (vertical dashed line), as was reported in [29]. The *right panel* shows the distribution of the off-source statistic over the sky in a North-polar stereographic projection, with the true location marked by the magenta crosshairs; the color of each hexagon gives the local average of $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}}$ over several sky bins.

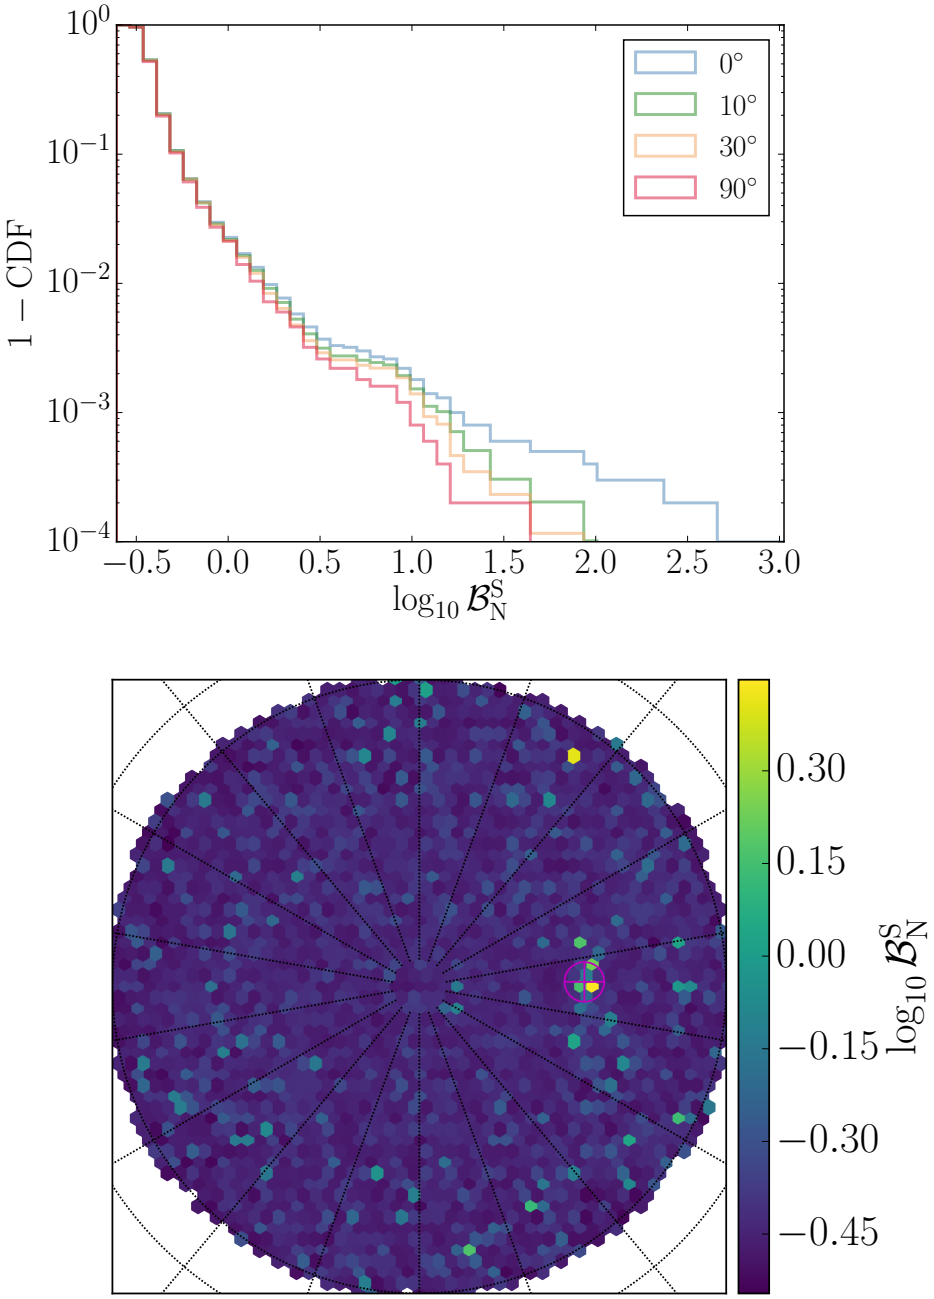

Figure 4.8: *Real O1 noise for a loud hardware injection (P03).* Off-sourced background produced from real O1 LIGO data prepared for hardware injection P03 [78], analyzed coherently with the Bayesian method of Sec. 4.2.3 (see Table 4.1). The *left panel* shows the survival function (1-CDF) of $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}}$, Eq. (4.14), for the off-sourced background for different excision areas around the source (different colors); the on-source statistic for this pulsar is $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}} = 504$ (off-scale). The *right panel* shows the distribution of the off-source statistic over the sky in a South-polar stereographic projection, with the true location marked by the magenta crosshairs; the color of each hexagon gives the local average of $\log_{10} \mathcal{B}^{\mathrm{S}}_{\mathrm{N}}$ over several sky bins.

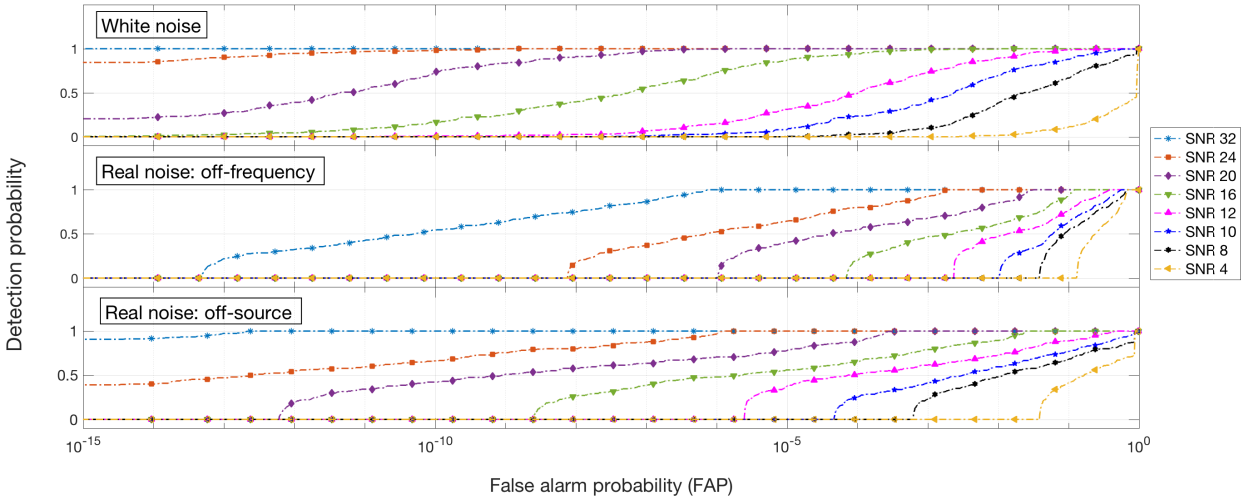

Figure 4.9: (Placeholder) Top panel: Outlier selection probability vs FAP threshold in Gaussian noise. Middle panel: Same but done in non Gaussian noise using the off-frequency method. Bottom: Same again but this time using the off-source method.

other. We then compute the number of detected signals as a function of false-alarm probability, i.e. the number of injections recovered with a detection statistic that is higher than or equal to the value corresponding to certain p-value, as established from the empirically-estimated background. For comparison, we repeat the above procedure but with a background generated via the standard "off-frequency" method (mentioned in Sec. 4.2.3 and described in detail in [80]), instead of off-sourcing. In both cases, the tails of the background distributions are extrapolated for very large values of the detection statistic using an exponential-decay fit.

The results of this study are summarized in the receiver operating characteristic (ROC) curves of Fig. 4.9. Each curve in that figure gives the normalized detection rate (detection probability) as a function of p-value (FAP) for different values of injected SNR (different colors and traces) and different methods used to estimate the FAP. The top two panels display results obtained by using off-sourcing to generate the background: the topmost one using simulated Gaussian noise, and the middle one with actual detector noise. On the other hand, the third panel shows results from the same set of actual detector noise, but with background estimated with the off-frequency method standard in 5-vector analyses.

It is clear from Fig. 4.9 that, at any fixed SNR and FAP, off-sourcing performs better in the presence of pure Gaussian noise (top panel) than actual LIGO noise

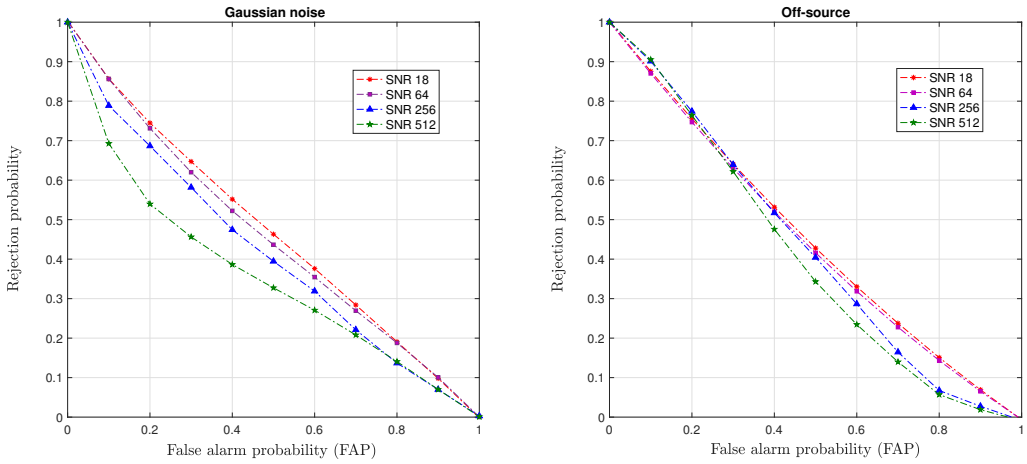

Figure 4.10: Rejection probability of an outlier born from a persistent monochromatic noise lines in raw data. Significance for selection estimated using a Gaussian model for the noise (bottom) and the off-source method (top).

(bottom panel). For example, with these settings we are only 50% likely to detect an SNR = 16 signal (green curve, down-pointing triangles) with FAP = 10^{-6} ($\sim 5\sigma$) in real O1 noise, but we are 75% likely to detect it with the same confidence in Gaussian noise. This is totally expected due to the fact that matched filtering is optimal for stationary Gaussian noise and agrees with the test cases shown in Sec. 4.4.1 (e.g. compare Fig. 4.5 to Fig. 4.8, with attention to the tails of the distributions). The artifacts in real noise produce tails in the background distribution at higher values of the detection statistic, hurting our chances to detect actual signals.

Most importantly, Fig. 4.9 also shows that off-sourcing outperforms traditional methods in the presence of actual instrumental noise, as can be seen by comparing the center and bottom panels. For instance, the off-frequency method (center panel) has essentially 0% chance of detecting an SNR = 16 signal at FAP = 10^{-6}, which is dramatically less than the 75% chance of detecting it via off-sourcing (center panel). In fact, for these settings, off-sourcing is consistently superior at all SNRs.

False-alarm rate

Besides assigning high significance to real signals, off-sourcing should also be able to robustly reject outliers arising from non-Gaussianities in the data. In other words, it should have a low false-alarm rate at any given level of confidence, rejecting artifacts with high probability. The study of noise features is a very general problem

due to the wide morphological range of non-Gaussianities that can be found in the data, which the same reason why modeling noise likelihoods from first principles is impossible in the first place. This makes the benchmarking of noise-rejection probability a challenging problem.

To address this, we use as proxy simulated monochromatic noise lines at frequencies close to the targeted pulsar frequency. The putative source for which we were looking was, again, a pulsar with the rotational parameters of P03 and the sky position $\alpha = 22^{\mathrm{h}}35^{\mathrm{m}}40.73^{\mathrm{s}}$, $\delta = 39°40'44.76''$. In particular, we produce 300 datasets with noise lines at fixed signal-to-noise ratios injected in Gaussian noise before performing the search with 4 months of observation time. Each noise line is injected with a frequency within 0.01 Hz of the targeted frequency, ensuring that the Rømer correction will cause it to contaminate the on-source analysis. This is due to the fact for a putative source at f 108.85 Hz the frequency shift due to the Rømer correction will be at maximum [72]:

$$\Delta f = f \cdot 10^{-4} \approx 0.01\mathrm{Hz}. \tag{4.27}$$

We pick the specific frequency and phase of the noise lines from a uniform distribution. After doing this, for every dataset we compute the significance of the outlier due to the injected noise line using off-sourcing with the 5-vector method (Sec. 4.2.3), producing 10^4 realizations of the noise background using sky patches separated with a minimum distance of 0.01 deg. We also evaluate the significance using the theoretical formula assuming pure Gaussian noise (cf. Fig. 4.4) [69].

We evaluate our method's ability to identify noise artifacts by studying the rejection probability as a function of the confidence threshold set to claim a detection—that is, how likely the analysis is to reject the artifact as we decrease our tolerance for false alarms (FAP). Figure Fig. 4.10 shows the results obtained empirically with off-sourcing (top panel) and analytically assuming Gaussian noise (bottom panel). The y-axis shows the percentage of noise lines which generate an outlier with a significance equal or higher than the detection-confidence threshold shown on the x-axis. In the ideal case, we would be able to perfectly measure the significance of an outlier and the curves in Fig. 4.10 would simply be a straight line with slope -1. As we can see, strong (SNR > 256) noise lines produce significant outliers if we assume the background to be Gaussian; however, this is no longer true for off-sourcing. The top plot clearly shows that the off-source can evaluate correctly the contribution of the noise to the noise-only distribution.

These tests can be extended to a general noise background. In principle, we can model coherent instrumental noise as a linear combination of monochromatic noise lines like those injected above. Every noise line will couple constructively or destructively with the other noise lines during the Doppler corrections, adding up to the overall contribution to the analysis. If the lines combined constructively, we will obtain a case very similar to the one presented in Fig. 4.10 but with a larger effective signal-to-noise ratio for the noise lines. While, if the noise lines combines destructively, then we will obtain a case very similar to the one presented in Fig. 4.10 but with a lower effective signal-to-noise ratio. The general case should lie somewhere in between.

4.5 Conclusion

Off-sourcing can provide a much-needed efficient and robust way to empirically estimate the background of searches for continuous gravitational waves targeted at known pulsars, enabling estimates of detection significance that are valid in actual (non-Gaussian) instrumental noise. This method has already been put into practice for diagnosing outliers in actual LIGO and Virgo searches [29, 53], but a systematic study of its performance was lacking from the literature. In this paper, we have filled in this gap by introducing the rationale behind this strategy, exploring its theoretical applicability and studying its performance in real and simulated data. We find that off-sourcing is an excellent tool, outperforming another standard method for estimating search backgrounds.

The procedure is simple: the original gravitational-wave data are time-corrected for multiple off-source sky locations to obtain as many instantiations of noise-only data, which are then analyzed by any of the usual searches with the same settings as the on-source search (Sec. 4.3.1). Under the right conditions, we show that the off-sourced data are blind to astrophysical signals while retaining the statistical properties of the noise. This allows for the direct empirical estimation of the background distribution of the different search statistics.

Two conditions need to be satisfied for off-sourcing to be effective: off-source sky locations must (i) be resolvably different and (ii) be drawn from the same hemisphere as the source. As long as this is true, off-sourcing will provide independent draws from the background distribution (Sec. 4.3.2). Furthermore, for realistic signal amplitudes, the distribution will be uncontaminated by the presence of a signal at the true on-source sky location. This is not true for extremely loud signals, but this is

not a problem because in those cases the on-source statistic is always louder than the background (Sec. 4.3.2). The phenomenon is analogous to that observed with strong signals in searches for compact binaries [3, 76].

We illustrate the efficacy of off-sourcing with several examples in real and fabricated data (Sec. 4.4.1). This includes simulated Gaussian noise in the absence of signal (Fig. 4.4), as well as in the presence of realistic (network-SNR = 5, Fig. 4.5) and strong (network-SNR = 700, Fig. 4.6) signals. We also demonstrate the method in the presence of real LIGO O1 noise with data prepared for PSR J1932+17 [29] and loud hardware injection [78]. Source parameters for all these case studies are summarized in Table 4.1.

Finally, we systematically study the performance of off-sourcing by looking at false-dismissal and false-alarm rates (Sec. 4.4.2). The former is quantified by the receiver operating curve of Fig. 4.9 and the latter by the rejection-probability vs confidence-level plot of Fig. 4.10. We find that off-sourcing outperforms the standard method for computing significances in the context of the 5-vector search.

Acknowledgments

The authors would like to thank Jonah Kanner and Alan Weinstein for insightful comments and suggestions regarding the method, as well as colleagues in the LIGO Scientific Collaboration Continuous Waves group for useful discussions. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. We are grateful for computational resources provided by Cardiff University and funded by an STFC grant supporting UK Involvement in the Operation of Advanced LIGO. MP is funded by the STFC under grant number ST/N005422/1. SM thanks the LIGO Visitor program. Plots produced using MATPLOTLIB [81].

Chapter 5

PROBING GRAVITATIONAL WAVE POLARIZATIONS WITH SIGNALS FROM COMPACT BINARY COALESCENCES

M. Isi and A. J. Weinstein, "Probing gravitational wave polarizations with signals from compact binary coalescences," (2017), M.I. carried out the analysis and wrote the manuscript, arXiv:1710.03794 [gr-qc] .

5.1 Introduction

The detection of gravitational waves (GWs) by the Advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) has enabled some of the first experimental studies of gravity in the highly dynamical and strong-field regimes [1–4, 82]. These first few detections have already been used to place some of the most stringent constraints on deviations from the general theory of relativity (GR) in this domain, which is inaccessible to laboratory, Solar System or cosmological tests of gravity.

However, it has not been possible to use LIGO signals to learn about the polarization content of GWs [82], a measurement highly relevant when comparing GR to many of its alternatives [83, 84]. In fact, all existing observations are so far consistent with the extreme case of purely non-GR polarizations. The reason for this is that the two LIGO instruments are nearly coaligned, meaning that they are sensitive to approximately the same linear combination of polarizations. This makes it nearly impossible to unequivocally characterize the polarization content of transient GW signals like the compact-binary coalescences (CBCs) observed so far, at least not without making assumptions about the way the signals were sourced [84, 85].

Existing observations that are usually taken to constrain the amount of allowed non-GR polarizations can do so only in an indirect manner. For example, measurements of the orbital decay of binary systems are sensitive to the total radiated GW power, but do not probe the geometric effect (namely, the directions in which space is stretched and squeezed) of the waves directly (see e.g. [86, 87], or [88, 89] for reviews). In the context of specific alternative theories (e.g. scalar-tensor) such observations can indeed constrain the power contained in extra polarizations. However, such measurements provide no direct, model-independent information on the actual

polarization content of the gravitational radiation. Thus, there may be multiple theories, with different polarization content, that still predict the correct observed GW emitted power.

To see that the above is the case, consider a scenario in which GWs are emitted precisely as in GR, but where the polarizations change during propagation: the phase evolution would be similar to GR, but the geometric effect of the wave would be completely different [90–93]. (This polarization mutation could take place if the linear polarization basis did not diagonalize the kinetic matrix of the theory, as is the case for neutrino oscillations [94, 95], or for the circular GW polarization states in dynamical Chern-Simons gravity [96].) Because the same limitations of pulsar binary analyses apply to studies of the details in the phasing of signals previously detected with LIGO and other traditional tests of GR (like Solar System tests) have no bearing on GWs, there currently exist no direct measurements of GW polarizations.

Prospects for the direct measurement of GW polarizations are improved by the addition of Advanced Virgo to the detector network. In principle, at least five noncoaligned differential-arm detectors would be needed to break *all* the degeneracies among the five nondegenerate polarizations allowed by generic metric theories of gravity [97, 98], if transient signals are used [18, 99]. However, as we will show, the current Advanced-LIGO–Advanced-Virgo network can already be used to distinguish between *some* of the possible combinations of polarizations without the need to use specific knowledge about the phase evolution of the source.

In this note, we present a simple Bayesian method to extract information about GW polarizations directly from strong CBC signals by using the relative amplitudes and timing at the different detectors.

5.2 Background

5.2.1 Polarizations

In all theories that respect Einstein's equivalence principle, including GR, gravitational interactions may be fully described via the universal coupling of matter to a metric tensor [83, 100]. Because of this, it may be shown that, in any such *metric theory*, a (nearly-)null plane GW may be encoded in at most six independent components of the Riemann tensor at any given point in spacetime [83, 97, 98]. These degrees of freedom give rise to six geometrically distinct polarizations, corresponding to the six linearly independent components of an arbitrary metric perturbation.

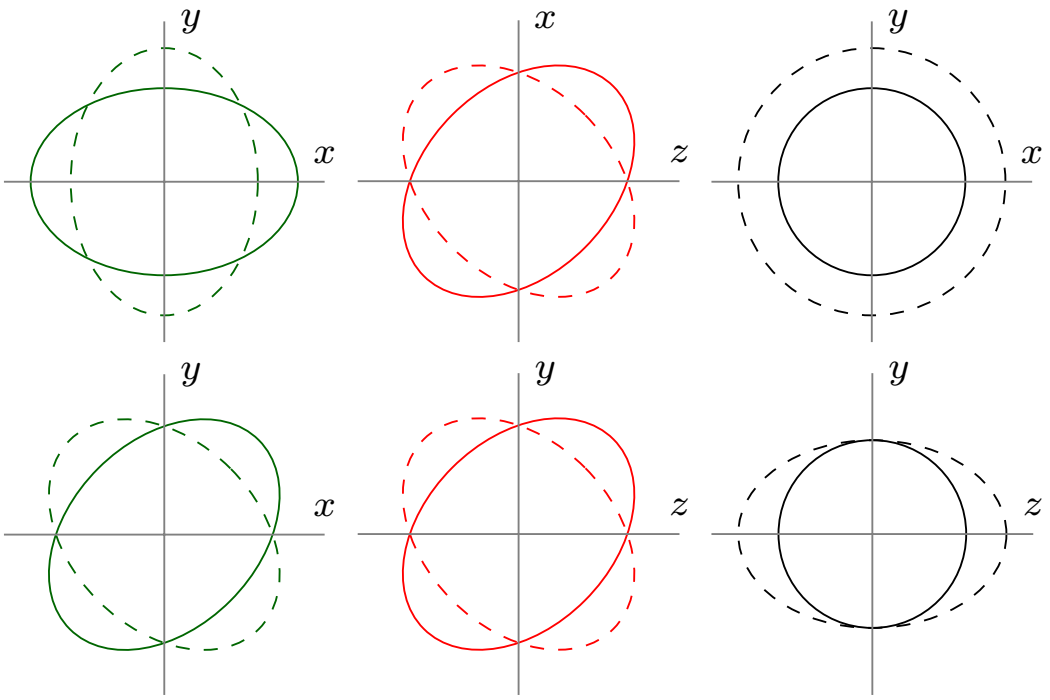

Figure 5.1: *Effect of different GW polarizations on a ring of free-falling test particles*. Plus (+) and cross (×) tensor modes (green); vector-x (x) and vector-y (y) modes (red); breathing (b) and longitudinal (l) scalar modes (black). In all of these diagrams the wave propagates in the z direction. This decomposition into polarizations was first proposed for generic metric theories in [98].

At any given spacetime point $\vec{x}$, the metric perturbation may thus be written as

$$h_{ab}(\vec{x}) = h_A(\vec{x})\, e^A{}_{ab}\,, \tag{5.1}$$

for six independent amplitudes, $h_A(\vec{x})$, and six polarization tensors $e^A{}_{ab}$ (implicit sum over polarizations A). For instance, letting $\mathbf{w}_z = \mathbf{w}_x \times \mathbf{w}_y$ be a spatial unit vector in the direction of propagation of the wave, we may consider the set of linear polarization tensors

$$\mathbf{e}^{+} = \mathbf{w}_x \otimes \mathbf{w}_x - \mathbf{w}_y \otimes \mathbf{w}_y\,, \tag{5.2}$$

$$\mathbf{e}^{\times} = \mathbf{w}_x \otimes \mathbf{w}_y + \mathbf{w}_y \otimes \mathbf{w}_x\,, \tag{5.3}$$

$$\mathbf{e}^{\mathrm{x}} = \mathbf{w}_x \otimes \mathbf{w}_z + \mathbf{w}_z \otimes \mathbf{w}_x\,, \tag{5.4}$$

$$\mathbf{e}^{\mathrm{y}} = \mathbf{w}_y \otimes \mathbf{w}_z + \mathbf{w}_z \otimes \mathbf{w}_y\,, \tag{5.5}$$

$$\mathbf{e}^{\mathrm{b}} = \mathbf{w}_x \otimes \mathbf{w}_x + \mathbf{w}_y \otimes \mathbf{w}_y\,, \tag{5.6}$$

$$\mathbf{e}^{\mathrm{l}} = \mathbf{w}_z \otimes \mathbf{w}_z\,. \tag{5.7}$$

Then Eq. (5.1) implies that there exists some gauge in which, in a local Lorentz frame with Cartesian coordinates along $(\mathbf{w}_x,\ \mathbf{w}_y,\ \mathbf{w}_z)$,

$$(h_{ij}) = \begin{pmatrix} h_{\rm b} + h_+ & h_\times & h_{\rm x} \\ h_\times & h_{\rm b} - h_+ & h_{\rm y} \\ h_{\rm x} & h_{\rm y} & h_{\rm l} \end{pmatrix}, \tag{5.8}$$

where the h_A's represent the amplitudes of the linear polarizations: plus (+), cross (×), vector x (x), vector y (y), breathing (b) and longitudinal (l). The effect of each of these modes on a ring of freely-falling particles is represented in Fig. 5.1.

Polarizations may be characterized by their behavior under Lorentz transformations, and different theories may be classified according to the polarizations they allow, as seen by different observers; this is known as the E(2) or *Eardley* classification [97, 98]. Five of these degrees of freedom (traceless and divergencessless) correspond to the helicity ±2, helicity ±1 and helicity 0 states of a massive spin-2 particle (the graviton); the remaining mode corresponds to a ghost-like degree of freedom associated with the trace of the perturbation. The correspondence between geometric (Eardly's classification) and field-theoretic (Wigner's classification) language is, however, limited because the the E(2) classification is only semi-Lorentz-invariant (although it is usually taken to hold, at least in the weak field regime) [98].

Einstein's theory only allows for the existence of linear combinations of the tensor + and × polarizations [83]. On the other hand, scalar-tensor theories famously predict the presence of some breathing component associated with the theory's extra scalar field [101], as do some theories with extra dimensions [102]. On top of tensor and scalar modes, bimetric theories, like Rosen or Lightman-Lee theories, may also predict vector modes [85, 103, 104]. The same is true in general for massive-graviton frameworks [105]. Furthermore, less conventional theories might, in principle, predict the existence of vector or scalar modes *only*, while still possibly being in agreement with all other non-GW tests of GR (see e.g. [106], for an unconventional example).

5.2.2 Antenna patterns

Because different polarizations have geometrically distinct effects, as illustrated in Fig. 5.1, GW detectors will react differently to each mode. The *strain* produced by a GW metric perturbation h_{ab} on certain detector I spatially located at $\mathbf{x}_I$, is given by

$$h_I(t) = D_I^{ab} h_{ab}(t, \mathbf{x}_I) = h_A(t, \mathbf{x}_I) D_I^{ab} e^A_{ab}. \tag{5.9}$$

The detector tensor, D^{ab}, encodes the geometry of the instrument and the measurement it makes; for differential-arm detectors (sometimes called *quadrupolar antennas*, because of the symetries of their angular response functions, cf. Fig. 5.2), like LIGO and Virgo, this is

$$D^{ab} = \frac{1}{2}\left(d_x{}^a d_x{}^b - d_y{}^a d_y{}^b\right), \tag{5.10}$$

where $\mathbf{d}_x$ and $\mathbf{d}_y$ are spatial unit vectors along the detector arms (with common origin at the vertex $\mathbf{x}_I$). Although D^{ab} is technically also a function of time due to the motion of Earth with respect to the fixed stars, in practice it can be taken as constant when treating short-lived CBC signals, as is done here.

The $h_A(t)$'s are determined by a nontrivial combination of the source dynamics, the details of the matter-gravity coupling, and the vacuum structure of the theory. However, the response (*antenna pattern*) of detector I to polarization A,

$$F^A \equiv D_I^{ab} e_{ab}^A, \tag{5.11}$$

depends *only* on the local geometry of the gravitational wave and the detector, irrespective of the properties of the source. This decoupling makes the antenna patterns a unique resource for studying GW polarizations directly.

The response functions, Eq. (5.11), encode the effect of a linearly A-polarized GW with unit amplitude, $h_A = 1$. Ground-based GW detectors, like LIGO and Virgo are quadrupolar antennas that perform low-noise measurements of the strain associated with the differential motion of two orthogonal arms. Their detector response functions can thus be written as [107–110]:

$$F_+ = \frac{1}{2}\left[(\mathbf{w}_x \cdot \mathbf{d}_x)^2 - (\mathbf{w}_x \cdot \mathbf{d}_y)^2 - (\mathbf{w}_y \cdot \mathbf{d}_x)^2 + (\mathbf{w}_y \cdot \mathbf{d}_y)^2\right], \tag{5.12}$$

$$F_\times = (\mathbf{w}_x \cdot \mathbf{d}_x)(\mathbf{w}_y \cdot \mathbf{d}_x) - (\mathbf{w}_x \cdot \mathbf{d}_y)(\mathbf{w}_y \cdot \mathbf{d}_y), \tag{5.13}$$

$$F_{\mathrm{x}} = (\mathbf{w}_x \cdot \mathbf{d}_x)(\mathbf{w}_z \cdot \mathbf{d}_x) - (\mathbf{w}_x \cdot \mathbf{d}_y)(\mathbf{w}_z \cdot \mathbf{d}_y), \tag{5.14}$$

$$F_{\mathrm{y}} = (\mathbf{w}_y \cdot \mathbf{d}_x)(\mathbf{w}_z \cdot \mathbf{d}_x) - (\mathbf{w}_y \cdot \mathbf{d}_y)(\mathbf{w}_z \cdot \mathbf{d}_y), \tag{5.15}$$

$$F_{\mathrm{b}} = \frac{1}{2}\left[(\mathbf{w}_x \cdot \mathbf{d}_x)^2 - (\mathbf{w}_x \cdot \mathbf{d}_y)^2 + (\mathbf{w}_y \cdot \mathbf{d}_x)^2 - (\mathbf{w}_y \cdot \mathbf{d}_y)^2\right], \tag{5.16}$$

$$F_{\mathrm{l}} = \frac{1}{2}\left[(\mathbf{w}_z \cdot \mathbf{d}_x)^2 - (\mathbf{w}_z \cdot \mathbf{d}_y)^2\right]. \tag{5.17}$$

Here, as before, the spatial vectors $\mathbf{d}_x$, $\mathbf{d}_y$ have unit norm and point along the detector arms such that $\mathbf{d}_z = \mathbf{d}_x \times \mathbf{d}_y$ is the local zenith; the direction of propagation of the wave from a source at known sky location (specified by right ascension α, and

declination δ) is given by $\mathbf{w}_z$, and $\mathbf{w}_x$, $\mathbf{w}_y$ are such that $\mathbf{w}_z = \mathbf{w}_x \times \mathbf{w}_y$. We choose $\mathbf{w}_x$ to lie along the intersection of the equatorial plane of the source with the plane of the sky, and let the angle between $\mathbf{w}_y$ and the celestial north be ψ, the *polarization angle*.

Because of their symmetries, the breathing and longitudinal modes are fully degenerate to networks of quadrupolar antennas (see e.g. Sec. VI of [85]). This means that no model-independent measurement with such a network can possibly distinguish between the two, so it is enough for us to consider just one of them explicitly; we will refer to the scalar modes jointly by the subscript "s". (This degeneracy may not be present for detectors with different geometries [111, 112].)

The response of a given differential-arm detector to signals of certain linear polarization and direction of propagation can be written, in the local Lorentz frame of the detector itself, as [see e.g. Eqs. (13.98) in [110] with $\psi \to -\psi - \pi/2$, to account for the different wave-frame definition]:

$$F_+(\vartheta, \varphi, \psi) = -\frac{1}{2}\left(1 + \cos^2\vartheta\right)\cos 2\varphi \cos 2\psi - \cos\vartheta \sin 2\varphi \sin 2\psi \,, \tag{5.18}$$

$$F_\times(\vartheta, \varphi, \psi) = \frac{1}{2}\left(1 + \cos^2\vartheta\right)\cos 2\varphi \sin 2\psi - \cos\vartheta \sin 2\varphi \cos 2\psi \,, \tag{5.19}$$

$$F_{\mathrm{x}}(\vartheta, \varphi, \psi) = -\sin\vartheta \sin 2\varphi \cos\psi + \sin\vartheta \cos\vartheta \cos 2\varphi \sin\psi \,, \tag{5.20}$$

$$F_{\mathrm{y}}(\vartheta, \varphi, \psi) = \sin\vartheta \sin 2\varphi \sin\psi + \sin\vartheta \cos\vartheta \cos 2\varphi \cos\psi \,, \tag{5.21}$$

$$F_{\mathrm{b/l}}(\vartheta, \varphi, \psi) = \mp\frac{1}{2}\sin^2\vartheta \cos 2\varphi \,, \tag{5.22}$$

where ϑ and φ are the polar an azimuthal coordinates of the source with respect to the antenna at any given time (with detector arms along the x and y-axes). The tensor, vector and scalar nature of the different polarizations is evident in this form, given how each mode depends on ψ (i.e. how it transforms under rotations around the direction of propagation).

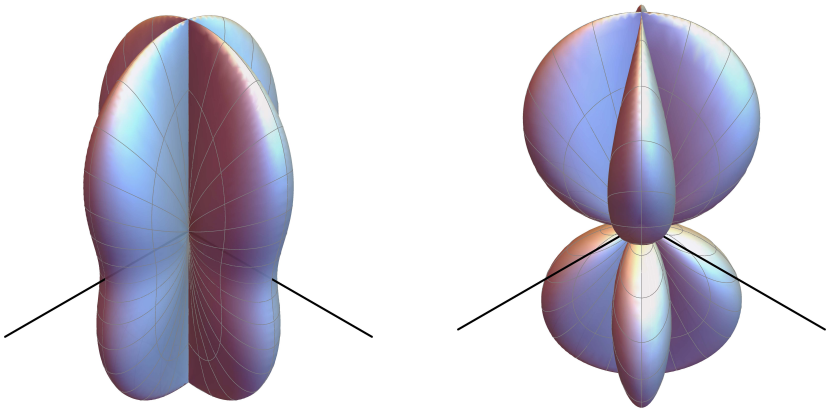

(a) Plus (+)

(b) Cross (×)

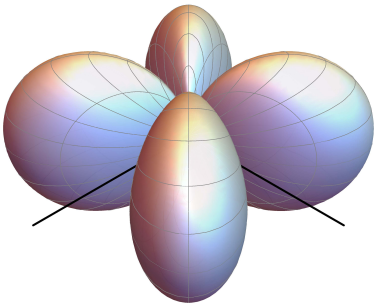

(c) Vector-x (x)

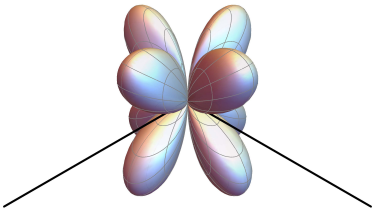

(d) Vector-y (y)

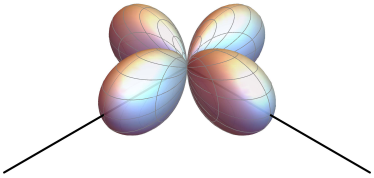

(e) Scalar (s)

Figure 5.2: *Angular response of a quadrupolar detector to each GW polarization*. The radial distance represents the response of a single quadrupolar antenna to a unit-amplitude gravitational signal of a tensor (top), vector (middle), or scalar (bottom) polarization, i.e. $|F_A|$ for each polarization A as given by Eqs. (5.18–5.22) for $\psi = 0$. The polar and azimuthal coordinates correspond to the source location with respect to the detector, which is to be imagined as placed with its vertex at the center of each plot and arms along the x and y-axes. The response is plotted to scale, such that the black lines representing the detector arms have unit length in all plots. The response to breathing and longitudinal modes is identical, so we only display it once and label it "scalar". (Reproduced from [18].)

Equations (5.18)–(5.22) are represented in Fig. 5.2 by a spherical polar plot in which the radial coordinate corresponds to the sensitivity given by the magnitude $|F_A|$, shown for $\psi = 0$. The angular response functions have quadrupolar symmetry around the detector's zenith, regardless of the helicitiy of the polarization itself. This figure also makes it clear that differential-arm detectors will generally be more sensitive to some polarizations than others, although this will vary with the sky location of the source. For example, for all but a few sky locations, quadrupolar antennas will respond significantly less to a breathing signal than a plus or cross signal.

Fig. 5.2 shows the response of a single differential-arm detector to waves coming from different directions in the local frame of the instrument. However, we are usually interested in the sensitivity of a *network* of detectors and its ability to distinguish the different polarizations. To visualize this, define the effective response to each of the helicities, for a given source sky-location $(\alpha,\ \delta)$ and detector I:

$$|F^I_{\mathrm{t}}(\alpha,\delta)| \equiv \sqrt{F^I_{+}(\alpha,\delta)^2 + F^I_{\times}(\alpha,\delta)^2}\,, \tag{5.23}$$

$$|F^I_{\mathrm{v}}(\alpha,\delta)| \equiv \sqrt{F^I_{\mathrm{x}}(\alpha,\delta)^2 + F^I_{\mathrm{y}}(\alpha,\delta)^2}\,, \tag{5.24}$$

$$\begin{aligned} |F^I_{\mathrm{s}}(\alpha,\delta)| &\equiv \sqrt{F^I_{\mathrm{b}}(\alpha,\delta)^2 + F^I_{\mathrm{l}}(\alpha,\delta)^2} \\ &= \sqrt{2}\,|F^I_{\mathrm{b}}(\alpha,\delta)|\,, \end{aligned} \tag{5.25}$$

for tensor, vector and scalar waves respectively. (Here, since we are not dealing with any specific source, we *define* our polarization frame letting $\psi = 0$.) For a network of N detectors, we may then construct an effective response vector for each of the polarization sets above,

$$\vec{F}_H(\alpha,\delta) \equiv \Big(|F^1_H(\alpha,\delta)|,\ \ldots, |F^N_H(\alpha,\delta)|\Big)\,, \tag{5.26}$$

for $H \in \{\mathrm{t},\ \mathrm{v},\ \mathrm{s}\}$. Finally, we may compare the overall sensitivity of the network to different polarizations by defining the *overlap*, as a normalized inner product between two of these vectors.

For instance, to compare the effective scalar or vector network sensitivity to the tensor one, we may look at the overlap factor:

$$\mathcal{F}_{H/\mathrm{t}}(\alpha,\delta) = \frac{\vec{F}_H(\alpha,\delta)\cdot\vec{F}_{\mathrm{t}}(\alpha,\delta)}{\vec{F}_{\mathrm{t}}(\alpha,\delta)\cdot\vec{F}_{\mathrm{t}}(\alpha,\delta)}, \tag{5.27}$$

which will take values greater (less) than unity if the response to polarizations H is better (worse) than to tensor, with $\mathcal{F}_{\mathrm{t}/\mathrm{t}}(\alpha,\delta) = 1$ by construction. The scalar and

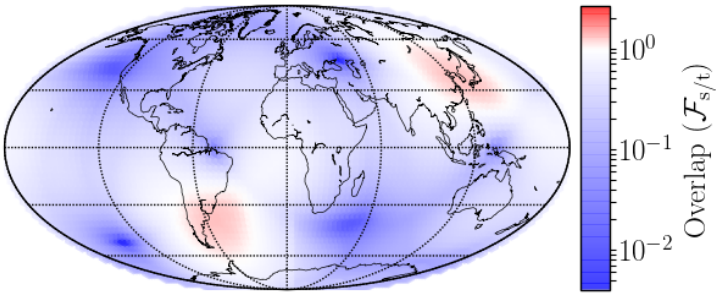

(a) Scalar

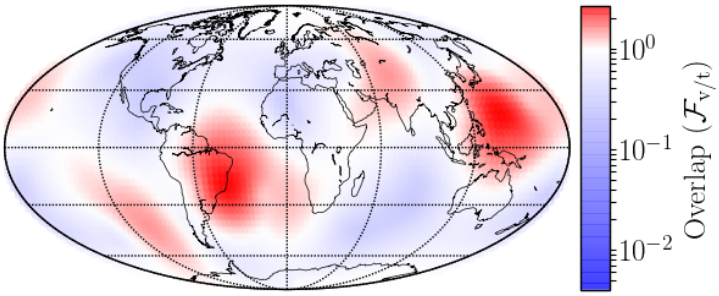

(b) Vector

Figure 5.3: *Overlaps of LIGO-Virgo network effective antenna patterns.* The normalized inner-products of Eq. (5.27) for the three-instrument network. The top plot compares scalar to tensor ($\mathcal{F}_{s/t}$), and the bottom one compares vector to tensor ($\mathcal{F}_{v/t}$). Blue (red) marks regions for which the effective nontentor response is greater (less) than tensor. A map of Earth is overlaid for reference.

vector overlaps with tensor are displayed for the LIGO-Virgo network in the skymap of Fig. 5.3, over a map of Earth for reference. Colored regions roughly correspond to areas in the sky for which the tensor and nontensor responses of the network are highly distinguishable. The patterns are anchored to angular locations with respect to Earth (not the fixed stars) and are determined by the specific location and orientation of the three detectors.

Averaged over all sky locations, the response of the network is worse for scalar signals than tensor ones, which is apparent from the top skymap in Fig. 5.3 and the distribution in Fig. 5.4. This is expected given that each interferometer is individually less sensitive to scalar waves, as seen in Fig. 5.2. On average, there is no significant difference between vector and tensor responses.

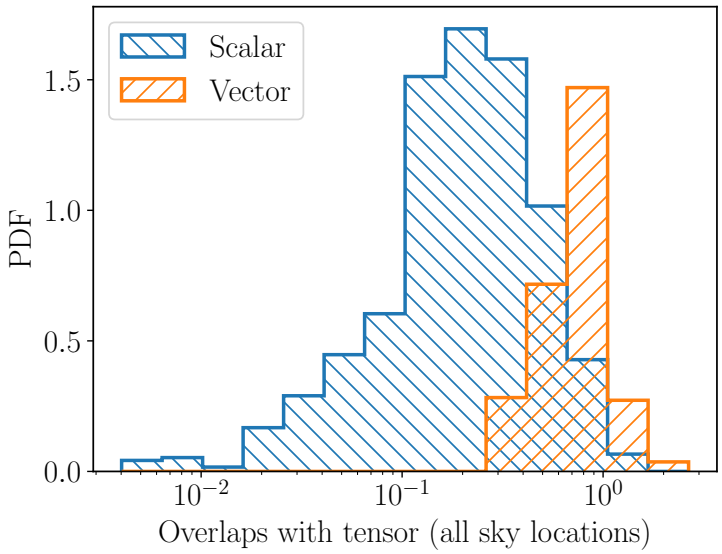

Figure 5.4: *Overlaps of LIGO-Virgo network effective antenna patterns.* The normalized inner-products of Eq. (5.27) for the three-instrument network. The blue (orange) histogram corresponds to the values shown in the top (bottom) panel of Fig. 5.3.

5.3 Method

Ideally, we would like to unequivocally measure the polarizations of the GW that produced a given transient strain signal in our detector network. Formally, this would mean finding which of the seven possible Bayesian hypotheses the data favor: pure tensor ($\mathcal{H}_{\rm t}$), pure vector ($\mathcal{H}_{\rm v}$), pure scalar ($\mathcal{H}_{\rm s}$), scalar-tensor ($\mathcal{H}_{\rm st}$), vector-tensor ($\mathcal{H}_{\rm tv}$), scalar-vector ($\mathcal{H}_{\rm sv}$), or scalar-vector-tensor ($\mathcal{H}_{\rm stv}$). A comprehensive Bayesian treatment of this polarization model-selection problem was presented in [18] for the case of continuous signals from known pulsars, later applied to stochastic GW backgrounds in [99], and could be easily by adapted to the case of transient signals considered here.

Yet, a simple counting argument is enough to show that three detectors are not sufficient to break *all* degeneracies between the five distinguishable GW polarizations using transient signals [83, 85]. Therefore, with the current LIGO-Virgo network, we expect the results of an all-encompassing model-selection analysis, as discussed above, to be inconclusive or dominated by priors. Nevertheless, we may still attempt to distinguish between *some* of the possible hypotheses.

As mentioned in the introduction, *all LIGO-only observations so far are consistent with the extreme scenario of GWs being composed of purely vector or purely scalar*

polarizations. Therefore, here we will focus on the problem of directly distinguishing between these theoretically far-fetched, yet phenomenologically valid, possibilities. That is, we will study our ability to choose between $\mathcal{H}_{\rm s}$ vs $\mathcal{H}_{\rm t}$, and between $\mathcal{H}_{\rm v}$ vs $\mathcal{H}_{\rm t}$. Importantly, this is qualitatively distinct from the more standard question about the presence of small nontensorial components in addition to the tensor wave predicted by GR. Although perhaps not as interesting as these "mixed" polarization studies (which, as explained above, will not fully succeed with current detectors), the problem of distinguishing between the "pure" polarization cases is well-defined and experimentally valuable.

We would like to ask the question: *is it geometrically possible that a given strain signal observed in the LIGO-Virgo network was produced by a GW with polarization other than GR's tensor + and ×?* The only way for us to answer this question is to probe the antenna patterns of our instruments, Eq. (5.11), which are a direct manifestation of local geometry only (polarizations and detector geometry), independent of source or the details of the underlying theory (see Sec. 5.2.2). We may thus exploit the difference in the response of the network to the different polarizations (Fig. 5.3).

One way to extract polarization information using the antenna patterns would be to construct linear combinations of the detector outputs that are guaranteed to contain no tensorial signal [85]. If coherent power (as seen by, e.g. a wavelet analysis) remains in such a *null-stream*, then that signal could not have been produced by a tensor (GR) wave. This approach has the strong advantage that it requires no knowledge of the spectral features of the signal whatsoever. However, to construct null-streams one needs to very accurately know the location of the source *a priori*, which is never the case without an electromagnetic counterpart (or more detectors).

Alternatively, one could carry out a morphology-independent sine-Gaussian analysis (e.g. using BayesWave [113, 114]) to reconstruct the best-fit unmodeled waveform from the data, and use that to extract information about times of arrival, phase offsets and relative amplitudes at different detectors. One could then just replace the tensor antenna patterns used in the signal reconstruction by their scalar or vector counterparts and see how well each case fits the data (as measured by a Bayes factor). In such test, *no polarization information is extracted from the phase evolution*. In particular, the waveform reconstruction is only used to infer the source location from the time lag between detectors, and the best-fitting combination of antenna patterns from the amplitudes and phases at peak energy. (See pedagogical example in Sec. 5.3.1 below.) An analysis like this was implemented for scalar modes and applied to

the GW150914 signal, yielding no conclusive results as mentioned above [82].

However, all signals observed by LIGO so far are exceptionally well described by GR CBC waveforms [2–4, 82]. This match is established on a case-by-case basis through comparisons between the GR templates and morphology independent burst reconstructions of the signal in the data and is largely independent of the polarization. In fact, for any of these confident detections, the waveform reconstructed from burst analyses is effectively identical to a GR template. As emphasized above, in the pure-polarization test ($\mathcal{H}_s$ vs $\mathcal{H}_t$, or $\mathcal{H}_v$ vs $\mathcal{H}_t$) all that matters is that most of the signal power is captured by the template, regardless of small potential mismatches in the phasing. Therefore, we may carry out the same study proposed in the previous paragraph using GR waveforms to fit the data, while replacing the tensor antenna patterns with those of different polarizations.

In other words, when the signal is clearly well-captured by a GR template, we may use that directly to extract polarization information from the antenna patterns in a model independent way, without implicitly assuming that the GW that caused it was tensor polarized as GR predicts. The waveform reconstruction will be dominated by the measurement at the most sensitive detector, while the amplitude information is encoded in the relations between measurements by different detectors.

Whether we use GR templates or a collection of sine-Gaussians to reconstruct the waveform, the effect of changing the antenna patterns will always result in different inferred sky location and orientation for the source. Yet, not all antenna patterns will be equally consistent with the observed relative amplitudes, phase offsets and delays between the signals in our three detectors—this will result in a poorer signal likelihood, and hence odds favoring tensor vs nontensor. Precisely because the waveform used to capture the signal is the same, we know that any difference between the tensor and nontensor results *must* come from the antenna patterns (polarizations).

This approach does not extract any information from the specific phase evolution of the signal and is insensitive to small changes in the waveform. Therefore, using a GR template to measure the signal power is justified, and does not imply a contradiction when testing for nontensorial polarizations. For the purpose of this study, the CBC signal is just probing the impulse response function of our network, and the same results would be obtained if the waveform were just a Delta function rather than a chirp.

5.3.1 Toy example

For concreteness, consider the example of an elliptically-polarized, two-component GW (e.g. two tensor modes or two vector modes) with waveform roughly described by a simple sine-Gaussian wavepacket, with some characteristic frequency Ω and relaxation time τ. Letting t be the time measured at Earth's center, then the strain measured by a given detector I will be:

$$h_I(t) = \Re\left[A\left(F_1^I + i\epsilon F_2^I\right)e^{i\Omega(t-t_0-\delta t_I)}\right]e^{-(t-t_0-\delta t_I)^2/\tau^2}, \tag{5.28}$$

where F_1^I and F_2^I are the responses of detector I to the two polarizations, $A \equiv |A|e^{i\phi_0}$ is a complex-valued amplitude, ϵ is an ellipticity parameter controlling the relative amounts of each polarization, and $\Re$ denotes the real part. Also, t_0 marks the time of arrival at Earth's center, which is delayed with respect to each interferometer by

$$\delta t_I = \hat{\mathbf{n}} \cdot \mathbf{x}_I / c\,, \tag{5.29}$$

where $\hat{\mathbf{n}}$ is a unit vector from Earth to the source, and $\mathbf{x}_I$ joins Earth's center to the detector (with magnitude equal to Earth's radius). Here we are assuming that the GW travels at the speed of light, c.

The signal of Eq. (5.28) may be written more simply as

$$h_I(t) = \mathcal{A}_I \cos[\Omega(t-\Delta t_I) + \Phi_I]\, e^{-(t-\Delta t_I)^2/\tau^2}\,, \tag{5.30}$$

after defining the three main observables at each detector:

$$\mathcal{A}_I \equiv |A|\left|F_1^I + i\epsilon F_2^I\right|\,, \tag{5.31}$$

$$\Phi_I \equiv \phi_0 + \arctan(\epsilon F_2^I / F_1^I)\,, \tag{5.32}$$

$$\Delta t_I \equiv t_0 + \delta t_I \tag{5.33}$$

From the output of three detectors (H, L, V), we may implement a simple inference analysis to extract these three numbers for the signal as seen by each instrument. The times at peak amplitude provide the three Δt_I's, while measurements of the phase and amplitude at peak itself give the Φ's and $\mathcal{A}_I$'s respectively. As always, recovery of all these parameters will be negatively affected by instrumental noise.

The three timing measurements alone suffice to recover the sky location of the source, $\hat{\mathbf{n}}$. With this knowledge, it is then possible to compute the values of all the corresponding antenna response functions, and thus obtain predictions for the

$(F_1^I + i\epsilon F_2^I)$ factors for any given ellipticity. Ratios of amplitudes and phase differences between detectors may then be used to infer measured values for these quantities, and to then find the best fitting polarization model. This may be achieved, for instance, via a maximum-likelihood analysis, effectively minimizing the distance between vectors like those of Eq. (5.26) and a similar one inferred from the data. (Note $|A|$, ϕ_0, and ϵ are nuisance parameters, and can be marginalized over.)

Although for this example we used a simple sine-Gaussian wavepacket to measure the signal, at no point did we make use of the specific details of this phase evolution. The only requirement is that the GW have a well-defined peak, in order to extract meaningful information about how the relative timing, phase and amplitude of this peak as seen by different detectors. In particular, this analysis would work precisely the same way if CBC-like chirp waveform were used, as long as most of the power in the actual signal was indeed captured by such a template.

This toy analysis makes the dependence on A_I, Φ_I, Δt_I explicit. In reality, when studying actual data, one would ideally implement a full Bayesian analysis, marginalizing over all parameters to compute evidences for the different polarization hypotheses ($\mathcal{H}_{\rm t}$, $\mathcal{H}_{\rm v}$, $\mathcal{H}_{\rm s}$) and corresponding Bayes factors (likelihood ratios). This can be achieved using a code like LALInference [20]. The polarization information extracted by this more rigorous analysis would still, nonetheless, effectively come from the values of A_I, Φ_I, Δt_I. As emphasized before, this is the case whether one uses GR templates or a collection of sine-Gaussians to capture the signal power.

5.4 Conclusion

By extracting polarization information from the antenna patterns we may directly probe the geometry of the GW metric perturbation (i.e. the directions along which space is stretched and squeezed by the passing wave) from its projection onto our detector network. With transient signals, instruments at five or more different orientations would be needed to break all degeneracies between the five independent (as seen by differential-arm detectors) polarizations allowed by generic metric theories of gravity. However, we may already distinguish between some of the possibilities using the current LIGO-Virgo network. How well we can do this will depend on the specific properties of each transient event (mainly, sky location).

The kind of geometric observational statement discussed in this note is independent of any theory or source model and is only possible with the addition of Virgo to the network. Although here we focused on the problem of distinguishing between

"pure" polarization states (tensor, vector or scalar), the case of "mixed" polarizations will be addressed in future work. More details and a demonstration of the analysis proposed here on simulated signals will be provided soon in an expanded version of this document.

Acknowledgments

LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058.

Chapter 6

CONSTRAINTS ON GRAVITATIONAL-WAVE POLARIZATIONS FROM COMPACT-BINARY COALESCENCES

B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **119**, 141101 (2017), M.I. carried out the polarization analysis, wrote the section describing it and provided feedback on the manuscript as a whole, arXiv:1709.09660 [gr-qc] .

B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), "Tests of General Relativity with GW170817," (2018), M.I. contributed the polarization analysis and provided feedback on the text, arXiv:1811.00364 [gr-qc] .

6.1 Introduction

Generic metric theories of gravity may predict up to six polarization modes for metric perturbations: two tensor (helicity ±2), two vector (helicity ±1), and two scalar (helicity 0) modes [97, 98]. In contrast, one of the key predictions of general relativity (GR) is that metric perturbations possess only two tensor degrees of freedom [83, 115]. Therefore, a detection of any nontensorial mode would be unambiguous indication of physics beyond GR.

Before the beginning of LIGO's second observation run, some evidence that gravitational waves (GWs) are described by the tensor metric perturbations of GR had been obtained from measurements of the rate of orbital decay of binary pulsars, in the context of specific beyond-GR theories (see e.g. [86, 87], or [88, 89] for reviews), and from the rapidly changing GW phase of binary black-hole mergers observed by LIGO, in the framework of parameterized models [1, 4, 82]. The addition of Advanced Virgo to the network in 2017 enabled another, more compelling, way of probing the nature of polarizations by studying GW geometry directly through the projection of the metric perturbation onto our detector network [18, 85, 99, 116].

The GW strain measured by a detector can be written in general as $h(t) = F^A h_A$, where h_A are the 6 independent polarization modes and F^A represent the detector responses to the different modes $A = (+, \times, \mathrm{x}, \mathrm{y}, \mathrm{b}, \mathrm{l})$. The antenna response functions depend only on the detector orientation and GW helicity, i.e. they are independent of the

Table 6.1: Polarization constraints from compact binaries (LIGO and Virgo).

	GW170814	GW170817	GW170817 (fixed sky)
$\log_{10}\mathcal{B}$(tensor vs vector)	2.5 ± 0.2	0.72 ± 0.09	20.81 ± 0.08
$\log_{10}\mathcal{B}$(tensor vs scalar)	3.1 ± 0.2	5.84 ± 0.09	23.09 ± 0.08

intrinsic properties of the source. We can therefore place bounds on the polarization content of a given GW by studying which combination of response functions is consistent with the signal observed [18, 79, 85, 99, 116, 117]. However, in order to break degeneracies and constrain all distinguishable polarization combinations with short compact-binary coalescences (CBCs), we would need at least 5 detectors, even if the location of the source was known *a priori*—a network with fewer detectors can only make partial statements about polarizations (see Chapter 5). Furthermore, as the two LIGO instruments have similar orientations, little information about polarizations can be obtained using the LIGO detectors alone.

The first test on the polarization of GWs was performed for GW150914 [82]. The number of GR polarization modes expected was equal to the number of detectors in the network that observed the signal, rendering this test inconclusive. The addition of Virgo to the network of GW detectors allowed for the first informative test of polarization for GW170814 [5]. As described in Sec. 6.2, this analysis established that the GW data was better described by pure tensor modes than pure vector or pure scalar modes. A similar analysis was carried out for GW170817, the binary neutron star detected shortly after [6]. In that case, and as described in Sec. 6.3, the identification of an electromagnetic counterpart allowed for much stronger results, effectively rejecting the two pure-nontensorial models to any reasonable doubt [118]. Results for both events are summarized in Table 6.1.

6.2 GW170814

On August 14, 2017, GWs from the coalescence of two black holes at a luminosity distance of 540^{+130}_{-210} Mpc, with masses of $30.5^{+5.7}_{-3.0}\,\mathrm{M}_\odot$ and $25.3^{+2.8}_{-4.2}\,\mathrm{M}_\odot$, were observed in all three LIGO and Virgo detectors. The signal was first observed at the LIGO Livingston detector at 10:30:43 UTC, and then at the LIGO Hanford and Virgo detectors with delays of ∼8 ms and ∼14 ms, respectively.

The signal-to-noise ratio (SNR) time series, the time-frequency representation of the strain data and the time series data of the three detectors together with the inferred GW waveform, are shown in Fig. 6.1. The different sensitivities and responses of the

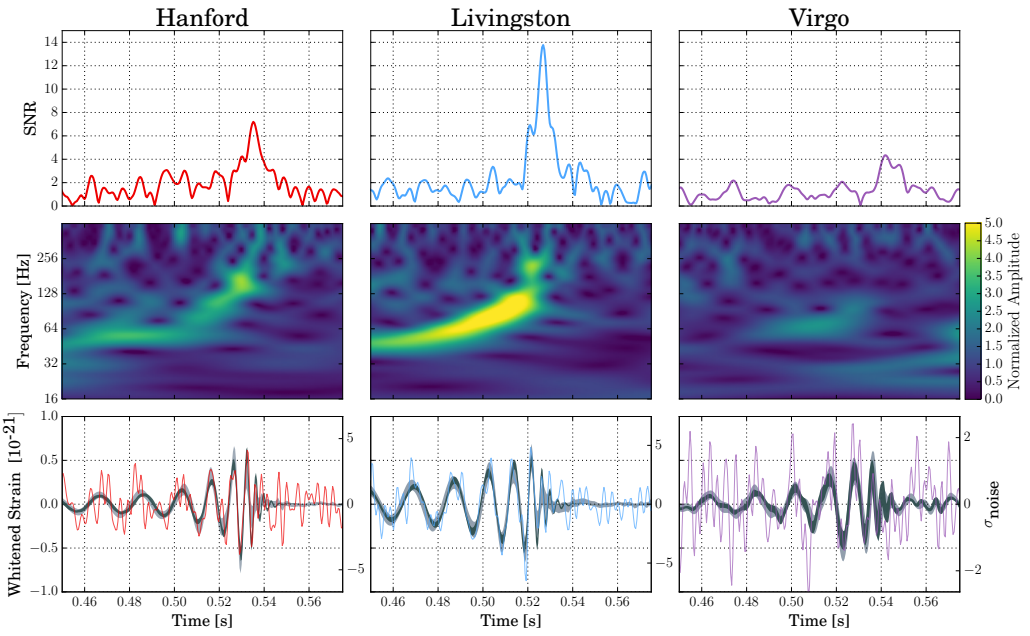

Figure 6.1: *GW170814 as observed by LIGO Hanford, LIGO Livingston and Virgo*. Times are shown from August 14, 2017, 10:30:43 UTC. *Top row*: SNR time series produced in low latency. The time series were produced by time-shifting the best-match template from the online analysis and computing the integrated SNR at each point in time. The single-detector SNRs in Hanford, Livingston and Virgo are 7.3, 13.7 and 4.4, respectively. *Second row*: Time-frequency representation of the strain data around the time of GW170814. *Bottom row*: Time-domain detector data (in color), and 90% credible intervals for waveforms reconstructed from a wavelet analysis [25] (light gray) and GR templated models (dark gray), whitened by each instrument's noise.(Reproduced from [5], see that Ref. for details on data manipulation for this figure.)

three detectors result in the GW producing different values of matched-filter SNR in each detector. A measurement of the signal amplitude by multiple (non-coaligned) detectors is the first ingredient needed to study polarizations.

Until Advanced Virgo became operational, typical GW position estimates were highly uncertain compared to the fields of view of most telescopes. The baseline formed by the two LIGO detectors allowed us to localize most mergers to roughly annular regions spanning hundreds to about a thousand square degrees at the 90% credible level [119–121]. Virgo adds additional independent baselines, which in cases such as GW170814 can reduce the positional uncertainty by an order of magnitude or more [120]. Good sky localization is the second ingredient needed to study polarizations.

With the addition of Advanced Virgo we can probe, for the first time, gravitational-

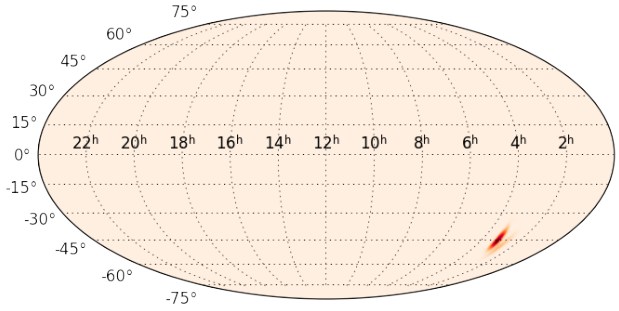

(a) Tensor

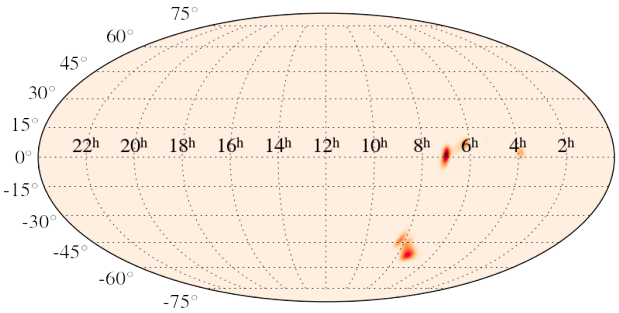

(b) Vector

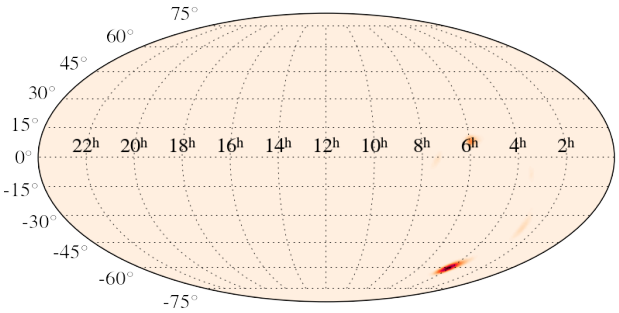

(c) Scalar

Figure 6.2: *GW170814 skymaps.* Reconstructed sky location under the assumption of the three different polarization hypotheses. Color represents probability density, as a function of equatorial coordinates in a Mollweide projection.

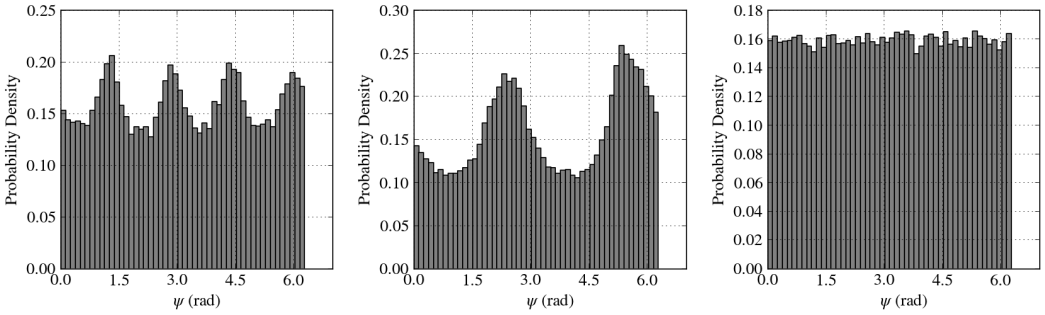

(a) Tensor (b) Vector (c) Scalar

Figure 6.3: *GW170814 polarization-angle posteriors.* Posterior probability density in the polarization angle under the three different polarization assumptions. Note the periodicities corresponding to the different helicities.

wave polarizations geometrically by projecting the wave's amplitude onto the three detectors. To do this, the coherent Bayesian analysis used to infer signal parameters is repeated after replacing the standard tensor antenna response functions with those appropriate for scalar or vector polarizations [116]. In our analysis, we are interested in the geometric projection of the GW onto the detector network, and therefore the details of the phase model itself are less relevant as long it is a faithful representation of the fit to the data in Fig. 6.1. Hence, we assume a GR phase model.

We find Bayes factors of more than 200 and 1000 in favor of the purely tensor polarization against purely vector and purely scalar, respectively. We also find that, as expected, the reconstructed sky location, distance and orientation change significantly depending on the polarization content of the source, with non-overlapping 90% credible regions for tensor, vector and scalar (Fig. 6.2). In particular, the posterior on the polarization angle ψ reveals the symmetries intrinsic to each helicity: periodic over $\pi/2$ for tensor, periodic over π for vector, and totally insensitive to changes in ψ for scalar (Fig. 6.3). The inferred detector-frame masses and spins are always the same, because that information is encoded in the signal phasing. An example of this is shown in Fig. 6.4 for the "chirp" mass, $\mathsf{M} = (m_1 m_2)^{3/5}(m_1 + m_2)^{-1/5}$. The most probable waveforms recovered under different hypotheses confirm that Virgo data is the key differentiating factor (Fig. 6.5).

6.3 GW170817

On August 17, 2017 at 12:41:04 UTC, the Advanced LIGO and Advanced Virgo gravitational-wave (GW) detectors made their first observation of a binary neutron

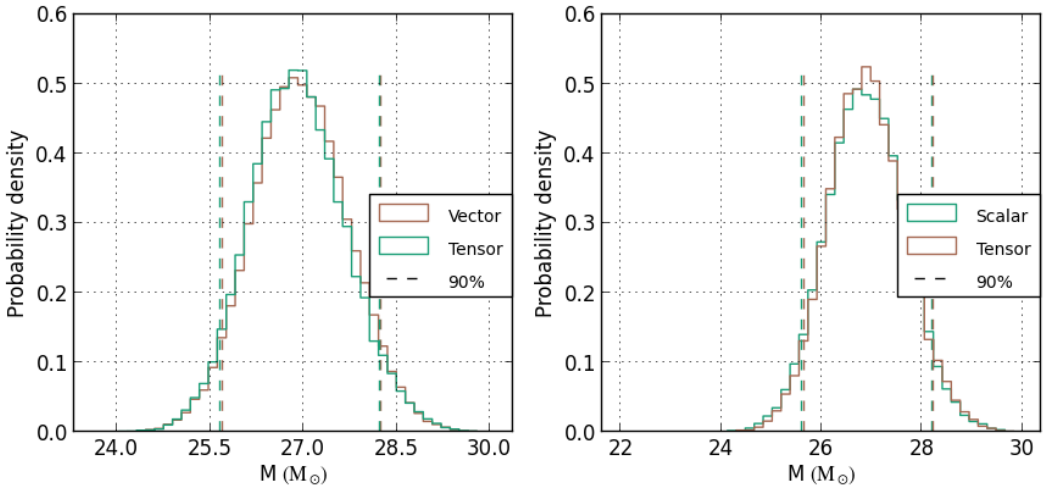

(a) Vector and tensor

(b) Scalar and tensor

Figure 6.4: *GW170814 chirp-mass comparison.* Posterior probability densities recovered for the detector-frame chirp mass $\mathsf{M} = (m_1 m_2)^{3/5}(m_1 + m_2)^{-1/5}$, where m_1 and m_2 are the two component masses. The different traces correspond to different polarization hypotheses: on the left, vector (brown) and tensor (green); on the right scalar (green) and tensor (brown). They are all equivalent.

star inspiral signal (GW170817) [6]. A representation of their data is given in Fig. 6.6. Associated with this event, a gamma ray burst [123] was independently observed, and an optical counterpart was later discovered [124]. The source was successively associated with the galaxy NGC4993.

In terms of fundamental physics, these coincident observations led to a stringent constraint on the difference between the speed of gravity and the speed of light, allowed new bounds to be placed on local Lorentz invariance violations and enabled a new test of the equivalence principle by bounding the Shapiro delay between gravitational and electromagnetic radiation [123]. These bounds, in turn, helped to strongly constrain the allowed parameter space of alternative theories of gravity that offered gravitational explanations for the origin of dark energy [125–131] or dark matter [132].

GW170817 also offers important clues about GR polarizations. We carry out a test similar to [5] (Sec. 6.2) by performing a coherent Bayesian analysis of the signal properties using either the tensor or the vector or the scalar response functions. We assume that the phase evolution of the GW can be described by GR templates, but the polarization content can vary [116]. The phase evolution is modeled with the GR waveform model IMRPhenomPv2 and the analysis is carried out with

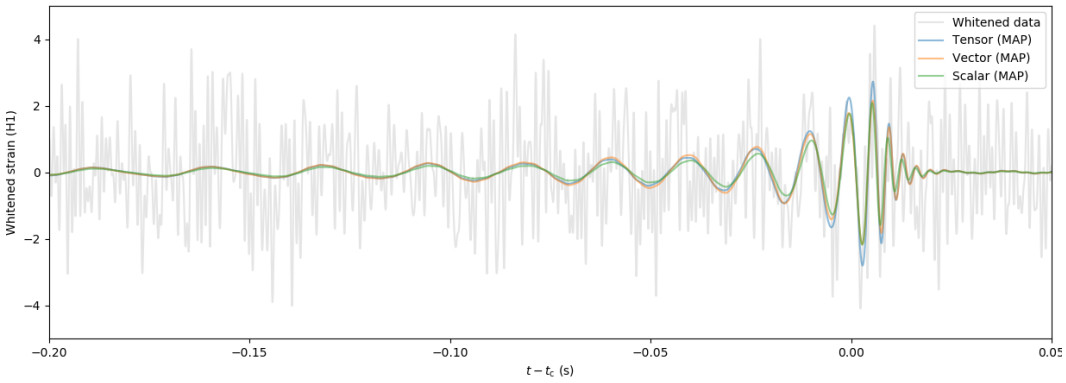

(a) LIGO Hanford

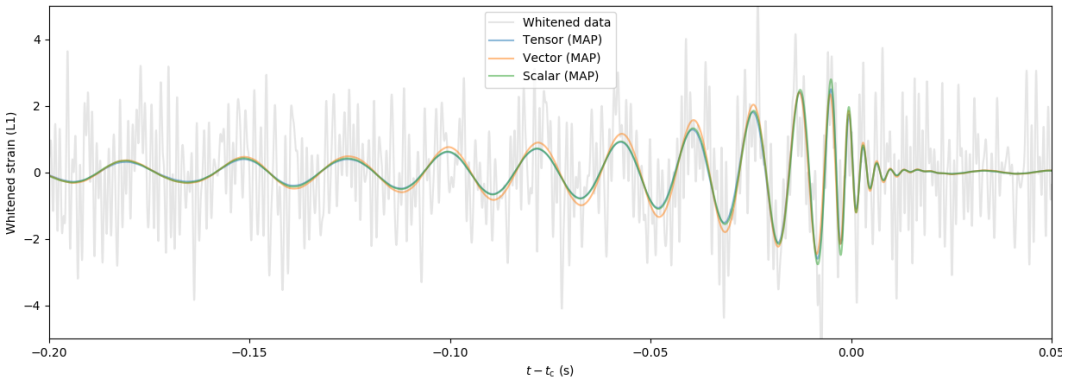

(b) LIGO Livingston

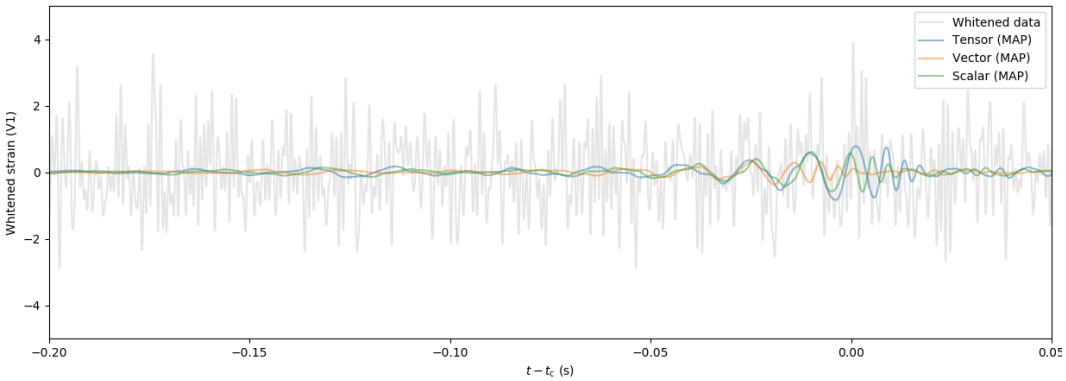

(c) Virgo

Figure 6.5: *GW170814 waveforms*. Reconstructed waveforms with the maximum *a posteriori* probability (MAP) under the assumption of fully tensor (blue), vector (orange) or scalar (green) polarizations, as seen by each of the three detectors in the network.

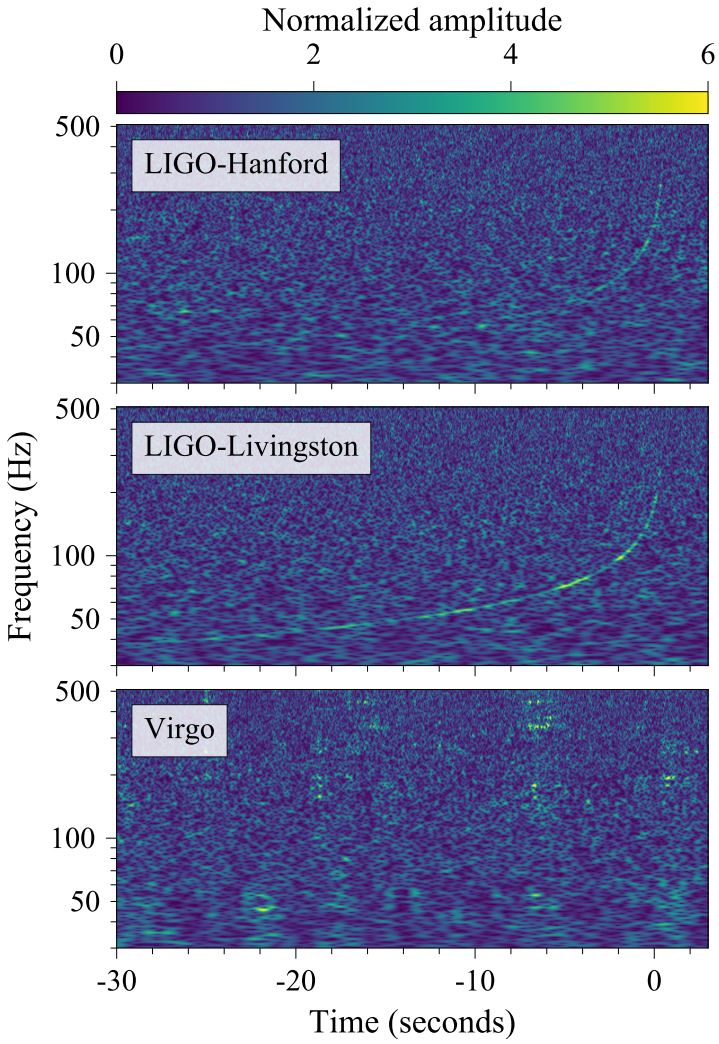

Figure 6.6: Time-frequency representations [122] of data containing the gravitational-wave event GW170817 , observed by the LIGO-Hanford (top), LIGO-Livingston (middle), and Virgo (bottom) detectors. Times are shown relative to August 17, 2017 12:41:04 UTC. The amplitude scale in each detector is normalized to that detector's noise amplitude spectral density. In the LIGO data, independently observable noise sources and a glitch that occurred in the LIGO-Livingston detector have been subtracted, as described in [6]. (Reproduced from [6].)

LALInference [20]. Tidal effects are not included in this waveform model, but this is not expected to affect the results presented below, since the polarization test is sensitive to the antenna pattern functions of the detectors and not the phase evolution of the signal, as argued above.

If the sky location of GW170817 is constrained to NGC 4993, we find overwhelming evidence in favor of pure tensor polarization modes in comparison to pure vector and pure scalar modes with a (base ten) logarithm of the Bayes factor of $+20.81 \pm 0.08$ and $+23.09 \pm 0.08$ respectively. This result is many orders of magnitudes stronger than the GW170814 case both due to the sky position of GW170817 relative to the detectors and the fact that the sky position is determined precisely by electromagnetic observations. Indeed if the sky location is unconstrained we find evidence against scalar modes with $+5.84 \pm 0.09$, while the test is inconclusive for vector modes with $+0.72 \pm 0.09$. From this analysis, we also see that only the tensor hypothesis is consistent with the location of the electromagnetic counterpart (Fig. 6.7).

The above results are to be expected given the marked difference in the network's sensitivity to signals with different polarizations in the direction of GW170817 (with respect to the network at the time of arrival). This is demonstrated by Fig. 6.8, which shows the location of NGC 4993 (cyan star) over each detector's effective response to tensor, vector or scalar signals (color). For the purpose of this figure, we define the "effective response" as the quadrature sum of the antenna patterns for each polarization of a given helicity evaluated at a given sky location [see Eqs. (5.23)–(5.25)]. This figure suggests that, had GW170817 been purely scalar or vector, the detectors would have measured drastically different relative amplitudes. In particular, the signal would not have been measured loudly in the LIGO detectors and not at all in Virgo, as it was.

6.4 Conclusion

The observations of GW170814 and GW170817 have allowed us to gather the first direct evidence that GWs come in the polarizations predicted by GR, as opposed to the extreme alternatives of purely vector or purely scalar polarizations. This is an important phenomenological check that the weak-field geometry of spacetime does not deviate *drastically* from what Einstein predicted. The log likelihood ratios (Bayes factors $\mathcal{B}$) encapsulating this conclusion are listed in Table 6.1.

However, we have not yet been able to address the case of mixed polarizations (e.g. GR plus some small nontensorial component), most interesting to theorists

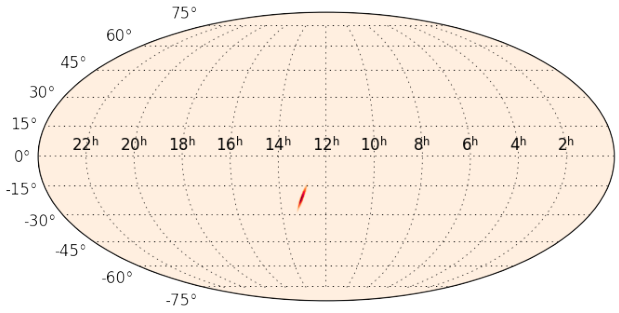

(a) Tensor

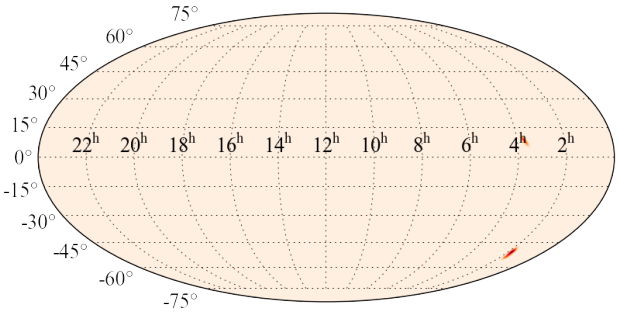

(b) Vector

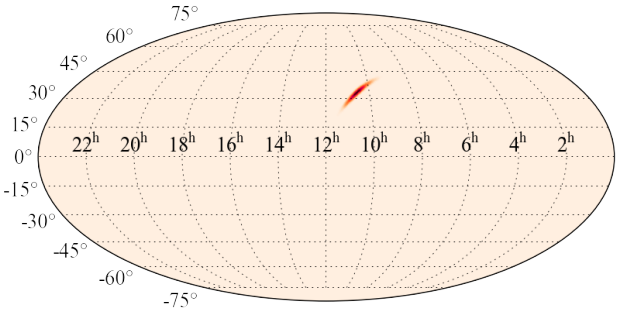

(c) Scalar

Figure 6.7: *GW170817 skymaps.* Reconstructed sky location under the assumption of the three different polarization hypotheses. Color represents probability density, as a function of equatorial coordinates in a Mollweide projection. Only the tensor case agrees with the EM counterpart.

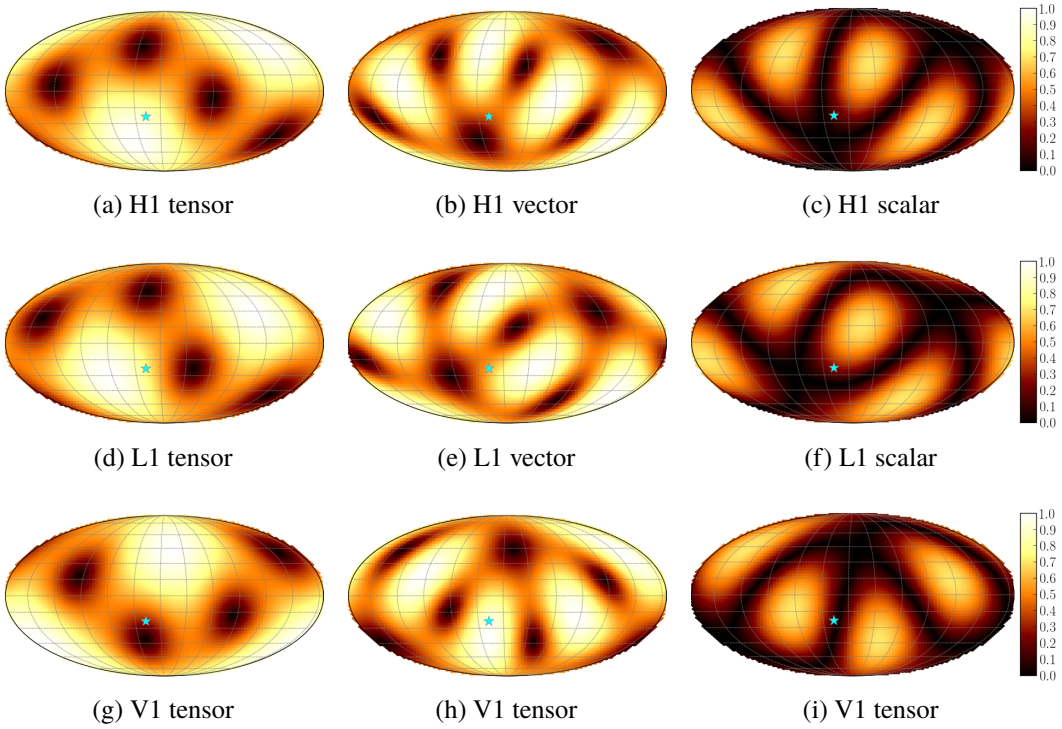

(a) H1 tensor (b) H1 vector (c) H1 scalar

(d) L1 tensor (e) L1 vector (f) L1 scalar

(g) V1 tensor (h) V1 tensor (i) V1 tensor

Figure 6.8: *Antenna patterns with GW170817 location.* Effective detector response to tensor (left), vector (center) and scalar (right) polarizations for the LIGO Hanford (top), LIGO Livingston (center) and Virgo (bottom) detectors, with the location of the electromagnetic counterpart of GW170817 (cyan star). The "effective antenna patterns" are defined as in Eqs. (5.23)–(5.25). The coordinates in this plot are the same as in Fig. 6.7.

developing extensions to GR. Even in the presence of an electromagnetic counterpart, as for GW170817, it is not currently possible to break all degeneracies between such mixed models by studying individual *transient* (i.e., short-lived) GW signals. However, we may, in the future, extract interesting polarization information by means of a rigorous Bayesian analysis of multiple CBC sources and, eventually, using extra detectors, LIGO India [133, 134] and KAGRA [135, 136]. Alternatively, when detected, *persistent* GW signals will allow us to thoroughly probe the polarization content of GWs, even with one or two detectors. This possibility is explored in the following few chapters.

Chapter 7

DETECTING BEYOND-EINSTEIN POLARIZATIONS OF CONTINUOUS GRAVITATIONAL WAVES

M. Isi, A. J. Weinstein, C. Mead, and M. Pitkin, Phys. Rev. **D91**, 082002 (2015), M.I. carried out the analysis and wrote the manuscript, arXiv:1502.00333 [gr-qc] .

7.1 Introduction

Since its introduction in 1915, Einstein's theory of General Relativity (GR) has been confirmed by experiment on every occasion [84]. However, GR has not yet been tested with great precision on scales larger than the solar system or for highly dynamical and strong gravitational fields [137]. Those kinds of rapidly changing fields give rise to gravitational waves (GWs)—self propagating stretching and squeezing of spacetime originating in the acceleration of massive objects, like spinning neutron stars with an asymmetry in their moment of inertia (e.g., see [138, 139]).

Although GWs are yet to be directly observed,[1] detectors such as the Laser Interferometer Gravitational Wave Observatory (LIGO) expect to do so in the coming years, giving us a chance to probe GR on new grounds [140, 141]. Because GR does not present any adjustable parameters, these tests have the potential to uncover new physics [84]. By the same token, LIGO data could also be used to test alternative theories of gravity that disagree with GR on the properties of GWs.

Furthermore, when looking for a weak signal in noisy LIGO data, certain physical models are used to target the search and are necessary to make any detection possible [137]. Because these are usually based on predictions from GR, assuming an incorrect model could yield a weak detection or no detection at all. Similarly, if GR is not a correct description for highly dynamical gravity, checking for patterns given by alternative models could result in detection where no signal had been seen before.

There exist efforts to test GR by looking at the deviations of the parametrized post-Newtonian coefficients extracted from the inspiral phase of compact binary

[1] This analysis was carried out before LIGO's first detection of gravitational waves on September 14, 2015. This notwithstanding, its conclusions remain valid and relevant to the extent that they were not superseded by the analysis in Chapter 8, which grew out of this study.

coalescence events [21, 142, 143]. Besides this, deviations from GR could be observed in generic GW properties such as polarization, wave propagation speed or parity violation [84, 85, 144]. Tests of these properties have been proposed which make use of GW burst search methods [145].

In this paper, we present methods to search LIGO-like detector data for continuous GW signals of any polarization mode, not just those allowed by GR. We also compare the relative sensitivity of different model-dependent and independent templates to certain kinds of signals. Furthermore, we provide expected sensitivity curves for GR and non-GR signals, obtained by means of blind searches over LIGO noise (not actual upper limits).

Section 7.2 provides the background behind GW polarizations and continuous waves, while Secs. 7.3 and 7.4 present search methods and the data analysis procedures used to evaluate sensitivity for detection. Results and final remarks are provided in Secs. 7.5 & 7.6 respectively.

7.2 Background

7.2.1 Polarizations

Just like electromagnetic waves, GWs can present different kinds of polarizations. Most generally, metric theories of gravity could allow six possible modes: plus (+), cross (×), vector x (x), vector y (y), breathing (b) and longitudinal (l). Their effects on a free-falling ring of particles are illustrated in Fig. 5.1. Transverse GWs (+, × and b) change the distance between particles separated in the plane perpendicular to the direction of propagation (taken to be the z-axis). Vector GWs are also transverse; but, because all particles in a plane perpendicular to the direction of propagation are equally accelerated, their relative separation is not changed. Nonetheless, particles farther from the source move at later times, hence varying their position relative to points with both different x-y coordinates and different z distance. Finally, longitudinal GWs change the distance between particles separated along the direction of propagation.

7.2.2 Signal

Because of their persistence, continuous gravitational waves (CWs) provide the means to study GW polarizations without the need for multiple detectors. For the same reason, continuous signals can be integrated over long periods of time, thus

Table 7.1: LIGO detectors [147][148]

	LHO	LLO	VIR
Latitude (λ)	46.45° N	30.56° N	43.63° N
Longitude (ϕ)	119.41° W	90.77° W	10.5° E
Orientation (γ)	125.99°	198.0°	71.5°

improving the likelihood of detection. Furthermore, these GWs are quasi-sinusoidal and present well-defined frequencies. This allows us to focus on the amplitude modulation, where the polarization information is contained.

CWs are produced by localized sources with periodic motion, such as binary systems or spinning neutron stars [146]. Throughout this paper, we target known pulsars (e.g., the Crab pulsar) and assume an asymmetry in their moment of inertia (rather than precession of the spin axis or other possible, but less likely, mechanisms) causes them to emit gravitational radiation. A source of this type can generate GWs only at multiples of its rotational frequency ν. In fact, it is expected that most power be radiated at twice this value [54]. For that reason, we take the GW frequency, $\nu_{\rm gw}$, to be 2ν. Moreover, the frequency evolution of these pulsars is well-known thanks to electromagnetic observations, mostly at radio wavelengths but also in gamma-rays.

Simulation of a CW from a triaxial neutron star is straightforward. The general form of a such a signal is:

$$h(t) = \sum_p A_p(t; \psi|\alpha, \delta, \lambda, \phi, \gamma, \xi)\, h_p(t; \iota, h_0, \phi_0, \nu, \dot{\nu}, \ddot{\nu}), \tag{7.1}$$

where, for each polarization p, A_p is the detector response (antenna pattern) and h_p a sinusoidal waveform of frequency $\nu_{\rm gw} = 2\nu$. The detector parameters are: λ, longitude; ϕ, latitude; γ, angle of the detector x-arm measured from East; and ξ, the angle between arms. Values for the LIGO Hanford Observatory (LHO), LIGO Livingston Observatory (LLO) and Virgo (VIR) detectors are presented in Table 7.1. The source parameters are: ψ, the signal polarization angle; ι, the inclination of the pulsar spin axis relative to the observer's line-of-sight; h_0, an overall amplitude factor; ϕ_0, a phase offset; and ν, the rotational frequency, with $\dot{\nu}$, $\ddot{\nu}$ its first and second derivatives. Also, α is the right ascension and δ the declination of the pulsar in celestial coordinates.

Note that the inclination angle ι is defined as is standard in astronomy, with $\iota = 0$ and $\iota = \pi$ respectively meaning that the angular momentum vector of the source points towards and opposite to the observer. The signal polarization angle ψ is related to

Table 7.2: Axis polarization (ψ) and inclination (ι) angles for known pulsars [149].

	ψ (deg)	ι (deg)		ψ (deg)	ι (deg)
Crab	124.0	61.3	J0205+6449	90.3	91.6
Vela	130.6	63.6	J0537−6910	131	92.8
J1930+1852	91	147	B0540−69	144.1	92.9
J2229+6114	103	46	J1124−5916	16	105
B1706−44	163.6	53.3	B1800−21	44	90
J2021+3651	45	79	J1833−1034	45	85.4

the position angle of the source, which is in turn defined to be the East angle of the projection of the source's spin axis onto the plane of the sky.

Although there are hundreds of pulsars in the LIGO band, in the majority of cases we lack accurate measurements of their inclination and polarization angles. The few exceptions, presented in Table 7.2, were obtained through the study of the pulsar spin nebula [149]. This process cannot determine the spin direction, only the orientation of the spin axis. Consequently, even for the best studied pulsars, ψ and ι are only known modulo a reflection: we are unable to distinguish between ψ and $-\psi$ or between ι and $\pi - \iota$. As will be discussed in Sec. 7.3, our ignorance of ψ and ι must be taken into account when searching for CWs.

Frequency evolution

In Eq. (7.1), $h_p(t)$ is a sinusoid carrying the frequency modulation of the signal:

$$h_p(t) = a_p \cos\left(\phi(t) + \phi_p + \phi_0^{\mathrm{gw}}\right) \tag{7.2}$$

$$\phi(t) = 4\pi\left(\nu t_{\mathrm{b}} + \frac{1}{2}\dot{\nu} t_{\mathrm{b}}^2 + \frac{1}{6}\ddot{\nu} t_{\mathrm{b}}^3\right) + \phi_0^{\mathrm{em}}, \tag{7.3}$$

where t_{b} is the Solar System barycentric arrival time, which is the local arrival time t modulated by the standard Rømer Δ_R, Einstein Δ_E and Shapiro Δ_S delays [56]:

$$t_{\mathrm{b}} = t + \Delta_R + \Delta_E + \Delta_S. \tag{7.4}$$

The leading factor of four in the r.h.s. of Eq. (7.3) comes from the substitution $\nu_{\mathrm{gw}} = 2\nu$. For known pulsars, ϕ_0^{em} is the phase of the radio pulse, while ϕ_0^{gw} is the phase difference between electromagnetic and gravitational waves. Both factors contribute to an overall phase offset of the signal ($\phi_0^{\mathrm{em}} + \phi_0^{\mathrm{gw}}$). This is of astrophysical significance since it may provide insights about the relation between EM & GW radiation and provide information about the physical structure of the source.

The a_p and ϕ_p coefficients in Eq. (7.2) respectively encode the relative amplitude and phase of each polarization. These values are determined by the physical model. For instance, GR predicts:

$$a_+ = h_0(1 + \cos^2 \iota)/2\,,\ \phi_+ = 0, \tag{7.5}$$

$$a_\times = h_0 \cos \iota\,,\ \phi_\times = -\pi/2, \tag{7.6}$$

while $a_\mathrm{x} = a_\mathrm{y} = a_\mathrm{b} = 0$. On the other hand, according to G4v [106]:

$$a_\mathrm{x} = h_0 \sin \iota\,,\ \phi_\mathrm{x} = -\pi/2, \tag{7.7}$$

$$a_\mathrm{y} = h_0 \sin \iota \cos \iota\,,\ \phi_\mathrm{x} = 0. \tag{7.8}$$

while $a_+ = a_\times = a_\mathrm{b} = 0$. In both cases, the overall amplitude h_0 can be characterized by [54, 57, 106]:

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{I_{zz} \nu^2}{r} \epsilon, \tag{7.9}$$

where r is the distance to the source, I_{zz} the pulsar's moment of inertia along the principal axis, $\epsilon = (I_{xx} - I_{yy})/I_{zz}$ its equatorial ellipticity and, as before, ν is the rotational frequency. Choosing some canonical values,

$$h_0 \approx 4.2 \times 10^{-26} \frac{I_{zz}}{10^{28}\ \mathrm{kg\ m^2}} \left[\frac{\nu}{100\ \mathrm{Hz}}\right]^2 \frac{1\ \mathrm{kpc}}{r} \frac{\epsilon}{10^{-6}}, \tag{7.10}$$

it is easy to see that GWs from triaxial neutron stars are expected to be relatively weak [150]. However, the sensitivity to these waves grows with the observation time because the signal can be integrated over long periods of time [57].

As indicated in the introduction to this section, we have assumed CWs are caused by an asymmetry in the moment of inertia of the pulsar. Other mechanisms, such as precession of the spin axis, are expected to produce waves of different strengths and with dominant components at frequencies other than 2ν. Furthermore, these effects vary between theories: for instance, in G4v, if the asymmetry is not perpendicular to the rotation axis, there can be a significant ν component as well as the 2ν component. In those cases, Eqs. (7.2, 7.9) do not hold (e.g., see [54] for precession models).

Amplitude modulation

At any given time, GW detectors are not equally sensitive to all polarizations. The response of a detector to a particular polarization p is encoded in a function $A_p(t)$

depending on the relative locations and orientations of the source and detector. As seen from Eq. (7.1), these functions provide the amplitude modulation of the signal.

A GW is best described in an orthogonal coordinate frame defined by wave vectors $(\mathbf{w}_x,\ \mathbf{w}_y,\ \mathbf{w}_z)$, with $\mathbf{w}_z = \mathbf{w}_x \times \mathbf{w}_y$ being the direction of propagation. Furthermore, the orientation of this wave-frame is fixed by requiring that the East angle between $\mathbf{w}_y$ and the celestial North be ψ. In this gauge, the different polarizations act through six orthogonal basis strain tensors [107, 108]:[2]

$$e^{+}_{jk} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \qquad e^{\times}_{jk} = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{2,3}$$

$$e^{\mathrm{x}}_{jk} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \qquad e^{\mathrm{y}}_{jk} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \tag{4,5}$$

$$e^{\mathrm{b}}_{jk} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \qquad e^{\mathrm{l}}_{jk} = \sqrt{2}\begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{6,7}$$

with j, k indexing x, y and z components. These tensors can be written in an equivalent, frame-independent form

$$\mathbf{e}^{+} = \mathbf{w}_x \otimes \mathbf{w}_x - \mathbf{w}_y \otimes \mathbf{w}_y, \tag{7.17}$$

$$\mathbf{e}^{\times} = \mathbf{w}_x \otimes \mathbf{w}_y + \mathbf{w}_y \otimes \mathbf{w}_x, \tag{7.18}$$

$$\mathbf{e}^{\mathrm{x}} = \mathbf{w}_x \otimes \mathbf{w}_z + \mathbf{w}_z \otimes \mathbf{w}_x, \tag{7.19}$$

$$\mathbf{e}^{\mathrm{y}} = \mathbf{w}_y \otimes \mathbf{w}_z + \mathbf{w}_z \otimes \mathbf{w}_y, \tag{7.20}$$

$$\mathbf{e}^{\mathrm{b}} = \mathbf{w}_x \otimes \mathbf{w}_x + \mathbf{w}_y \otimes \mathbf{w}_y, \tag{7.21}$$

$$\mathbf{e}^{\mathrm{l}} = \sqrt{2}\,(\mathbf{w}_z \otimes \mathbf{w}_z)\,. \tag{7.22}$$

If a detector is characterized by its unit arm-direction vectors ($\mathbf{d}_x$ and $\mathbf{d}_y$, with $\mathbf{d}_z$ the detector zenith), its differential-arm response A_p to a wave of polarization p is:

$$A_p = \frac{1}{2}\left(\mathbf{d}_x \otimes \mathbf{d}_x - \mathbf{d}_y \otimes \mathbf{d}_y\right) : \mathbf{e}^{p}, \tag{7.23}$$

[2] The normalization for $\mathbf{e}^{\mathrm{l}}$ used here is unique to this chapter: in the rest of this thesis we define $\mathbf{e}^{\mathrm{l}} = \mathrm{diag}(0, 0, 0, 1)$, i.e. without the $\sqrt{2}$ prefactor.

where the colon indicates double contraction. As a result, Eqs. (2-13) imply:

$$A_+ = \frac{1}{2}\left[(\mathbf{w}_x \cdot \mathbf{d}_x)^2 - (\mathbf{w}_x \cdot \mathbf{d}_y)^2 - (\mathbf{w}_y \cdot \mathbf{d}_x)^2 + (\mathbf{w}_y \cdot \mathbf{d}_y)^2\right], \tag{7.24}$$

$$A_\times = (\mathbf{w}_x \cdot \mathbf{d}_x)(\mathbf{w}_y \cdot \mathbf{d}_x) - (\mathbf{w}_x \cdot \mathbf{d}_y)(\mathbf{w}_y \cdot \mathbf{d}_y), \tag{7.25}$$

$$A_\mathrm{x} = (\mathbf{w}_x \cdot \mathbf{d}_x)(\mathbf{w}_z \cdot \mathbf{d}_x) - (\mathbf{w}_x \cdot \mathbf{d}_y)(\mathbf{w}_z \cdot \mathbf{d}_y), \tag{7.26}$$

$$A_\mathrm{y} = (\mathbf{w}_y \cdot \mathbf{d}_x)(\mathbf{w}_z \cdot \mathbf{d}_x) - (\mathbf{w}_y \cdot \mathbf{d}_y)(\mathbf{w}_z \cdot \mathbf{d}_y), \tag{7.27}$$

$$A_\mathrm{b} = \frac{1}{2}\left[(\mathbf{w}_x \cdot \mathbf{d}_x)^2 - (\mathbf{w}_x \cdot \mathbf{d}_y)^2 + (\mathbf{w}_y \cdot \mathbf{d}_x)^2 - (\mathbf{w}_y \cdot \mathbf{d}_y)^2\right], \tag{7.28}$$

$$A_\mathrm{l} = \frac{1}{\sqrt{2}}\left[(\mathbf{w}_z \cdot \mathbf{d}_x)^2 - (\mathbf{w}_z \cdot \mathbf{d}_y)^2\right]. \tag{7.29}$$

Accounting for the time dependence of the arm vectors due to the rotation of the Earth, Eqs. (7.24-7.29) can be used to compute $A_p(t)$ for any value of t. In Fig. 7.1 we plot these responses for the LIGO Hanford Observatory (LHO) observing the Crab pulsar, over a sidereal day (the pattern repeats itself every day). Note that the b and l patterns are degenerate ($A_\mathrm{b} = -\sqrt{2}A_\mathrm{l}{}^3$), which means they are indistinguishable up to an overall constant.

Although the antenna patterns are ψ-dependent, a change in this angle amounts to a rotation of A_+ into $A_\times$ or of A_x into A_y, and vice-versa. If the orientation of the source is changed such that the new polarization is $\psi' = \psi + \Delta\psi$, where ψ is the original polarization angle and $\Delta\psi \in [0, 2\pi]$, it is easy to check that the new antenna patterns can be written [108]:

$$A'_+ = A_+ \cos 2\Delta\psi + A_\times \sin 2\Delta\psi, \tag{7.30}$$

$$A'_\times = A_\times \cos 2\Delta\psi - A_+ \sin 2\Delta\psi, \tag{7.31}$$

$$A'_x = A_x \cos \Delta\psi + A_y \sin \Delta\psi, \tag{7.32}$$

$$A'_y = A_y \cos \Delta\psi - A_x \sin \Delta\psi, \tag{7.33}$$

$$A'_b = A_b, \tag{7.34}$$

$$A'_l = A_l, \tag{7.35}$$

and the tensor, vector and scalar nature of each polarization becomes evident from the ψ dependence.

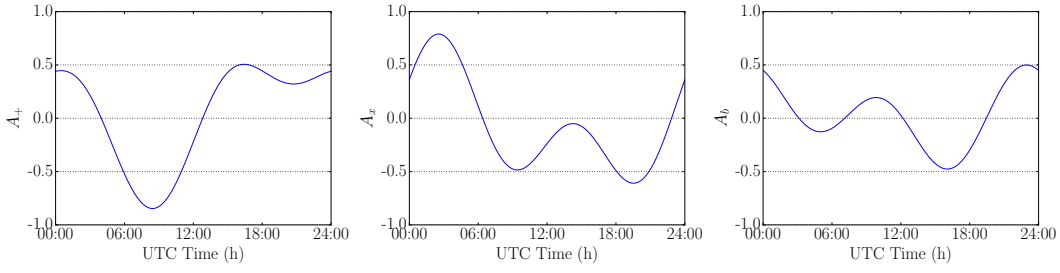

(a) Plus (+) (b) Vector x (x) (c) Breathing (b)

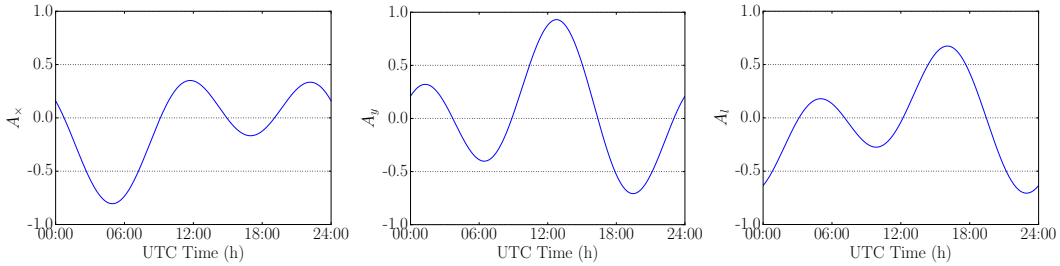

(d) Cross (×) (e) Vector y (y) (f) Longitudinal (l)

Figure 7.1: LHO response $A_p(t)$, Eq. (7.23), to different polarizations from the Crab (PSR J0534+2200), from 00:00 UTC to 24:00 UTC.

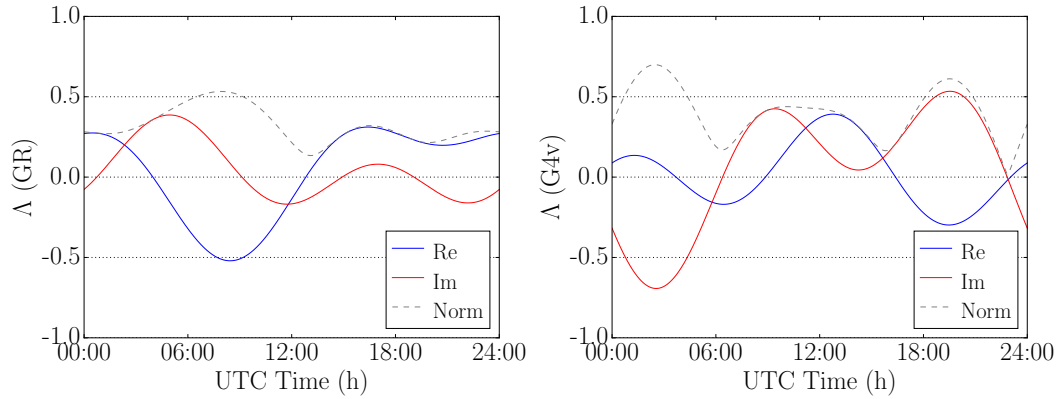

Figure 7.2: *Simulated GR (left) and G4v (right) heterodyned Crab signals as seen by LHO*. The templates are generated from Eq. (7.37) with the model parameters given in Eqs. (7.5–7.8) and setting $h_0 = 1$, $\phi_0 = 0$. The solid curves represent the real (blue) and imaginary (red) parts, while the dashed curve corresponds to the complex norm.

7.3 Method

7.3.1 Data reduction

For some set of interferometric data, we would like to detect CW signals from a given source, regardless of their polarization, and to reliably distinguish between the different modes. Because detector response is the only factor distinguishing CW polarizations, all the relevant information is encoded in the amplitude modulation of the signal. As a result, it suffices to consider a narrow frequency band around the GW frequency and the data can be considerably reduced following the complex heterodyne method developed in [66] and [57].

A signal of the form of Eq. (7.1) can be rewritten as

$$h(t) = \Lambda(t)e^{i\phi(t)} + \Lambda^*(t)e^{-i\phi(t)}, \tag{7.36}$$

$$\Lambda(t) = \frac{1}{2}\sum_{p=1}^{5} a_p e^{i\phi_p + i\phi_0} A_p(t), \tag{7.37}$$

with $*$ indicating complex conjugation and $\phi(t)$ as given in Eq. (7.3). Note that we have slightly simplified the notation in Eq. (7.37) by renaming $\phi_0^{\rm gw} \rightarrow \phi_0$. Also, the summation is over only five values of p because the breathing and longitudinal polarizations are indistinguishable to the detectors.

The key of the heterodyne method is that, since we can assume the phase evolution is well-known from electromagnetic observations (ephemerides obtained through the pulsar timing package TEMPO2 [56]), we can multiply our data by $\exp\left[-i\phi(t)\right]$ (heterodyning) so that the signal therein becomes

$$h'(t) \equiv h(t)e^{-i\phi(t)} = \Lambda(t) + \Lambda^*(t)e^{-i2\phi(t)} \tag{7.38}$$

and the frequency modulation of the first term is removed, while that of the second term is doubled. A series of low-pass filters can then be used to remove the quickly-varying term, which enables the down-sampling of the data by averaging over minute-long time bins. As a result, we are left with $\Lambda(t)$ only and Eq. (7.37) becomes the template of our complex-valued signal. One period of such GR and G4v signals coming from the Crab are presented as seen by LHO in Fig. 7.2.

From Eq. (7.38) we see that, in the presence of a signal, the heterodyned and down-sampled noisy detector strain data B_k for the k^{th} minute-long time bin (which

[3]For the definitions used in all other chapters, the relation is $A_{\rm b} = -A_{\rm l}$ (see previous footnote).

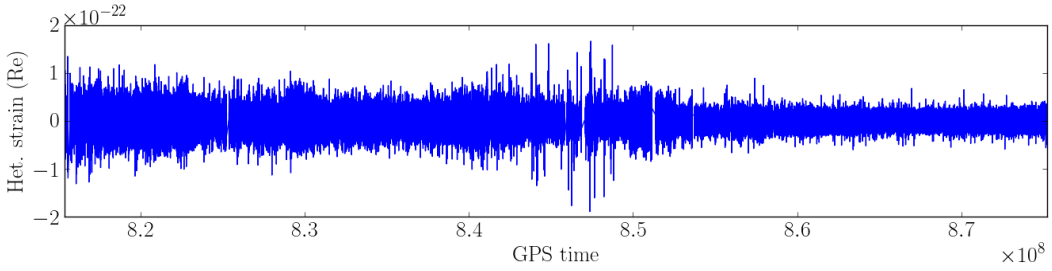

Figure 7.3: Real part of LIGO Science Run 5 Hanford 4km detector (H1) minute-sampled data prepared for the Crab spanning approximately two years. A signal in these data would be described by Eq. (7.39).

can be labeled by GPS time of arrival) are expected to be of the form:

$$B_{\text{expected}}(t_k) = \frac{1}{2}\sum_{p=1}^{5} a_p(t_k) e^{i\phi_p + i\phi_0} A_p(t_k) + n(t_k), \tag{7.39}$$

where $n(t_k)$ is the heterodyned, averaged complex noise in bin k, which carries no information about the GW signal. As an example, Fig. 7.3 presents the real part of actual data heterodyned and filtered for the Crab pulsar. We can clearly see already that the data are non-stationary, an issue addressed in the Sec. 7.3.2 and Appendix 7.7.

7.3.2 Search

Given data in this form, we analyze it to obtain the parameters of a signal that would best fit the data and then incorporate the results into the frequentist analysis described in Sec. 7.4. Regressions are performed by minimizing the χ^2 of the system (same as a matched-filter). For certain template $T(t_k)$, this is:

$$\chi^2 = \sum_{k=0}^{N} \left[T(t_k) - B(t_k)\right]^2 / \sigma_k^2, \tag{7.40}$$

where σ_k is the estimate standard deviation of the noise in the data at time t_k. In the presence of Gaussian noise, the χ^2 minimization is equivalent to a maximum likelihood analysis.

Any linear template T can be written as a linear combination of certain basis functions f_i, so that $T(t) = \sum_i \tilde{a}_i f_i(t)$ and each $\tilde{a}_i$ is found as a result of minimizing (7.40). For instance, $T(t_k)$ could be constructed in the from of Eq. (7.37). In such model-dependent searches, the antenna patterns are the basis set, i.e. $\{f_i\} = \{A_p\}$,

and the $\tilde{a}_i$ weights correspond to the $a_p \exp\left(i\phi_p\right)$ prefactors. (From here on, the tilde denotes the coefficient that is fitted for, rather than its predicted value.)

The regression returns a vector $\tilde{\mathbf{a}}$ containing the values of the $\tilde{a}_i$'s that minimize Eq. (7.40). These quantities are complex-valued and encode the relative amplitude and phase of each contributing basis. From their magnitude, we define the overall *recovered signal strength* to be:

$$h_{\mathrm{rec}} = |\tilde{\mathbf{a}}|. \tag{7.41}$$

The significance of the fit is evaluated through the covariance matrix C. This can be computed by taking the inverse of $A^T A$, where A is the design matrix of the system (built from the f_i set). In particular, we define the *significance* of the resulting fit (signal SNR) as

$$s = \sqrt{\tilde{\mathbf{a}}^\dagger C^{-1} \tilde{\mathbf{a}}}, \tag{7.42}$$

where † indicates Hermitian conjugation.

χ^2-minimizations have optimal performances when the noise is Gaussian. However, although the central limit theorem implies that the averaged noise in (7.39) should be normally distributed, actual data are far from this ideal (see Fig. 7.3). In fact, the quality of the data changes over time, as it is contingent on various instrumental factors. The time series is plagued with gaps and is highly nonstationary. This makes estimating σ_k nontrivial.

As done in regular CW searches [150], we address this problem by computing the standard deviation for the data corresponding to each sidereal day throughout the data run, rather than for the series as a whole. This method improves the analysis because the data remains relatively stable over the course of a single day, but not throughout longer periods of time (see Appendix 7.7). Furthermore, noisier days have less impact on the fit, because σ_k in Eq. (7.40) will be larger. The evolution of the daily value of the standard deviation for H1 data heterodyned for the Crab pulsar is presented in Fig. 7.4.

Model-dependent

In a model-dependent search, a particular physical model is assumed in order to create a template based on Eq. (7.37). In the case of GR, if ψ and ι are known, it is

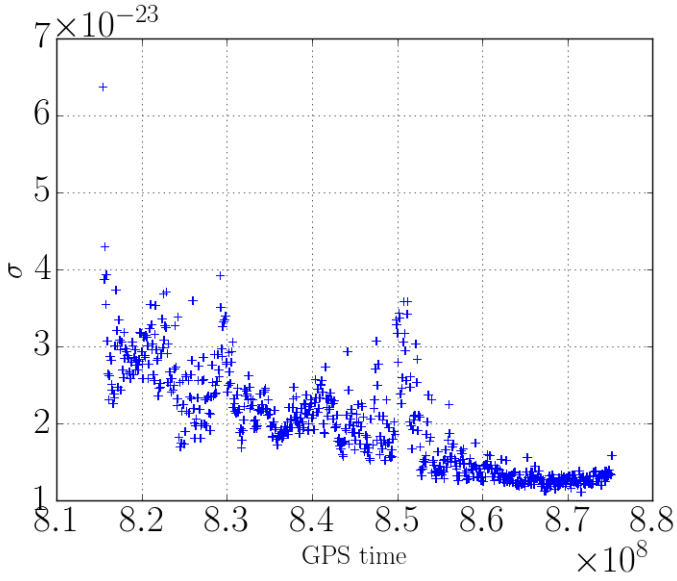

Figure 7.4: Daily standard deviation of S5 H1 data heterodyned for the Crab pulsar (Fig. 7.3).

possible to construct a template with only one complex-valued free parameter $\tilde{h}_0$:

$$T_{\mathrm{GR}}(t) = \tilde{h}_0 \frac{1}{2}\left[\frac{1}{2}(1+\cos^2\iota)A_+(t;\psi)+ \right. \\ \left. + \cos\iota A_\times(t;\psi)e^{-i\pi/2}\right], \tag{7.43}$$

where the factor of 2 comes from the heterodyne, cf. Eq. (7.37). Similarly for G4v:

$$T_{\mathrm{G4v}}(t) = \tilde{h}_0 \frac{1}{2}\left[\sin\iota\, e^{-i\pi/2}A_{\mathrm{x}}(t;\psi) + \sin\iota\cos\iota A_{\mathrm{y}}(t;\psi)\right], \tag{7.44}$$

Analogous templates could be constructed for scalar-tensor theories, or any other model. In the former case, there would be a second free parameter to represent the unknown scalar contribution.

However, as mentioned in Sec. 7.2, even in the case of the best-studied pulsars we know ι only in absolute value. This ambiguity creates the need to use two model-dependent templates like Eqs. (7.43, 7.44): one corresponding to ι and one to $\pi - \iota$. Note that the indeterminacy of ψ is absorbed by the overall phase of $\tilde{h}_0$, so it has no effect on the template. Thus, if the ambiguity in ι is accounted for, the overall signal strength h_0 and the angle ϕ_0 can be inferred directly from the angle and phase of $h_{\mathrm{rec}} = \tilde{h}_0$.

In most cases, ψ and ι are completely unknown. It is then convenient to regress to each antenna pattern independently, allowing for two free parameters. This can be done

by computing the antenna patterns assuming any arbitrary value of the polarization angle, say $\psi = 0$. Indeed, Eqs. (7.30–7.35) guarantee that the subspace of tensor, vector or scalar antenna patterns for *all* ψ is spanned by a pair of corresponding tensor, vector or scalar antenna patterns assuming any *particular* ψ.

In the case of GR, this means we can use a template

$$T_{\rm GR}(t) = \tilde{\alpha}_+ A_+(t;\psi=0)/2 + \tilde{\alpha}_\times A_\times(t;\psi=0)/2 \tag{7.45}$$

with two complex weights $\tilde{\alpha}$'s to be determined by the minimization. In the presence of a signal and in the absence of noise, Eqs. (7.30, 7.31) indicate that the values returned by the fit would be a function of the *actual*, unknown ψ and ι:

$$\alpha_+ = a_+(\iota)e^{i\phi_0}\cos 2\psi - a_\times(\iota)e^{i\phi_0 - i\pi/2}\sin 2\psi, \tag{7.46}$$

$$\alpha_\times = a_\times(\iota)e^{i\phi_0 - i\pi/2}\cos 2\psi + a_+(\iota)e^{i\phi_0}\sin 2\psi, \tag{7.47}$$

with the $\alpha(\iota)$'s as given in Eqs. (7.5, 7.6).

Again, a (semi-) model-dependent template, like Eq. (7.45), can be constructed for any given theory by selecting the corresponding antenna patterns to be used as basis for the regression. For G4v, this would be:

$$T_{\rm G4v}(t) = \tilde{\alpha}_x A_x(t;\psi=0)/2 + \tilde{\alpha}_y A_y(t;\psi=0)/2 \tag{7.48}$$

with two complex weights $\tilde{\alpha}$'s to be determined by the minimization. As before, in the presence of a signal and in the absence of noise, Eqs. (7.32, 7.33) indicate that the values returned by the fit would be a function of the *actual*, unknown ψ and ι:

$$\alpha_x = a_x(\iota)e^{i\phi_0 - i\pi/2}\cos\psi - a_y(\iota)e^{i\phi_0}\sin\psi, \tag{7.49}$$

$$\alpha_y = a_y(\iota)e^{i\phi_0}\cos\psi + a_x(\iota)e^{i\phi_0 - i\pi/2}\sin\psi. \tag{7.50}$$

In this case, we cannot directly relate our recovered strength to h_0 and the framework does not allow to carry out parameter estimation. The proper way to do that is using Bayesian statistics, marginalizing over the orientation parameters. Since we are mostly interested in quantifying our ability to detect alternative signals rather than estimating source parameters, we do not cover such methods here. However, it would be straightforward to incorporate our generalized likelihoods (as given by our templates) into a full Bayesian analysis (cf. [57]).

Model-independent

In a model-independent search, the regression is performed using all five nondegenerate antenna patterns and the phases between the A_p's are not constrained. Thus,

$$T_{\text{indep}}(t) = \sum_{p=1}^{5} \tilde{a}_p A_p(t). \tag{7.51}$$

Because we do not consider any particular model, there is no information about the relative strength of each polarization; hence, the $\tilde{a}_p$'s are unconstrained. Again, Eqs. (7.30–7.35) enable us to compute the antenna patterns for any value of ψ.

By calculating the necessary inner products, it can be shown that a regression to the *antenna pattern basis*,

$$\left\{A_+,\ A_\times,\ A_{\text{x}},\ A_{\text{y}},\ A_{\text{b}}\right\}, \tag{7.52}$$

is equivalent to a regression to the *sidereal basis*,

$$\{1,\ \cos\omega t,\ \cos 2\omega t,\ \sin\omega t,\ \sin 2\omega t\}, \tag{7.53}$$

where $\omega = 2\pi/(86164\ \text{s})$ is the sidereal rotational frequency of the Earth. This is an orthogonal basis which spans the space of the antenna patterns. In this basis,

$$T_{\text{indep}}(t) = \sum_{i=1}^{5} \tilde{a}_i f_i(t). \tag{7.54}$$

with f_i representing the set in (7.53). This is the same basis set used in so-called *5-vector searches* [69].

Because they span the same space, using either basis set yields the same results with the exact same significance, as defined in Eq. (7.42). Furthermore, the weights obtained as results of the fit can be converted back and forth between the two bases by means of a time-independent coordinate transformation matrix.

A model-independent search is sensitive to all polarizations but is prone to error due to noise when distinguishing between them. It also has more degrees of freedom (compared with a pure-GR template) that can respond to noise fluctuations, resulting in a search that is less sensitive to pure-GR signals. However, the analysis can be followed by model-dependent searches to clarify which theory fits with most significance.

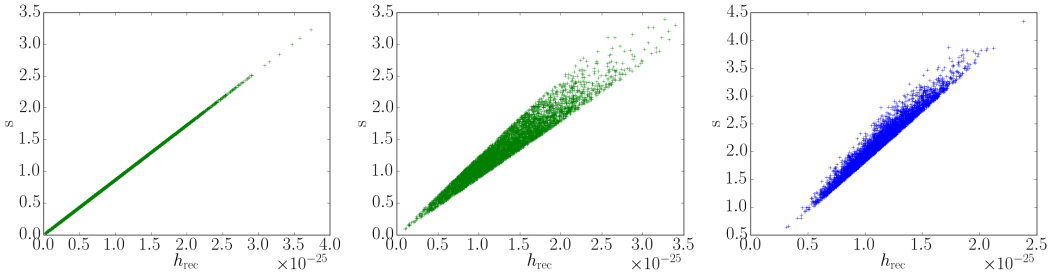

Figure 7.5: Significance, Eq. (7.42), vs recovered strength, Eq. (7.41), for searches over 5000 noise-only H1 S5 Crab instantiations using model-dependent Eq. (7.43) (left), semi-dependent Eq. (7.45) (center), and independent Eq. (7.54) (right) templates. The model-dependent case assumes *fully* known ι and ψ. Note that the number of degrees of freedom in the regression is manifested in the spread, which is due to noise: templates with a single degree of freedom are less susceptible to noise and the spread is minimal. The two plots on the left were generated using a GR template, but similar results are obtained for G4v.

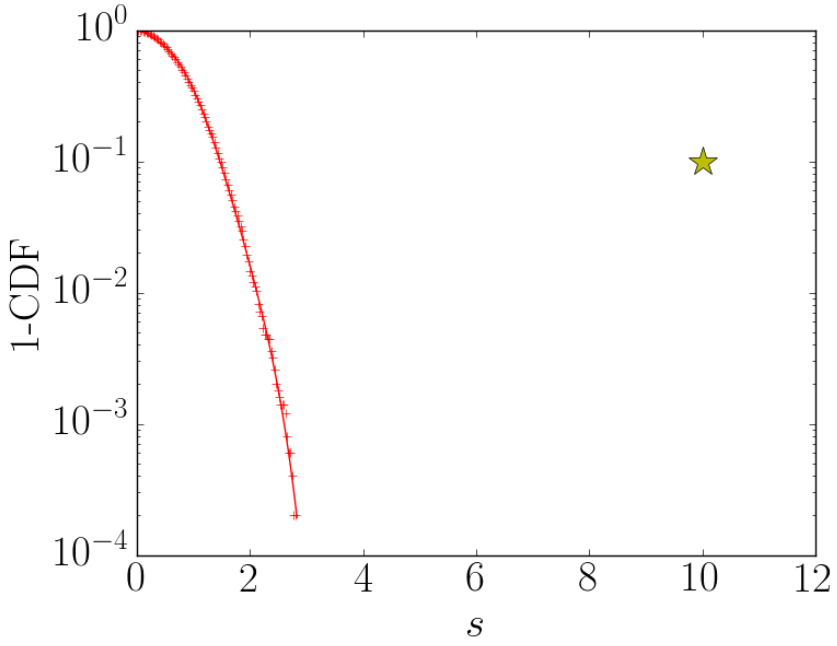

Figure 7.6: Example plot of $p = 1 - \text{CDF}$ vs the recovery significance for a particular template. A loud injection in noise is manifested as an outlier (star) over the noise-only background (red). Note that the injection is plotted arbitrarily at $p = 10^{-1}$.

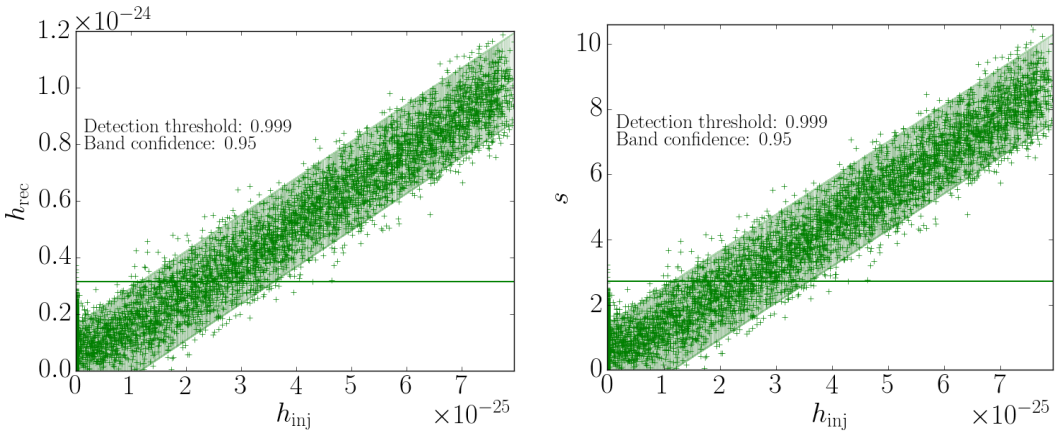

Figure 7.7: Neyman plot of recovered signal strength $h_{\rm rec}$ (left) and significance s (right) vs injected strength $h_{\rm inj}$. In this case, GR signals are recovered with GR templates, but results are qualitatively the same with G4v injections recovered with G4v templates, or either kind of injection recovered with model-independent templates. The collection of points at $h_{\rm inj} = 0$ is noise-only and the detection threshold (horizontal line) is placed above $\alpha_{\rm n} = 99.9\%$ of them. The shaded band includes $\alpha_{\rm b} = 95\%$ of the data points above the threshold and it is centered on their best-fit line. The fit forced null y-intersect.

7.4 Analysis

We wish to detect any CW signal originating in a given pulsar, regardless of its polarization in a model-independent way. We can then determine whether the measured polarization content agrees with theoretical predictions. This information can be used to obtain frequentist confidence levels for a potential detection and to generate upper limits for the strength of signals of any polarization potentially buried in the data.

In order to test the statistical properties of the noisy data filtered through our templates, we produce numerous instantiations of detector noise by taking actual data processed as outlined in Sec. 7.3 and reheterodyning over a small band close to the frequency of the original heteredoyne. Any true signal in the data stream is scrambled in the process and what remains is a good estimate of the noise. This allows us to perform searches under realistic conditions with or without injections of simulated signals, while remaining blind to the presence of a true signal.

By heterodyning at different frequencies, we are able to generate a large number of instantiations of the data. Because our S5 datasets span roughly 1.9 years and are sampled once per minute, our bandwidth is 8.3×10^{-3} Hz with a lowest resolvable

frequency of 1.7×10^{-8} Hz. This means we could theoretically reheterodyne our data at a maximum of $8.3 \times 10^{-3}/1.7 \times 10^{-8} \approx 4.9 \times 10^{5}$ independent frequencies. In our study, we picked 10^4 frequencies in the $10^{-7} - 10^{-3}$ Hz range, avoiding the expected signal frequency of $\sim 10^{-5}$ Hz (period of a sidereal day) and its multiples.

We quantify the results of a particular search by looking at the obtained recovered signal strength, Eq. (7.41), and significance, Eq. (7.42). As expected, these two parameters are strongly correlated (Fig. 7.5). However, the significance is, in the presence of Gaussian noise, a direct indicator of goodness-of-fit and can be used to compare results from templates with different numbers of degrees of freedom.

By performing searches on multiple instantiations of noise-only data, we construct cumulative distribution function (CDF) probability plots showing the distribution of recovered signal strength, Eq. (7.41), and significance, Eq. (7.42), corresponding to a given template. Such plots give the probability that the outcome of the regression is consistent with noise (i.e. provide p-values). As shown in Fig. 7.6, an instantiation that contains a loud injected signal becomes manifest in this plot as an outlier. This sort of plot can also be used when searching for an actual signal in the data—namely, when looking at the original, non-reheterodyned series. In that case, the $1 -$ CDF curve can be extrapolated or interpolated to find the p-value corresponding to the significance with which the injection was recovered.

After injecting and retrieving increasingly loud signals with a given polarization content in different background instantiations, we produce plots of recovered strength vs injected strength ($h_{\rm rec}$ vs $h_{\rm inj}$) and significance vs injected strength (s vs $h_{\rm inj}$). Recall that injections are of the form of Eqs. (7.43, 7.44). Examples of such plots are presented in Fig. 7.7. These plots, and corresponding fits, can be used to assess the sensitivity of a template to certain type of signal, define thresholds for detection and produce confidence bands for recovered parameters. (In the frequentist literature, these plots are sometimes referred to as *Neyman constructions* [151].)

We define a horizontal *detection threshold* line above an arbitrary fraction $\alpha_{\rm n}$ (e.g., $\alpha_{\rm n} = 99.9\%$) of noise-only points (i.e. points with $h_{\rm inj} = 0$, but $h_{\rm rec} \neq 0$), so that data points above this line can be considered detected with a p-value of $p = 1 - \alpha_{\rm n}$ (e.g., $p = 0.1\%$). For a particular template, this fractional threshold can be directly translated into a significance value $s_{\alpha_{\rm n}}$ (e.g., $s_{99.9\%} = 2.5$). The sensitivity of the template is related to the number of injections recovered with a significance higher than $s_{\alpha_{\rm n}}$. Therefore, for a given $\alpha_{\rm n}$, a lower $s_{\alpha_{\rm n}}$ means higher sensitivity to true signals.

For the results of each template, the fractional threshold $\alpha_{\rm n}$ can also be associated to a strain value. We define this to be the loudness of the minimum injection detected above this threshold with some arbitrary *upper-limit confidence* $\alpha_{\rm up}$. This value can be determined from the s vs $h_{\rm inj}$ plot by placing a line parallel to the best fit but to the right of a fraction $\alpha_{\rm up}$ of all data points satisfying $0 < h_{\rm inj}$. The intersection of this line with the $\alpha_{\rm n}$ line occurs at $h_{\rm inj} = h_{\rm min}^{\alpha_{\rm up}}$, which is the strain value above which we can have $\alpha_{\rm up}$ confidence that a signal will be detected (i.e. recovered with significance $s > s_{\alpha_{\rm n}}$).

We refer to $h_{\rm min}^{\alpha_{\rm up}}$ as the *expected sensitivity* or *strain detection threshold at* $\alpha_{\rm n}$. This value allows not only for the definition of upper limits for the presence of signals, but also the comparison of different model dependent and independent templates. See Fig. 7.8b for a juxtaposition of the results of matching and nonmatching model-dependent templates for the case of the Crab pulsar.

The efficiency of a template is also quantified by the slope of the $h_{\rm rec}$ vs $h_{\rm inj}$ best-fit line, which should be close to 1 for a template that matches the signal. We perform this fit by taking into account only points above the $\alpha_{\rm n}$ line and forcing the y-intersect to be null. The deviations from this fit are used to produce confidence intervals for the recovered strength. This is done by defining a band centered on the best-fit line and enclosing an arbitrary fraction $\alpha_{\rm b}$ (e.g., $\alpha_{\rm b} = 95\%$) of the data points, corresponding to the confidence band placed around best-fit line. The intersection between this band and a horizontal line at some value of $h_{\rm rec}$ yields a confidence interval for the true strength with $\alpha_{\rm b}$ confidence. Note that deviations above and below the best-fit line are taken independently to obtain asymmetric confidence intervals. The same analysis can be done on the s vs $h_{\rm inj}$ plots, taking into account proper scaling of the best-fit slope.

In general, when performing injections we pick parameters with a uniform distribution over the uncertainty ranges of location and orientation values obtained from the ATNF Pulsar Catalog [152]. When there is no orientation information, we must draw ψ and ι from the ranges $[-\pi/2, \pi/2]$ and $[0, 2\pi]$ respectively. Note that standard searches consider tensor signals (2ψ-dependent) only and therefore assume $\psi \in [-\pi/4, \pi/4]$; however, a bigger range must be used when taking into account vector signals (ψ-dependent). The reason these ranges need not cover the full $[0, \pi]$ range is that a change in ψ of $\pi/2$ for tensor and π for vector signals is equivalent to a change of signal sign. Therefore, this is taken care of by varying the overall phase $\phi_0 \in [0, \pi]$.

We tested the aforementioned methods on LIGO data taken by the Hanford and

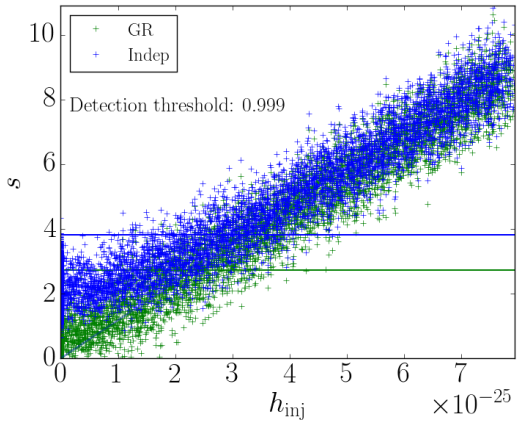

(a) GR injections; GR and independent recovery

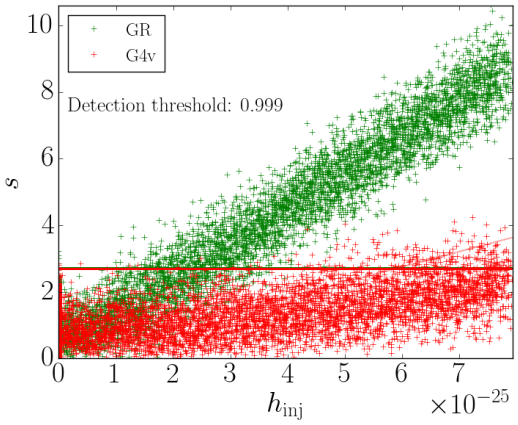

(b) GR injections; GR and G4v recovery

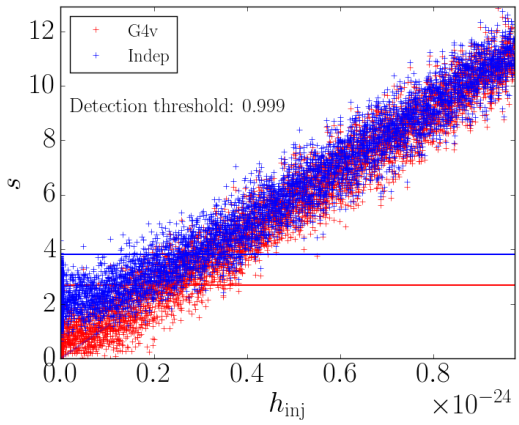

(c) G4v injections; G4v and independent recovery

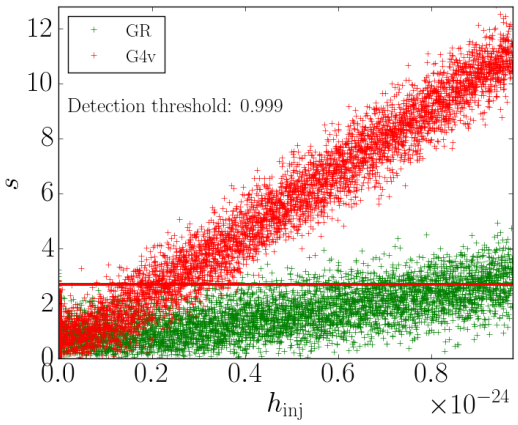

(d) G4v injections; G4v and GR recovery

Figure 7.8: GR (top) and G4v (bottom) injection results of search over LIGO S5 H1 data heterodyned for the Crab pulsar. Plots show significance, Eq. (7.42), vs injected strength. Color corresponds to the template used for recovery: GR, green; G4v, red; model-independent, blue. This particular search was performed using 10^4 instantiations, half of which contained injections using the values of ι and ψ given in Table 7.1. The model-dependent templates assumed the same same ι as the injections. Horizontal lines correspond to a detection threshold $\alpha_{\rm n} = 99.9\%$.

Livingston detectors over LIGO Science Run 5 (S5). During this run, which took place from November 2005 through September 2007 (GPS times 815155213 - 875232014), the three LIGO detectors operated in data-taking mode at design sensitivity, collecting a year of coincident detector data. The root-mean-square strain noise of the instrument reached values as low as 3×10^{-22} for bands of 100 Hz over the most sensitive frequencies [153]. LIGO S5 data has been recently released to the public and is accessible online through the LIGO Open Science Center [154].

In particular, we looked at data for 115 pulsars, obtained by reducing S5 H1, H2 and L1 strain data as outlined in Sec. 7.3.2. But for the inclusion of PSR J0024-72040 and the exclusion of PSR J2033+17 and Vela, these are the same heterodyned time series analyzed in reference [150]. However, that study presented Bayesian upper limits to the presence of GR signals and did not consider alternative polarizations.

7.5 Results

Here we present the results of a study of the signal sensitivity of the analysis procedure described in Sec. 7.4, using the data described at the end of Sec. 7.4. We perform a "closed box" analysis, using only reheterodyned data, which are insensitive to the presence of actual signals, and simulated signal injections. A full "open box" analysis, using Bayesian methods to produce model-dependent and model-independent signal detection confidence bands or upper limits, is in preparation.

In particular, we produced 10^4 reheterodyned instantiations of data for each pulsar by picking linearly spaced frequencies in the $10^{-7} - 10^{-3}$ Hz range (cf. sec. 7.4). Half of those were injected with simulated signals of increasing strength. The data were then analyzed with each template (GR, G4v and model-independent), producing plots like those in Fig. 7.7. For the Crab pulsar, since the source orientation information is known, the full model-dependent templates, Eqs. (7.43, 7.44), were used; otherwise, the semi-model-dependent templates, Eqs. (7.45, 7.48), were used. The whole process was carried out for both GR and G4v injections. In all cases, we set $\alpha_{\rm n} = 99.9\%$ and $\alpha_{\rm up} = 95.0\%$.

7.5.1 Crab pulsar

Results for searches over H1 S5 data prepared for the Crab pulsar ($\nu = 30.22$ Hz, $\nu_{\rm GW} = 60.44$ Hz) are presented in Fig. 7.8. The results using templates matched to the injections are compared to those of the model-independent (left) and nonmatching

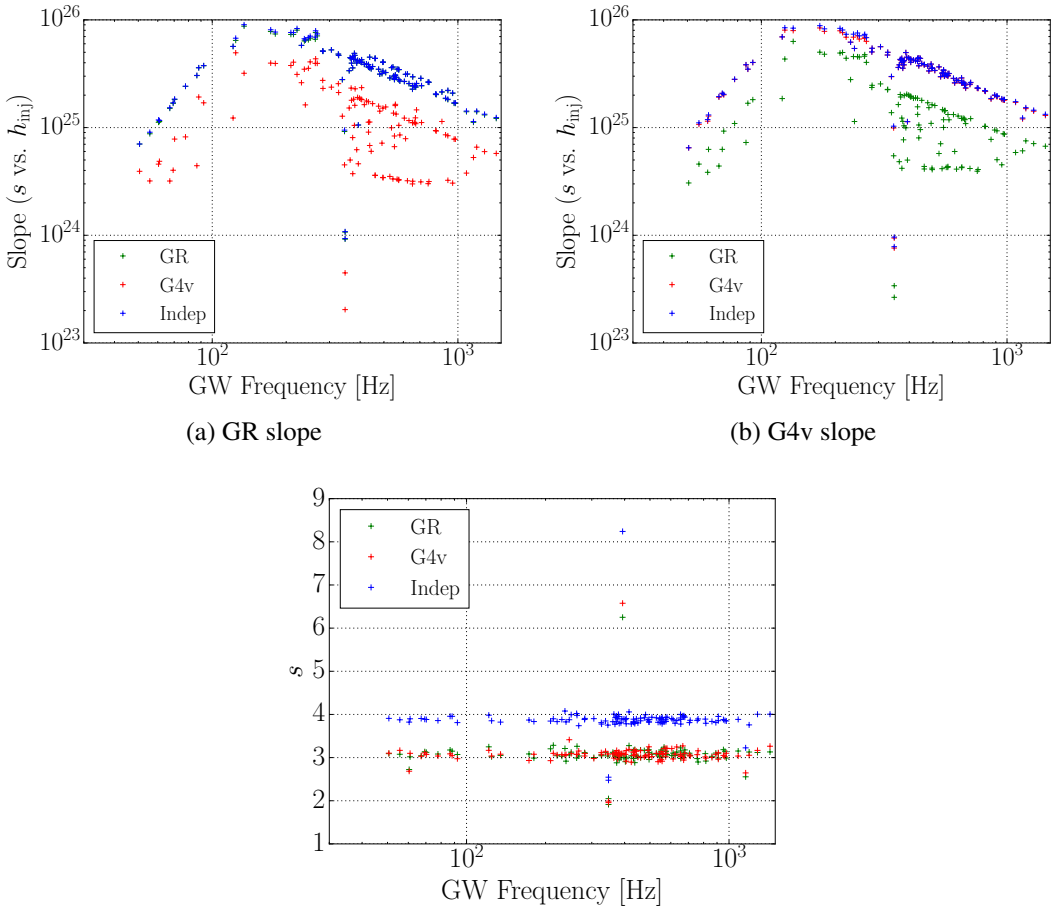

(a) GR slope

(b) G4v slope

(c) Detection threshold

Figure 7.9: Slope of the s vs $h_{\rm inj}$ best-fit-line (left and center) and significance detection threshold at $\alpha_{\rm n} = 99.9\%$ (right) vs GW frequency and for GR and G4v injections on S5 H1 data for 115 pulsars. Color corresponds to search template: GR, green; G4v, red; and model-independent, blue. Note that for both kinds of injections, the model-independent points overlap the matching template.

templates (right). The expected sensitivities, as defined in Sec. 7.4, for each injection template and search model are provided in Table 7.3. Recall that the Crab is a special case, since its orientation in the sky is well-known, which enables us to use full model-dependent templates, Eqs. (7.43, 7.44). However, searches for actual signals would still have to make use of two templates for each theoretical model because of the ambiguity in ι described in Sec. 7.2.2. In order to avoid doing this, a semi-model-dependent or model-independent search could be carried out instead.

A number of interesting observations can be drawn from Fig. 7.8 and Table 7.3. As inferred from the values of $h_{\rm min}$, the model-independent template is roughly 25%

Table 7.3: Summary of expected sensitivity for the Crab pulsar S5 H1 searches ($\alpha_{\rm n}$ = 99.9%, $\alpha_{\rm up}$ = 95.0%). Rows correspond to injection type and columns to search template. The rotational frequency of the Crab is ν = 30.22 Hz and, therefore, $\nu_{\rm GW}$ = 60.44 Hz.

	GR	G4v	Independent
GR	3.41×10^{-25}	7.49×10^{-25}	4.20×10^{-25}
G4v	8.90×10^{-25}	3.30×10^{-25}	4.15×10^{-25}

less sensitive than the matching one, regardless of the theory assumed when making injections. This is understood by the presence of four extra degrees of freedom in the model-independent template, compared to the single tunable coefficient in the full model-dependent one. If instead the semi-model-dependent template with two degrees of freedom is used, the improvement with respect to the model-independent search goes down to 15%. In any case, the accuracy of matching and model-independent searches, given by the width of the confidence bands, are almost identical.

Model dependent templates are significantly less sensitive to nonmatching signals. Table 7.3 indicates that model-dependent templates are 120-170% less sensitive to nonmatching signals than their matching counterpart. A consequence of this is the existence of a range of signals which would be detected by templates of one theory, but not the other (see Figs. 7.8b & 7.8d). This is particularly interesting given that previous LIGO searches assume GR to be valid and use a template equivalent to Eq. (7.43). Therefore, our results suggest it is possible that those searches might have missed fully-non-GR signals buried in the data (see Sec. 7.6 for further discussion).

7.5.2 All pulsars

The Crab pulsar is only one of the 115 sources we analyzed. The results, presented in Figs. 7.9a & 7.9b generally confirm the observations anticipated from the Crab. While model-independent searches are of the same accuracy as matching semi-model-dependent ones, their strain detection threshold is louder due to the extra degrees of freedom (Fig. 7.9c). Consequently, model-independent templates demand a higher significance to be able to distinguish a signal from noise. The detection thresholds for GR and G4v templates are of the same magnitude, since both have the same number of degrees of freedom. Among all the 115 pulsars, the sources with best expected sensitivities to GR and G4v signals were PSR J1603-7202 and PSR J1748-2446A respectively (see Table 7.4).

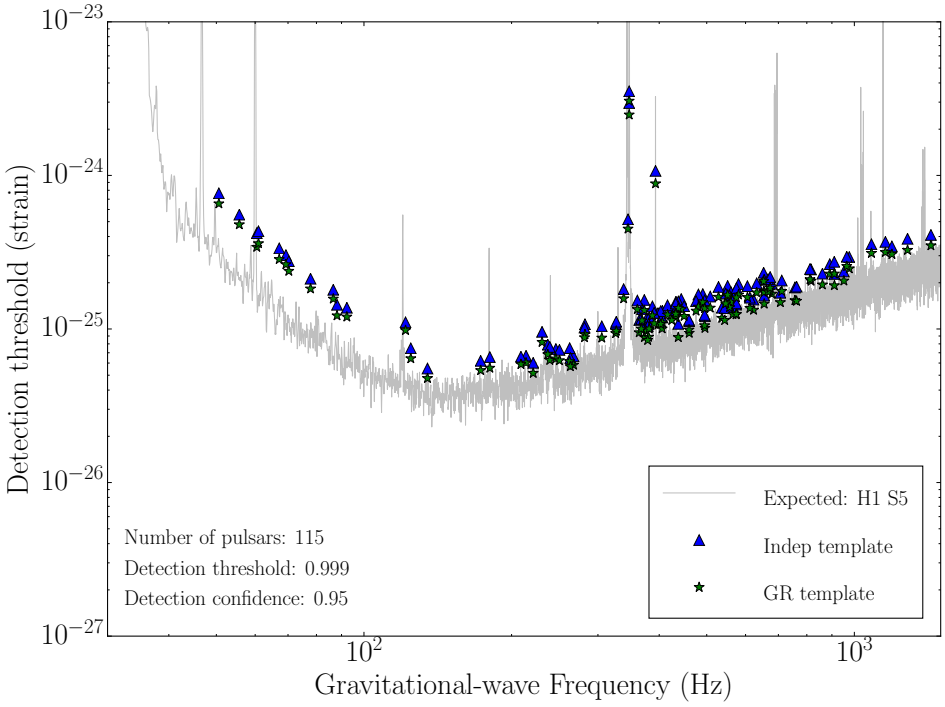

(a) GR injections

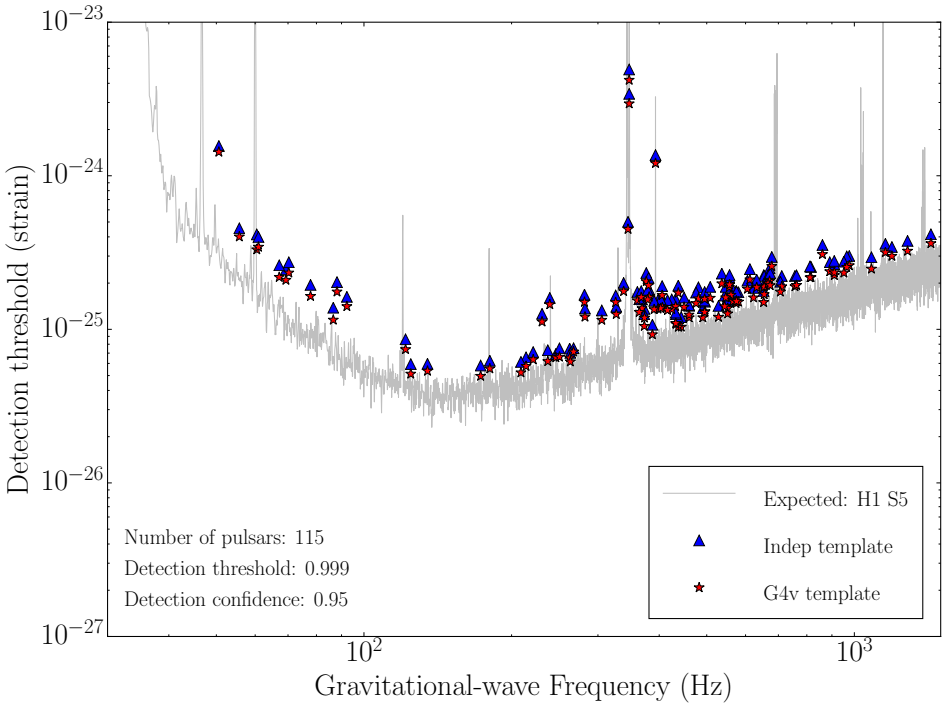

(b) G4v injections

Figure 7.10: S5 H1 expected sensitivity (strain detection threshold at $\alpha_{\rm n} = 99.9\%$ with $\alpha_{\rm up} = 95.0\%$ confidence) vs GW frequency for 115 pulsars. Color corresponds to search template: GR, green; G4v, red; and model-independent, blue. The gray line is the anticipated sensitivity of a standard Bayesian search, Eq. (7.55).

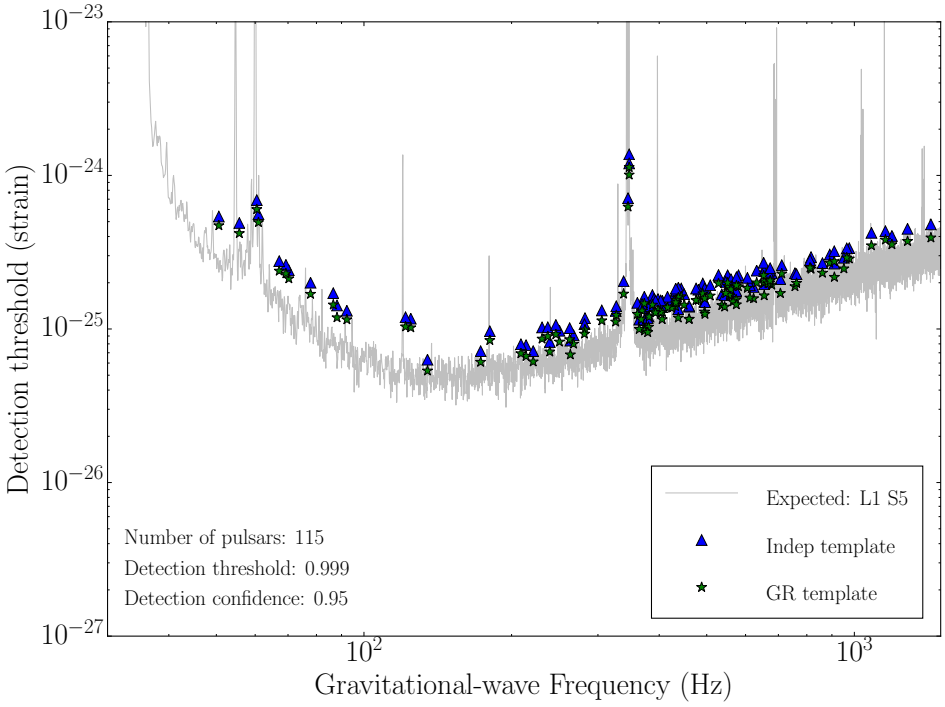

(a) GR injections

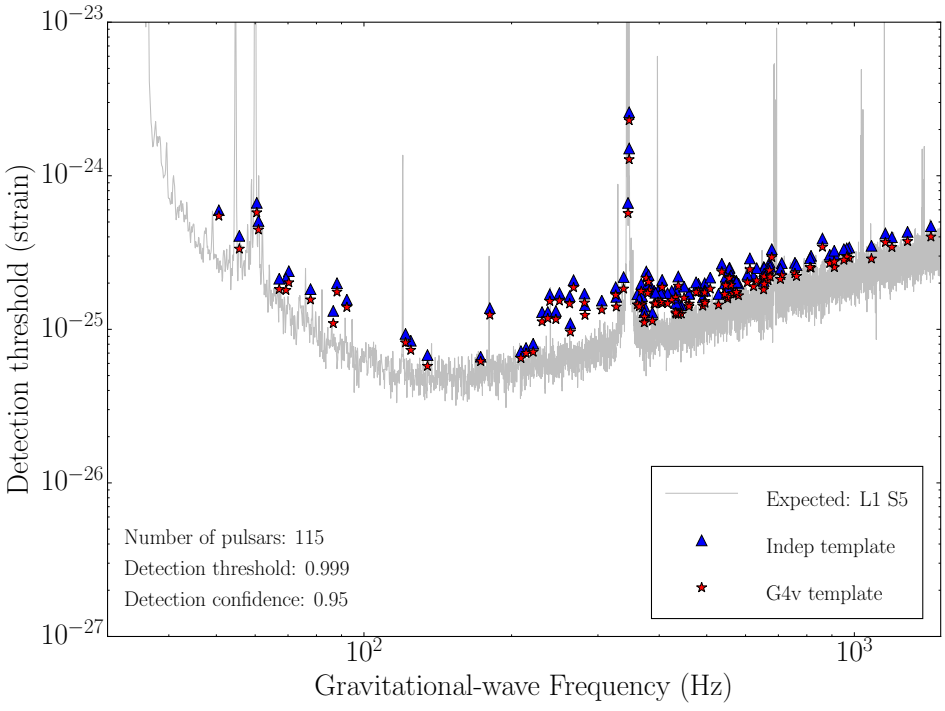

(b) G4v injections

Figure 7.11: S5 L1 expected sensitivity (strain detection threshold at $\alpha_{\rm n} = 99.9\%$ with $\alpha_{\rm up} = 95.0\%$ confidence) vs GW frequency for 115 pulsars. Color corresponds to search template: GR, green; G4v, red; and model-independent, blue. The gray line is the anticipated sensitivity of a standard Bayesian search, Eq. (7.55).

Table 7.4: Best expected sensitivities for S5 H1 searches ($\alpha_n = 99.9\%$, $\alpha_{up} = 95.0\%$). Rows correspond to injection type and columns to pulsar name (PSR), rotation frequency (ν) and strain detection threshold for matching dependent (h_{dep}) and independent (h_{indep}) templates.

	PSR	ν (Hz)	h_{dep}	h_{indep}
GR	J1603-7202	67.38	4.77×10^{-26}	5.53×10^{-26}
G4v	J1748-2446A	86.48	4.96×10^{-26}	5.81×10^{-26}

The key results of our study are summarized in Fig. 7.10 for H1 and Fig. 7.11 for L1. These plots present the expected sensitivity (strain detection threshold at $\alpha_n = 99.9\%$ with $\alpha_{up} = 95.0\%$ confidence) vs GW frequency ($\nu_{GW} = 2\nu$). The outliers seen in Figs. 7.9-7.11 correspond to pulsars whose value of ν_{GW} are very close to instrumental noise spectral lines associated with violin resonances of the detectors' test mass pendulum suspensions.

For the matching or model-independent templates, the resulting data points trace the noise curve of the instrument; however, due to the long integration time, we are able to detect signals below LIGO's standard strain noise. The gray curve shown in Figs. 7.10, 7.11 represents the expected sensitivity of a regular Bayesian GR search (e.g., [150]). This is proportional to the amplitude spectral density of the detector and inversely proportional to the square root of the observation time. The particular empirical relationship used to generate the curve in Figs. 7.10 & 7.11 is:

$$\langle h_{min} \rangle = 10.8\sqrt{S_n(f)/T}, \tag{7.55}$$

with $S_n(f)$ the noise power spectral density and T the total observation time (527 days for S5 H1 and 405 days for S5 L1) [57]. This formula enables the comparison of the methods presented here with the expected performance of standard Bayesian searches.

By the same token, we can define a figure of merit ρ for our searches by the ratio:

$$\rho\,(\nu_{GW}) = h_{min}/\sqrt{S_n(\nu_{GW})/T}. \tag{7.56}$$

The average of this value over all pulsars, $\langle\rho\rangle$, can be semiquantitatively compared to the 10.8 prefactor in Eq. (7.55). The equivalence is not direct because, besides the intrinsic differences between Bayesian and frequentist approaches, Eq. (7.55) was obtained by averaging the results of 4000 simulated searches [57], while we include just the 115 pulsars at hand. The values of $\langle\rho\rangle$ for our S5 H1 & L1 analyses are

Table 7.5: Average sensitivity ratios $\langle\rho\rangle$, Eq. (7.56), for S5 H1 (first value) and S5 L1 (second value) searches. Rows correspond to injection type and columns to search template.

	GR		G4v		Independent	
GR	16.11	14.65	58.53	51.89	18.83	17.15
G4v	61.21	55.06	18.42	16.76	21.24	19.32

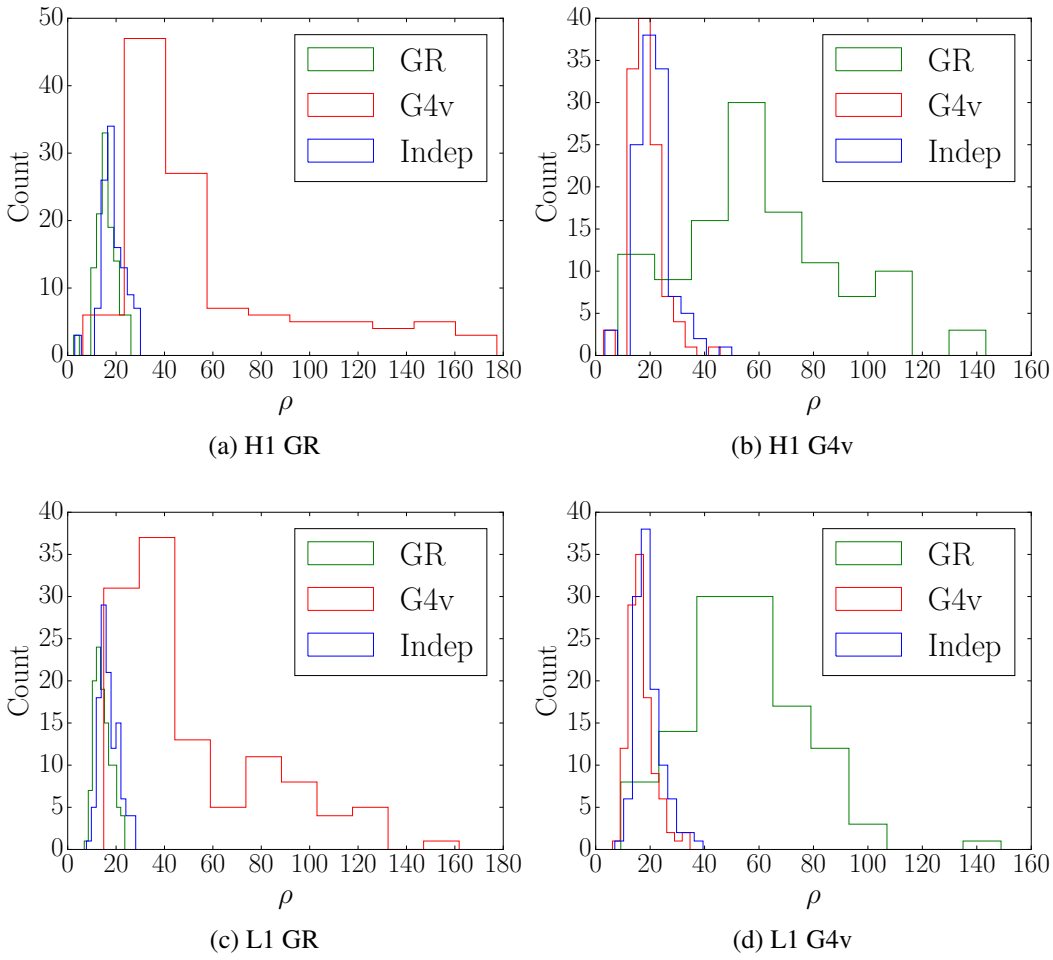

(a) H1 GR

(b) H1 G4v

(c) L1 GR

(d) L1 G4v

Figure 7.12: Histograms of the figure of merit ρ, Eq. (7.56), for our searches over S5 H1 (top) and L1 (bottom) data sets with GR (left) and G4v (right) injections, corresponding to 115 pulsars. Color corresponds to search template: GR, green; G4v, red; and model-independent, blue.

Table 7.6: Sensitivity ratio ρ, Eq. (7.56) evaluated at the Crab's GW frequency, for S5 H1 (first value) and S5 L1 (second value) searches. Rows correspond to injection type and columns to search template.

	GR		G4v		Independent	
GR	20.75	10.40	45.52	27.15	25.54	11.94
G4v	54.06	20.30	20.07	9.96	25.21	11.52

presented in Table 7.5 and Fig. 7.12. The specific values for the Crab pulsar are shown in Table 7.6. A smaller ρ indicates better performance.

As mentioned above, the remarks made about the Crab pulsar hold for most other sources, except that detectability is slightly lower because orientation parameters are unknown. In all cases, the matching template is the best at recovering signals, followed closely by the model-independent one. Searches that assume the incorrect model are substantially less efficient and their $h_{\rm min}$ vs $\nu_{\rm GM}$ curves do not follow the instrumental noise line. This is reflected, for instance, by the figures of merit presented in Table 7.5.

7.6 Conclusions

We have developed novel model-independent methods to search for CW signals coming from targeted sources in LIGO-like interferometric data. These searches are able to detect signals of *any* polarization content with high significance.

In order to test our methods in the presence of realistic noise conditions, we implemented a procedure to produce thousands of noise-only instantiations from actual data. We then proceeded by injecting and retrieving increasingly loud signals of different polarization content.

We studied 115 pulsars using S5 data from the LIGO Hanford and Livingston detectors. Although the methods are general, we restricted our study to two theories that predict starkly different GW polarization contents (GR and G4v).

Our results indicate that assuming the wrong theoretical model greatly reduces the sensitivity of a search to signals buried in the data. Yet, our model-independent searches are almost as effective as the model-dependent templates that match the kind of signal injected (i.e. when the models used for injection and search are the same). This means that our model-independent templates can be used to find signals of any polarization without additional computational requirements.

We are able to reach sensitivities comparable to previous studies, although slightly worse than those presented in [150]. This is probably due to our making use of a single detector and to differences between frequentist and Bayesian approaches.

We have shown that, for some combinations of detectors, sources, and signal strengths, G4v signals are invisible to GR templates and viceversa. Therefore, it is possible that, if GWs are composed uniquely of vector modes, previous LIGO searches, which assume GR, may have missed their signals.

It is clear that the next step in this study consists of incorporating our model-independent templates into the Bayeasian machinery used in standard LIGO Scientific Collaboration searches. This will allow us to properly marginalize over all nuisance parameters and to produce multi-detector model-dependent and model-independent signal detection confidence bands or upper limits. We will also employ methods to constrain other theories (e.g., scalar-tensor) in the event of a model-independent detection.

Acknowledgments

The authors would like to thank Holger Pletsch for helpful discussions. M. Pitkin is funded by the STFC through grant number ST/L000946/1. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. This paper carries LIGO Document Number LIGO-P1400169.

7.7 Appendix: Statistical properties of LIGO data

The χ^2 minimization is equivalent to a maximum likelihood procedure only in the presence of Gaussian noise. When this requirement is not satisfied, the regression is still valid, but the χ^2 values resulting from the fit will be distributed in a nontrivial way, rather than the χ^2 distribution expected in the case of Gaussian noise. Furthermore, the relationship between the covariance matrix of the system and the standard uncertainties of the recovered coefficients becomes unclear. Therefore, it is important to statistically characterize the data and understand the limitations of our assumption of Gaussianity.

When taken as a whole, LIGO detector noise does not conform to a *stationary* Gaussian distribution. This can be visually confirmed by means of a histogram, as shown in Fig. 7.13 for the case of S5 H1 data prepared for the Crab. The divergence

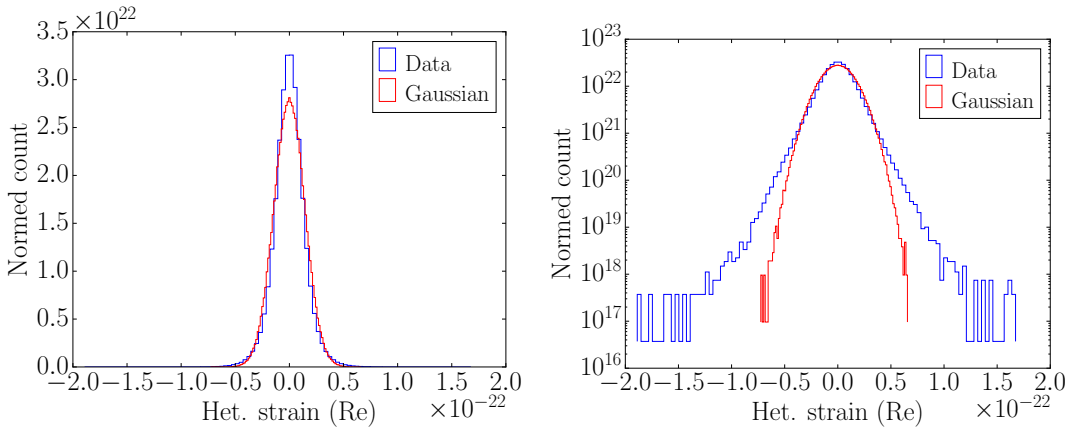

Figure 7.13: Normalized histogram of the real part of S5 H1 data heterodyned for the Crab in linear (left) and logarithmic y scales. A Gaussian curve with the same standard deviation is plotted in red for comparison.

from Gaussianity is evident from the long tails, seen most clearly in the log-y version of the plot. As expected, the data fail more rigorous standard Gaussianity tests, such as the Kolmogorov-Smirnov (KS) or the Anderson-Darling (AD) tests.

However, it is possible to split up the data into day-long (or shorter) segments, as was described in Sec. 7.3.2, so as to study the Gaussianity of the data on a day-to-day basis. The results of the KS and AD tests for each day-segment, together with those for reference Gaussian noise series, are presented in Figs. 7.14a and 7.14b respectively. The KS test returns the p-value for a null hypothesis that assumes the data are normally distributed; therefore, a lower p-value implies a higher probability that the data are not Gaussian [155]. The AS test returns a figure of merit which is indirectly proportional to the significance with which the hypothesis of Gaussianity can be rejected; therefore a higher AS statistic implies a higher probability that the data are not Gaussian [156].

It can be seen from the results of these tests that the statistical properties of the segments vary considerably from day to day. This could have been guessed from the non-stationarity of the data in Fig. 7.3, the daily variation of the standard deviation (Fig. 7.4) and other irregularities of the data. Nonetheless, most of the segments seem to pass the Gaussianity tests, with some remarkable exceptions around the days 250–400 of the run. This corresponds to the spiking observed in the heterodyned data (GPS times $8.4 \times 10^8 - 8.5 \times 10^8$ in Fig. 7.3).

In order to confirm that our assumption of Gaussianity is not too far from reality, we repeated our analysis (see Sec. 7.5) on sets of synthetic Gaussian noise. In order to do

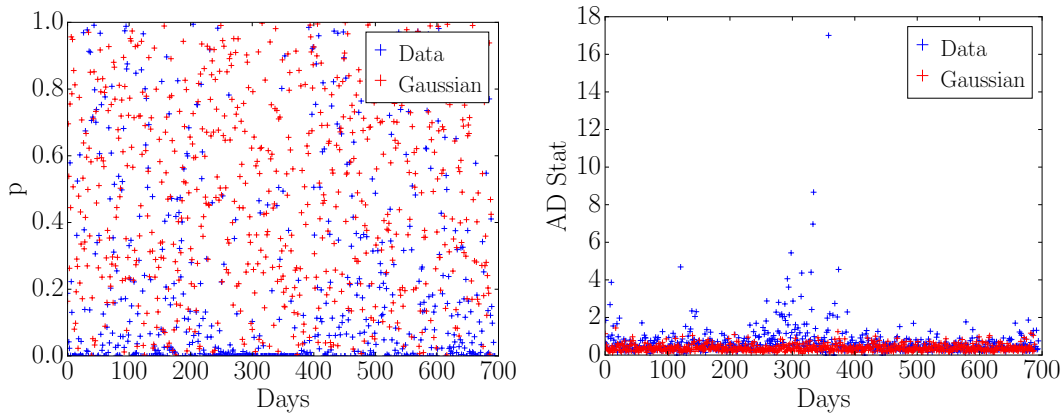

(a) Kolmogorov-Smirnov

(b) Anderson-Darling

Figure 7.14: Results of Gaussianity tests for daily segments of S5 H1 data prepared for the Crab (blue). The results for corresponding sets of Gaussian noise are presented for comparison (red).

this, for each pulsar we generated streams of complex-valued data randomly selected from a normal distribution with the same standard deviation as the corresponding original LIGO data set. These series replaced the instantiations of reheterodyned data, but the search process was otherwise unchanged. The results of this comparison for S5 H1 are shown in Figs. 7.15, where we juxtaposed expected sensitivities obtained using Gaussian noise and actual LIGO noise (cf. Sec. 7.5). These plots confirm that, indeed, we obtain qualitatively the same results with Gaussian noise as with actual LIGO data.

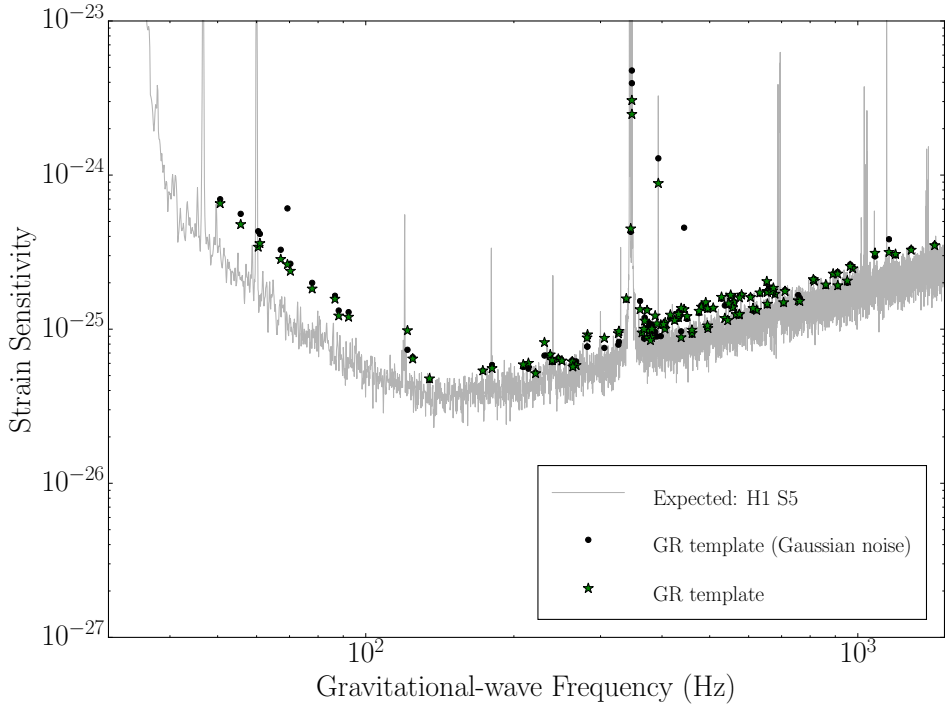

(a) GR injections in Gaussian noise

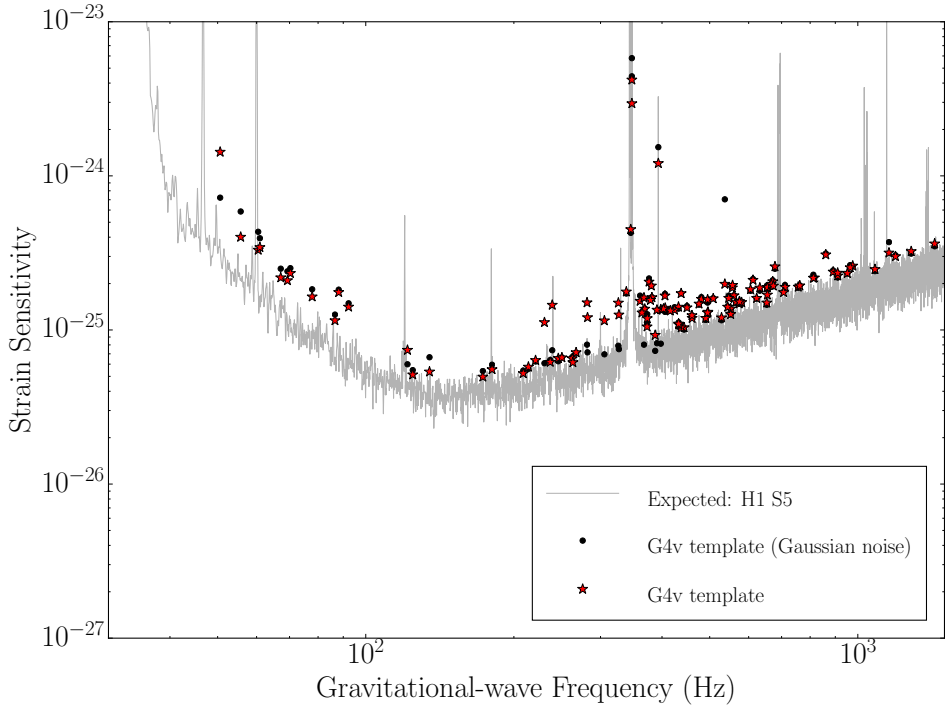

(b) G4v injections in Gaussian noise

Figure 7.15: Expected sensitivity ($\alpha_{\mathrm{n}} = 99.9\%$, $\alpha_{\mathrm{up}} = 95.0\%$) vs GW frequency. Comparison between fabricated Gaussian noise and actual LIGO noise. Searches were made with semi-model-dependent templates, Eqs. (7.45, 7.48). The colored stars correspond to actual LIGO H1 noise (cf. Fig. 7.10), while the black dots correspond to fabricated Gaussian noise.

Chapter 8

PROBING DYNAMICAL GRAVITY WITH THE POLARIZATION OF CONTINUOUS GRAVITATIONAL WAVES

M. Isi, M. Pitkin, and A. J. Weinstein, Phys. Rev. **D96**, 042001 (2017), M.I. carried out the analysis and wrote the manuscript, arXiv:1703.07530 [gr-qc] .

8.1 Introduction

The recent detection of gravitational waves (GWs) by the advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) heralds the beginning of the long-awaited era of GW astronomy [1, 2]. One of the main goals of this field is to use GWs as a probe of fundamental physics in the highly dynamical and strong-field regimes of gravity, as predicted by the general theory of relativity (GR). The first few GW detections have already been used to place some of the most stringent constraints on deviations from GR in this domain, which is inaccessible to laboratory, Solar System or cosmological tests of gravity.

However, it has not been possible to use LIGO signals to learn about the polarization content of GWs [82], a measurement highly relevant when comparing GR to many of its alternatives [83, 84]. The reason for this is that the relative orientation of the two LIGO detectors makes it nearly impossible to unequivocally characterize the polarizations of transient GW signals like the compact-binary coalescences (CBCs) observed so far. In fact, at least five noncoaligned quadrupolar detectors would be needed to break the degeneracies of all five nondegenerate polarizations allowed by generic metric theories of gravity [97, 98].

Existing observations that are usually taken to constrain the amount of allowed non-GR polarizations can do so only in an indirect and strongly model-dependent manner. For example, measurements of the orbital decay of binary systems are sensitive to the total radiated GW power but do not probe the waves directly (see e.g. [86, 87], or [88, 89] for reviews). In the context of specific alternative theories (e.g. scalar-tensor) such observations can indeed constrain the power in extra polarizations; however, they provide no direct, model-independent information on the actual polarization content of the gravitational radiation. Thus, there may be multiple theories, with

different polarization content, that still predict the correct observed GW emitted power. Because other traditional tests of GR (like Solar System tests) have no bearing on GWs, there currently exist no direct measurements of GW polarizations.

Unlike CBC transients, continuous gravitational waves (CWs) are, by definition, long-lasting narrow-band signals. Although they have not yet been observed [29, 43–47], CWs are expected to be emitted by stable systems, like spinning neutron stars with an asymmetric moment of inertia [48]. If detected, such signals would allow for tests of gravity complementary to those achievable with transients, including the study of GW polarizations [109].

In [109] we showed that it is possible to search for CWs in a polarization-agnostic way and to disentangle the polarization content if a signal is present. However, the data analysis methods proposed were based on a frequentist approach to statistics and suffered from the associated limitations. In this paper, we reframe the ideas of [109] in a more sophisticated Bayesian framework that allows us to achieve the following novel goals:

1. *Model-independent detection:* determine whether a set of GW detector data, prepared for any given known pulsar and from one or multiple detectors, provides evidence for the presence of an astrophysical signal of any polarization content.

2. *Model selection:* in the presence of a signal, determine whether the data favor GR or a generic non-GR model, as well as comparing specific alternative theories among themselves and to GR; combine data for multiple sources into a single statement about the validity of GR.

3. *Inference:* if the data favor the presence of a GR signal, place constraints on specific alternative theories using the tools of Bayesian parameter estimation.

Furthermore, while [109] treated only the case of a single detector, we are now able to consider the generic case of a network of detectors.

We present Bayesian methods to achieve the three goals above in the context of searches targeted to known pulsars and present sensitivity estimates for the advanced detector era, including the first generic estimates of sensitivity to nontensorial CW polarizations ever published. In Sec. 8.2, we review the basics of beyond-Einstein polarizations and the targeted pulsar CW search. In Sec. 8.3, we phrase our problem

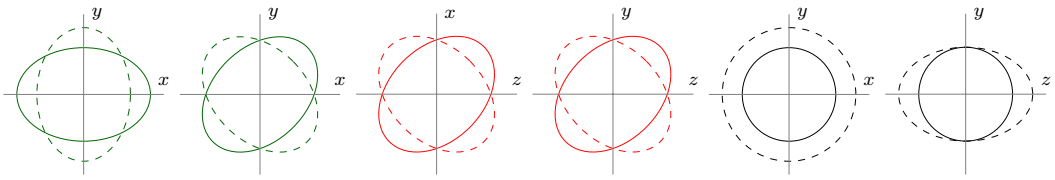

Figure 8.1: *Effect of different GW polarizations on a ring of free-falling test particles*. Plus (+) and cross (×) tensor modes (green); vector-x (x) and vector-y (y) modes (red); breathing (b) and longitudinal (l) scalar modes (black). In all of these diagrams the wave propagates in the z direction. This decomposition into polarizations was first proposed for generic metric theories in [98].

in the language of model selection and explain the construction of hypotheses that will allow us to distinguish GR from non-GR signals. In Sec. 8.4 we specify the details of our analysis, and we explain our results in Sec. 8.5. Finally, we summarize our findings and explain caveats in Sec. 8.6.

8.2 Background

8.2.1 Polarizations

GWs can be decomposed into different polarizations, which arise from the linearly independent components of the three-dimensional, rank-two tensor representing the spatial metric perturbation [98]. A generic metric theory of gravity may thus allow any combination of up to six independent modes: plus (+), cross (×), vector x (x), vector y (y), breathing (b) and longitudinal (l). The effect of each of these modes is represented in Fig. 8.1. The rotational properties of the fields underlying any given theory determines which polarizations the theory supports: + and × correspond to tensor fields (helicity ±2), x and y to vector fields (helicity ±1), and b and l to scalar fields (helicity 0).

The components of the tensor and vector pairs are not separable, in the sense that a signal model that includes one element of the group must also include the other (e.g. it is not possible to have a model that allows plus + but not ×), because the distinction between + and ×, or x and y, is contingent on the frame of reference (e.g. relative orientation of source and detector).

Einstein's theory only allows the existence of the + and × polarizations. On the other hand, scalar-tensor and massive-graviton theories may also predict the presence of some b and/or l component associated to the theory's extra scalar field [84]. On top of tensor and scalar modes, bimetric theories, like Rosen or Lightman-Lee theories,

also predict vector modes [85]. Furthermore, less conventional theories might, in principle, predict the existence of vector or scalar modes *only*, while still possibly being in agreement with all other non-GW tests of GR (see [106] for an example). Although all these different theoretical frameworks serve as motivation for our study, our approach to the measurement of GW polarizations is phenomenological and, thus, theory-agnostic (Sec. 8.3). It is important to underscore that the detection of a GW signal with a non-GR polarization, no matter how small, is sufficient to falsify GR (note the converse is not true, however).

Because different polarizations have geometrically distinct effects, GW detectors will react differently to each mode. This is manifested in the detector response function F_p for each polarization p, which encodes the effect of a linearly p-polarized GW with unit amplitude, $h_p = 1$. Ground-based GW detectors, like LIGO and Virgo are quadrupolar antennas that perform low-noise measurements of the strain associated with the differential motion of two orthogonal arms. Their detector response function can thus be written as [107–110]:

$$F_+ = \frac{1}{2}\left[(\mathbf{w}_x \cdot \mathbf{d}_x)^2 - (\mathbf{w}_x \cdot \mathbf{d}_y)^2 - (\mathbf{w}_y \cdot \mathbf{d}_x)^2 + (\mathbf{w}_y \cdot \mathbf{d}_y)^2\right], \tag{8.1}$$

$$F_\times = (\mathbf{w}_x \cdot \mathbf{d}_x)(\mathbf{w}_y \cdot \mathbf{d}_x) - (\mathbf{w}_x \cdot \mathbf{d}_y)(\mathbf{w}_y \cdot \mathbf{d}_y), \tag{8.2}$$

$$F_{\mathrm{x}} = (\mathbf{w}_x \cdot \mathbf{d}_x)(\mathbf{w}_z \cdot \mathbf{d}_x) - (\mathbf{w}_x \cdot \mathbf{d}_y)(\mathbf{w}_z \cdot \mathbf{d}_y), \tag{8.3}$$

$$F_{\mathrm{y}} = (\mathbf{w}_y \cdot \mathbf{d}_x)(\mathbf{w}_z \cdot \mathbf{d}_x) - (\mathbf{w}_y \cdot \mathbf{d}_y)(\mathbf{w}_z \cdot \mathbf{d}_y), \tag{8.4}$$

$$F_{\mathrm{b}} = \frac{1}{2}\left[(\mathbf{w}_x \cdot \mathbf{d}_x)^2 - (\mathbf{w}_x \cdot \mathbf{d}_y)^2 + (\mathbf{w}_y \cdot \mathbf{d}_x)^2 - (\mathbf{w}_y \cdot \mathbf{d}_y)^2\right], \tag{8.5}$$

$$F_{\mathrm{l}} = \frac{1}{2}\left[(\mathbf{w}_z \cdot \mathbf{d}_x)^2 - (\mathbf{w}_z \cdot \mathbf{d}_y)^2\right]. \tag{8.6}$$

Here, the spatial vectors $\mathbf{d}_x$, $\mathbf{d}_y$ have unit norms and point along the detector arms such that $\mathbf{d}_z = \mathbf{d}_x \times \mathbf{d}_y$ is the local zenith; the direction of propagation of the wave from a source at known sky location (specified by right ascension α, and declination δ) is given by $\mathbf{w}_z$, and $\mathbf{w}_x$, $\mathbf{w}_y$ are such that $\mathbf{w}_z = \mathbf{w}_x \times \mathbf{w}_y$. We choose $\mathbf{w}_x$ to lie along the intersection of the equatorial plane of the source with the plane of the sky and let the angle between $\mathbf{w}_y$ and the celestial north be ψ, the *polarization angle*.

Because of their symmetries, the breathing and longitudinal modes are fully degenerate to networks of quadrupolar antennas (see e.g. Sec. VI of [85]). This means that no model-independent measurement with such a network can possibly distinguish between the two, so it is enough for us to consider just one of them explicitly. We will refer to the scalar modes jointly by the subscript "s".

The response of gravitational detectors to signals of a given polarization and direction of propagation can be represented, as in Fig. 8.2, by a spherical polar plot in which the radial coordinate corresponds to the sensitivity given by Eqs. (8.1–8.6). In the frame of a given detector, this can be written as [see e.g. Eqs. (13.98) in [110] with $\psi = -\pi/2$, to account for the different wave-frame definition]:

$$F_+(\vartheta, \varphi\,; \psi = 0) = -\frac{1}{2}\left(1 + \cos^2\vartheta\right)\cos 2\varphi\,, \tag{8.7}$$

$$F_\times(\vartheta, \varphi\,; \psi = 0) = -\cos\vartheta \sin 2\varphi\,, \tag{8.8}$$

$$F_{\rm x}(\vartheta, \varphi\,; \psi = 0) = -\sin\vartheta \sin 2\varphi\,, \tag{8.9}$$

$$F_{\rm y}(\vartheta, \varphi\,; \psi = 0) = \sin\vartheta \cos\vartheta \cos 2\varphi\,, \tag{8.10}$$

$$F_{\rm b/l}(\vartheta, \varphi\,; \psi = 0) = \mp\frac{1}{2}\sin^2\vartheta \cos 2\varphi\,, \tag{8.11}$$

where ϑ and φ are the polar an azimuthal coordinates of the source with respect to the antenna at any given time (with detector arms along the x and y-axes), and we have fixed the wave frame so that $\psi = 0$. The representation of Fig. 8.2 makes it clear that quadrupolar detectors will generally be more sensitive to some polarizations than others, although this will vary with the sky location of the source. For example, for all but a few sky locations, quadrupolar antennas will respond significantly less to a breathing signal than a plus or cross signal.

For a given detector, polarization angle and sky location, the antenna patterns of Eqs. (8.1–8.6) become simple, distinct functions of time determined by the rotation of the Earth. This can be pictured by noting that, as the Earth spins on its axis, the angular location of the source with respect to detector will change, tracing an arc on the surfaces of Fig. 8.2 with varying radial distance. As we explain in Sec. 8.3.1, the F_p's of polarizations with different rotational properties can be distinguished even in the absence of information on the source orientation; for the minority of cases in which such information exists, it can be taken into account to better distinguish among specific signal models (see Sec. 8.4).

Because their characteristic period (a sidereal day) is much longer than the CBC timescale (order of minutes or less), the F_p's are treated as constants in transient searches; however, this simplification is not valid for CW searches, since their coherent-integration time can be of the order of months to years. As we have pointed out before, this can be used to distinguish the polarization content of a signal [109]. Assuming wave frequency and speed are the same for all modes, the only differences between CWs of different polarizations arise from the sidereal-period amplitude modulations caused by each antenna pattern.

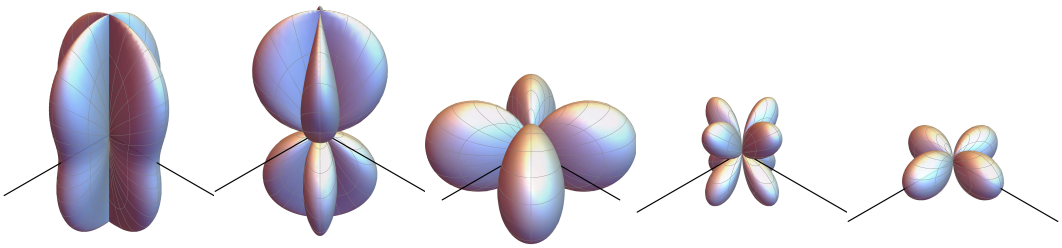

(a) Plus (+) (b) Cross (×) (c) Vector-x (x) (d) Vector-y (y) (e) Scalar (s)

Figure 8.2: *Angular response of a quadrupolar detector to each GW polarization*. The radial distance represents the response of a single quadrupolar antenna to a unit-amplitude gravitational signal of a tensor (top), vector (middle), or scalar (bottom) polarization, i.e. $|F_p|$ for each polarization p as given by Eqs. (8.7–8.11). The polar and azimuthal coordinates correspond to the source location with respect to the detector, which is to be imagined as placed with its vertex at the center of each plot and arms along the x and y-axes. The response is plotted to scale, such that the black lines representing the detector arms have unit length in all plots. The response to breathing and longitudinal modes is identical, so we only display it once and label it "scalar".

8.2.2 Continuous waves

Signal

A CW is an almost-monochromatic gravitational perturbation with constant intrinsic amplitude and phase evolution $\phi(t)$. For arbitrary polarization content, such a GW will induce a strain in a quadrupolar detector which can be written as:

$$h(t) = \sum_p F_p(t) h_p(t), \tag{8.12}$$

where the sum is over the five independent polarizations, $p \in \{+, \times, \mathrm{x}, \mathrm{y}, \mathrm{s}\}$; the F_p's are those of Eqs. (8.1–8.5) and thus implicitly depend on the relative location and orientation of source and detector by means of ψ, α and δ; the h_p term encodes the amplitude and phase of the wave before being projected onto the frame of the detector:

$$h_p(t) = a_p \cos\left(\phi(t) + \phi_p\right), \tag{8.13}$$

where a_p is a time-independent amplitude with a functional dependence on source parameters determined by each particular theory of gravity; $\phi(t)$ the phase evolution, a consequence of the dynamics of the source in that theory; and ϕ_p a phase offset for each polarization. The polarization amplitudes a_p and phases ϕ_p may take arbitrary values depending on the specific theory of gravity and emission mechanism.

In GR, there are several ways in which a neutron star could emit CWs, but the most likely is the presence of a nonaxisymmetry in the star's moment of inertia [54]. For this type of *triaxial*, nonprecessing source, GR predicts:

$$h_+(t) = h_0 \frac{1}{2}(1 + \cos^2 \iota) \cos \phi(t), \tag{8.14}$$

$$h_\times(t) = h_0 \cos \iota \sin \phi(t), \tag{8.15}$$

$$h_\mathrm{x} = h_\mathrm{y} = h_\mathrm{s} = h_\mathrm{l} = 0, \tag{8.16}$$

where ι is the inclination angle between the spin axis of the source and the observing line-of-sight, and h_0 is an overall amplitude given by:

$$h_0 = \frac{16\pi^2 G}{c^4} \frac{\epsilon I_{zz} f_\mathrm{rot}^2}{r}, \tag{8.17}$$

where r is the source distance, f_rot its rotation frequency around the principal axis z, $\mathbf{I}$ the moment-of-inertia tensor and $\epsilon \equiv (I_{xx} - I_{yy})/I_{zz}$ the equatorial ellipticity. For the triaxial case, the GW frequency f is twice the rotational value f_rot, so that we can write:

$$\phi(t) = 2\phi_\mathrm{rot}(t) + \phi_\mathrm{(GW-EM)}, \tag{8.18}$$

where ϕ_rot is the rotational phase as measured via electromagnetic (EM) observations and $\phi_\mathrm{(GW-EM)}$ is a potential, constant phase offset between the GW and EM signals that can be absorbed into the definition of the ϕ_p's in Eq. (8.13).

Note that other emission mechanisms may result in GW radiation at $f = f_\mathrm{rot}$ [146], or even noninteger powers f_rot [157–159]. Furthermore, alternative theories of gravity may (and, in general, will) support signals at any harmonic. Although in this paper we only consider the case in which only the second rotational harmonic appears in the GW phase, the analysis can be easily generalized to also include contributions from the fundamental and other multiples of f_rot (see Sec. 8.6).

Targeted search

We would like to search a given set of data (from one or more detectors) for CW signals coming from a specific candidate pulsar which has already been observed and timed electromagnetically. Timing solutions are obtained through the pulsar timing package TEMPO2 [55, 56]. We want to achieve this regardless of polarization content, and to reliably distinguish between the different modes present.

If we assume all polarizations share the same phase evolution, then detector response is the only factor distinguishing CW polarizations and, thus, all the relevant information is encoded in the sidereal-day-period amplitude modulation of the signal. This allows us to focus on a narrow frequency band around the expected GW frequency by processing the data following the complex-heterodyne method presented in [66] and [57]. This procedure is summarized below.

A signal like Eqs. (8.12, 8.13) can be rewritten in the form:

$$h(t) = \Lambda(t)e^{i\phi(t)} + \Lambda^*(t)e^{-i\phi(t)}, \tag{8.19}$$

$$\Lambda(t) \equiv \frac{1}{2}\sum_{p=1}^{5} a_p e^{i\phi_p} F_p(t_k; \psi, \alpha, \delta), \tag{8.20}$$

with $*$ indicating complex conjugation and $\phi(t)$ given by a Taylor expansion around $f = 2f_{\rm rot}$:

$$\phi(t) = 2\pi\left(2f_{\rm rot}\tau + \dot{f}_{\rm rot}\tau^2 + ...\right), \tag{8.21}$$

where τ is itself a function of time given by:

$$\tau(t) = t + \Delta_{\rm R} + \Delta_{\rm E} + \Delta_{\rm S} + \Delta_{\rm binary}\ . \tag{8.22}$$

Above, τ is the time measured by a clock inertial with respect to the pulsar; t is the time as measured at a given detector; $\Delta_{\rm R}$ is the Roemer delay; $\Delta_{\rm E}$ is the Solar-System Einstein delay; $\Delta_{\rm S}$ is the Solar-System Shapiro delay; $\Delta_{\rm binary}$ is the delay originating from the motion of the pulsar in its binary (a term that vanishes for isolated sources) [57].

It is important to remember that, the F_p's are functions of the source orientation and sky location relative to the detector, so we have made this dependence explicit in Eq. (8.20) by writing $F_p(t_k)$ as $F_p(t_k; \psi, \alpha, \delta)$. Also, recall that these functions have a characteristic period of a sidereal day ($\sim 10^{-5}$ Hz).

Because the phase evolution $\phi(t)$, including all corrections from Eq. (8.22), is known (with known uncertainties) from electromagnetic observations, we can digitally heterodyne the data by multiplying by $\exp[-i\phi(t)]$ so that the signal therein becomes:

$$h'(t) \equiv h(t)e^{-i\phi(t)} = \Lambda(t) + \Lambda^*(t)e^{-i2\phi(t)} \tag{8.23}$$

and the frequency modulation of the first term is removed, while that of the second term is doubled. A series of low-pass filters can then be used to remove the quickly varying term, which enables the down-sampling of the data by averaging over

minute-long time bins. As a result, $\Lambda(t)$ is the only contribution from the original signal left in our data, and hence we can use Eq. (8.20) as the template for our search. Note that, although we started with real-valued data, after this process the data are now complex.

From Eq. (8.23) we see that, in the presence of a signal, the heterodyned and down-sampled noisy detector strain data B_k for the k^{th} time bin (which can also be labeled by the Earth-frame GPS time-of-arrival at the detector, t_k) are expected to be of the form:

$$B_{\text{expected}}(t_k) = \Lambda(t_k) + n(t_k), \tag{8.24}$$

where $n(t_k)$ is the heterodyned, filtered and downsampled noise in bin k, which carries no information about the GW signal. Note then that $B_k(t_k) - \Lambda(t_k)$ should be expected to have the statistical properties of noise, a fact that will be used below in defining likelihoods.

8.3 Method

8.3.1 Model selection

We use the tools of Bayesian model selection (also known as *second-level inference*) to determine whether the data contain a signal and, if so, whether that signal agrees with the GR prediction or not. Our procedure is hierarchical and consists of the following stages:

1. *detection*: select between signal and noise models;
2. *test of GR*: if a signal is present, select between GR and non-GR models;
3. *upper limits*: if GR is favored, place upper limits on nontensorial strain amplitudes, in the context of specific alternative polarization models.

This subsection covers only the first two items in this list, since the placement of upper limits belongs in the section on parameter estimation. We treat the case of a single data set in 8.3.1 and 8.3.1, and we show how to combine results from multiple analyses in 8.3.1; we offer some considerations about how to approach the problem of non-Gaussian noise in 8.3.1.

Hypotheses

For any given pulsar, we would first like to use *reduced* (i.e. heterodyned, filtered and downsampled) GW data to decide between the following two logically disjoint hypotheses:

1. *noise* ($\mathcal{H}_\mathrm{N}$): no signal, the data are drawn from a Gaussian distribution of zero mean and some (possibly slowly varying) standard deviation;
2. *signal* ($\mathcal{H}_\mathrm{S}$): the data contain noise drawn from a Gaussian distribution and a signal with the assumed phase evolution and *any* polarization content.

In order to perform model selection, we need to translate these hypotheses into the corresponding Bayesian models; this means setting a likelihood function derived from the expected noise properties and picking a multidimensional prior distribution over all parameters. It is important to underscore that a Bayesian model is *defined* by the choice of these two probability distributions.

For $\mathcal{H}_\mathrm{N}$, the construction of the likelihood is straightforward. First, let σ be the standard deviation of the detector noise at or near the expected GW frequency; then, for each complex-valued data point B_k, Gaussianity implies:

$$p(B_k \mid \sigma, \mathcal{H}_\mathrm{N}) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{|B_k|^2}{2\sigma^2}\right). \tag{8.25}$$

Here, and throughout this document, a lower-case p is used for probability densities, while an uppercase P is used for discrete probabilities.

If the data are split into N_S segments of lengths s_j ($j = 1, \ldots, N_\mathrm{S}$) over which the standard deviation σ_j is assumed to remain constant, we can analytically marginalize over this parameter to obtain a likelihood for the entire data set $\mathbf{B}$ in the form of a Student's t-distribution [51, 57]:

$$P(\mathbf{B} \mid \mathcal{H}_\mathrm{N}) = \prod_{j=1}^{N_\mathrm{S}} A_j \left(\sum_{k=\kappa_j}^{K_j} |B_k|^2\right)^{-s_j}, \tag{8.26}$$

with $A_j = (s_j - 1)!/2\pi^{s_j}$, $\kappa_j = 1 + \sum_{n=1}^{j} s_{n-1}$, $K_j = \kappa_j + s_j - 1$ and $s_0 = 0$. Data streams from N_D detectors can be analyzed coherently by generalizing this to:

$$P(\mathbf{B} \mid \mathcal{H}_\mathrm{N}) = \prod_{i=1}^{N_\mathrm{D}} \prod_{j=1}^{N_{\mathrm{S}i}} A_{i,j} \left(\sum_{k=\kappa_{i,j}}^{K_{i,j}} \left|B_{i,k}\right|^2\right)^{-s_{i,j}}, \tag{8.27}$$

where i indexes detectors, $B_{i,k} \equiv B_i(t_k)$ is the datum corresponding to the i^{th} detector at time t_k, and $A_{i,j}$, $\kappa_{i,j}$ and $K_{i,j}$ are defined analogously to s_j, κ_j above. The splitting of the data into segments of constant standard deviation may be achieved with a strategy similar to the Bayesian-blocks algorithm of [160], and explained in detail in [51].

Note that the likelihood $p(\mathbf{B} \mid \vec{\theta}, \mathcal{H})$ of some hypothesis $\mathcal{H}$, is the probability of observing the data $\mathbf{B}$ assuming $\mathcal{H}$ is true and given a specific choice of free parameters $\vec{\theta}$ from the model's parameter space Θ. However, in the case of the noise ("null") hypothesis, as defined by the Student's t likelihood above, there are no free parameters. Consequently, $\Theta = \varnothing$ and $p(\mathbf{B} \mid \vec{\theta}, \mathcal{H}_{\text{N}}) = P(\mathbf{B} \mid \mathcal{H}_{\text{N}})$.

The case of $\mathcal{H}_{\text{S}}$ requires more careful attention. One could be tempted to use Eq. (8.24) to define a likelihood like Eq. (8.27) with the substitution $|B_k| \to |B_k - \Lambda_k|$, for $\Lambda_k \equiv \Lambda(t_k)$ including all polarizations like in Eq. (8.20); the priors would reflect uncertainties in measured source parameters and extend over reasonable ranges for a_p and ϕ_p. However, for most realistic prior choices, that would correspond to a hypothesis that assigns most of the prior probability to regions of parameter space for which $a_p \neq 0$ for all p, thus downweighting more conservative models (including GR) that we would like to prioritize. This is simply because the subspace in parameter space corresponding to any of these smaller subhypotheses (which, for example, fix one of the a_p's to be zero) has infinitely less volume (i.e. it offers infinitesimally less support) than its complement; hence any practical choice of prior probability density will also assign this subspace infinitely less weight, and so the prior for the corresponding subhypothesis will be vanishingly small.

Formally, the inadequacy of the naive construction of $\mathcal{H}_{\text{S}}$ as proposed in the previous paragraph is related to the logical independence of nested hypotheses. We refer to this important point multiple times in the following sections; in particular, we discuss it in the context of odds computations in the text surrounding Eq. (8.38). We refer readers not familiar with this line of reasoning to a similar discussion in [21], or, more generally, to Ch. 4 in [22] or Ch. 28 in [23].

Instead, we will construct $\mathcal{H}_{\text{S}}$ from two logically disjoint component hypotheses:

1. *GR signal* ($\mathcal{H}_{\text{GR}}$ or $\mathcal{H}_{\text{t}}$): the data contain Gaussian noise and a tensorial signal with the assumed $\phi(t)$;

2. *non-GR signal* ($\mathcal{H}_{\text{nGR}}$): the data contain Gaussian noise and a signal with non-GR polarization content, but with the assumed $\phi(t)$.

The tensorial hypothesis is embodied most generally by a signal model such that

$$\Lambda_{\rm t}(t) = \frac{1}{2}\left[a_+ e^{i\phi_+} F_+(t;\psi=0) + a_\times e^{i\phi_\times} F_\times(t;\psi=0)\right], \tag{8.28}$$

where a_+, $a_\times$, ϕ_+ and $\phi_\times$ are free parameters, and we pick a specific polarization frame by setting $\psi = 0$ (we are allowed to do this because of a degeneracy between ψ and a_+, $a_\times$ explained in Appendix 8.7). An alternative parametrization can be derived from the triaxial emission model of Eqs. (8.14–8.16), namely

$$\Lambda_{\rm GR}(t) = \frac{1}{2} h_0 e^{i\phi_0}\left[\frac{1}{2}(1+\cos^2\iota)F_+(t;\psi) - i\cos\iota F_\times(t;\psi)\right], \tag{8.29}$$

where the free parameters are now h_0, ϕ_0, ι and ψ [in the notation of Eqs. (8.13, 8.20), $\phi_+ = \phi_0$ and $\phi_\times = \phi_0 - \pi/2$]. This is the parametrization used in most traditional GR-only searches (see e.g. [29, 51]).

The templates of Eq. (8.28) and Eq. (8.29) span the same signal space; therefore, if we pick parameter priors properly related by their Jacobian, the respective hypotheses ($\mathcal{H}_{\rm t}$ and $\mathcal{H}_{\rm GR}$) will be logically equivalent (i.e. $\mathcal{H}_{\rm t} \equiv \mathcal{H}_{\rm GR}$). However, we will sometimes want to restrict ψ or ι in Eq. (8.29) to incorporate measurements of the source orientation (see Table 3 in [49]) and compare those results to the unconstrained model of Eq. (8.28). In such cases, $\mathcal{H}_{\rm t}$ and $\mathcal{H}_{\rm GR}$ are no longer equivalent: the former corresponds to a *free-tensor* signal, while the latter now corresponds to a GR *triaxial* signal for some given source orientation [i.e. a signal with the functional dependence on ι and ψ of Eq. (8.29)]. Because of lack of any orientation information, this is a distinction without a difference for most pulsars. (See Appendix 8.7 for more details.)

The non-GR hypothesis, $\mathcal{H}_{\rm nGR}$, can itself be seen as a composite hypothesis encompassing all the signal models that depart from GR in some way, i.e. models that include polarizations other than + and ×. We denote such subhypotheses with a subscript listing the polarizations included in the signal. For example, "st" (meaning "scalar plus tensor") corresponds to a model with unrestricted scalar and tensor contributions:

$$\begin{aligned}\Lambda_{\rm st}(t) = \frac{1}{2}\Big[&a_+ e^{i\phi_+} F_+(t;\psi=0) + a_\times e^{i\phi_\times} F_\times(t;\psi=0) \\ &+ a_s e^{i\phi_s} F_s(t;\psi=0)\Big].\end{aligned} \tag{8.30}$$

With this notation extended to the names of the relevant hypotheses, we may then

write $\mathcal{H}_{\rm nGR}$ as the logical union ("or" junction, $\vee$)

$$\begin{aligned}\mathcal{H}_{\rm nGR} &\equiv \mathcal{H}_{\rm s} \vee \mathcal{H}_{\rm v} \vee \mathcal{H}_{\rm st} \vee \mathcal{H}_{\rm sv} \vee \mathcal{H}_{\rm tv} \vee \mathcal{H}_{\rm stv} \\ &= \bigvee_{m\in\tilde{M}} \mathcal{H}_m, \end{aligned} \tag{8.31}$$

where, for convenience, we have defined the non-GR subscript set $\tilde{M}$:

$$\tilde{M} \equiv \{\rm s,\ v,\ st,\ sv,\ tv,\ stv\}\,. \tag{8.32}$$

Just as before, we may equivalently use the triaxial parametrization, Eq. (8.29), for the tensor modes in the non-GR hypotheses by instead defining $\tilde{M}$ as

$$\tilde{M} = \{\rm s,\ v,\ sv,\ GR + s,\ GR + v,\ GR + sv\}, \tag{8.33}$$

where, for example, GR+s denotes a signal template like

$$\begin{aligned}\Lambda_{\rm GR+s}(t) = \frac{h_0}{2} e^{i\phi_0} \left[\frac{1}{2}(1+\cos^2\iota)F_+(t;\psi) - i\cos\iota F_\times(t;\psi)\right] \\ + \frac{1}{2} a_b e^{i\phi_b} F_b(t;\psi), \end{aligned} \tag{8.34}$$

and similarly for GR+v and GR+sv, with the added vector modes. Again, the two definitions of $\tilde{M}$, Eqs. (8.32, 8.33), are equivalent unless orientation information is incorporated in the way explained above.

By the same token, the signal hypothesis can be built from the logical union of $\mathcal{H}_{\rm GR}$ or $\mathcal{H}_{\rm t}$, and $\mathcal{H}_{\rm nGR}$:

$$\mathcal{H}_{\rm S} \equiv \mathcal{H}_{\rm GR/t} \vee \mathcal{H}_{\rm nGR} = \bigvee_{m\in M} \mathcal{H}_m, \tag{8.35}$$

with M defined similarly to $\tilde{M}$, but also including the tensor-only hypothesis, $\mathcal{H}_{\rm GR}$ or $\mathcal{H}_{\rm t}$:

$$M \equiv \tilde{M} \cup \{\rm GR/t\}. \tag{8.36}$$

The validity of Eqs. (8.31, 8.35) is contingent on the mutual logical independence of all the $\mathcal{H}_m$'s. This requirement is satisfied by construction, since each of the $\mathcal{H}_m$'s is defined to exclude regions of parameter space that would correspond to other hypotheses nested within it (e.g. $\mathcal{H}_{\rm GR+s}$ is defined over all values of the scalar amplitude except $a_{\rm s} = 0$, to avoid including $\mathcal{H}_{\rm GR}$). In practice, however, it is not necessary to explicitly exclude these infinitesimal regions of parameter space, as will be explained in the following section.

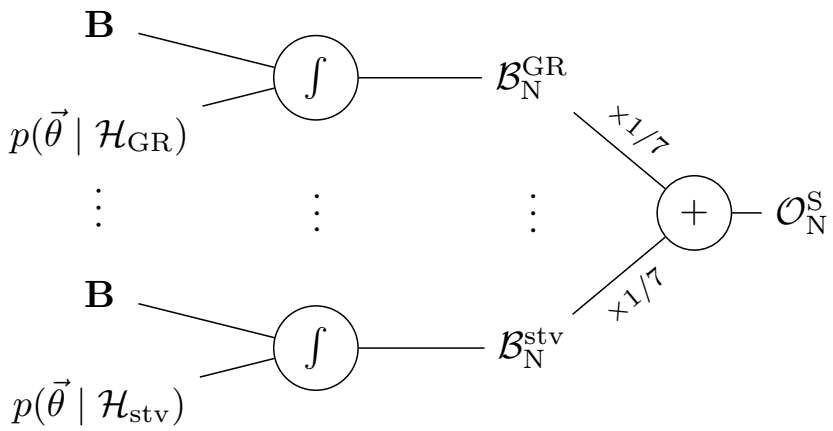

Figure 8.3: *Computation of* $O^{\rm S}_{\rm N}$. First, the Bayes factor $\mathcal{B}^m_{\rm N}$ is obtained from the data **B** and corresponding priors $p(\theta|\mathcal{H}_m)$ for each model $m \in M$, by evaluating the integral of Eq. (8.40) using a nested sampling algorithm that samples over $\vec{\theta}$ (step indicated by integral sign); these values are then added and multiplied by $P(\mathcal{H}_m)/P(\mathcal{H}_{\rm N})$ to obtain $O^{\rm S}_{\rm N}$, as in Eq. (8.43). (Note that here we have set $P(\mathcal{H}_m)/P(\mathcal{H}_{\rm N}) = 1/7$, as explained Sec. 8.4.) The computation of $O^{\rm nGR}_{\rm GR}$ is analogous.

Odds

We can construct a Bayesian model for $\mathcal{H}_{\rm S}$ starting from its components: for each subhypothesis $\mathcal{H}_m$ for $m \in M$, we use a likelihood function like Eq. (8.27) with the substitution $|B_{i,k}| \to |B_{i,k} - \Lambda_{m,i,k}|$, i.e.

$$p(\mathbf{B} \mid \vec{\theta}, \mathcal{H}_m) = \prod_{i=1}^{N_{\rm D}} \prod_{j=1}^{N_{{\rm S}i}} A_{i,j} \left(\sum_{k=\kappa_{i,j}}^{K_{i,j}} \left| B_{i,k} - \Lambda_{m,i,k} \right|^2 \right)^{-s_{i,j}} \tag{8.37}$$

(where $\Lambda_{m,i,k}$ is the template corresponding to model m, for detector i and time-bin k), and suitable priors on the model parameters $\vec{\theta}_m \in \Theta_m$; then, we combine the posteriors with priors on the models themselves to obtain the posterior for $\mathcal{H}_{\rm S}$. This last step allows us to incorporate our *a priori* beliefs about the validity of each of the components. This procedure is represented schematically in Fig. 8.3 and fleshed out below.

The choice of model priors can be made clearer by considering the posterior probability for the signal model. Given some set of detector data **B** and underlying assumptions I (suppressed from the following expressions), the posterior probability for $\mathcal{H}_{\rm S}$ is:

$$P(\mathcal{H}_{\rm S} \mid \mathbf{B}) = \sum_{m \in M} P(\mathcal{H}_m \mid \mathbf{B}) \tag{8.38}$$

by Eq. (8.35) and because the components are all logically independent [i.e. $\mathcal{H}_{m_1} \wedge \mathcal{H}_{m_2} = \text{False}$, hence $P(\mathcal{H}_{m_1} \wedge \mathcal{H}_{m_2} \mid \mathbf{B}) = 0$ for any $m_1,\ m_2 \in M$ such that $m_1 \neq m_2$]. Note that this is true even for hypotheses that may contain each other as special cases. For instance, even though the GR template can be obtained from GR+s by setting the scalar amplitude to $a_\mathrm{s} = 0$, the points in the GR+s parameter space satisfying this condition define an infinitesimally-thin slice in parameter space that offers no support to the prior distribution and is thus assigned no weight (see similar discussion in [21]).

We can expand each term on the RHS of Eq. (8.38) using Bayes' theorem:

$$P(\mathcal{H}_\mathrm{S} \mid \mathbf{B}) = \sum_{m \in M} P(\mathcal{H}_m) P(\mathbf{B} \mid \mathcal{H}_m) / P(\mathbf{B}). \tag{8.39}$$

Each of the terms on the RHS is made up of three factors: a marginalized likelihood $P(\mathbf{B} \mid \mathcal{H}_m)$, a prior on the model $P(\mathcal{H}_m)$, and a normalization constant $P(\mathbf{B})$.

The marginalized likelihood (also known as *evidence*) is computed from the data:

$$P(\mathbf{B} \mid \mathcal{H}_m) = \int_{\Theta_m} p(\mathbf{B} \mid \vec{\theta}_m, \mathcal{H}_m)\ p(\vec{\theta}_m \mid \mathcal{H}_m)\ \mathrm{d}\vec{\theta}_m, \tag{8.40}$$

where $p(\mathbf{B} \mid \vec{\theta}_m, \mathcal{H}_m)$ is itself the likelihood. The evaluation of the multidimensional integral of Eq. (8.40) is the most computationally intensive part of our analysis (see Sec. 8.4 for details).

We are free to choose the model priors (discussed in Sec. 8.4), as long as we satisfy the constraint:

$$P(\mathcal{H}_\mathrm{N}) + \sum_{m \in M} P(\mathcal{H}_m) = 1. \tag{8.41}$$

This is a statement about the exhaustiveness and disjointedness of the hypotheses we are considering: we assume that reality will agree with *one and only one* of the hypotheses at hand. (As we will see in Sec. 8.6, this assumption might not hold; for example, the noise may not be Gaussian.) The particular choice of prior for each model will encode our expectations about the corresponding theory (before seeing the data), and thus allow for some degree of subjectivity.

Note that we cannot directly compute $P(\mathbf{B})$ in a straightforward manner and without assuming that our hypothesis set is indeed exhaustive (which is not the case for non-Gaussian detector noise, see Sec. 8.3.1). However, the need for this computation can be avoided by looking at relative probabilities, i.e. *odds*. The odds for $\mathcal{H}_\mathrm{S}$ vs $\mathcal{H}_\mathrm{N}$ is defined as:

$$\mathcal{O}^\mathrm{S}_\mathrm{N} \equiv \frac{P(\mathcal{H}_\mathrm{S} \mid \mathbf{B})}{P(\mathcal{H}_\mathrm{N} \mid \mathbf{B})}. \tag{8.42}$$

Using Bayes' theorem again and canceling the $P(\mathbf{B})$ factors, this simplifies to:

$$\mathcal{O}^{\mathrm{S}}_{\mathrm{N}} = \frac{\sum P(\mathcal{H}_m)P(\mathbf{B} \mid \mathcal{H}_m)}{P(\mathcal{H}_{\mathrm{N}})P(\mathbf{B} \mid \mathcal{H}_{\mathrm{N}})} = \sum_{m\in M} \frac{P(\mathcal{H}_m)}{P(\mathcal{H}_{\mathrm{N}})}\mathcal{B}^m_{\mathrm{N}}, \tag{8.43}$$

where, in the second equality, we have used the definition of the *Bayes factor*:

$$\mathcal{B}^i_j \equiv \frac{P(\mathbf{B} \mid \mathcal{H}_i)}{P(\mathbf{B} \mid \mathcal{H}_j)}, \tag{8.44}$$

for any two hypotheses $\mathcal{H}_i, \mathcal{H}_j$.

The odds in Eq. (8.43) can be used as a detection statistic to determine whether it is likely that the data contain a signal (of any polarization) or not. Once the presence of a signal has been established, a similar ratio can be constructed to assess agreement with GR:

$$\mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}} = \frac{P(\mathcal{H}_{\mathrm{nGR}} \mid \mathbf{B})}{P(\mathcal{H}_{\mathrm{GR}} \mid \mathbf{B})} = \sum_{m\neq GR} \frac{P(\mathcal{H}_m)}{P(\mathcal{H}_{\mathrm{GR}})}\mathcal{B}^m_{\mathrm{GR}}. \tag{8.45}$$

This ratio encodes the relative probability that there is a GR violation. Because it is now assumed that there is a signal in the data, $P(\mathcal{H}_{\mathrm{N}}) = 0$ and the model priors must instead satisfy:

$$\sum_{m\in M} P(\mathcal{H}_m) = 1. \tag{8.46}$$

We can reduce the number of computations needed to obtain $\mathcal{O}^{\mathrm{S}}_{\mathrm{N}}$ and $\mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ by using the fact that:

$$\mathcal{B}^i_j = \frac{P(\mathbf{B} \mid \mathcal{H}_i)}{P(\mathbf{B} \mid \mathcal{H}_j)} = \frac{P(\mathbf{B} \mid \mathcal{H}_i)}{P(\mathbf{B} \mid \mathcal{H}_{\mathrm{N}})}\frac{P(\mathbf{B} \mid \mathcal{H}_{\mathrm{N}})}{P(\mathbf{B} \mid \mathcal{H}_j)} = \frac{\mathcal{B}^i_{\mathrm{N}}}{\mathcal{B}^j_{\mathrm{N}}}. \tag{8.47}$$

This means that we need to evaluate an integral like Eq. (8.40) seven times per set of data, to compute $\mathcal{B}^m_{\mathrm{N}}$ for each m in M. Those seven numbers, together with the evidence for $\mathcal{H}_{\mathrm{N}}$, are enough to compute all the quantities of interest.

Instead of asking about a generic deviation from GR, we may also compare GR to a particular alternative theory. For such purpose, we will usually assign equal prior weight to GR and its alternative to compute:

$$\mathcal{O}^j_{\mathrm{GR}} = \frac{P(\mathcal{H}_j)}{P(\mathcal{H}_{\mathrm{GR}})}\mathcal{B}^j_{\mathrm{GR}} = \mathcal{B}^j_{\mathrm{GR}}, \tag{8.48}$$

where $\mathcal{H}_j$ may be any of the hypotheses in $\tilde{M}$ or an even more specific hypothesis. (The latter case demands an extra execution of the inference code.)

Multiple data sets

So far we have assumed that the data **B**, corresponding to one or more GW detectors, can be analyzed coherently; however, there are cases in which we would like to combine results from sets of data analyzed incoherently. Examples are data sets corresponding to different sources or observation periods. Our Bayesian framework makes it possible to combine the respective odds in order to make an overall model selection statement (in our case, about the presence of signal or the validity of GR).

For instance, we may analyze data for $N_{\rm P}$ pulsars and ask about the probability that any of them contain a signal; treating each as an independent observation, the combined probability can be constructed from the odds above. Letting $\mathcal{H}_{\rm S_{\it i}}$, $\mathcal{H}_{\rm N_{\it i}}$ respectively denote signal and noise hypotheses for the $i^{\rm th}$ source, while $\mathcal{H}_{\rm S_{any}}$ corresponds to a signal being present in *any* of the sources and $\mathcal{H}_{\rm N_{all}}$ corresponds to Gaussian noise in data for *all* sources:

$$
\begin{aligned}
{}^{(N_{\rm P})}O^{\rm S_{any}}_{\rm N_{all}} &= \frac{P(\mathcal{H}_{\rm S_{any}} \mid \mathbf{B})}{P(\mathcal{H}_{\rm N_{all}} \mid \mathbf{B})} = \frac{1 - P(\mathcal{H}_{\rm N_{all}} \mid \mathbf{B})}{P(\mathcal{H}_{\rm N_{all}} \mid \mathbf{B})} \\
&= \frac{1}{P(\bigwedge_i \mathcal{H}_{\rm N_{\it i}} \mid \mathbf{B})} - 1 = \left[\prod_{i=1}^{N_{\rm P}} \frac{1}{P(\mathcal{H}_{\rm N_{\it i}} \mid \mathbf{B}_i)}\right] - 1 \\
&= \left[\prod_{i=1}^{N_{\rm P}} \frac{P(\mathcal{H}_{\rm S_{\it i}} \mid \mathbf{B}_i) + P(\mathcal{H}_{\rm N_{\it i}} \mid \mathbf{B}_i)}{P(\mathcal{H}_{\rm N_{\it i}} \mid \mathbf{B}_i)}\right] - 1 \\
&= \left[\prod_{i=1}^{N_{\rm P}} \left(O^{\rm S_{\it i}}_{\rm N_{\it i}} + 1\right)\right] - 1,
\end{aligned}
\tag{8.49}
$$

where we have used the exclusivity and exhaustiveness of the signal and noise hypotheses, i.e.

$$P(\mathcal{H}_{\rm S_{any}} \mid \mathbf{B}) + P(\mathcal{H}_{\rm N_{all}} \mid \mathbf{B}) = 1, \tag{8.50}$$

$$P(\mathcal{H}_{\rm S_{\it i}} \mid \mathbf{B}_i) + P(\mathcal{H}_{\rm N_{\it i}} \mid \mathbf{B}_i) = 1, \tag{8.51}$$

with i indexing data sets. Note that the data sets for different sources ($\mathbf{B}_i$'s) are *not* conditionally independent under $\mathcal{H}_{\rm S_{any}}$ or $\mathcal{H}_{\rm N_{all}}$. Also, Eq. (8.49) does not enforce the requirement that, if signals are present in multiple sources, they all correspond to the same model from Eq. (8.36); such a constraint could be implemented at this stage, but is more easily enforced by examining individual values of $O^m_{\rm N}$ when necessary.

The construction of Eq. (8.49) implicitly assigns model priors to each of the meta-

hypotheses $\mathcal{H}_{\mathrm{S_{any}}}$ and $\mathcal{H}_{\mathrm{N_{all}}}$ such that:

$$\frac{P(\mathcal{H}_{\mathrm{S_{any}}})}{P(\mathcal{H}_{\mathrm{N_{all}}})} = \left[\frac{P(\mathcal{H}_{\mathrm{S}})}{P(\mathcal{H}_{\mathrm{N}})} + 1\right]^{N_{\mathrm{P}}} - 1, \tag{8.52}$$

where we have assumed the priors for signal vs noise are equal for all sources, i.e. $P(\mathcal{H}_{\mathrm{S}_i}) = P(\mathcal{H}_{\mathrm{S}})$ and $P(\mathcal{H}_{\mathrm{N}_i}) = P(\mathcal{H}_{\mathrm{N}})$ for all i. When making combined statements for multiple sources, we may wish to choose $P(\mathcal{H}_{\mathrm{S}})/P(\mathcal{H}_{\mathrm{N}})$ such as to produce any desired value of $P(\mathcal{H}_{\mathrm{S_{any}}})/P(\mathcal{H}_{\mathrm{N_{all}}})$, say $P(\mathcal{H}_{\mathrm{S_{any}}}) = P(\mathcal{H}_{\mathrm{N_{all}}})$. Furthermore, one may wish to weight each pulsar differently within $\mathcal{H}_{\mathrm{S_{any}}}$ by incorporating information about the source distance (or other parameters) into the priors via a parametrization like Eq. (8.17); this may improve the sensitivity of the ensemble odds to weak signals in the set, as suggested in [161]. However, using such a parametrization generally implies committing to a specific gravitational theory (or family of theories). We choose not to take such approach in this study.

Besides combining data for multiple pulsars, for a given source, we could also (incoherently) combine the results of analyses using data from different observation periods. Since the astrophysical CWs we are considering should either be present in all N_{R} observation runs or in none of them, the relevant odds, generalizing Eq. (8.43), are:

$$\begin{aligned}
{}^{[N_{\mathrm{R}}]}\mathcal{O}^{\mathrm{S}}_{\mathrm{N}} &= \frac{P(\mathcal{H}_{\mathrm{S}} \mid \mathbf{B})}{P(\mathcal{H}_{\mathrm{N}} \mid \mathbf{B})} = \sum_{m\in M} \frac{P(\mathcal{H}_m \mid \mathbf{B})}{P(\mathcal{H}_{\mathrm{N}} \mid \mathbf{B})} \\
&= \sum_{m\in M} \frac{P(\mathbf{B} \mid \mathcal{H}_m)P(\mathcal{H}_m)}{P(\mathbf{B} \mid \mathcal{H}_{\mathrm{N}})P(\mathcal{H}_{\mathrm{N}})} \\
&= \sum_{m\in M} \frac{P(\mathcal{H}_m)}{P(\mathcal{H}_{\mathrm{N}})} \prod_{j=1}^{N_{\mathrm{R}}} \left(\mathcal{B}^m_{\mathrm{N}}\right)_j ,
\end{aligned} \tag{8.53}$$

where we have again used $\mathbf{B} = \{\mathbf{B}_j\}_{j=1}^{N_{\mathrm{R}}}$ to refer to the totality of data, with j indexing observation runs. The independence of the $\mathbf{B}_j$'s, conditional on $\mathcal{H}_m$ and $\mathcal{H}_{\mathrm{N}}$, is applied on the last line to write the result in terms of the individual Bayes factors for each run, $\left(\mathcal{B}^m_{\mathrm{N}}\right)_j$.

Similarly, we can use multiple data sets to make a single statement about deviations from GR. Once we have made N_{P} detections from different sources, the odds for a GR violation is:

$${}^{(N_{\mathrm{P}})}\mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}} = \sum_{m\in \tilde{M}} \frac{P(\mathcal{H}_m)}{P(\mathcal{H}_{\mathrm{GR}})} \prod_{i=1}^{N_{\mathrm{P}}} \left(\mathcal{B}^m_{\mathrm{GR}}\right)_i , \tag{8.54}$$

where, again, i indexes sources; this is a generalization of Eq. (8.45). (See Sec. IIID of [21] for an analogous derivation.)

Non-Gaussian noise

Up to this point, like most other CW studies, we have assumed that the detector noise is Gaussian. However, although previous work has indicated that this is generally a very good approximation [29, 109], it is not exactly true for actual detector noise (for some frequencies more so than others). Happily, most of the model selection statements expounded so far are valid also in the presence of non-Gaussian instrumental noise, after some light reinterpretation.

If the assumption of Gaussianity does not hold, the hypotheses constructed in Sec. 8.3.1 are no longer exhaustive: the data may not only be explained by Gaussian noise or a signal (GR or otherwise), but also by non-Gaussian artifacts that are impossible to satisfactorily model. Nevertheless, the computation and interpretation of evidences and odds remain unchanged for all the hypotheses under consideration.

Because “noise" no longer just means “Gaussian noise", $O^{\mathrm{S}}_{\mathrm{N}}$ (which compares the signal model vs *Gaussian* noise) has to be treated more carefully for detection purposes. Indeed, instrumental features that are clearly non-Gaussian (e.g. a loud, narrow-band artifact wandering across the frequency of interest) will generally result in a relatively large value of $O^{\mathrm{S}}_{\mathrm{N}}$, even if there is no detectable astrophysical signal in the data. This issue affects the standard GR searches as well [29], although perhaps to a lesser degree due to the reduced signal parameter space.

It is possible to mitigate this problem by constructing a hypothesis that captures some key characteristic of instrumental features and helps discriminate those from real astrophysical signals. Perhaps the best way to do this is to take advantage of the fact that an astrophysical CW must manifest itself coherently across detectors, while the same is not true for detector artifacts [27]. We can thus define an *instrumental feature hypothesis* ($\mathcal{H}_{\mathrm{I}}$) to encompass the cases in which the data are composed of Gaussian noise or features that look like astrophysical signals but are not coherent across detectors (viz. they do not have a consistent phase evolution and they are best described by different waveform parameters).

Formally, we define $\mathcal{H}_{\mathrm{I}}$ by:

$$\mathcal{H}_{\mathrm{I}} \equiv \bigwedge_{d=1}^{N_{\mathrm{D}}} \left(\mathcal{H}_{\mathrm{S}_d} \vee \mathcal{H}_{\mathrm{N}_d}\right), \tag{8.55}$$

where the subscript d identifies detectors, and $\wedge$ is the logical “and” junction. This definition does not explicitly encompass instrumental features that are coherent across

some subset of the detectors. Also, note that Eq. (8.55) implicitly contains a term equivalent to the usual noise hypothesis $\mathcal{H}_\mathrm{N} = \bigwedge_d \mathcal{H}_{\mathrm{N}_d}$. Similarly, it also contains a term corresponding to the presence of signals in all detectors ($\bigwedge_d \mathcal{H}_{\mathrm{S}d}$). Importantly, such an *incoherent* term is not equivalent to the *coherent* signal hypothesis $\mathcal{H}_\mathrm{S}$, as given by the multidetector likelihood of Eq. (8.37):

$$\mathcal{H}_\mathrm{S} \neq \bigwedge_{d=1}^{N_\mathrm{D}} \mathcal{H}_{\mathrm{S}d}. \tag{8.56}$$

While the evidence integral of Eq. (8.40) factorizes into single-detector terms for $\mathcal{H}_\mathrm{N}$ (due to the null parameter space), the same is not true for $\mathcal{H}_\mathrm{S}$. Furthermore, because it does not demand detector coherence, the RHS of Eq. (8.56) is associated with a considerably larger parameter space than the LHS. Thus, in the presence of an astrophysical signal, model selection will favor $\mathcal{H}_\mathrm{S}$ due to its smaller Occam's penalty. The same is true, of course, when comparing $\mathcal{H}_\mathrm{S}$ to $\mathcal{H}_\mathrm{I}$ as a whole.

From Eq. (8.55), it is straightforward to write the evidence for $\mathcal{H}_\mathrm{I}$ as

$$\begin{aligned} P(\mathbf{B} \mid \mathcal{H}_\mathrm{I}) = \prod_{d=1}^{N_\mathrm{D}} \Big[& P(\mathbf{B}_d \mid \mathcal{H}_{\mathrm{S}_d}) P(\mathcal{H}_{\mathrm{S}_d} \mid \mathcal{H}_\mathrm{I}) \\ & + P(\mathbf{B}_d \mid \mathcal{H}_{\mathrm{N}_d}) P(\mathcal{H}_{\mathrm{N}_d} \mid \mathcal{H}_\mathrm{I}) \Big] \end{aligned} \tag{8.57}$$

and use this to construct the odds comparing against $\mathcal{H}_\mathrm{S}$:

$$O^\mathrm{S}_\mathrm{I} = \frac{P(\mathcal{H}_\mathrm{S})}{P(\mathcal{H}_\mathrm{I})} \frac{\mathcal{B}^\mathrm{S}_\mathrm{N}}{\prod_{d=1}^{N_\mathrm{D}} \left[P(\mathcal{H}_{\mathrm{S}_d} \mid \mathcal{H}_\mathrm{I})(\mathcal{B}^{\mathrm{S}_d}_{\mathrm{N}_d} - 1) + 1 \right]}. \tag{8.58}$$

Here we have used Eq. (8.57), together with the fact that $P(\mathcal{H}_{\mathrm{S}_d} \mid \mathcal{H}_\mathrm{I}) + P(\mathcal{H}_{\mathrm{N}_d} \mid \mathcal{H}_\mathrm{I}) = 1$ and $P(\mathbf{B} \mid \mathcal{H}_\mathrm{N}) = \prod_d P(\mathbf{B} \mid \mathcal{H}_{\mathrm{N}d})$, to write O^S_I as a function of the detector-coherent signal vs noise Bayes factor $\mathcal{B}^\mathrm{S}_\mathrm{N}$, the single-detector signal vs noise Bayes factors $\mathcal{B}^{\mathrm{S}_d}_\mathrm{N}$, and model priors $P(\mathcal{H}_\mathrm{S})$, $P(\mathcal{H}_\mathrm{I})$ and $P(\mathcal{H}_{\mathrm{S}_d} \mid \mathcal{H}_\mathrm{I})$.

As usual, we are free to choose the model priors to give more or less weight to different hypotheses. For example, we recover the choice of [29] (Appendix A3) by setting $P(\mathcal{H}_{\mathrm{S}_d} \mid \mathcal{H}_\mathrm{I}) = 0.5$ for all d and $P(\mathcal{H}_\mathrm{S}) = P(\mathcal{H}_\mathrm{I}) \times 0.5^{N_\mathrm{D}}$ such that:

$$\ln O^\mathrm{S}_\mathrm{I} = \ln \mathcal{B}^\mathrm{S}_\mathrm{N} - \sum_{d=1}^{N_\mathrm{D}} \ln \left(\mathcal{B}^{\mathrm{S}_d}_{\mathrm{N}_d} + 1 \right). \tag{8.59}$$

(When comparing to Appendix A3 of [29], however, note that in that work "I" is used to denote both the background information and the "incoherent-signal-or-noise" hypothesis, which can be identified with our $\mathcal{H}_\mathrm{I}$.)

There is reason to believe that $\ln \mathcal{O}^{\mathrm{S}}_{\mathrm{I}}$, with model priors as in Eq. (8.59), is quite good at picking out instrumental features, even for data from just two instruments [29]. (Note that we would expect the discriminatory power of $\ln \mathcal{O}^{\mathrm{S}}_{\mathrm{I}}$ to grow with the number of detectors available.) However, at the end of the day, we can never be fully confident that $\mathcal{H}_{\mathrm{I}}$ will indeed capture all nonastrophysical disturbances. To address this, we may always treat $\ln \mathcal{O}^{\mathrm{S}}_{\mathrm{N}}$ and $\ln \mathcal{O}^{\mathrm{S}}_{\mathrm{I}}$ as any generic detection statistic and use estimates of the background distribution to establish significance.

8.3.2 Parameter estimation

Besides choosing between different models, we can use Bayesian statistics to obtain posterior probability density functions (PDFs) on the parameters of a given template (*first-level inference*). In the absence of a loud signal, this can be used to obtain credible intervals that yield upper-limits on the amplitudes of GR deviations.

For a model $\mathcal{H}$ with N parameters, an N-dimensional posterior PDF covering the parameter space Θ can be obtained from Bayes' theorem:

$$p(\vec{\theta} \mid \mathbf{B}, \mathcal{H}) = \frac{p(\mathbf{B} \mid \vec{\theta}, \mathcal{H})\, p(\vec{\theta} \mid \mathcal{H})}{P(\mathbf{B} \mid \mathcal{H})}, \tag{8.60}$$

for $\vec{\theta}$ in Θ, and with $p(\vec{\theta} \mid \mathcal{H})$ the prior over Θ. To obtain a one-dimensional PDF for a single parameter (call it θ_i), the N-dimensional distribution must be marginalized over all nuisance parameters (viz. all parameters except θ_i):

$$\begin{aligned} p(\theta_i \mid \mathbf{B}, \mathcal{H}) &= \int_{\Theta'} p(\vec{\theta} \mid \mathbf{B}, \mathcal{H})\, \mathrm{d}^{N-1}\theta_j \\ &\propto \int_{\Theta'} p(\mathbf{B} \mid \vec{\theta}, \mathcal{H})\, p(\vec{\theta} \mid \mathcal{H})\, \mathrm{d}^{N-1}\theta_j, \end{aligned} \tag{8.61}$$

where $0 < j \leq N$, such that $j \neq i$, and Θ' denotes the parameter space Θ with the i^{th} dimension removed. Note that the equality has been replaced by a relation of proportionality because we have excluded the evidence $P(\mathbf{B} \mid \mathcal{H})$ from the expression. (Although of great importance for model selection, this quantity is uninteresting for the purposes of parameter estimation and can be treated as a simple normalization constant.) As discussed in Sec. 8.4, we evaluate Eq. (8.61) with the same algorithm used to compute the evidence.

Eq. (8.61) can be used to place upper limits on model parameters; in particular, we will use it to place limits on the amplitude of GR deviations. Consider, for instance, the case of a scalar-tensor theory that can be encapsulated by our GR+s model as

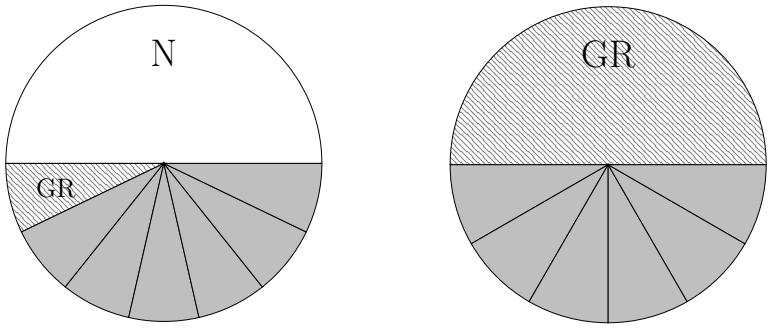

Figure 8.4: *Model priors*. Distribution of prior probability over subhypotheses for the construction of $O^{\mathrm{S}}_{\mathrm{N}}$ (left) and $O^{\mathrm{nGR}}_{\mathrm{GR}}$ (right), according to Eqs. (8.43) and (8.45) respectively. For $O^{\mathrm{S}}_{\mathrm{N}}$, we assign equal weight to the $\mathcal{H}_{\mathrm{N}}$ (white) and $\mathcal{H}_{\mathrm{S}}$ (gray); as in Eq. (8.64), we make no *a priori* distinction between non-GR models (solid) and GR (hatched). For $O^{\mathrm{nGR}}_{\mathrm{GR}}$, we set equal prior probability for $\mathcal{H}_{\mathrm{GR}}$ and $\mathcal{H}_{\mathrm{nGR}}$, distributing the prior equally among non-GR models, as in Eq. (8.66).

described in the previous section; the 95%-credible upper limit on the strength of the breathing mode is $h_{\mathrm{s}}^{95\%}$, defined by:

$$0.95 = \int_{\min(h_{\mathrm{s}})}^{h_{\mathrm{s}}^{95\%}} p(h_{\mathrm{s}} \mid \mathbf{B}, \mathcal{H}_{\mathrm{GR+s}})\, \mathrm{d}h_{\mathrm{s}}, \tag{8.62}$$

where $\min(h_{\mathrm{s}})$ is the minimum value of h_{s} allowed by the prior.

Note that there may be reasons to compute posteriors under different priors than when computing Bayes factors. In particular, it is conventional to present upper limits obtained using a uniform prior over some broad range of the amplitude parameters. With a uniform prior, the posterior is trivially related to the likelihood. This approach produces a more conservative upper limit than other choices, e.g. a Jeffreys prior (see Appendix 8.8).

8.4 Analysis

We quantify our ability to use Bayesian model selection to detect CW signals and determine their polarization content as described above. To do this, we use one year of simulated data from three advanced interferometric detectors at design sensitivity: LIGO Hanford (H1), LIGO Livingston (L1) and Virgo (V1). Detector noise is simulated by drawing from a Gaussian distribution with zero mean and variance corresponding to the power spectral density (PSD) of each detector at the GW frequency of the pulsar. (Previous work has shown that these are good assumptions for actual reduced detector data [29, 109].)

As mentioned in the previous section, the key step in our analysis is the computation of the evidence integral of Eq. (8.40) for the hypotheses under consideration (one noise model, plus seven signal submodels; see Sec. 8.3.1). We carry this out using a version of the Bayesian inference code used for the targeted pulsar search in [28, 51], which we modified to handle signals from theories other than GR. This inference code is itself built on the implementation of Skilling's nested-sampling algorithm [19] in the LALInference package [20], part of the LIGO Algorithm Library Suite [162]. This is the same inference software used for CBC analyses, including GW150914 [31].

In computing likelihoods, we take source location, frequency and frequency derivatives as known quantities (relevant uncertainties are negligible for this analysis). Unless otherwise stated, priors uniform in the logarithm are used for amplitude parameters (h_0 or h_p's), since these are the least informative priors for scaling coefficients (also known as "Jeffreys priors") [163]; we make the somewhat arbitrary choice of restricting the strain amplitudes to the 10^{-28}–10^{-24} range (this is of little consequence for model selection, as explained in Appendix 8.8). Flat priors are placed over all phase offsets (ϕ_0 and all the ϕ_p's).

All plots for the Crab pulsar (PSR J0534+2200) in Sec. 8.5 are produced using known values of its orientation parameters, $\cos\iota$ and ψ, and with the triaxial parametrization of tensor modes; for other pulsars, however, the free-tensor parametrization is used instead. (See Sec. 8.3.1 and Appendix 8.7.)

We follow common practice by adopting the principle of indifference (see e.g. Ch. 5 of [22]) in assigning equal prior probability to the signal and noise models, i.e. we let

$$P(\mathcal{H}_{\rm S}) = P(\mathcal{H}_{\rm N}) = 1/2\ . \tag{8.63}$$

We must also decide how to split the prior among the different $\mathcal{H}_m$'s when computing $O^{\rm S}_{\rm N}$ and $O^{\rm nGR}_{\rm GR}$. In the former case we choose to distribute the prior weight uniformly among all signal models, so that:

$$P(\mathcal{H}_m) = |M|^{-1}/2 = 1/14, \tag{8.64}$$

with $|M| = 7$ the cardinality of M [i.e. the number of signal models that go into the construction of $\mathcal{H}_{\rm S}$, see Eq. (8.36)]. In the latter, however, we prioritize GR by setting:

$$P(\mathcal{H}_{\rm GR} \mid \mathcal{H}_{\rm S}) = 1/2, \tag{8.65}$$

$$P(\mathcal{H}_m \mid \mathcal{H}_{\rm S}) = |\tilde{M}|^{-1}/2 = 1/12. \tag{8.66}$$

This distribution is illustrated schematically in Fig. 8.4. Note that these are not the only justifiable options; for example, we might want to prioritize $\mathcal{H}_{\rm GR}$ when constructing $\mathcal{H}_{\rm S}$ in order to better handle a noise background that does not conform to our assumption of Gaussianity. (Other strategies to tackle non-Gaussian noise are discussed in Sec. 8.3.1.) In any case, the code is sufficiently flexible to make different choices for the model priors if desired.

To study our method in the presence of signal, we perform several injections of scalar, vector and tensor polarizations (and combinations thereof) for all the 200 pulsars analyzed in [29]. The simulated signals have a range of signal-to-noise ratios (SNRs), which we proxy below by their *effective strain amplitudes*. We define these in terms of the a_p's from Eq. (8.13) by:

$$h_{\rm t} \equiv \sqrt{a_+^2 + a_\times^2}, \tag{8.67}$$

$$h_{\rm v} \equiv \sqrt{a_{\rm x}^2 + a_{\rm y}^2}, \tag{8.68}$$

$$h_{\rm s} \equiv a_{\rm s}, \tag{8.69}$$

for tensor, vector and scalar signals respectively. Each simulated signal is generated with a random value of the nuisance phase parameters (ϕ_0 or ϕ_p's). GR injections are always carried out using the triaxial template of Eq. (8.29), with random orientation parameters (ψ and ι) when those are not known. Location is always taken to be fixed at the known value for each pulsar.

8.5 Results

8.5.1 Model selection

Signal vs noise

We first show that $\mathcal{O}_{\rm N}^{\rm S}$, as defined in Eq. (8.43), can be used to discriminate signals of any polarization from Gaussian noise, without significant loss of sensitivity to GR signals. The black histogram in Fig. 8.5 shows the distribution of the natural logarithm of this quantity (henceforth, "log-odds"), obtained from the analysis of an ensemble of noise instantiations corresponding to a single source—in this case, the Crab pulsar. For comparison, the gray (unhatched) histogram in Fig. 8.5 is the analogous distribution for $\ln \mathcal{B}_{\rm N}^{\rm GR}$ [note that $\mathcal{B}_{\rm N}^{\rm GR} = \mathcal{O}_{\rm N}^{\rm GR}$ if we assign equal priors to the GR and Gaussian noise models, cf. Eq. (8.43) with m = GR]; this is the value computed in regular, GR-only targeted pulsar searches, although with different signal

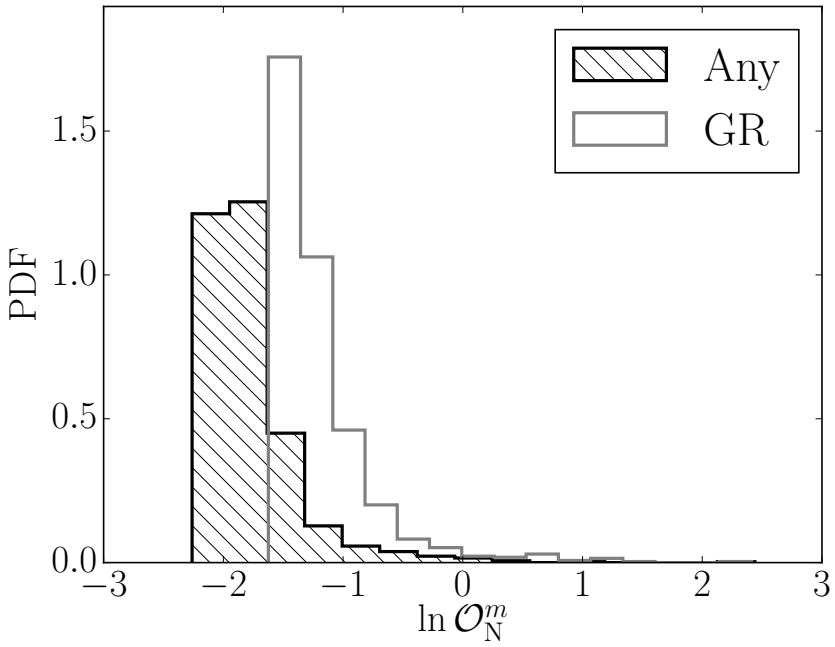

Figure 8.5: *Signal vs noise log-odds background distributions for any-signal and GR hypotheses*. Histograms of $\ln \mathcal{O}^{\rm S}_{\rm N}$ (black line, hatched) and $\ln \mathcal{O}^{\rm GR}_{\rm N}$ (gray line) over an ensemble of 1000 simulated noise instantiations corresponding to the Crab pulsar. For each instantiation, three time series of Gaussian noise were produced using the design noise spectra of H1, L1 and V1, as outlined in Sec. 8.4; the data are analyzed coherently across detectors. (Note that here $\ln \mathcal{O}^{\rm GR}_{\rm N} = \ln \mathcal{B}^{\rm GR}_{\rm N}$, since we assign equal weight to both models.)

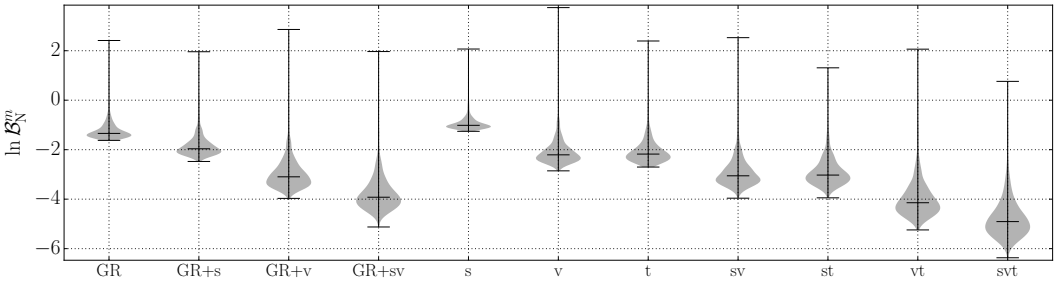

Figure 8.6: *Signal vs noise log-Bayes background distributions for all subhypotheses*. Violin plots representing histograms of the log-Bayes of several models vs noise, computed over an ensemble of 1000 simulated noise instantiations each corresponding to H1, L1 and V1 design data prepared for the Crab pulsar; the data are analyzed coherently across detectors. The labels on the x-axis indicate which hypothesis is being compared against noise; the "GR" label indicates tensor modes parametrized by Eq. (8.29) with fixed ψ and ι. Black lines mark the range and median of each distribution. (The gray histogram in Fig. 8.5 corresponds to the leftmost distribution here.)

amplitude priors [29]. Note that odds carry an intrinsic probabilistic meaning in terms of gambling probabilities, and a background histogram like this is not required to interpret their value (see e.g. [22]).

For both quantities shown in Fig. 8.5, a negative value marks a preference for the noise model ($\mathcal{H}_{\rm N}$, as defined at the beginning of section 8.3.1). However, note that a conservative (as determined by the priors) analysis should not be expected to strongly favor $\mathcal{H}_{\rm N}$, since the presence of a weak signal below the noise threshold cannot be discarded; this explains why the ranges in the plots of Fig. 8.5 do not extend to more negative values. Generally speaking, the magnitude of the signal prior volume (viz. the volume of parameter space allowed by the signal model, weighted by the prior function) will determine the mean of background distributions like Fig. 8.5, which will be more negative the greater the signal volume. This is a manifestation of an implicit Occam's penalty.

The relationship between the Bayes factors for different signal hypotheses vs noise is illustrated in Fig. 8.6, which shows violin plots representing the noise-ensemble distributions of $\ln \mathcal{B}^m_{\rm N}$ for all models discussed in 8.3.1. The values for $m \in \{\rm s, v, sv, GR, GR + s, GR + t, GR + sv\}$ are combined to produce $\ln \mathcal{O}^{\rm S}_{\rm N}$ in Fig. 8.5. As explained above, the "GR" label indicates that the tensor modes have been parametrized using the triaxial model of Eq. (8.29), with orientation parameters fixed at the known values for the Crab pulsar; on the other hand, the "t" label corresponds to the free-tensor template of Eq. (8.28). We include both parametrizations to demonstrate the effect of assuming a triaxial emission mechanism and restricting the orientation of the source (see also Appendix 8.7).

Interestingly, Fig. 8.6 reveals the relationship between $\ln \mathcal{B}^m_{\rm N}$ and the number of degrees of freedom (a proxy for the prior volume) of model m: models with more degrees of freedom have a greater prior volume and are correspondingly downweighted, resulting in more negative values of $\ln \mathcal{B}^m_{\rm N}$; this is a manifestation of the Occam's penalty automatically applied by the Bayesian analysis (see e.g. Ch. 28 in [23]). We underscore that this feature arises naturally from the computation of the evidence integral, and not from manually downweighting either model *a priori*.

If the data contain a sufficiently loud signal of any polarization, the evidence for $\mathcal{H}_{\rm S}$ will surpass that for $\mathcal{H}_{\rm N}$, and this can be used to establish a detection. Fig. 8.7 shows the response of $\ln \mathcal{O}^{\rm S}_{\rm N}$ and $\ln \mathcal{B}^{\rm GR}_{\rm N}$ to the presence of GR and non-GR signals. In particular, the second panel in Fig. 8.7 shows results for injected signals of the vector-only model of [106], but the behavior would be the same for scalar-only

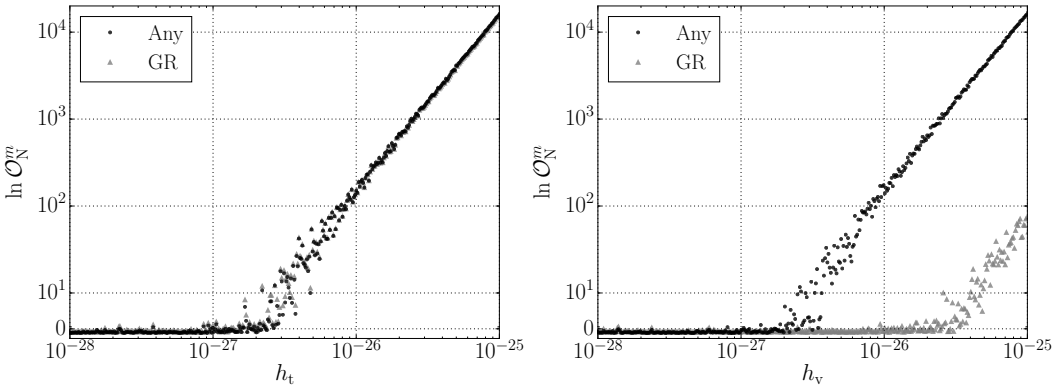

Figure 8.7: *Expected sensitivity to GR and vector injections*. Log-odds of any-signal ($\mathcal{H}_{\rm S}$, black circles) and GR ($\mathcal{H}_{\rm GR}$, gray triangles) vs noise ($\mathcal{H}_{\rm N}$) hypotheses, as a function of injection amplitude, for signals corresponding to both GR (left) and the vector-only model from [106] (right). The any-signal odds are defined in Eq. (8.43). Each of the 500 points corresponds to a data instantiation (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar signal of the indicated strength. The injections were performed with random values of the nuisance phase parameters, and the data were analyzed coherently across detectors. A logarithmic scale is used for the y-axis, except for a linear stretch corresponding to the first decade.

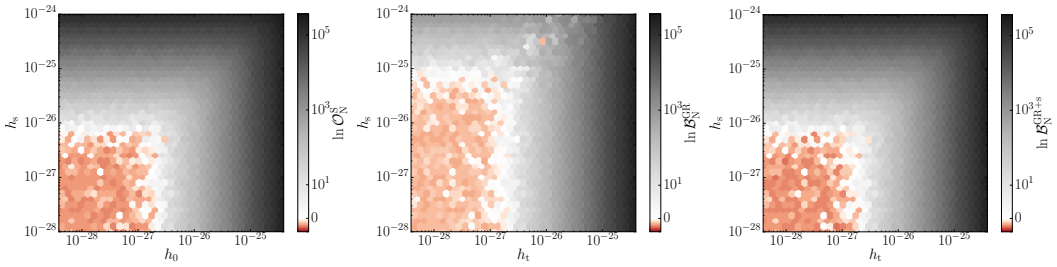

(a) Any (b) GR (c) GR+s

Figure 8.8: *Expected sensitivity to scalar-tensor injections*. Log-odds of any-signal ($\mathcal{H}_{\rm S}$, left), GR ($\mathcal{H}_{\rm GR}$, center) and GR+s ($\mathcal{H}_{\rm GR+s}$, right) hypotheses vs noise. The any-signal odds are defined in Eq. (8.43). Each plot was produced by analyzing 2500 instantiations of data (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar GR+s signal of the indicated tensor (x-axis) and scalar (y-axis) amplitudes. The color of each hexagon represents the average value of the log-odds in that region of parameter space; color is normalized logarithmically, except for a linear stretch in the $(-1, 1)$ range.

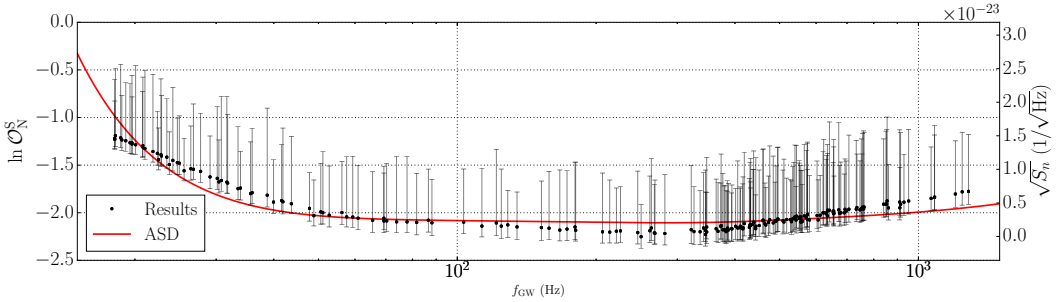

Figure 8.9: *Signal log-odds vs GW frequency for noise-only data*. Circles mark the mean of the distribution of $\ln \mathcal{O}_N^S$, as a function of the expected GW frequency for each pulsar in our set; vertical lines indicate one-sided standard deviations for each source. Each data point and corresponding bars summarize the shape of a distribution like Fig. 8.5 for each of the pulsars, but produced from only 100 runs per source. The effective noise amplitude spectral density $\sqrt{S_n}$ (ASD, red curve), corresponding to the harmonic mean of each detector PSD, is overlaid for comparison (scaling obtained from a linear regression).

signals. The general features of these plots confirm our expectations that for weak, subthreshold signals, the analysis should not be able to distinguish between the signal and noise models, yielding a Bayes factor close to unity (more precisely, a value of $\ln \mathcal{O}_N^S$ consistent with the background distributions of Fig. 8.5). Note that, in agreement with Fig. 8.5, the noise baseline for $\ln \mathcal{O}_N^S$ lies below that of $\ln \mathcal{B}_N^{GR}$, due to its greater prior volume.

For stronger (detectable) signals, the basic form of our likelihood functions, Eq. (8.25), leads us to expect $\ln \mathcal{O}_N^S$ to scale linearly with the square of the signal-to-noise ratio (SNR):

$$\ln \mathcal{B}_N^m \sim \left(\mathbf{B} \cdot \mathbf{\Lambda}_{MP} - |\mathbf{\Lambda}_{MP}|^2/2\right)/\sigma^2 \propto \left(h_{inj}/\sigma\right)^2, \tag{8.70}$$

where the variance σ^2 proxies the PSD and we let $\mathbf{\Lambda}_{MP}$ be the time-series vector corresponding to the maximum probability template; for a stationary PSD, this implies $\ln \mathcal{B}_N^m \propto h_{inj}^2$, as observed in Fig. 8.7. The spread around the trendline is due to the individual features of each noise instantiation and (much less so) to numerical errors in the computation of the evidence, Eq. (8.40). For details on numerical uncertainty, see Appendix 8.9.

From the left panel of Fig. 8.7, we conclude that $\ln \mathcal{O}_N^S$ can be as good an indicator of the presence of GR signals as $\ln \mathcal{B}_N^{GR}$ itself; this implies that we may include non-GR polarizations in our search without significantly sacrificing sensitivity to GR signals. However, the power of $\ln \mathcal{O}_N^S$ lies in responding also to non-GR signals. As

an example of this, the right panel in Fig. 8.7 shows $\ln \mathcal{O}^{\mathrm{S}}_{\mathrm{N}}$ and $\ln \mathcal{B}^{\mathrm{GR}}_{\mathrm{N}}$ as a function of the amplitude of a fully non-GR injection. Here, we have chosen to inject a particular model of vector signal developed in [106], but the results are generic.

Note that, for sufficiently loud signals, $\mathcal{H}_{\mathrm{GR}}$ becomes preferable over $\mathcal{H}_{\mathrm{N}}$ (hence $\ln \mathcal{B}^{\mathrm{GR}}_{\mathrm{N}} > 0$), even when the injection model does not match the search; this is because the noise evidence drops faster than GR's and becomes very small (i.e. the data do not look at all like Gaussian noise, although they do not match the expected GR signal well either). The particular SNR at which this occurs will depend on the overlap between the antenna patterns of the injection and those of GR, and will consequently vary among sources.

For the interesting case of scalar-tensor theories (here, templates composed of GR plus an extra breathing component, and denoted "GR+s"), the behavior is slightly different. This is both because GR+s has an extra amplitude degree of freedom (a_{s}) and, as discussed in Sec. 8.3.1, because $\mathcal{H}_{\mathrm{GR}}$ can be recovered as a special case of $\mathcal{H}_{\mathrm{GR+s}}$ (namely, when $a_{\mathrm{s}} \to 0$). In Fig. 8.8, we present the log-odds of signal vs noise hypotheses as a function of injected GR (x-axis) and scalar (y-axis) strengths. These plots divide the h_{s}–h_{t} plane in roughly two regions where the associated signal model ($\mathcal{H}_{\mathrm{S}}$, $\mathcal{H}_{\mathrm{GR}}$ or $\mathcal{H}_{\mathrm{GR+s}}$) is preferred (black) and where it is not (red). The latter corresponds to the area of parameter space associated with subthreshold signals that cannot be detected.

As expected, the best coverage is obtained when analyzing the data using the model matching the injection, GR+s, (rightmost plot) or the all-signal model (leftmost plot). In both these cases, the results improve with either scalar or tensor SNR. In contrast, the GR analysis (center plot) is sensitive to tensor strain, but, as evidenced by the extended red region in the central plot, it misidentifies strong scalar signals as noise. Nevertheless, if the scalar component is larger than $\sim 5 \times 10^{-26}$, the GR analysis will disfavor the noise hypothesis, even for a small tensor component, as in the right panel of Fig. 8.7; this is the same behavior observed in Fig. 8.7. In contrast, the any-signal analysis is sensitive to the total power of the injected signal, regardless of polarization.

We have produced distributions of background $\ln \mathcal{O}^{\mathrm{S}}_{\mathrm{N}}$, like those of Fig. 8.5, for all 200 known pulsars in the sensitive band of the three detectors under consideration (same set analyzed in [29]. In Fig. 8.9, these are represented by their respective means and one-sided standard deviations as a function of the pulsar's GW frequency. The frequency dependence is explained by variations in the instrumental noise spectra.

This is explained by the fact that, for a particular prior choice, more information is gained from the data if the noise floor is lower: with less noise it is possible to discard the presence of weaker signals, so the value of $\ln \mathcal{O}^{\mathrm{S}}_{\mathrm{N}}$ decreases.

GR vs non-GR

In the presence of a signal, $\mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$, as defined by Eq. (8.45), indicates whether there is reason to believe there is a GR violation or not. Because there could always be an unresolvably small departure from GR, we do not expect our analysis (with priors as chosen) to ever strongly favor the GR hypothesis; rather, in the presence of a GR signal we will find that $\ln \mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ remains relatively close to zero, simply meaning that there is no strong evidence for or against non-GR features. This is indeed the behavior observed in the left panel of Fig. 8.10, where $\ln \mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ is shown to be roughly insensitive to tensor injection amplitude. For values of h_{t} below certain threshold (which, in this case, is around 3×10^{-27}), the search does not detect a signal and, consequently, no information is gained for or against $\mathcal{H}_{\mathrm{GR}}$, i.e. $\ln \mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}} \sim 0$. The difference between the two populations (below and above threshold) is determined mainly by the choice of amplitude priors.

The behavior of $\mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ is less ambiguous in the presence of a non-GR signal. For instance, if the data contain a detectable signal that completely lacks tensor components, then $\mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ will unequivocally reflect this. This is evidenced by the growth of $\ln \mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ with injected nontensorial SNR in the right panel of Fig. 8.10. In other words, while the analysis is inconclusive for GR injections because it cannot discard the presence of subthreshold non-GR components hidden by the noise, vector signals are are clearly identified as not conforming to GR. This is a reflection of the fact that, as mentioned in the introduction, any evidence of a nontensorial component is fatal for GR, but absence of non-GR components does not mean Einstein's theory is necessarily correct.

As might be expected, $\mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ responds to non-GR signals that include a tensor component with a combination of features from both panels of Fig. 8.10. As an example, the left plot of Fig. 8.11 shows $\mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ in the presence of GR+s injections, as a function of injected tensor and scalar amplitudes. This plot can be split into three clearly demarcated regions: one in which the signal is not detected (light red, bottom left), one in which the signal is detected and the non-GR model is preferred (black, top), and one which the signal is detected but where the evidence for a deviation

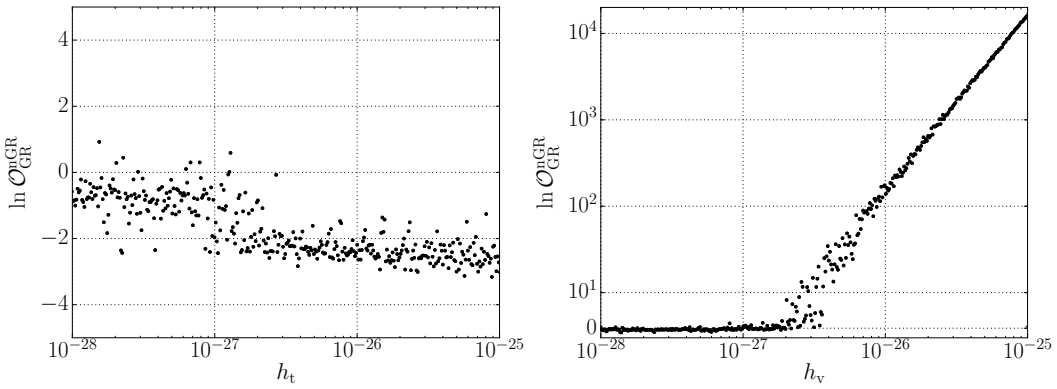

Figure 8.10: *Categorizing tensor and vector injections* ($\mathcal{H}_{\mathrm{nGR}}$ vs $\mathcal{H}_{\mathrm{GR}}$). Non-GR vs GR log-odds, as a function of effective injection amplitude, for both GR (left) and the vector-only model from [106] (right). Each of the 500 points corresponds to a data instantiation (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar signal of the indicated strength. The injections were performed with random values of the nuisance phase parameters, and the data were analyzed coherently across detectors. Note that, on the right, a logarithmic scale is used for the y-axis, except for a linear stretch corresponding to the first decade.

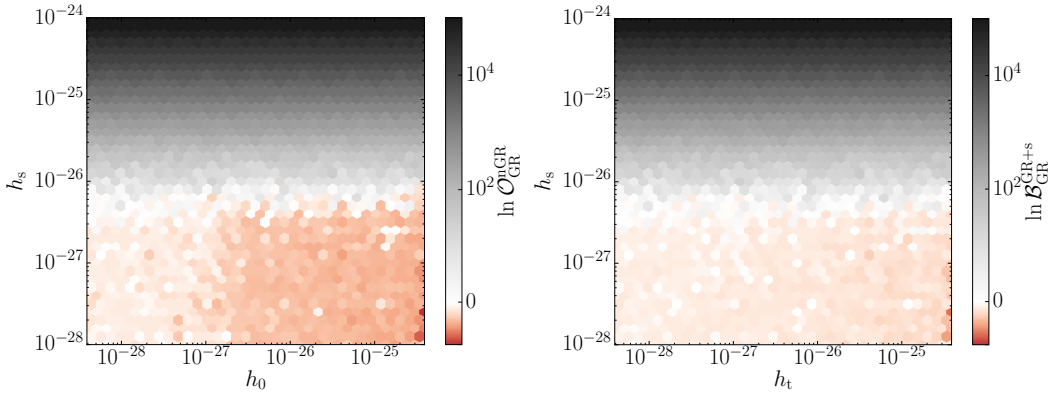

Figure 8.11: *Categorizing scalar-tensor injections* ($\mathcal{H}_{\mathrm{nGR}}$ & $\mathcal{H}_{\mathrm{GR+s}}$ vs $\mathcal{H}_{\mathrm{GR}}$). Log-odds comparing the non-GR and GR+s hypotheses to GR. The non-GR odds is defined in Eq. (8.45). Each plot was produced by analyzing 2500 instantiations of data (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar GR+s signal of the indicated tensor (x-axis) and scalar (y-axis) amplitudes. The color of each hexagon represents the average value of the log-odds in that region of parameter space; color is normalized logarithmically, except for a linear stretch in the $(-1, 1)$ range.

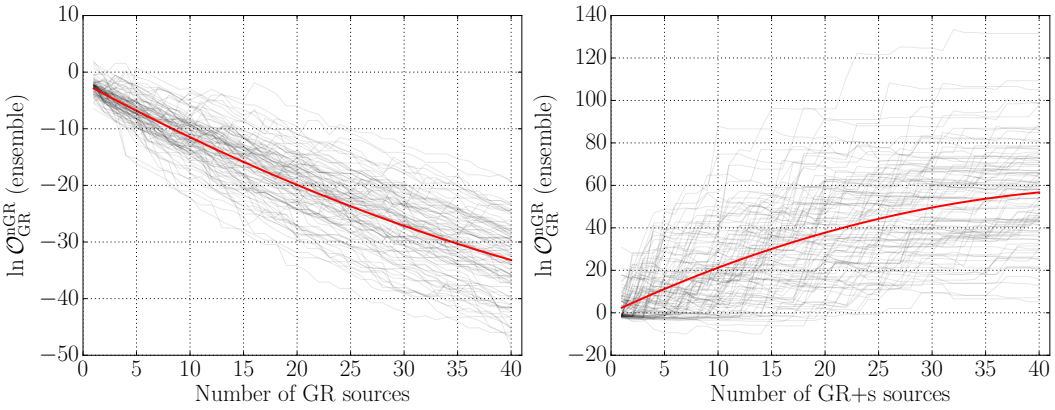

Figure 8.12: *Ensemble non-GR vs GR log-odds*. Non-GR vs GR log-odds computed from data for multiple sources vs the number of sources in the set. Each light-gray trace marks a possible progression of the ensemble log-odds as new sources are added; the red line corresponds to the best quadratic fit. For each pulsar, we chose an arbitrary data instantiation containing a GR (left) or GR+s (right); GR signals are restricted to $10^{-27} < h_{\mathrm{t}} < 10^{-26}$, while GR+s signals also satisfy $0.3 < h_{\mathrm{s}}/h_{\mathrm{t}} < 1$. We compute the value of $\ln \mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$ for each signal in the set and combine them according to Eq. (8.54) to obtain the ensemble value plotted in the y-axis.

from GR is not clear due to the predominance of the tensorial component (darker red, bottom right). The first corresponds to the subthreshold population on either side of Fig. 8.10, while the second and third correspond to the above-threshold populations on the right and left sides of Fig. 8.10 respectively; indeed, note that a horizontal slice taken over the red region of the left plot produces a series of points like those in the left panel of Fig. 8.10. For reference, Fig. 8.11 also includes the direct comparison of GR+s and GR on the right.

We can make a stronger statement about the agreement of the data with GR by making use of signals from multiple sources, as discussed in Sec. 8.3.1. The power of combining multiple signals is illustrated in Fig. 8.12, where $\ln \mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$, as defined in Eq. (8.54), is plotted vs number of GR (left) and GR+s (right) signals detected. Note that this presumes that, for each source, the presence of a signal has already been established from the value of $\ln \mathcal{O}^{\mathrm{S}}_{\mathrm{N}}$. Computing the ensemble $\ln \mathcal{O}^{\mathrm{nGR}}_{\mathrm{GR}}$, as done here, is a good way of summarizing the information contained in the data about the relative likelihoods between the two models, but it provides no information not already present in the set of individual single-source odds.

8.5.2 Parameter estimation

When no conclusive evidence for a CW is found in the data, we are still interested in placing upper limits on the strength of possible signals (up to some credibility), and this is done as explained in Sec. 8.3.2. By the same token, if a signal consistent with GR is detected, we can always place an upper limit on the amplitude of non-GR modes, even if the odds indicate there is no clear sign of a GR violation.

For instance, we can get a quantitative estimate of our sensitivity to scalar modes from a given source by looking at the distribution of $h_{\rm s}^{95\%}$, defined in Eq. (8.62), computed for a set of noise-only data instantiations. Such distribution for the Crab pulsar is presented in the left panel of Fig. 8.13. Similarly, the right panel presents estimates for the sensitivity to vector modes coming from the Crab pulsar, assuming a vector-tensor model. In this case, however, the quantity plotted is the upper limit on total, effective vector strain amplitude $h_{\rm v}$, Eq. (8.68). These plots include distributions produced using the same log-uniform prior used to obtain Bayes factors, as well as more conservative ones obtained using uniform amplitude priors (see Appendix 8.8). In either case, the magnitude of $h_{\rm v}^{95\%}$ is comparable to that of $h_{\rm s}^{95\%}$.

Interestingly, our ability to measure scalar and vector amplitudes is unaffected by the presence of other modes. We illustrate this for the Crab pulsar in Fig. 8.14, which results from analyzing data with GR+s (left) and GR+v (right) injections. There we plot $h_{\rm s}^{95\%}$ as a function of scalar and tensor injection amplitudes on the left, and $h_{\rm v}^{95\%}$ as a function of vector and tensor injection amplitudes on the right. From these plots, one can conclude that $h_{\rm s}^{95\%}$ and $h_{\rm v}^{95\%}$ are sensitive only to the corresponding scalar and vector components, and not by $h_{\rm t}$. (It is worth emphasizing that the upper limits, $h_{\rm s}^{95\%}$ and $h_{\rm v}^{95\%}$, are well-defined even when the non-GR component is strong enough to be detected, as is the case for the darker-colored regions.)

As shown previously in the literature, the mean of distributions like those of Fig. 8.13 will scale with $\sqrt{S_{\rm n}(f)/T}$, where $S_{\rm n}(f)$ is the effective PSD of the detector noise at the expected GW frequency f, and T is the integration time (cf. Eq. (26) and Fig. 1 of [57]). Because of this, the mean of this distribution will vary with the source's expected GW frequency, as shown in Fig. 8.15. Following convention, these upper limits are computed using uniform amplitude priors, which means they are a factor of a few less stringent than those obtained with a log-uniform prior (see Fig. 8.13 and Appendix 8.8). Also, for completeness, Fig. 8.15 also includes the expected tensor upper limits, $h_{\rm t}^{95\%}$. Note that those values are *not* the same as would be obtained by the standard GR-only search, because that analysis looks at the triaxial h_0 of Eq.

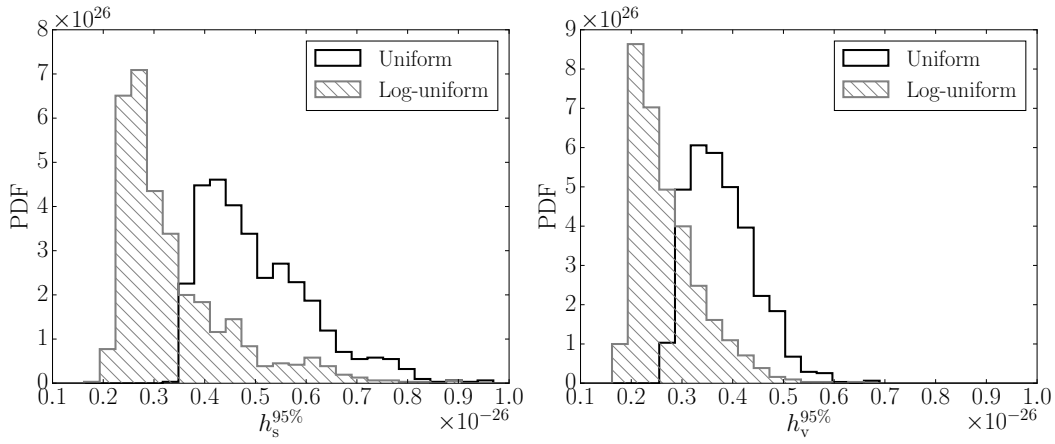

Figure 8.13: *Expected Crab non-GR upper limits in absence of signal*. Histogram of 95%-credible upper limits for the scalar (left) and vector (right) amplitudes, for a set of 1000 noise-only data sets, computed using priors uniform in the amplitude (black) or uniform in the logarithm of the amplitude (hatched gray); the differences between these two priors are discussed in detail in Appendix 8.8. Each instantiation (one time series for each detector: H1, L1 and V1) is made up of simulated Gaussian noise with standard deviation given by the advanced design PSDs. Scalar and vector upper limits are produced using the GR+s and GR+v models respectively.

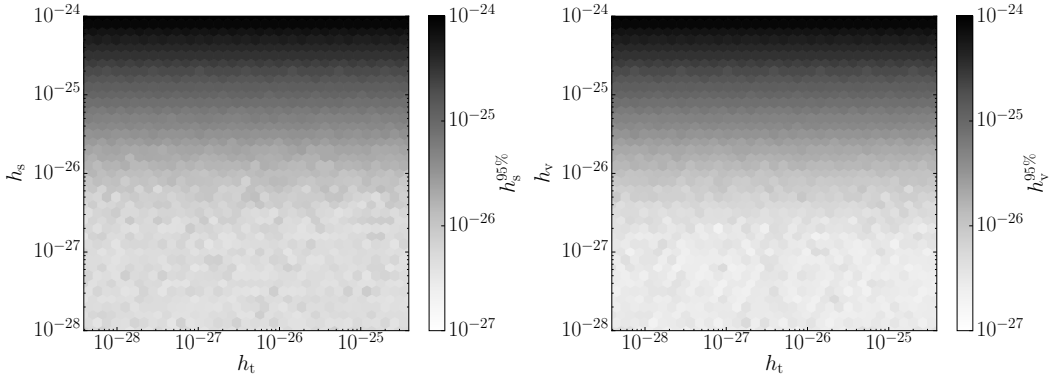

Figure 8.14: *Expected Crab scalar and vector upper limits in presence of GR+s and GR+v signals*. Shading represents the 95%-credible upper limit for the scalar ($h_{\rm s}^{95\%}$, left) and vector ($h_{\rm v}^{95\%}$, right) amplitudes vs the amplitude of injected GR (x-axis) and corresponding non-GR (y-axis) components. Each plot was produced by analyzing 2500 instantiations of data (one time series for each detector: H1, L1 and V1) made up of Gaussian noise plus a simulated Crab-pulsar GR+s (left) or GR+v (right) signal with indicated strains. The color of each hexagon represents the average value of the upper limit in that region of parameter space.

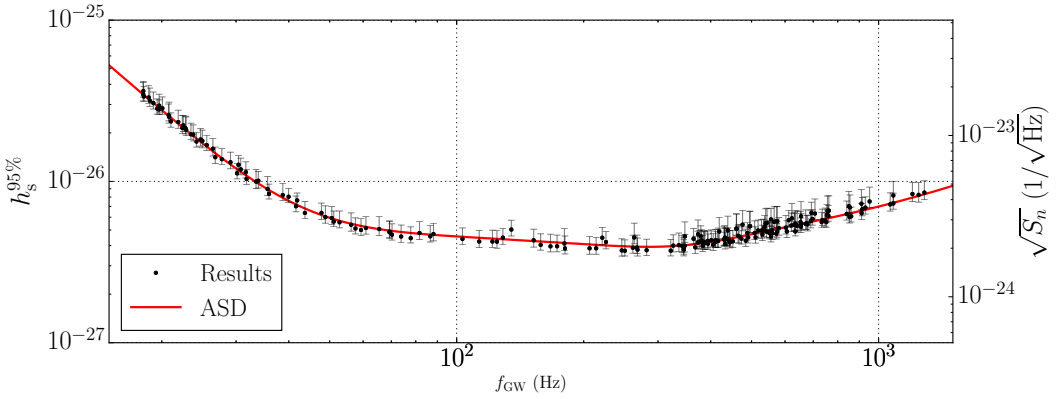

(a) Scalar-tensor

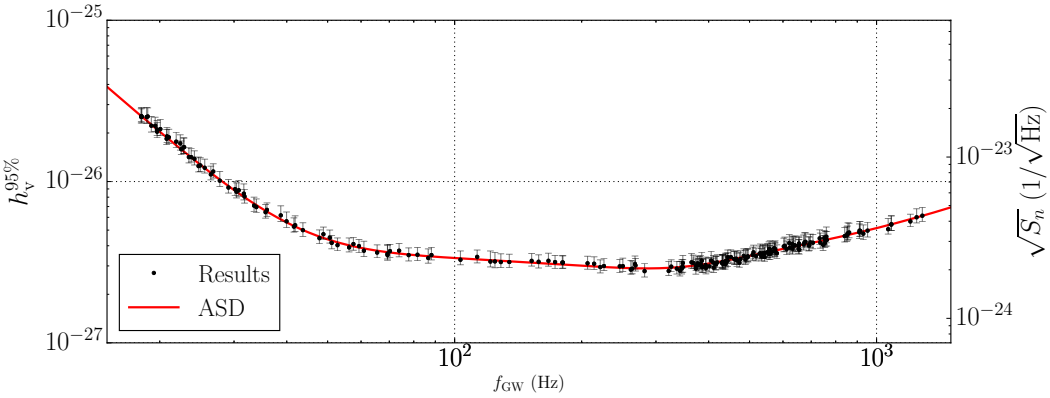

(b) Vector-tensor

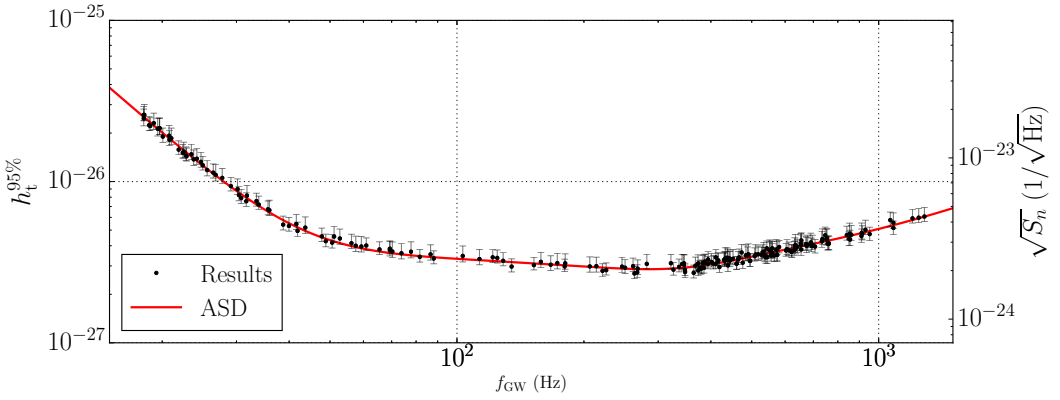

(c) Tensor

Figure 8.15: *Expected upper limits in absence of signal vs GW frequency*. Circles mark the mean of the distribution of $h_s^{95\%}$ (top), $h_v^{95\%}$ (middle) and $h_t^{95\%}$ (bottom), as a function of expected GW frequency for each pulsar in our set; vertical lines mark one-sided standard deviations for each source. Each data point and corresponding bars summarize the shape of a distribution like those of Fig. 8.13, but produced from 100 noise instantiations each. The scalar, vector and tensor upper limits were produced assuming st, vt and t models respectively. We use uniform priors in all amplitude parameters (see Fig. 8.13 and Appendix 8.8).

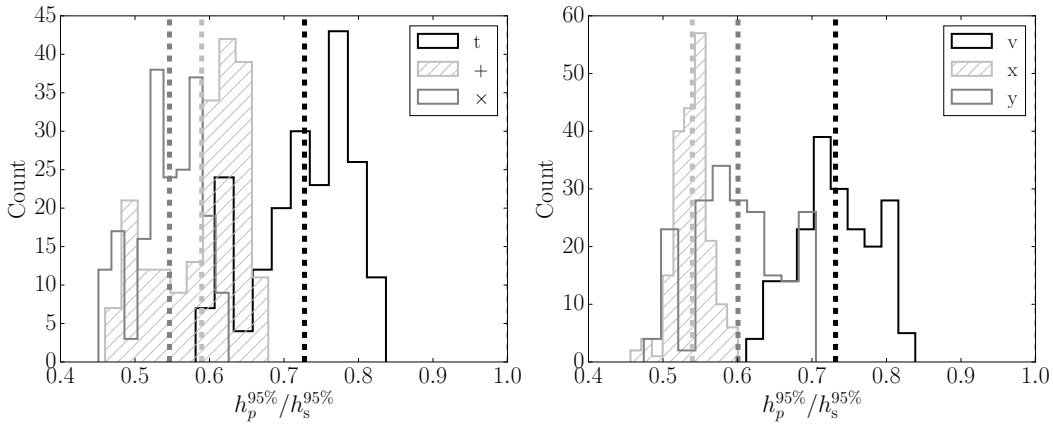

Figure 8.16: *Tensor and vector upper limits as a ratio of scalar upper limits*. Histogram of tensor (top) and vector (bottom) upper limits divided by the scalar upper limit for each pulsar. The top plot shows ratios for $h_t^{95\%}$ (black), $h_+^{95\%}$ (light gray, hatched), and $h_\times^{95\%}$ (dark gray); the bottom plot shows ratios for $h_v^{95\%}$ (black), $h_x^{95\%}$ (light gray, hatched), and $h_x^{95\%}$ (dark gray). Vertical dashed lines mark the mean of each distribution.

(8.29), rather than h_t.

In order to compare our sensitivity to the different polarizations, in Fig. 8.16 we histogram the tensor and vector upper limits as a ratio of the scalar upper limits—this includes the $h_t^{95\%}$ and $h_v^{95\%}$ values shown in Fig. 8.15, as well as the limits on the individual amplitudes from which they are constructed ($h_+^{95\%}$, $h_\times^{95\%}$, $h_x^{95\%}$ and $h_y^{95\%}$). The mean of these distributions (vertical dashed lines in Fig. 8.16) indicate that, for most pulsars, the scalar upper limit is slightly larger in magnitude than those for the +, ×, x or y modes; this systematic effect is a manifestation of the decreased sensitivity of quadrupolar detectors to scalar waves, which was discussed in Sec. 8.2.1 (see, in particular, Fig. 8.2). The fact that the difference between $h_s^{95\%}$ and $h_t^{95\%}$, or $h_v^{95\%}$, is less pronounced can be easily be explained as a statistical factor arising from the definitions of h_t and h_v as square-roots of sums of squares, Eqs. (8.67, 8.68). Both these scalings are discussed in more detail in Appendix 8.10.

8.6 Summary

We have developed a Bayesian framework to detect CW signals from known sources regardless of polarization content, to disentangle the modes present in a given signal, and to constrain the amplitudes of extra polarizations that may be hiding under the noise. We have implemented this as an extension of LIGO's Bayesian targeted CW search pipeline [51] and thus benefit from the power of the nested sampling algorithm

on which it is based.

We have tested our methods on one year of simulated noise for three advanced-era detectors at design sensitivity (H1, L1, V1), and prepared for a set of multiple known sources in their frequency band. This allows us to estimate our future sensitivity to CW polarizations, in this most optimistic case. Under these conditions and for the Crab pulsar in particular, we expect signals of any polarization to become detectable for characteristic strain amplitudes $h \gtrsim 3 \times 10^{-27}$ (Figs. 8.7 and 8.8); this threshold will vary among sources, due to differences in position (sky location and orientation) and detector PSD at the expected GW frequency (cf. e.g. Fig. 8.9). Furthermore, the value of this threshold will decrease linearly with the square-root of the observation time [57].

A signal louder than the detection threshold will allow us to determine whether its polarization content is consistent with GR or not, and the strength of this statement will depend almost exclusively on the power of the non-GR component (Figs. 8.10 and 8.11). In other words, from a model-selection standpoint, the non-GR hypothesis will only be unequivocally favored if the total power in non-GR modes is greater than the threshold value, regardless of the strength of the GR modes. However, for signals that do not satisfy this, we may always place upper limits on nontensorial amplitudes and thus constrain deviations from GR; for instance, Fig. 8.15 presents the most optimistic projections for 95%-credible upper limits for scalar and vector amplitudes of CW signals from all pulsars in our set ($h_{\mathrm{s}}^{95\%} \sim 4 \times 10^{-27}$ and $h_{\mathrm{v}}^{95\%} \sim 3 \times 10^{-27}$, in the best case). As far as we are aware, these are the first generic estimates of sensitivity to scalar and vector polarizations ever published [1].

From our projected upper limits, we have found that, at design sensitivity, the LIGO-Virgo network will be generally less sensitive to continuous scalar signals than to the individual vector or tensor modes by factors of 0.45–0.7, depending on the location of the source (Fig. 8.16); this diminished sensitivity to scalar modes stems from the quadrupolar nature of the detector antenna patterns (Fig. 8.2 and Appendix 8.10). Also, our injection studies indicate that the upper limits on the amplitudes of nontensorial modes will be roughly unaffected by the presence or absence of a tensor signal in the data (Fig. 8.14).

Although the results presented here made use of simulated Gaussian noise, the procedure is identical for actual detector data. Furthermore, the assumption of

[1] Note that sensitivity estimates presented in [109] were restricted to the specific vector-only model of [106]

Gaussianity has been shown to hold relatively well for real CW data [109], so the actual sensitivity limits should not be far from those presented here. If the data are strongly non-Gaussian, however, one must be careful in using $\ln O_{\mathrm{N}}^{\mathrm{S}}$ for detection purposes and may instead wish to adopt one of the strategies suggested in Sec. 8.3.1.

Another important limitation of our results is that here we only consider CW signals emitted at $f = 2f_{\mathrm{rot}}$, while it is to be expected that other mechanisms (within GR or not) allow emission at other harmonics, $f = f_{\mathrm{rot}}$ in particular. Yet, the only change required to account for this is to modify the template in Eq. (8.12) to include terms at different harmonics; the ability to do this already exists within our current infrastructure. We also assume that other aspects of the waves, like their speed, remain in agreement with the GR prediction, an assumption that will be relaxed in a future study.

Acknowledgments

The authors would like to thank Ian Jones and Walter Del Pozzo for carefully reading this manuscript and providing insightful suggestions; we also thank Tjonnie Li and Carver Mead, as well as many colleagues in the LIGO Scientific Collaboration Continuous Waves group, for many useful comments. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. We are grateful for computational resources provided by Cardiff University, and funded by an STFC grant supporting UK Involvement in the Operation of Advanced LIGO. MP is funded by the STFC under grant number ST/N005422/1. Plots produced using Matplotlib [81]. This paper carries LIGO Document Number LIGO-P1600305.

8.7 Appendix: Tensor models

A conceptual distinction can be drawn between the triaxial GR model and a free-tensor model that includes + and × but does not restrict their relative amplitudes (denoted "t"). The former has four free parameters (overall amplitude, h_0; overall phase, ϕ_0; inclination, ι; polarization, ψ) and corresponds to a signal template of the form [same as Eq. (8.29)]:

$$\Lambda_{\mathrm{GR}}(t) = \frac{1}{2}h_0 e^{i\phi_0}\left[\frac{1}{2}(1+\cos^2\iota)F_+(t;\psi) - i\cos\iota F_\times(t;\psi)\right]. \tag{8.71}$$

This is a reparametrization of the free-tensor model, which also has four parameters (plus amplitude, a_+; cross amplitude, $a_\times$; plus phase, ϕ_+; cross phase, $\phi_\times$) and whose

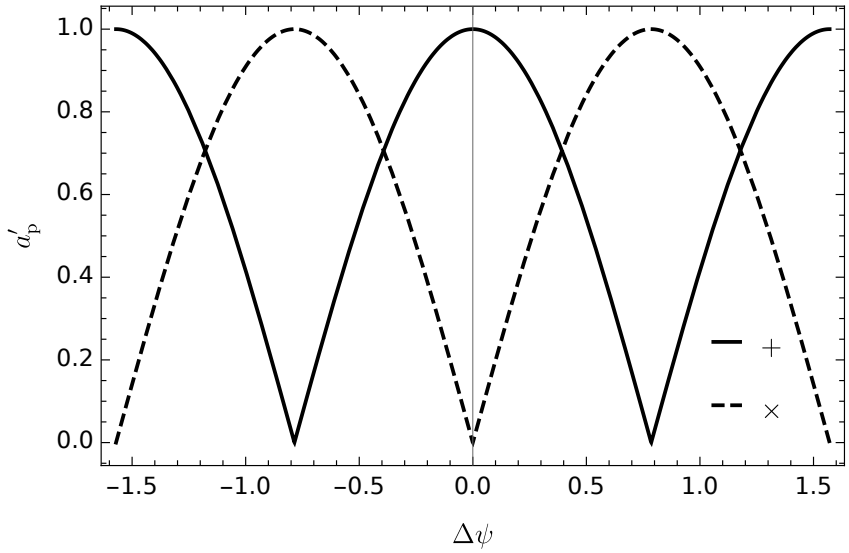

Figure 8.17: *Effect of changing polarization angle*. Norm of the complex plus (a'_+, solid line) and cross ($a'_\times$, dashed line) weights after rotating the source by $\Delta\psi$ in the plane of the sky, i.e. letting $\psi \to \psi' = \psi + \Delta\psi$; this transformation is expressed in Eqs. (8.75) and (8.76). In this case, we start from $a_+ = 1$, $a_\times = 0$ and $\psi = 0$.

template is [same as Eq. (8.28)]:

$$\Lambda_{\rm t}(t) = \frac{1}{2}\left[a_+ e^{i\phi_+} F_+(t;\psi=0) + a_\times e^{i\phi_\times} F_\times(t;\psi=0)\right]. \tag{8.72}$$

If ψ and ι are known, it is clear that the two models are different, since $\mathcal{H}_{\rm GR}$ has two free parameters (h_0, ϕ_0) and $\mathcal{H}_{\rm t}$ has four (a_+, $a_\times$, ϕ_+, $\phi_\times$). If the orientation is *not* fixed, however, the two models span the same signal space. This is because there is a degeneracy between ψ and a_+, $a_\times$ due to the way the antenna patterns transform under changes in ψ:

$$F_+(t;\psi') = F_+(t;\psi)\cos 2\Delta\psi + F_\times(t;\psi)\sin 2\Delta\psi, \tag{8.73}$$

$$F_\times(t;\psi') = F_\times(t;\psi)\cos 2\Delta\psi - F_+(t;\psi)\sin 2\Delta\psi, \tag{8.74}$$

with $\psi' = \psi + \Delta\psi$. Eqs. (8.73) and (8.74) can be derived from Eqs. (8.1) and (8.2) respectively, as in [108] (or see, e.g., Sec. 9.2.2 in [48]). Consequently, changing $\psi \to \psi'$ in Eq. (8.72) is equivalent to leaving ψ fixed [at, say, $\psi = 0$ as in Eq. (8.72)] while replacing the plus and cross complex amplitudes by:

$$a'_+ e^{i\phi'_+} = a_+ e^{i\phi_+}\cos 2\Delta\psi - a_\times e^{i\phi_\times}\sin 2\Delta\psi, \tag{8.75}$$

$$a'_\times e^{i\phi'_\times} = a_\times e^{i\phi_\times}\cos 2\Delta\psi + a_+ e^{i\phi_+}\sin 2\Delta\psi. \tag{8.76}$$

This is illustrated in Fig. 8.17.

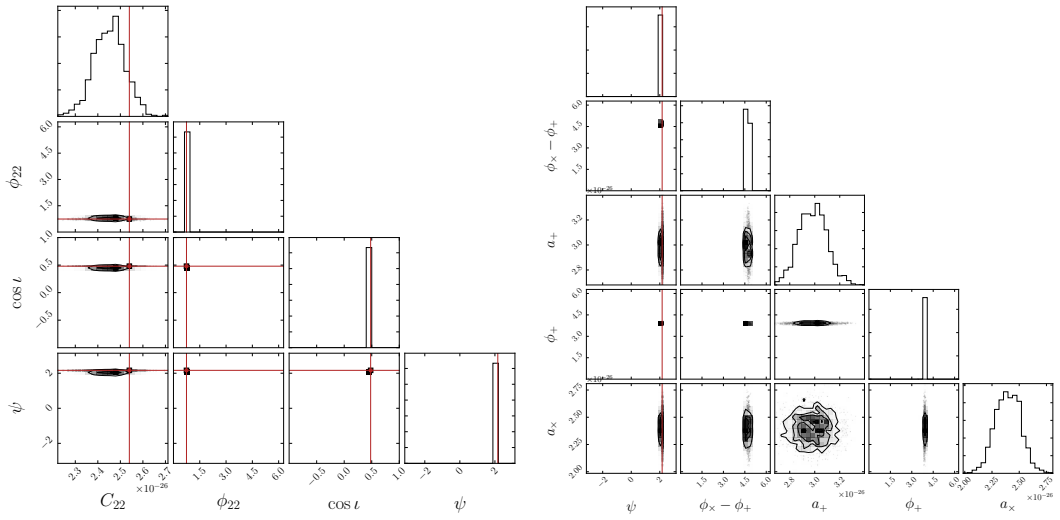

(a) GR, fixed orientation

(b) Free tensor, fixed orientation

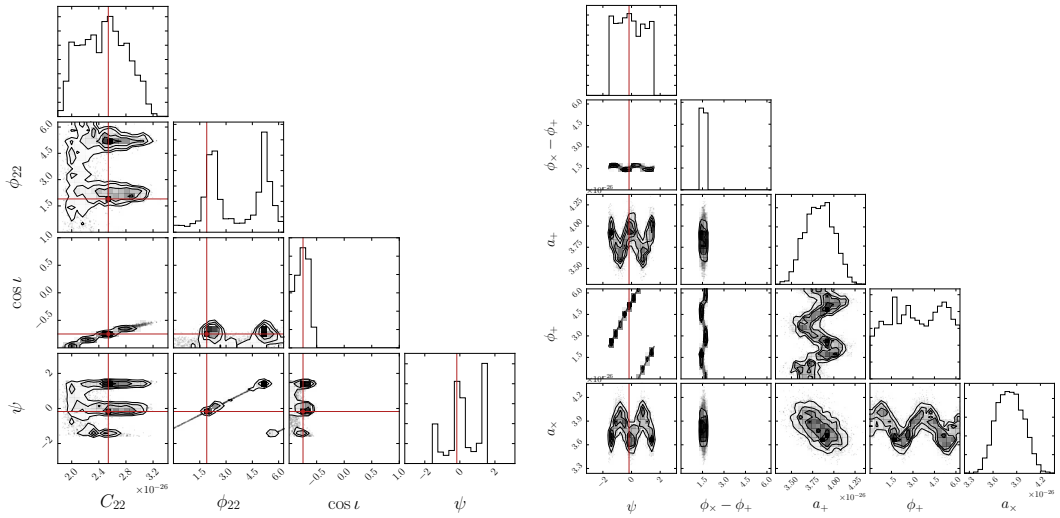

(c) GR, unfixed orientation

(d) Free tensor, unfixed orientation

Figure 8.18: *Tensor posteriors in presence of signal*. Posterior PDFs for parameters of $\mathcal{H}_{\rm GR}$ (left) and $\mathcal{H}_{\rm t}$ (right) with fixed (top) and unfixed (bottom) source orientation (ψ, ι). Each panel consists of a *corner plot* displaying the two-dimensional posteriors for each pair of parameters as indicated by the x and y labels, with the diagonals showing a histogram of the one-dimensional PDF for each parameter [i.e. the 1D PDF obtained after marginalization of the multidimensional posterior PDF all other quantities, as in Eq. (8.61)]. The data analyzed contain signals with parameters indicated by the red lines; note that $C_{22} = h_0/2$ is the quantity that was actually used to parametrize GR triaxial amplitudes in the code [51]. In both (a) and (b), $\cos\iota$ and ψ are fully known, and their resolution in these plots is limited by binning only. These plots were produced using the CORNER.PY package [164].

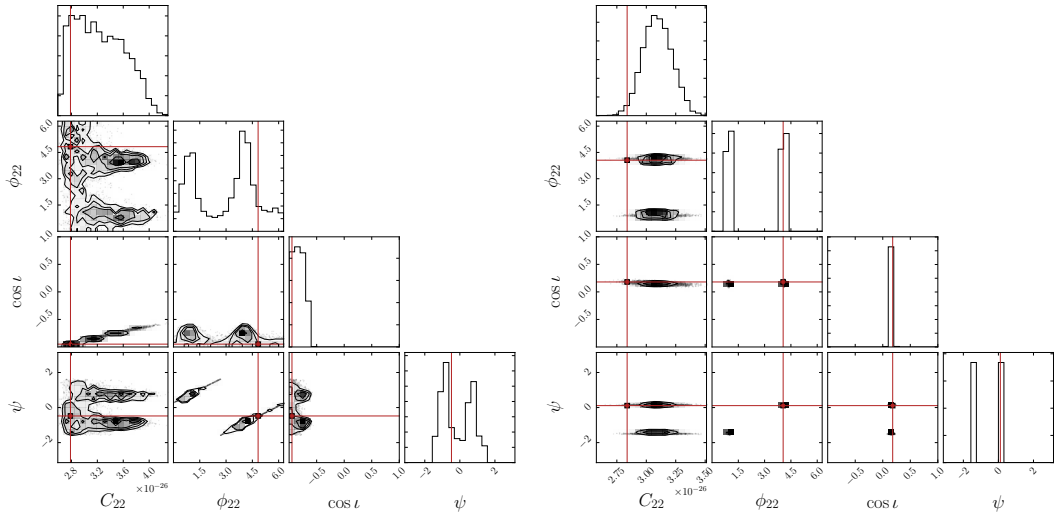

Figure 8.19: *Effect of inclination.* Posterior PDFs for parameters of $\mathcal{H}_{\rm GR}$ with unfixed source orientation (ψ, ι). Each panel consists of a *corner plot* displaying the two-dimensional posteriors for each pair of parameters as indicated by the x and y labels, with the diagonals showing a histogram of the one-dimensional PDF for each parameter. The data sets analyzed contain signals with parameters indicated by the red lines; note that $C_{22} = h_0/2$ is the quantity that was actually used to parametrize GR triaxial amplitudes in the code [51]. On the left, the injected signal corresponds to a face-off source ($\cos\iota \approx -1$), making it difficult to constrain the polarization angle ψ; on the right, the injection has similar amplitude but corresponds to an edge-on source ($\cos\iota \approx 0$), making it easy to constrain ψ [modulo $\pi/2$ due to the $2\Delta\psi$ dependence of Eqs. (8.73) and (8.74)]. These plots were produced using the CORNER.PY package [164].

These rotational properties are easily understood by recalling that GW polarizations can be defined in *any* frame, although a given signal might look more or less simple given the choice of frame. Eqs. (8.73) and (8.74) provide the transformation between frames that are coaligned except for a rotation of $\Delta\psi$ around their z-axes. Because waveform predictions, e.g. Eq. (8.71), are made in *specific* frames, it is important to orient the wave frame appropriately when working with a given theory and emission mechanism. However, if the signal parametrization is independent of any theory, e.g. Eq. (8.72), one is free to pick any frame (i.e. any ψ).

The relationship between the different tensor model parameters is reflected in the posterior probability plots of Fig. 8.18. For fixed orientation, both the triaxial (a) and free-tensor (b) analyses accurately determine the amplitude and phase of the injected signal. In panel (b), a_+ and $a_\times$ are constrained to lie within a region consistent with $h_{\rm inj}^2 = a_+^2 + a_\times^2$ and $a_+/a_\times = (1+\cos^2\iota)/(2\cos\iota)$, for an effective injection amplitude

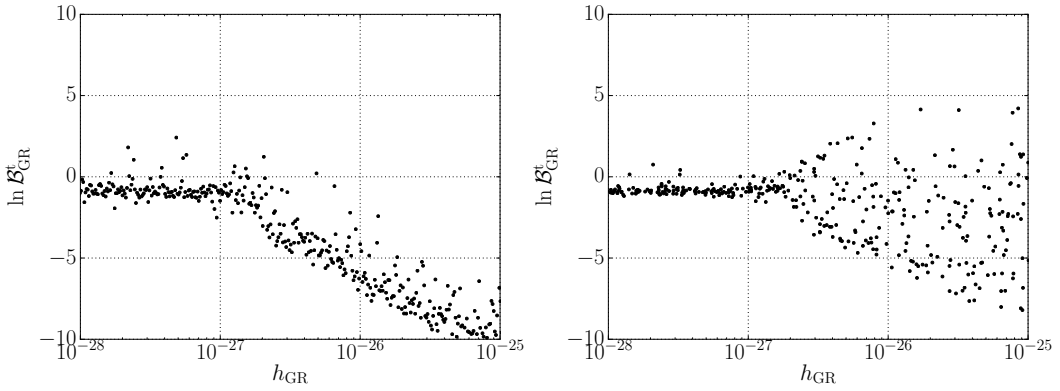

(a) Fixed orientation

(b) Unfixed orientation

Figure 8.20: *Free-tensor vs GR*. Natural logarithm of the Bayes factor comparing $\mathcal{H}_{\mathrm{t}}$ to $\mathcal{H}_{\mathrm{GR}}$, as a function of GR injection amplitude for fixed (left) and unfixed (right) source orientation. On the left, the analysis correctly gives preference to $\mathcal{H}_{\mathrm{GR}}$ for signals above the detection threshold; on the right, however, the analysis is unable to satisfactorily distinguish between $\mathcal{H}_{\mathrm{t}}$ and $\mathcal{H}_{\mathrm{GR}}$, due to the orientation degeneracies discussed in Appendix 8.7.

given by $h^2_{\mathrm{inj}} = h^2_0(1+\cos^2\iota)^2/4 + h^2_0\cos^2\iota$, as in Eq. (8.67). When the orientation is allowed to vary, we observe the expected correlations between the recovered triaxial amplitude and the orientation parameters in panel (c); in this case, ψ and $\cos\iota$ will also become correlated, as better shown in Fig. 8.19.

The degeneracy between ψ and a_+, $a_\times$ is particularly evident in Fig. 8.18d, where the one-dimensional PDF for ψ shows that this parameter cannot be constrained, even for a loud signal. Furthermore, joint posteriors between ψ and a_+ & $a_\times$ confirm that this is due to the degeneracy from Eqs. (8.75) and (8.76), as seen by comparing these two-dimensional PDFs to Fig. 8.17. Physically, this is a consequence of the fact that we are free to orient the polarization frame as we wish.

Because their signal templates are degenerate when ψ and ι are allowed to vary, the distinction between $\mathcal{H}_{\mathrm{GR}}$ and $\mathcal{H}_{\mathrm{t}}$ is not really meaningful for unfixed orientation. This can be seen from the values of $\ln \mathcal{B}^{\mathrm{t}}_{\mathrm{GR}}$ in the cases of known and unknown orientations, as in Fig. 8.20. On the left panel, $\mathcal{H}_{\mathrm{GR}}$ is defined with specific values of ψ and $\cos\iota$ that match the injections; on the right, the $\mathcal{H}_{\mathrm{GR}}$ priors allow ψ and $\cos\iota$ to range over their full ranges, and the injections are performed with random values of both. When the orientation is fixed, $\mathcal{H}_{\mathrm{GR}}$ will always be preferred to $\mathcal{H}_{\mathrm{t}}$ for resolvable signals because of its lower Occam's penalty; however, that is not true for unfixed orientation. Note that, in the strictest sense, the two hypotheses are not

logically equivalent unless their parameter priors are related by the Jacobian of the coordinate transformation between the two parametrizations, Eqs. (8.71, 8.72); only in that case all regions of signal space are treated equally by $\mathcal{H}_{\rm GR}$ and $\mathcal{H}_{\rm t}$. This explains the variation in values of $\ln \mathcal{B}^{\rm t}_{\rm GR}$ on the right of Fig. 8.20.

If one knew the source orientation and one believed that the *only* viable mechanism for producing CWs at the assumed frequency in GR is the triaxial model embodied by Eq. (8.71), then one could include the free-tensor hypothesis and all of its derivatives (i.e. t, st, vt, stv) in the non-GR set $\tilde{M}$, on top of {s, v, sv, GR + s, GR + v, GR + sv}. Doing so would mean treating a tensor-only signal that does not conform to Eq. (8.71) as evidence of a GR violation, rather than of a different emission mechanism within GR. Given the many simplifications intrinsic to the triaxial model, however, having that much confidence in its validity seems unwarranted; hence we choose to not take that approach.

8.8 Appendix: Amplitude priors

Previous CW Bayesian searches targeted to known pulsars have always applied a flat prior on the signal amplitude parameter [29]. This is because flat priors, if wide enough, cause the posterior to be only determined by the likelihood (up to normalization), yielding more conservative upper limits on the signal strength. Furthermore, unlike with priors uniform in the logarithm of the quantity, upper limits derived with flat priors will generally not depend on the limits set by the prior (again, assuming the range allowed extends from zero amplitude to some large value that does not truncate the likelihood).

Upper limits obtained using log-uniform priors (uniform in the logarithm of the quantity) will, generally, be dependent on the range of the prior, although not strongly. For example, consider a one-dimensional problem on some positive parameter x. For simplicity, further assume we have a flat likelihood between $x = 0$ and an upper cutoff at $x = x_{\rm max}$; then, $x_{\rm max}$ will necessarily also be an upper bound for the posterior. Because the likelihood is uniform, below the cutoff the posterior will be determined, up to normalization, by the prior only, i.e. for $x < x_{\rm max}$,

$$p(x \mid \mathbf{B}, \mathcal{H}) \propto p(x \mid \mathcal{H}). \tag{8.77}$$

Now consider a log-uniform prior $p(x \mid \mathcal{H}) \propto \mathrm{d}(\log x) \propto 1/x$, with a lower bound $x_{\rm min}$, such that $0 < x_{\rm min} < x_{\rm max}$. Because such prior is uniform in the $\log x$, this

implies that the 95%-credible upper limit on x will be given by:

$$\begin{aligned}\log x^{95\%} &= \log x_{\min} + 0.95(\log x_{\max} - \log x_{\min}) \\ &= \log\left(x_{\max}^{0.95}/x_{\min}^{0.95-1}\right).\end{aligned} \tag{8.78}$$

Since $x_{\max}$ is set by the likelihood (by construction), if the prior is changed by rescaling $x_{\min}$ by a factor α,

$$x_{\min} \rightarrow x'_{\min} = \alpha x_{\min}, \tag{8.79}$$

then, for a given set of data, the upper limit becomes $x_\alpha^{95\%}$, satisfying:

$$x_\alpha^{95\%}/x^{95\%} = \alpha^{0.05}. \tag{8.80}$$

Thus, the dependence of the upper limit on the range defined by the log-uniform prior is quite weak, as illustrated in Fig. 8.21. This explains why upper limits obtained with a log-uniform prior differ only by a factor of a few from those obtained with a flat one, as seen in Fig. 8.13.

However, the flat priors do not properly represent our ignorance of the scale of the signal amplitude. This problem manifests itself in negative Bayes factors that too quickly favor the noise hypothesis if no loud signal is clearly present, rather than reflecting our expectation that a signal might be hiding under the noise. This can be seen in Fig. 8.22, where we show the distributions of $\ln \mathcal{B}_{\mathrm{N}}^{\mathrm{GR}}$, obtained for several noise-only data instantiations for the Crab pulsar, corresponding to flat and log-uniform priors in the GR amplitude parameter, h_0; a uniform prior results in lower values of $\ln \mathcal{B}_{\mathrm{N}}^{\mathrm{GR}}$ that strongly favor $\mathcal{H}_{\mathrm{N}}$. This behavior is not specific to the GR model.

For most of our analysis, we choose to apply priors uniform in the logarithm of all amplitude quantities. However, for the sake of consistency with previous searches and in order to make our limits more conservative, we also present upper limits produced using flat amplitude priors, as shown in Fig. 8.13.

8.9 Appendix: Numerical error

The fractional numerical error in the computation of the natural logarithm of the evidence by nested sampling is usually estimated by:

$$\delta\left[\ln P\left(\mathbf{B} \mid \mathcal{H}\right)\right] \sim \sqrt{H/N_{\mathrm{live}}}, \tag{8.81}$$

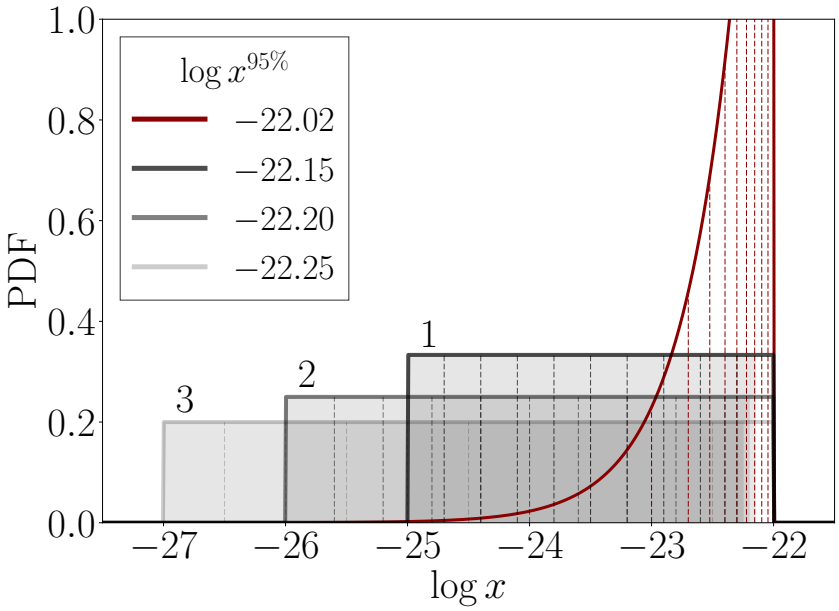

Figure 8.21: *Log-uniform prior and upper limits*. For a 1D random variable x, we show the probability densities corresponding to a uniform likelihood with upper cutoff $\log x_{\rm max} = -22$ (red) and log-uniform priors with different lower cutoffs ($\log x_{\rm min} = -25$ for box 1, $\log x_{\rm min} = -26$ for box 2 and $\log x_{\rm min} = -27$ for box 3). Vertical dashed lines mark areas of equal probability mass for each distribution. The combined effect of the likelihood and each of the prior distributions is to produce 95%-credible upper limits on x with values shown in the legend. The value obtained using only the likelihood corresponds to that obtained with a uniform prior with a broad enough range. As expected from Eq. (8.80), the upper limit is not very sensitive to the lower bound set by the prior.

where $N_{\rm live}$ is the number of of live points and H is the information gained in the analysis:

$$H \equiv \int_{\Theta} p(\vec{\theta} \mid \mathbf{B}, \mathcal{H}) \, \ln \frac{p(\vec{\theta} \mid \mathbf{B}, \mathcal{H})}{p(\vec{\theta} \mid \mathcal{H})} \mathrm{d}\vec{\theta}, \tag{8.82}$$

a quantity that is easy to estimate from the output of the nested sampling code [19, 165].

An example of the actual statistical error as function of SNR is presented in Figs. 8.23 and 8.24, where the injected GR signal amplitude serves as proxy for ρ (for fixed PSD). From these plots it becomes apparent that, although the actual error might exceed the estimator of Eq. (8.81), its absolute magnitude is quite small and should not affect our results. In any case, Eq. (8.81) indicates that any level of accuracy may be achieved by increasing the number of live points (at the cost of increased computational burden). For more details on the numerical error of the

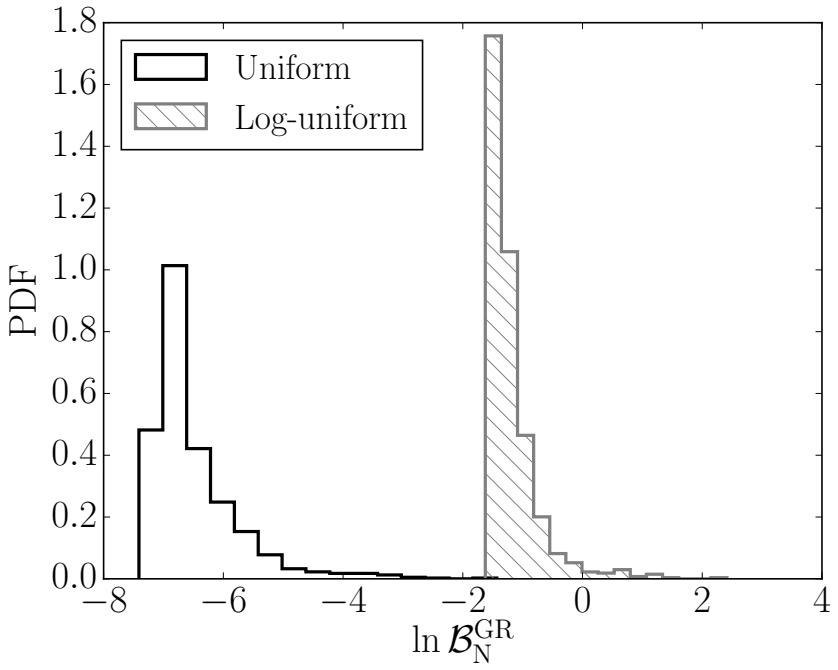

Figure 8.22: *Log-uniform vs flat amplitude priors*. The logarithm of the GR vs noise Bayes factor is computed for 1000 instantiations of Crab pulsar noise. For the GR amplitude h_0, we apply priors uniform in the quantity (black) and uniform in the logarithm of the quantity (hatched gray). The flat prior causes one to more strongly favor the noise model, due to a larger implicit Occam's penalty.

nested sampling algorithm in LALInference, we refer the reader to Sec. IVB of [11].

8.10 Appendix: Upper-limit ratios

When comparing upper limits for the different modes, as in Fig. 8.16, two scalings become apparent: first, the +, ×, x, and y upper limits are, on average, more stringent than those for the scalar polarization by a factor of ∼1.8; second, the upper limits on h_{t} (h_{t}) are a factor of ∼1.3 larger than those on the individual + and × (x and y) amplitudes.

The scaling between the scalar upper limit and those for the other individual strain amplitudes can be accounted for by the decreased sensitivity of quadrupolar GW detectors to scalar waves. For a single instrument (that is, not a network), this can be appreciated visually from Fig. 8.2, by noting that for most sky locations the magnitude of the scalar response is considerably less than for the other modes.

To properly evaluate the effect of the detector geometry on the analysis, however,

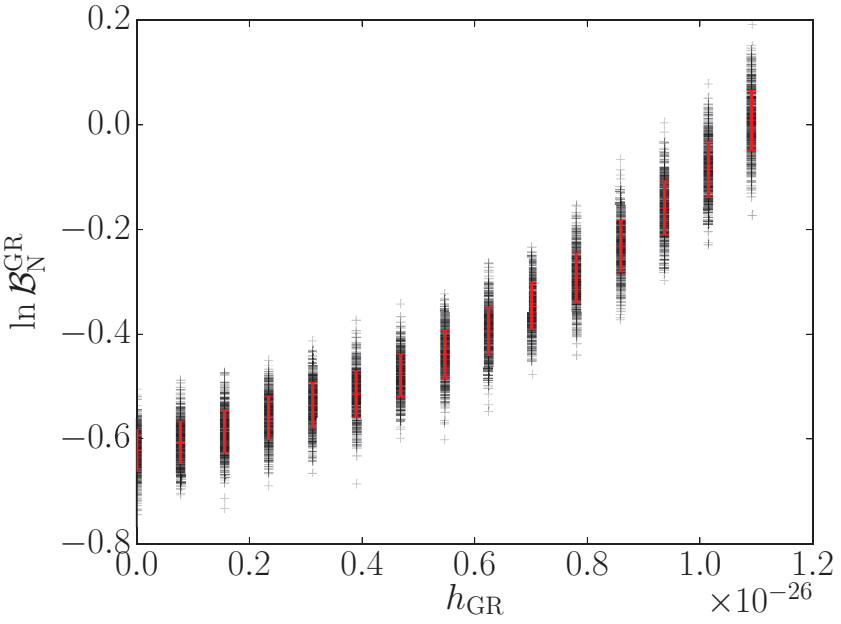

Figure 8.23: *Numerical error in Bayes factor computation*. The logarithm of the GR vs noise Bayes factor is computed 500 times for different values of injected GR signal amplitude. The noise realization is not varied between computations with the same injection strength, only the seed for the random number generator used by the nested sampling algorithm. The red bars mark one standard deviation around the mean.

it is necessary to look at the relative SNRs of unit-amplitude scalar, vector and tensor GWs from a given source, as they are received by the detector network under consideration (H1, L1, V1) after some fixed observation time. Assuming all detectors have comparable noise levels, the network SNR can be proxied by the root-mean-square (RMS) amplitude of the effective network antenna patterns, defined by

$$F^{\mathrm{rms}}_{p,\,\mathrm{net}} = \sqrt{\frac{1}{T}\int_0^T \sum_d F^2_{p,d}(t)}, \tag{8.83}$$

for each polarization p, some long observation time T, and where the sum is over detectors d. [Here we have fixed the source and detector parameters so that the F_p's of Eqs. (8.1–8.5) are now just simple functions of time.] We may then compute this for all five polarizations and for multiple sources to obtain a sky-average of the ratio of the scalar RMS antenna pattern to those of the other polarizations. We find this ratio to be roughly ∼0.55 for all polarizations, in agreement with Fig. 8.16, since we should expect

$$\left\langle \frac{F^{\mathrm{rms}}_{\mathrm{s,\,net}}}{F^{\mathrm{rms}}_{p,\,\mathrm{net}}} \right\rangle \sim \left\langle \frac{h^{95\%}_p}{h^{95\%}_{\mathrm{s}}} \right\rangle, \tag{8.84}$$

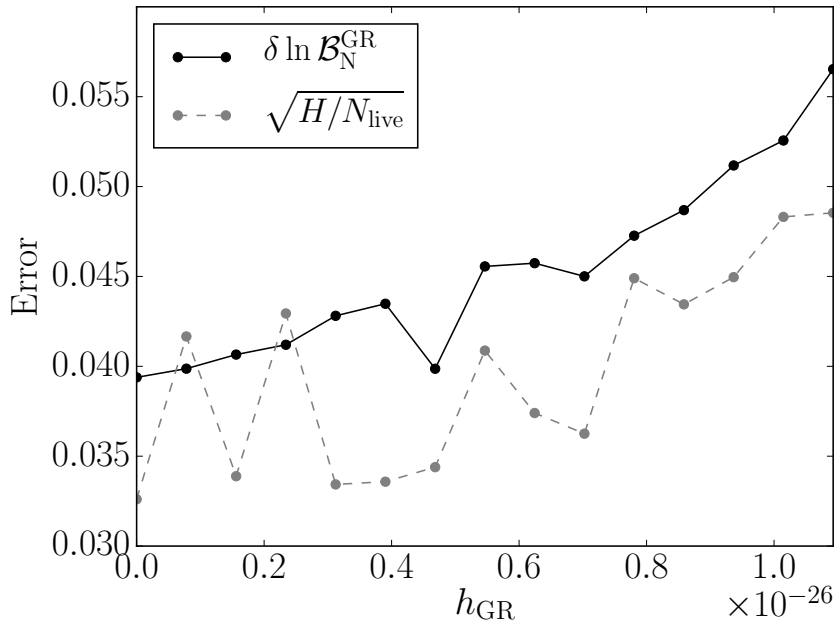

Figure 8.24: *Observed error and prediction*. Error in the computation of the logarithm of the GR vs noise Bayes factor as a function of injected GR signal amplitude. The solid, black curve corresponds to measured standard deviations from the computation of $\ln \mathcal{B}^{\rm GR}_{\rm N}$ 500 times per injection strength (red bars in Fig. 8.23). The dashed, gray curve shows the theoretical prediction for the error in the logarithm of the evidence, Eq. (8.81).

where the average $\langle \cdot \rangle$ is taken over multiple sources distributed across the sky.

The relation between the $h_{\rm t}$ ($h_{\rm v}$) upper limits and those for their component amplitudes, + and × (x and y), can be easily understood by noting that, if using flat priors and in the absence of signal, the marginalized posteriors for each of the component amplitudes (h_+, $h_\times$, $h_{\rm x}$, $h_{\rm y}$) will roughly be described by a one-sided normal distribution. Consequently, it can be shown that posterior for the square-root of the sum of the squares of two of these quantities will be given by a *chi distribution* with two degrees of freedom. Considering the definitions of Eqs. (8.67, 8.68), it is straightforward to show (numerically or analytically) that this explains the observed factor of ∼1.3 difference between $h_{\rm t}^{95\%}$ ($h_{\rm v}^{95\%}$) and $h_+^{95\%}$ or $h_\times^{95\%}$ ($h_{\rm x}^{95\%}$ or $h_{\rm y}^{95\%}$).

Chapter 9

FIRST SEARCH FOR NONTENSORIAL GRAVITATIONAL WAVES FROM KNOWN PULSARS

B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **120**, 031104 (2018), M.I. carried out the analysis and wrote the manuscript, arXiv:1709.09203 [gr-qc] .

9.1 Introduction

The first gravitational waves (GWs) detected by the Advanced Laser Interferometer Gravitational-wave Observatory (aLIGO) and Virgo have already been used to place some of the most stringent constraints on deviations from the general theory of relativity (GR) in the highly-dynamical and strong-field regimes of gravity [1–4]. However, even though some partial progress has been made with the observation of GW170814 [5, 116] and in spite of the wealth of new information provided by GW170817 [6, 123], it has not yet been possible to unambiguously confirm GR's prediction that the associated metric perturbations are of a tensor nature (helicity ± 2), rather than vector (helicity ± 1), or scalar (helicity 0) [82]. This is unfortunate, since the presence of nontensorial modes is a key prediction of many extensions to GR [83, 84, 97, 98, 166]. Most importantly, the detection of a scalar or vector component, no matter how small, would automatically point to physics beyond Einstein's theory [83, 84].

In order to experimentally study GW polarizations directly, one needs a local measurement of their geometric effect (i.e. which directions are streched and squeezed) that breaks degeneracies between the five distinguishable (to differential-arm instruments) modes supported by a generic metric theory of gravity [97, 98]. For transient waves like those detected so far, this cannot be fully achieved with the LIGO-Virgo network, as at least five noncooriented differential-arm antennas are required to break *all* such degeneracies [84, 85]. Constraints on the magnitude of non-GR polarizations inferred from indirect measurements, like the rate of orbital decay of binary pulsars, are only meaningful in the context of specific theories (see e.g. [86, 87], or [88, 89] for reviews).

Theory-independent polarization measurements could instead be carried out with current detectors in the presence of signals sufficiently long to probe the detector antenna patterns, which are themselves polarization-sensitive [18, 99, 107, 109]. Such is the case, for instance, for the continuous, almost-monochromatic waves expected from spinning neutron stars with an asymmetric moment of inertia [48]. Known galactic pulsars are one of the main candidates for searches for such signals in data from ground-based detectors, and analyses targeting them have already achieved sensitivities that are comparable to, or even surpass, their canonical spin-down limit (i.e. the strain that would be produced if the observed slowdown in the pulsar's rotation was completely due to gravitational radiation) [29].

However, all previous targeted searches have been, by design, restricted to tensorial gravitational polarizations *only*. This leaves open the possibility that, due to a departure from GR, the neutron stars targeted in previous searches may indeed be emitting strong continuous waves with nontensorial content, in spite of the null results of standard searches.

In this paper, we present results from a search for continuous GWs in aLIGO data that makes no assumptions about how the gravitational field transforms under spatial rotations and is thus sensitive to any of the five measurable polarizations allowed by a generic metric theory of gravity. We targeted 200 known pulsars using data from aLIGO's first observation run (O1) and assumed GW emission at twice the rotational frequency of the source.

Our data provide no evidence for the emission of gravitational signals of tensorial or nontensorial polarization from any of the pulsars targeted. For sources in the most sensitive band of our detectors, we constrain the strain of the scalar and vector modes to be below 1.5×10^{-26} at 95% credibility. These are the first direct upper limits for scalar and vector strain ever published and may be used to constrain beyond-GR theories of gravity.

9.2 Analysis

We search aLIGO O1 data from the Hanford (H1) and Livingston (L1) detectors for continuous waves of any polarization (tensor, scalar or vector) by applying the Bayesian time-domain method of [57], generalized to non-GR modes as described in [18] and summarized below. Our analysis follows closely that of [29] and uses the exact same interferometric data.

Calibrated detector data are heterodyned and filtered using the timing solutions for each pulsar obtained from electromagnetic observations. The maximum calibration uncertainties estimated over the whole run give a limit on the combined H1 and L1 amplitude uncertainties of 14%—this is the conservative level of uncertainty on the strain upper limits [29, 167].

The data streams start on 2015 Sep 11 at 01:25:03 UTC for H1 and 18:29:03 UTC for L1 and finish on 2016 Jan 19 at 17:07:59 UTC at both sites. The pulsar timing solutions used are also the same as in [29] and were obtained from the 42-ft telescope and Lovell telescope at Jodrell Bank (UK), the 26-m telescope at Hartebeesthoek (South Africa), the Parkes radio telescope (Australia), the Nancay Decimetric Radio Telescope (France), the Arecibo Observatory (Puerto Rico) and the Fermi Large Area Telescope (LAT).

As described in detail in [18], we construct a Bayesian hypothesis that captures signals of any polarization content (our *any-signal* hypothesis, $\mathcal{H}_{\rm S}$) by combining the sub-hypotheses corresponding to the signal being composed of tensor, vector, scalar modes, or any combination thereof. Each of these sub-hypotheses corresponds to a different signal model; in particular, the least restrictive template includes contributions from all polarizations and can be written as:

$$h(t) = \sum_{p} F_p(t; \alpha, \delta, \psi) h_p(t), \tag{9.1}$$

where the sum is over the five independent polarizations: plus (+), cross (×), vector-x (x), vector-y (y) and scalar (s) [98]. The two scalar modes in the most common strain-basis, breathing and longitudinal, are degenerate for networks of quadrupolar antennas [84], so we do not make a distinction between them.

Each term in Eq. (9.1) is the product of an antenna pattern function F_p and an intrinsic strain function h_p. We define the different polarizations in a wave-frame such that the z-axis points in the direction of propagation, x lies in the plane of the sky along the line of nodes (here defined to be the intersection of the equatorial plane of the source with the plane of the sky) and y completes the right-handed system, such that the polarization angle ψ is the angle between the y-axis and the projection of the celestial North onto the plane of the sky (see e.g. [168]). We can thus write the F_p's as implicit functions of the source's right ascension α, declination δ and polarization ψ. (For the sources targeted here, α and δ are always known to high accuracy, while ψ is usually unknown.) The antenna patterns acquire their time

dependence from the sidereal rotation of the Earth; explicit expressions for the F_p's are given in [18, 107–110].

For a continuous wave, the polarizations take the simple form:

$$h_p(t) = a_p \cos(\phi(t) + \phi_p), \tag{9.2}$$

where a_p is a time-independent strain amplitude, $\phi(t)$ is the intrinsic phase evolution and ϕ_p a phase offset for each polarization. The nature of these three quantities depends on the specifics of the underlying theory of gravity and the associated emission mechanism (for different emission mechanisms within GR, see e.g. [146, 157, 158]). While we treat a_p and ϕ_p as free parameters, we take $\phi(t)$ to be the same as in the traditional GR analysis [29]:

$$\phi(t) = 2\pi \sum_{j=0}^{N} \frac{\partial_t^{(j)} f_0}{(j+1)!} \left[t - T_0 + \delta t(t)\right]^{(j+1)}, \tag{9.3}$$

where $\partial_t^{(j)} f_0$ is the j^{th} time derivative of the GW frequency measured at the fiducial time T_0; $\delta t(t)$ is the time delay from the observatory to the solar system barycenter (including the known Rømer, Shapiro and Einstein delays) and can also include binary system corrections to transform to the time coordinate to a frame approximately inertial with respect to the source; N is the order of the series expansion (1 or 2 for most sources).

The GW frequency f is related to the rotational frequency of the source f_{rot}, which is in turn known from electromagnetic observations. Although arbitrary theories of gravity and emission mechanisms may predict gravitational emission at any multiple of the rotational frequency, here we assume $f = 2f_{\text{rot}}$, in accordance with the most favored emission model in GR [48]. This restriction arises from practical considerations affecting our specific implementation and will be relaxd in future studies.

For convenience, we define *effective strain amplitudes* for tensor, vector and scalar modes respectively by

$$h_{\text{t}} \equiv \sqrt{a_+^2 + a_\times^2}, \tag{9.4}$$

$$h_{\text{v}} \equiv \sqrt{a_{\text{x}}^2 + a_{\text{y}}^2}, \tag{9.5}$$

$$h_{\text{s}} \equiv a_{\text{s}}, \tag{9.6}$$

in terms of the intrinsic a_p amplitudes of Eq. (9.2). These quantities may serve as proxy for the total power in each polarization group.

Table 9.1: Existing orientation information for pulsars in our band, obtained from observations of the pulsar wind nebulae (see Table 3 in [49], and [149, 169] for measurement details).

	ι	ψ
J0534+2200	$62°.2 \pm 1°.9$	$35°.2 \pm 1°.5$
J0537–6910	$92°.8 \pm 0°.9$	$41°.0 \pm 2°.2$
J0835–4510	$63°.6 \pm 0°.6$	$40°.6 \pm 0°.1$
J1833–1034	$85°.4 \pm 0°.3$	$45° \pm 1°$
J1952+3252	N/A	$-11°.5 \pm 8°.6$

One may recover the GR hypothesis considered in previous analysis by setting:

$$a_+ = h_0(1 + \cos^2 \iota)/2\,,\ \phi_+ = \phi_0, \tag{9.7}$$

$$a_\times = h_0 \cos \iota\,,\ \phi_\times = \phi_0 - \pi/2, \tag{9.8}$$

$$a_\mathrm{x} = a_\mathrm{y} = a_\mathrm{s} = 0, \tag{9.9}$$

where ι is the inclination (angle between the line of sight and the spin axis of the source) and h_0, ϕ_0 are free parameters. (As with ψ, ι is unknown for most pulsars.) This corresponds to the standard triaxial-star emission mechanism (see e.g. [54]). We use this parameterization only when we wish to incorporate known orientation information as explained below; otherwise, we parametrize the tensor polarizations directly in terms of a_+, $a_\times$, ϕ_+ and $\phi_\times$.

Templates of the form of Eq. (9.1), together with appropriate priors, allow us to compute Bayes factors (marginalized likelihood ratios) for the presence of signals in the data vs Gaussian noise. We do this using an extension of the nested sampling implementation presented in [28] (see [18] for details specific to the non-GR polarizations). The Bayes factors corresponding to each signal model may be combined into the odds O^S_N that the data contain a continuous signal of any polarization vs Gaussian noise:

$$O^\mathrm{S}_\mathrm{N} = P(\mathcal{H}_\mathrm{S} \mid \mathbf{B})/P(\mathcal{H}_\mathrm{N} \mid \mathbf{B})\,, \tag{9.10}$$

i.e. the ratio of the posteriors probabilities that the data $\mathbf{B}$ contain a signal of any polarizations ($\mathcal{H}_\mathrm{S}$) vs just Gaussian noise ($\mathcal{H}_\mathrm{N}$). We compute these odds by setting model priors such that $P(\mathcal{H}_\mathrm{S}) = P(\mathcal{H}_\mathrm{N})$; then, by Bayes' theorem, $O^\mathrm{S}_\mathrm{N} = \mathcal{B}^\mathrm{S}_\mathrm{N}$, with the Bayes factor

$$\mathcal{B}^\mathrm{S}_\mathrm{N} \equiv P(\mathbf{B} \mid \mathcal{H}_\mathrm{S})/P(\mathbf{B} \mid \mathcal{H}_\mathrm{N}). \tag{9.11}$$

Built into the astrophysical signal hypothesis, $\mathcal{H}_{\rm S}$, is the requirement of coherence across detectors, which must be satisfied by a real GW signal. In order to make the analysis more robust against non-Gaussian instrumental features in the data, we also define an *instrumental feature* hypothesis, $\mathcal{H}_{\rm I}$, that identifies non-Gaussian noise artifacts by their lack of coherence across detectors [27, 29]. In particular, we define $\mathcal{H}_{\rm I}$ to capture Gaussian noise *or* a detector-incoherent signal (i.e. a feature that mimics an astrophysical signal in a single instrument but is not recovered consistently across the network) in each detector [18]. We may then compare this to $\mathcal{H}_{\rm S}$ by means of the odds $\mathcal{O}^{\rm S}_{\rm I}$. For D detectors, this is given by:

$$\log \mathcal{O}^{\rm S}_{\rm I} = \log \mathcal{B}^{\rm S}_{\rm N} - \sum_{d=1}^{D} \log\left(\mathcal{B}^{{\rm S}_d}_{{\rm N}_d} + 1\right), \tag{9.12}$$

where $\mathcal{B}^{{\rm S}_d}_{{\rm N}_d}$ is the signal vs noise Bayes factor computed only from data from the $d^{\rm th}$ detector. This choice implicitly assigns prior weight to the models such that $P(\mathcal{H}_{\rm S}) = P(\mathcal{H}_{\rm I}) \times 0.5^D$ [18]. For an in depth analysis of the behavior of the different Bayesian hypotheses considered here, in the presence absence of simulated signals of all polarizations, we again refer the reader to the paper methods [18].

We compute likelihoods by taking source location, frequency and frequency derivatives as known quantities. In computing Bayes factors, we employ priors uniform in the logarithm of amplitude parameters (h_0 or h_p's), since these are the least informative priors for scaling coefficients [163]; we bound these amplitudes to the 10^{-28}–10^{-24} range [1]. On the other hand, flat amplitude priors are used to compute upper limits, in order to facilitate comparison with published GR results in [29]. In all cases, flat priors are placed over all phase offsets (ϕ_0 and all the ϕ_p's).

For those few cases in which some orientation information exists (see Table 9.1 in Supplementary Material), we analyze the data a second time using the triaxial parametrization of tensor modes, Eqs. (9.7) and (9.8), taking that information into account by marginalizing over ranges of $\cos\iota$ and ψ in agreement with measurement uncertainties. Following previous work [29], we only consider orientation constraints obtained from pulsar wind nebulae. However, pulsar orientations can also be inferred from other measurements, especially if the object is in a binary (e.g. [170–172]). We will consider incorporating such constraints in future searches.

[1] The specific range chosen for the amplitude priors has little effect on our results, as explained in Appendix B of [18]

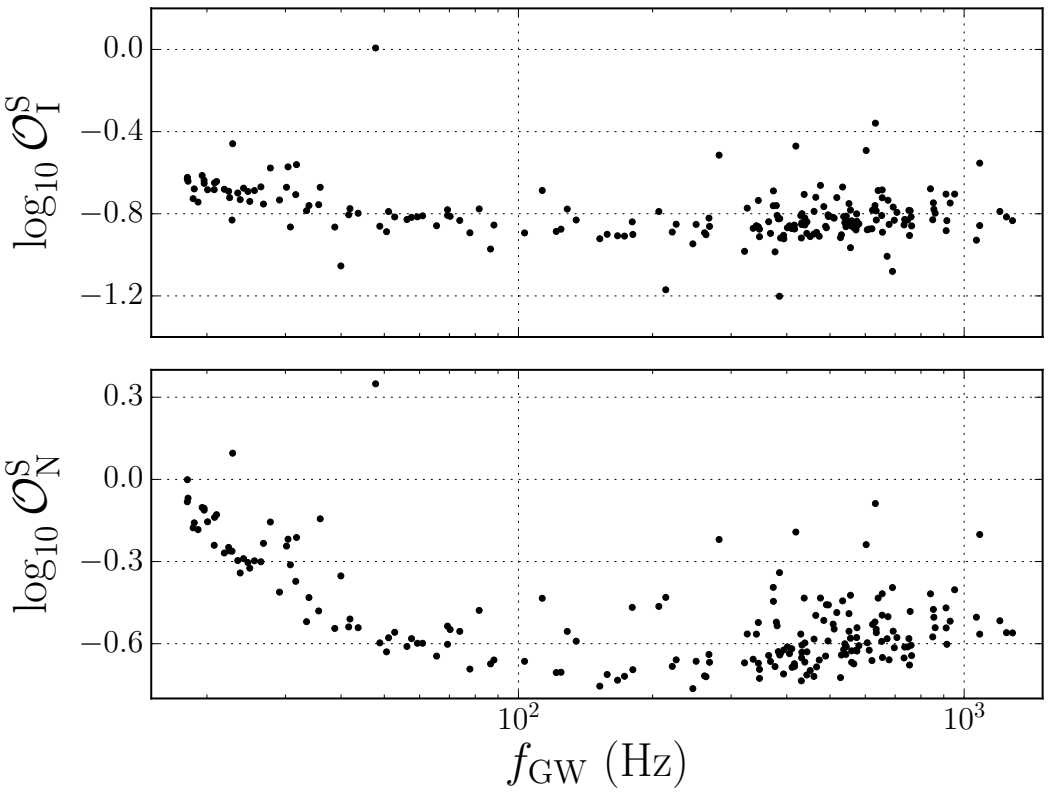

Figure 9.1: *Log-odds vs GW frequency*. Log-odds comparing the any-signal hypothesis to the instrumental (top) and Gaussian noise (bottom) hypotheses, as a function of assumed GW frequency, $f = 2f_{\mathrm{rot}}$, for each pulsar. Looking at the top plot for $\log_{10} O^{\mathrm{S}}_{\mathrm{I}}$, notice that the instrumental noise hypothesis is clearly favored for all pulsars except one, for which the analysis is inconclusive. (This is J1932+17, the same non-significant outlier identified in [29].) These results were obtained without incorporating any information on the source orientation and are tabulated in Table B.2 in the Supplementary Material. Expressions for both odds are given in Eq. (9.10) and Eq. (9.12).

9.3 Results

We find no evidence of continuous-wave signals of any polarization, tensorial or otherwise, from any of the 200 pulsars analyzed. Odds and 95%-credible upper limits are summarized in the Supplementary Material: Table B.1, for pulsars with measured orientations (using the triaxial parameterization of tensor modes) and in Table B.2, for all pulsars without incorporating any orientation information (using the unconstrained parameterization of tensor modes). Odds values are reported with an error of 5% at 90% confidence; errors on the upper limits due to the use of finite samples in estimating posterior probability distributions are at most 10% at 90% confidence, which is slightly less than the 15% error expected from calibration

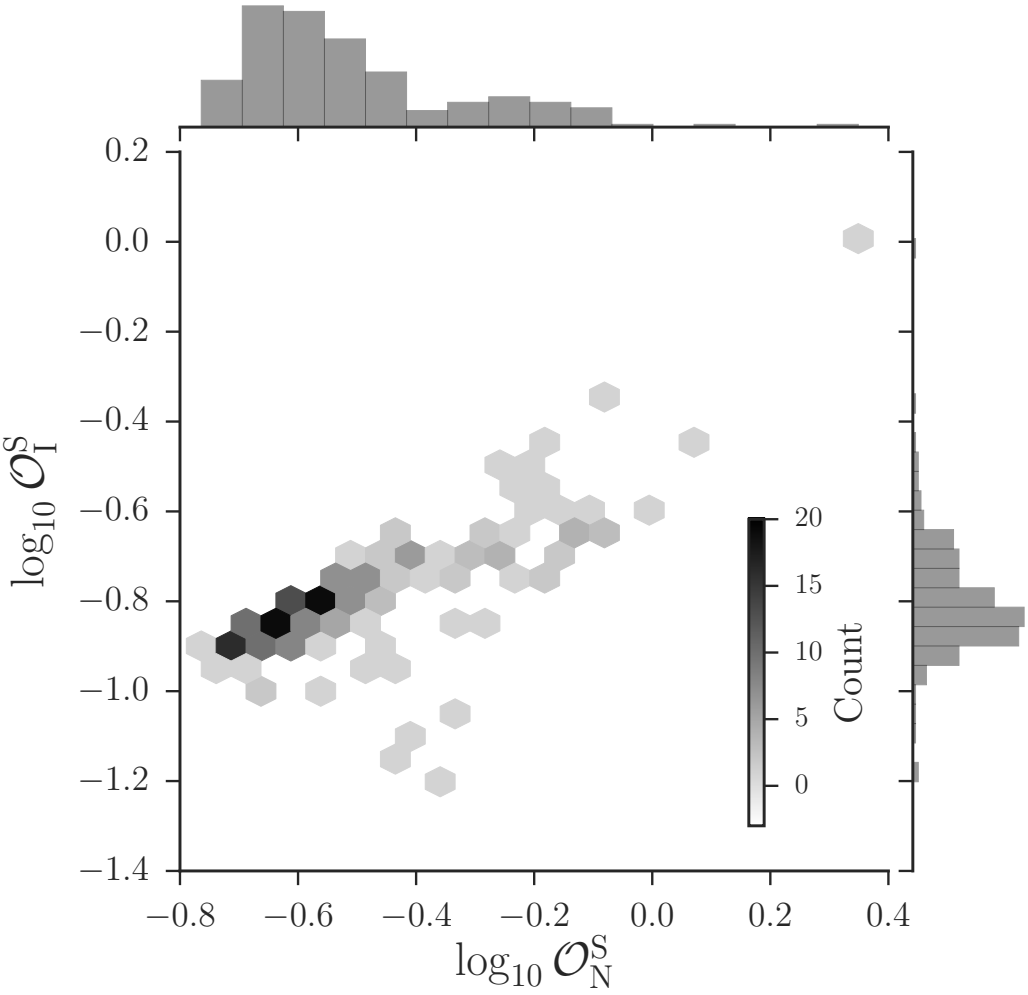

Figure 9.2: *Log-odds distributions*. Distributions of log-odds comparing the any-signal hypothesis to the instrumental (ordinate axis, right) and Gaussian noise (abscissa axis, top) hypotheses for all pulsars. This plot contains the same information as Fig. 9.1 and displays the same non-significant outlier. These results were obtained without incorporating any information on the source orientation and are tabulated in Table B.2 in the Supplementary Material. Expressions for both odds in this plot are given in Eq. (9.10) and Eq. (9.12). We underscore that, although this plot looks similar to Fig. 2 in [29], the signal hypothesis here incorporates scalar, vector and tensor modes, in all their combinations.

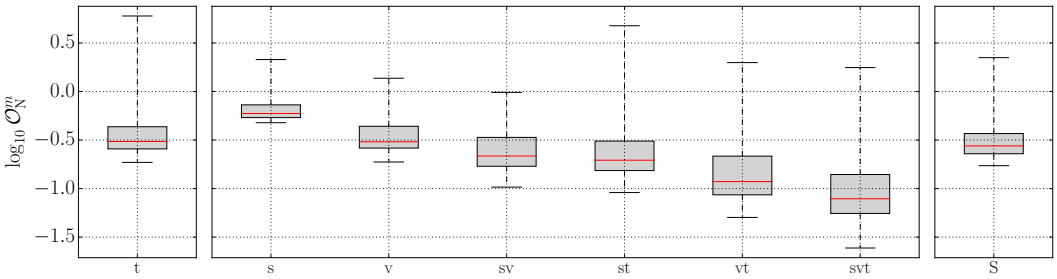

Figure 9.3: *Sub-hypothesis odds*. Box plots for the distribution of the signal vs noise log-odds for each of the sub-hypotheses considered, for all of the pulsars analyzed. The sub-hypotheses are: (st), vector-tensor (tv), scalar-vector-tensor tensor-only (t), scalar-only (s), vector-only (v), scalar-vector (sv), scalar-tensor (st), vector-tensor (vt), and scalar-vector-tensor (stv); these are all combined into the signal hypothesis (S). The quantity represented is $\log_{10} \mathcal{B}^m_{\rm N}$, which is the same as $\log_{10} \mathcal{O}^m_{\rm N}$ if neither $\mathcal{H}_m$ nor $\mathcal{H}_{\rm N}$ are favored *a priori* (hence the label on the ordinate axis). The horizontal red line marks the median of the distribution, while each gray box extends from the lower to upper quartile, and the whiskers mark the full range of the distribution of $\log_{10} \mathcal{O}^m_{\rm N}$ for the 200 pulsars analyzed. These results were produced without incorporating any information on the source orientation, and are tabulated in Table B.2 in the Supplementary Material.

uncertainties.

The main quantity of interest is $\log_{10} \mathcal{O}^{\rm S}_{\rm I}$, defined in Eq. (9.12), since it encodes the probability that the data contain a signal vs just instrumental noise (Gaussian or otherwise). This quantity, together with the log-odds for signal vs Gaussian noise, is presented as a function of assumed GW frequency for each pulsar in Fig. 9.1, and histogrammed in Fig. 9.2. Importantly, the outliers in Fig. 9.1 lose significance once $\log_{10} \mathcal{O}^{\rm S}_{\rm I}$ is taken into account; indeed, Figs. 9.1 and 9.2 reveal the usefulness of $\log_{10} \mathcal{O}^{\rm S}_{\rm I}$ in increasing the robustness of the search against non-Gaussian instrumental artifacts.

Based on the intrinsic probabilistic meaning of $\log_{10} \mathcal{O}^{\rm S}_{\rm I}$ in terms of betting odds, it is standard to demand at least $\log_{10} \mathcal{O}^{\rm S}_{\rm I} > 1$ to conclude that the signal model is favored (see e.g. the table in Sec. 3.2 of [173], or Jeffrey's original criteria in [174] or [175]). Since none of the odds obtained meet this criterion, we conclude that there is no evidence for signals from any of the pulsars targeted. In most cases, $\log_{10} \mathcal{O}^{\rm S}_{\rm I} < 0$ and the noise model is clearly favored; the single exception is J1932+17, for which $\log_{10} \mathcal{O}^{\rm S}_{\rm I} \sim 0$, so that we can make no conclusive statement about which hypothesis is preferred. (The presence of this non-significant outlier is to be expected, as it was already identified in [29].)

The distribution of the odds corresponding to the subhypotheses making up $\mathcal{H}_S$ is summarized in the box plots of Fig. 9.3. These correspond to tensor-only (t), scalar-only (s), vector-only (v), scalar-vector (sv), scalar-tensor (st), vector-tensor (vt), and scalar-vector-tensor (stv) models. The mean of these distributions decreases with the number of degrees of freedom in the model, which is to be expected from the associated Occam penalties [18]. The right-most panel in Fig. 9.3 shows the distribution of $\log_{10} \mathcal{O}^S_N$, which results from the combination of all the other odds; this is the same quantity histogrammed on the abscissa of Fig. 9.2.

In the absence of any discernible signals, we produce upper limits for the magnitude of scalar, vector and tensor polarizations, with a 95% credibility. As usual in Bayesian analyses, upper limits are obtained by integrating posterior probability distributions for the relevant parameters up to the desired credibility (see e.g. [18]). Using the effective amplitude definitions of Eqs. (9.4)–(9.6), these quantities are presented in Fig. 9.4 as a function of assumed GW frequency and in explicitly in the Supplementary Material. The plotted upper limits are computed under the assumption of a signal model that includes all five independent polarizations ($\mathcal{H}_{svt}$); limits obtained assuming other signal models may be found online in [2]. Previous work has demonstrated that the presence or absence of a GR component does not affect the non-GR upper limits (Fig. 13 in [18]).

As expected, the upper limits presented here are comparable in magnitude to the upper limits on the GR strain obtained by the traditional searches [29]. However, constraints on the scalar amplitude are, on average, around 20% less stringent than those on the vector or tensor amplitudes. This is a consequence of the fact that, for most source locations in the sky, the LIGO detectors are intrinsically less sensitive to continuous waves of scalar polarization [18].

Technically, traditional all-sky searches for continuous GWs are also sensitive to nontensorial modes, because they are generally designed to look for any signal of sidereal and half-sidereal periodicities in the data, without assuming knowledge of phase evolution or source sky-location [43, 45, 46, 176, 177]. However, as can be seen by comparing the magnitude of all-sky upper limits (e.g. Fig. 9 in [43]) to those in shown here in Fig. 9.4, the sensitivity of these searches would be substantially poorer than that of a targeted search like this one—if only because they are not targeted to a specific source. This is especially true if the search is optimized for a given signal polarization (e.g. circular combination of plus and cross).

[2] https://dcc.ligo.org/LIGO-P1700009/public

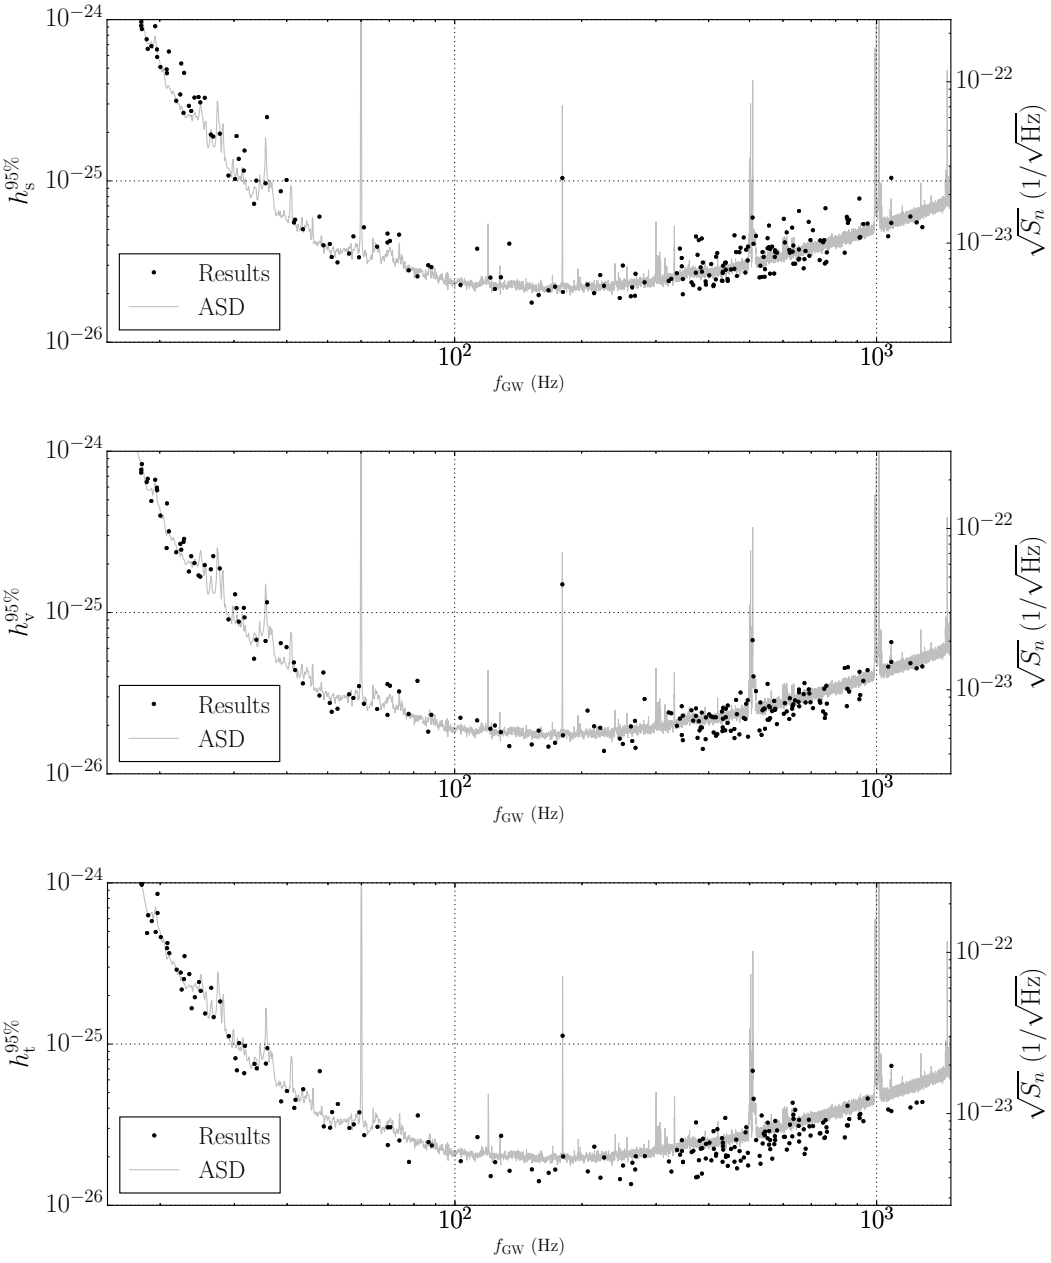

Figure 9.4: *Non-GR upper limits vs GW frequency*. Circles mark the 95%-credible upper limit on the scalar, $h_{\rm s}^{95\%}$ (top), and the effective vector, $h_{\rm v}^{95\%}$ (middle), and tensor $h_{\rm t}^{95\%}$ (bottom) strain amplitudes as a function of assumed GW frequency for each of the 200 pulsars in our set. The upper limits are obtained assuming a signal model including all five independent polarizations ($\mathcal{H}_{\rm stv}$), and incorporating no information on the orientation of the source (Table B.2 in Supplementary Material). The effective amplitude spectral density (ASD) of the detector noise is also displayed for reference; this is the harmonic mean of the H1 and L1 spectra; the scaling is obtained from linear regression to the upper limits.

9.4 Conclusion.

We have presented the results of the first direct search for nontensorial gravitational waves. This is also the first search for GWs targeted at known pulsars that is sensitive to any of the five measurable polarizations of the gravitational perturbation allowed by a generic metric theory of gravity. From the analysis of O1 data from both aLIGO observatories, we have found no evidence of signals from any of the 200 pulsars targeted.

In the absence of a clear signal, we have produced the first direct upper limits for scalar and vector strains (Fig. 9.4, and tables in the Supplementary Material). The values of the 95%-credible upper limits are comparable in magnitude to previously-published GR constraints, reaching $h \sim 1.5 \times 10^{-26}$ for pulsars whose frequency is in the most sensitive band of our instruments.

Our results have been obtained in a theory-independent fashion. However, our upper limits on nontensorial strain can be translated into model-dependent constraints by picking a specific alternative theory and emission mechanism. To do so, one should use the upper limits produced under the assumption of a signal model that incorporates the polarizations matching those allowed by the theory one wishes to constrain; these may not necessarily be those in Fig. 9.4 (e.g. for limits on a scalar-tensor theory, one needs upper limits from $\mathcal{H}_{\rm st}$). However, this also requires non-trivial knowledge of the dynamics of spinning neutron stars under the theory of interest.

While it is conventional to compare the sensitivity of continuous wave searches to the canonical spin-down limit for each pulsar, it is not possible to do so here without committing to a specific theory of gravity. This is because doing so would require specific knowledge of how each polarization contributes to the effective GW stress-energy, how matter couples to the gravitational field, how the waves propagate (dispersion and dissipation), and what the angular dependence of the emission pattern is. However, analogues of the canonical spin-down limit for specific theories may be obtained from the results presented here by using the strain upper limits obtained assuming the sub-hypotheses with polarizations corresponding to that theory, as mentioned above.

We have demonstrated the robustness of searches for generalized polarization states (tensor, vector, or scalar) in GWs from spinning neutron stars. Furthermore, even in the absence of a detection, we were able to obtain novel constraints on the strain

amplitude of nontensorial polarizations. In the future, once a signal is detected, similar methods will allow us to characterize the gravitational polarization content and, in so doing, perform novel tests of general relativity. Although this search assumed a GW frequency of twice the rotational frequency of the source, this restriction will be relaxed in future analyses.

Acknowledgments

The authors gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as 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. The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Foundation for Fundamental Research on Matter supported by the Netherlands Organisation for Scientific Research, for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Council of Scientific and Industrial Research of India, Department of Science and Technology, India, Science & Engineering Research Board (SERB), India, Ministry of Human Resource Development, India, the Spanish Ministerio de Economía y Competitividad, the Conselleria d'Economia i Competitivitat and Conselleria d'Educació, Cultura i Universitats of the Govern de les Illes Balears, the National Science Centre of Poland, the European Commission, the Royal Society, the Scottish Funding Council, the Scottish Universities Physics Alliance, the Hungarian Scientific Research Fund (OTKA), the Lyon Institute of Origins (LIO), the National Research Foundation of Korea, Industry Canada and the Province of Ontario through the Ministry of Economic Development and Innovation, the Natural Science and Engineering Research Council Canada, Canadian Institute for Advanced Research, the Brazilian Ministry of Science, Technology, and Innovation, Fundaçao de Amparo à Pesquisa do Estado de São Paulo (FAPESP), Russian Foundation for Basic Research, the Leverhulme Trust, the Research Corporation, Ministry of Science and Technology (MOST), Taiwan and the Kavli Foundation. The authors gratefully acknowledge the support of the NSF, STFC, MPS, INFN, CNRS and the State of Niedersachsen/Germany for provision of computational resources. This paper carries LIGO Document Number LIGO–P1700009.

Chapter 10

MEASURING STOCHASTIC GRAVITATIONAL-WAVE ENERGY BEYOND GENERAL RELATIVITY

M. Isi and L. C. Stein, "Measuring stochastic gravitational-wave energy beyond general relativity," (2018), M.I. conceived the project, carried out most calculations and led the writing of the manuscript, arXiv:1807.02123 [gr-qc] .

10.1 Introduction

Besides transient signals, like those detected so far [1–6, 178] by the Advanced Laser Interferometer Gravitational-wave Observatory (aLIGO) [8] and Virgo [9], gravitational-wave (GW) detectors are also expected to be sensitive to a persistent stochastic background [179–186]. This background signal is expected from primordial cosmological processes [187–197] or the incoherent addition of myriad individually-unresolvable astrophysical sources, like compact binary coalescences [198–205] or exotic topological defects [206–209]. Among many other rich scientific goals (see [210] for a review) detection of a stochastic background would provide an invaluable opportunity to study the fundamental nature of gravitational waves as they propagate over cosmological distances.

In the past decade or so, the formalism underlying stochastic GW searches has been extended to theories of gravity beyond general relativity (GR), primarily to account for the potential presence of nontensorial polarizations. Generic metric theories of gravity allow for up to six polarizations, corresponding to scalar (helicity 0), vector (helicity ± 1) and tensor (helicity ± 2) metric perturbations [97, 98]. The effect of these extra polarizations on the stochastic background has been studied in particular for theories with scalar modes [211–213], and in general for all possible modes in a theory-agnostic way [107, 214]. The problem of detecting nontensorial modes in the background has been studied in the context pulsar timing [111, 112, 215, 216] and GW measurements using astrometry [217]. Beyond these proposals, a comprehensive data analysis framework has been recently implemented to search LIGO and Virgo data for GWs of *any* polarization, tensorial or otherwise, and some first upper limits have been placed on their amplitudes [99, 117].

The goal of searches for stochastic backgrounds, within GR or beyond, is to measure the amount of energy that the Universe contains in the form of gravitational waves. Consequently, treatments of stochastic GW signals are predominantly parametrized in terms of their *effective energy-density spectrum* [$\Omega_{\rm GW}$, defined in Eq. (10.29) below]. Such parametrization is only possible thanks to a standard set of assumptions about the properties of gravitational waves, the detectors, and the statistics of the background itself. Although generally justified within GR, the fundamental structure of beyond-GR theories may not always warrant all (or any) of those standard assumptions—even without considering modifications to specific emission mechanisms, or expected source populations. One must therefore be careful in applying the usual premises to searches for stochastic waves that aim to be theory agnostic, and should be aware that adopting any of these assumptions may come with additional observational restrictions.

Perhaps the most important example of an assumption that has been dubiously applied beyond GR concerns the form of the effective stress-energy of GWs. Multiple studies of stochastic signals beyond GR assume that the fractional energy density spectrum in GWs is related to the wave amplitudes in the same way as it is in GR [107, 112, 213, 214, 216–218]. Yet, as pointed out in [219], the expression for the effective GW stress-energy need not be the same in all theories of gravity. This means that it is inadvisable to parametrize putatively model-independent searches for beyond-GR backgrounds assuming the GW energy density has the same functional form as in GR: doing so will result in the use of a quantity that should *not* generally be interpreted as the energy density in GWs. This is not only misleading but, most importantly, can lead to incorrect comparisons between observational limits and theoretical predictions.

Besides this, some of the simplifying assumptions about the properties of the stochastic background that are usually justified in GR are not acceptable in general, and should not be extended to model-independent analyses. This is the case even without considering changes to the potential sources of the background in beyond-GR theories, which may themselves break more of the assumed symmetries. For instance, it is not reasonable to always assume that the usual linear GW polarization amplitudes will be statistically independent, as this will not be true unless the chosen polarization basis diagonalizes the kinetic matrix of the underlying theory of gravity. Similar arguments can be made about the assumptions that the polarizations are equipartitioned, or even that they have well defined phase velocities—let alone that

they propagate at the speed of light.

In view of this, our goal is to straighten out the framework underlying searches for stochastic gravitational backgrounds, to make it formally valid and easily applicable to a large family of theories beyond GR. In Sec. 10.2, we lay out a generic formalism for such searches, review the most commonplace assumptions in standard analyses, and evaluate their degree of applicability to other frameworks; along the way, we also clarify some relevant differences in conventions used by the theory and data analysis literatures. In Sec. 10.3, we provide a series of examples of theories that break the premises behind one or more of these assumptions and show the impact this has on the analysis—we focus on differences in the predicted form of the effective GW stress energy, but also discuss other problematic points. In particular, we use these examples to show how to go from the action defining a theory to *(1)* a relation between the fractional GW energy density spectrum and the correlation of polarization amplitudes, and *(2)* to the cross-correlation of GW detector outputs—which is the relevant observable for ground-based detectors. We review the derivation for general relativity in Sec. 10.3.1, and then move on to Chern-Simons gravity in Sec. 10.3.2, scalar-tensor theories in Sec. 10.3.3, and Fierz-Pauli massive gravity in Sec. 10.3.4. Finally, we offer a summary and conclusions in Sec. 10.4.

10.2 Formalism

In this section, we provide the framework required to search for stochastic GW backgrounds without assuming GR is correct. In Sec. 10.2.1, we review the four-dimensional Fourier transform of a generic GW, lay out its decomposition into polarizations, and provide some useful identities for later use in Sec. 10.3. In Sec. 10.2.2, we focus on the properties of stochastic backgrounds, carefully reviewing the assumptions made in traditional analyses to determine whether they hold in theories beyond GR. In Sec. 10.2.3, we describe the measurement process, including complications that may arise in generic theories. Finally, in Sec. 10.2.4, we sketch the calculations needed to relate the effective stochastic GW energy in any given theory to the polarization amplitudes measurable by a detector.

Here, and throughout this Chapter, spatial three-vectors are identified by an arrow (e.g. $\vec{k}$), or a circumflex accent if they have unit norm (e.g. $\hat{k}$). Four-vectors and higher-rank tensors are denoted by boldface, or abstract index notation (e.g. $\mathbf{k}$ or k_a). For tensor coordinate components, spacetime Greek indices ($\alpha, \beta, \gamma, \ldots$) take values in the range 0–3, while spatial Latin indices ($i, j, k, \ldots$) span 1–3. We use metric

signature +2, using g_{ab} for generic background metrics and η_{ab} for the Minkowski metric. Our conventions for the Levi-Civita tensor follow [220]: $\epsilon_{abcd} = \sqrt{-g}[abcd]$ where g is the determinant of the metric, and $[abcd]$ is the Levi-Civita *symbol*, with $[0123] = +1$; similarly, $\epsilon_{ijk} = \sqrt{\gamma}[ijk]$ where γ is the determinant of the spatial metric, and $[123] = +1$. We normalize (anti-)symmetrizations as idempotent projection operations, e.g., $T_{(ab)} = (T_{ab} + T_{ba})/2$ and $T_{[ab]} = (T_{ab} - T_{ba})/2$.

10.2.1 Decomposition of the metric perturbation

In any metric theory of gravity, as long as the observation region is small compared to the curvature radius, an arbitrary GW metric perturbation $h_{ab}(\mathbf{x})$ at a spacetime point $\mathbf{x}$ may be expressed as a plane-wave expansion by the compact expression:

$$h_{ab}(\mathbf{x}) = \frac{1}{2\pi}\int \tilde{h}_{ab}(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}}\, \widetilde{\mathrm{d}\mathbf{k}}, \tag{10.1}$$

integrating over all directions of propagation, and over both positive and negative frequencies. Here $\tilde{h}_{ab}(\mathbf{k})$ is the complex-valued Fourier amplitude for the wave-vector $\mathbf{k} \equiv (\omega/c, \vec{k})$; we let $\omega = 2\pi f$ be the angular frequency, and $\vec{k} = |\vec{k}|\,\hat{k} \equiv -|\vec{k}|\,\hat{n}$ the spatial wave-vector, implicitly defining $\hat{n}$ as the sky location of the source. To simplify our notation in Eq. (10.1), we have defined the four-dimensional integral over the measure

$$\widetilde{\mathrm{d}\mathbf{k}} \equiv 2c\,\delta(|\vec{k}|^2 - |\vec{k}_\omega|^2)\,|\vec{k}|^{-1}\mathrm{d}\mathbf{k} = \mathrm{d}\omega\,\mathrm{d}\hat{n}, \tag{10.2}$$

where $\delta(x)$ is the Dirac delta function, and the last equality assumes an implicit integration over the magnitude of $\vec{k}$. In order to write this, we assume that there is just *one* dispersion relation, $\omega = \omega(\vec{k}) \equiv \omega_k$, that determines the modulus of $\mathbf{k}$ and implicitly defines $|\vec{k}_\omega| \equiv |\vec{k}|(\omega)$.[1] The dispersion relation is specific to the theory of gravity: for example, $\omega_k = c|\vec{k}|$ and $|\vec{k}_\omega| = \omega/c$ in GR.

With the integration measure defined as in Eq. (10.2), in a local Lorentz frame (so that $\mathbf{x}\cdot\mathbf{y} = \vec{x}\cdot\vec{y} - x_0 y_0$), Eq. (10.1) can be recast in a form most common in stochastic GW literature (see, e.g., [183, 185, 222]):

$$h_{ab}(t,\vec{x}) = \int_{-\infty}^{\infty}\int_{\mathrm{sky}} \tilde{h}_{ab}(f,\hat{n})\, e^{-2\pi i f(t+\hat{n}\cdot\vec{x}/v_{\mathrm{p}})}\,\mathrm{d}\hat{n}\,\mathrm{d}f\,, \tag{10.3}$$

where $v_{\mathrm{p}} \equiv |\vec{k}/\omega|^{-1}$ is the (potentially frequency-dependent) phase velocity of the wave ($v_{\mathrm{p}} = c$ in GR). Finally, to guarantee that $h_{ab}(\mathbf{x})$ be real, we must necessarily

[1] Some theories violate this assumption; for example, bimetric gravity [221] has one massless and one massive gravitational wave mode—we will allow for this briefly in Sec. 10.2.3 only.

have

$$\tilde{h}^{*}_{ab}(f,\hat{n}) = \tilde{h}_{ab}(-f,\hat{n}), \tag{10.4}$$

where the asterisk indicates complex conjugation. In Appendix 10.5, we elucidate the equivalence between Eqs. (10.1) and (10.3), derive the second equality in Eq. (10.2) and discuss differences between our Fourier conventions and those from the field theory literature.

For any given frequency and direction of propagation, the Fourier amplitudes may be written as a linear combination of at most six tensors corresponding to the six polarizations supported by generic metric theories of gravity [97, 98], even if the wave speed is slightly different from the speed of light [83]. Therefore, the most generic gravitational wave in this large category of theories may be written as a function of six independent amplitudes. We may, therefore, define six orthogonal polarization tensors, e^{A}_{ab}, such that

$$\tilde{h}_{ab}(\mathbf{k}) = \tilde{h}_{A}(\mathbf{k})\, e^{A}_{ab}(\hat{n})\,, \tag{10.5}$$

where the sum is over six polarizations indexed by A, and the $\tilde{h}_{A}(\mathbf{k})$'s are the Fourier transforms of the six scalar fields, $h_{A}(\mathbf{x})$, encoding the amplitude of each mode, as defined by means of Eq. (10.1)

In order to study interactions between waves and detectors, it is usually convenient to pick a "synchronous" gauge[2] such that the perturbation is purely spatial in the frame of interest ($h_{0\nu} = 0$), and correspondingly so are the polarization tensors. For instance, in an orthogonal frame in which the z-axis is aligned with the direction of propagation (so that $\hat{k}_i = \delta^3_i$ in that frame), we may write the six degrees of freedom as

$$(h_{ij}) = \begin{pmatrix} h_{\mathrm{b}} + h_{+} & h_{\times} & h_{\mathrm{x}} \\ h_{\times} & h_{\mathrm{b}} - h_{+} & h_{\mathrm{y}} \\ h_{\mathrm{x}} & h_{\mathrm{y}} & h_{\mathrm{l}} \end{pmatrix}, \tag{10.6}$$

in terms of the linear tensor polarizations (+, ×), linear vector polarizations (x, y), and scalar breathing (b) and longitudinal (l) modes.

For the purpose of analyzing the output of multiple GW detectors, it is often convenient to write the polarization tensors in terms of unit vectors tangent and

[2] In a diffeomorphism invariant theory, one may always gauge transform into synchronous gauge by solving an initial value problem. If the theory is not diff-invariant, the Stückelberg trick can be used to restore the symmetry and then gauge transform. We provide an example of this in Sec. 10.3.4.

normal to the celestial sphere at each sky location. A standard linear polarization basis is given by

$$e^{+}_{ab}(\hat{n}) = \hat{\phi}_a(\hat{n})\,\hat{\phi}_b(\hat{n}) - \hat{\theta}_a(\hat{n})\,\hat{\theta}_b(\hat{n}), \tag{10.7a}$$

$$e^{\times}_{ab}(\hat{n}) = \hat{\phi}_a(\hat{n})\,\hat{\theta}_b(\hat{n}) + \hat{\theta}_a(\hat{n})\,\hat{\phi}_b(\hat{n}), \tag{10.7b}$$

$$e^{\mathrm{x}}_{ab}(\hat{n}) = \hat{\phi}_a(\hat{n})\,\hat{k}_b(\hat{n}) + \hat{k}_a(\hat{n})\,\hat{\phi}_b(\hat{n}), \tag{10.7c}$$

$$e^{\mathrm{y}}_{ab}(\hat{n}) = \hat{\theta}_a(\hat{n})\,\hat{k}_b(\hat{n}) + \hat{k}_a(\hat{n})\,\hat{\theta}_b(\hat{n}), \tag{10.7d}$$

$$e^{\mathrm{b}}_{ab}(\hat{n}) = \hat{\phi}_a(\hat{n})\,\hat{\phi}_b(\hat{n}) + \hat{\theta}_a(\hat{n})\,\hat{\theta}_b(\hat{n}), \tag{10.7e}$$

$$e^{\mathrm{l}}_{ab}(\hat{n}) = \hat{k}_a(\hat{n})\,\hat{k}_b(\hat{n}), \tag{10.7f}$$

where $\hat{\theta}(\hat{n})$ and $\hat{\phi}(\hat{n})$ are respectively the celestial polar and azimuthal coordinate vectors for a given source sky location determined by $\hat{n}$; by design, these vectors satisfy $\hat{\theta}(\hat{n}) \times \hat{\phi}(\hat{n}) = -\hat{k}(\hat{n}) = \hat{n}$. Other frame choices are possible, and multiple conventions abound in the literature.

As an example of the polarization decomposition of Eq. (10.5), consider theories in which gravitational perturbations carry spin-weight 2, like GR. In that case, we may choose to work with the two transverse-traceless linear polarization tensors corresponding to the plus (+) and cross (×) amplitudes shown in Eq. (10.6), and Eq. (10.5) becomes simply:

$$\tilde{h}_{ab}(\mathbf{k}) = \tilde{h}_{+}(\mathbf{k})\,e^{+}_{ab}(\hat{n}) + \tilde{h}_{\times}(\mathbf{k})\,e^{\times}_{ab}(\hat{n})\,. \tag{10.8}$$

Because the linear polarization tensors are real-valued by definition [cf. Eq. (10.7)], the reality condition for the amplitudes, Eq. (10.4), implies

$$\tilde{h}_{+/\times}(-f, \hat{n}) = \tilde{h}^{*}_{+/\times}(f, \hat{n})\,. \tag{10.9}$$

Alternatively, instead of the linear modes of Eq. (10.6), we could choose to work with eigenmodes of the helicity operator, i.e. the right- and left-handed circular polarization tensors (denoted "R" and "L" respectively). These modes satisfy an eigenvalue equation

$$\epsilon^{ijk}\hat{k}_k e^{\hat{A}}{}_{\ell j} = i\,\varepsilon_{\hat{A}}\,e^{\hat{A}}{}_{\ell}{}^{i}\,, \tag{10.10}$$

for $\hat{A} \in \{\mathrm{R}, \mathrm{L}\}$ (not summed on the RHS), where we have defined the factor $\varepsilon_{\mathrm{R/L}} = \pm 1$, with the plus (minus) sign corresponding to the R (L) mode. Then, the circular polarization tensors can be written in terms of the ones for plus and cross as

$$\mathbf{e}_{\mathrm{R/L}} = \frac{1}{\sqrt{2}}\left(\mathbf{e}_{+} + i\,\varepsilon_{\mathrm{R/L}}\mathbf{e}_{\times}\right). \tag{10.11}$$

Using the circular tensors as a basis, we would write, instead of Eq. (10.8),

$$\tilde{h}_{ab}(\mathbf{k}) = \tilde{h}_{\rm R}(\mathbf{k})\, e^{\rm R}_{ab}(\hat{n}) + \tilde{h}_{\rm L}(\mathbf{k})\, e^{\rm L}_{ab}(\hat{n})\,, \tag{10.12}$$

and the reality condition, Eq. (10.4), would now imply (note the "L/R" subscript on the right hand side)

$$\tilde{h}_{\rm R/L}(-f,\hat{n}) = \tilde{h}^{*}_{\rm L/R}(f,\hat{n})\,, \tag{10.13}$$

instead of Eq. (10.9). The circular polarization modes can be similarly defined for the vector polarizations to obtain eigenmodes of helicity ± 1. On the other hand, the scalar modes have helicity 0, so in a sense are already circular.

For future reference, note that the spin-weight 2, spin-weight 1 and the transverse spin-weight 0 linear polarization tensors are normalized as usual such that, for a given direction of propagation,

$$e^{A\,ij} e^{A'}{}_{ij} = 2\,\delta^{AA'}\,, \tag{10.14}$$

for $A \in \{+, \times, \mathrm{x}, \mathrm{y}, \mathrm{b}\}$, and $\delta^{AA'}$ the Kronecker delta; on the other hand, the longitudinal tensor satisfies $(e^{\rm l})^{ij}(e^{A})_{ij} = \delta^{{\rm l}A}$. Similarly, the spin-weight 2 circular polarization tensors of Eq. (10.11) satisfy

$$(e^{\hat{A}\,ij})^{*} e^{\hat{A}'}{}_{ij} = 2\,\delta^{\hat{A}\hat{A}'}\,. \tag{10.15}$$

The basis tensors for the circularly-polarized vector modes also satisfy Eq. (10.15).

Although the two linear and circular bases discussed above are probably the most common in the GW literature (modulo normalizations), we are of course free to pick any other. For instance, in the analysis of differential-arm instruments, it is generally convenient to instead work with the traceless linear combination of $h_{\rm b}$ and $h_{\rm l}$, since that is what such detectors can measure. Similarly, different theories may also define their own preferred polarization bases, given by the choice that diagonalizes their kinetic matrices.

10.2.2 Stochastic signals

In the case of stochastic signals, the Fourier amplitudes, $\tilde{h}_{ij}(\mathbf{k})$, are, by definition, random variables and, as such, can be fully characterized by the moments of some (multivariate) probability distribution. Most standard searches for a stochastic GW background make the following assumptions about the random process that produced these amplitudes (see, e.g., [222] for a review): the random process is (*i*) Gaussian,

(*ii*) ergodic, and (*iii*) stationary, with no correlation between amplitudes from different (*iv*) sky locations or (*v*) polarizations, and with (*vi*) equipartition of power across polarizations; furthermore, the process is commonly (although not universally) assumed to be (*vii*) isotropic. We break down these assumptions below and introduce some important definitions along the way.

Stochastic backgrounds are expected to arise from primordial cosmological processes [187–197], or by the incoherent superposition of a great number of signals from contemporary astrophysical events [198–210]. The assumption (*i*) that the astrophysical background is produced by a Gaussian random process is motivated by the central limit theorem—this guarantees that the properties of any large number of incoherently-added GW signals will be normally distributed, regardless of the specific characteristics of any given source. A similar argument can be applied to primordial signals by considering the independent evolution of waves from causally-disconnected regions [184]. Although waves from inflation will technically have non-Gaussianities, they will be small as long as inflation satisfied the slow-roll approximation [184, 197, 223].

For Gaussian processes, all properties of the probability distribution are determined by its first two moments (correlation functions)—namely, the mean and power spectrum (respectively, the one- and two-point correlation functions). The first moment of the distribution, the mean $\langle \tilde{h}(f) \rangle$, will not appear explicitly in any of the expressions below, so we ignore it.[3] Here and below, the expectation value, denoted by angle brackets $\langle \cdot \rangle$, corresponds to ensemble averages, as well as space/time-averages by assumption (*ii*) of ergodicity. The expectation of ergodicity itself comes from the assumption that the Universe is homogeneous (for more discussion on this topic, see [197]).

The second moment of the distribution will end up being an important observable. In order to write down an expression for it, we can make use of assumptions (*iii*) and (*iv*). First, stationarity (*iii*) is motivated by the fact that observation times (order of months to years) are extremely small relative to the dynamical timescales intrinsic to the cosmological processes that could change the properties of the background (order of billions of years); therefore, any changes in the stochastic background would be unnoticeable to us. Formally, stationarity means that the first moment is constant,

[3] Some authors explicitly set this value to zero because the contribution from a nonvanishing mean would take the form of a *coherent* offset in the Fourier amplitudes as a function of frequency, which not only would be hard to justify physically, but would also hardly classify as "stochastic" (see, e.g., [185, 222]).

while the second moment depends only on time *differences* (see, e.g., [224]). As a consequence, the Fourier transform of a stationary random variable can be shown to be such that amplitudes at different frequencies will be statistically independent and, therefore, uncorrelated (Appendix 10.6).

Next, the assumption (*iv*) that amplitudes from different sky locations will be uncorrelated is justified for primordial waves because signals from different points in the sky are only coming into causal contact now at Earth, under ordinary topological assumptions. One could potentially search for nonstandard spatial topologies in a sufficiently "small" universe through angular correlations in gravitational waves [225], in much the same way as in the cosmic microwave background (CMB) [226, 227]. A small universe with nonstandard spatial topology would induce circles of excess correlation in both the CMB and gravitational-wave background. As there has been no evidence of this phenomenon in the CMB, in this Chapter we consider primordial signals from different sky directions to be uncorrelated.

For contemporary ("astrophysical") backgrounds, (*iv*) comes from the assumption that the contributions from multiple sources throughout the sky (say, binary systems) are added "incoherently"—that is, sources are not perfectly aligned and timed as would be needed for signals from different directions to reach us with matching phase and amplitude evolution. Even though such astrophysical sources were in causal contact at some point in the past, they are embedded in chaotic astrophysical environments (with e.g. turbulent magnetohydrodynamics) with Lyapunov times sufficiently short that in practice, they can be treated as uncorrelated. In principle, strong gravitational lensing may introduce correlations between sky bins into the stochastic background, whether primordial or contemporary, but we can expect this effect to be negligible in practice [228].

With assumptions (*iii*) and (*iv*) in place, we may write the second moment of the amplitude distribution in the form (Appendix 10.6):

$$\left\langle \tilde{h}^*_A(\mathbf{k})\, \tilde{h}_{A'}(\mathbf{k}') \right\rangle = \frac{1}{2}\delta(f - f')\, \delta(\hat{n} - \hat{n}')\, S_{AA'}(\mathbf{k})\,. \tag{10.16}$$

This equation defines the *one-sided cross-power spectral density*, $S_{AA'}(\mathbf{k}) \equiv S_{AA'}(f, \hat{n})$, for two signals, $\tilde{h}_{A/A'}(\mathbf{k})$, sharing a wave-vector $\mathbf{k}$ but with potentially different polarizations A and A'. For linear polarizations, this quantity satisfies $S_{AA'}(f, \hat{n}) = S_{A'A}(-f, \hat{n})$, because of the reality condition of Eq. (10.4). Since we are usually interested in the total measured power at a given frequency, regardless of sky direction,

we also define the integral of $S_{AA'}(\mathbf{k})$ over the sky,

$$S_{AA'}(f) \equiv \int_{\text{sky}} S_{AA'}(f, \hat{n})\, \mathrm{d}\hat{n}\,, \tag{10.17}$$

which carries units of strain^2/Hz. For $A = A'$, this is nothing more than the one-sided *power spectral density* (PSD) in polarization A, which we denote $S_A(f) \equiv S_{AA}(f)$. In general, for any real-valued random variable $X(t)$, the PSD can be approximated as twice the square of the band-limited Fourier transform [224],

$$S_X(f) = \lim_{T\to\infty} \frac{2}{T} \left| \int_{T/2}^{-T/2} X(t)\, e^{2\pi i f t} \mathrm{d}t \right|^2 , \tag{10.18}$$

in practice always computed for some long but finite integration time, T, on the order of months to years for observations of the stochastic background. As usual, the factor of 2 in Eq. (10.18) accounts for the fact that this is the *one-sided* PSD, $S(f) \equiv S(|f|)$.

Assumption (*v*) that the different polarizations are statistically independent may be used to discard off-diagonal terms in the cross-power spectrum, so that $S_{AA'}(f) = \delta_{AA'} S_A(f)$. However, one must be careful with this simplification: the assumption is valid *if and only if* one works in a polarization basis that diagonalizes the kinetic matrix of the theory. Importantly, as we will show with specific examples, such a basis need *not* be the linear polarization basis used in most GR analyses. Even when working within GR, it is generally better, from a theoretical standpoint, to work in terms of the circular modes, as they are eigenstates of the helicity operator, and they might be produced with different intensities in the early universe [229–231].

Besides assuming that the polarizations are uncorrelated, it is also common to assume that there is equipartition of power between them, assumption (*vi*) in our list above. Under this presumption, the background is said to be unpolarized and the polarization PSDs may be written in terms of the *total* GW spectral density, $S(f) = \sum S_A(f)$, such that $S_A = S(f)/N$, where N is the number of polarizations allowed to propagate in a given theory. In general, this assumption is only justified if the polarizations both diagonalize the kinetic matrix *and* interact similarly with matter, so that they are sourced in equal amounts. This not always the case: for example, in both massive gravity [105, 232] and dynamical Chern-Simons gravity [96, 233, 234], different polarizations couple to sources with different strengths.

Finally, the simplest searches for a stochastic background also adopt assumption (*vii*) of isotropy, in which case $S(f, \hat{n}) = S(f)/4\pi$, by Eq. (10.17). In GR, if one

disregards the proper motion of the solar system, this assumption is expected to hold well for most foreseeable sources of a stochastic background detectable by existing ground-based observatories, since they are expected to originate from cosmological distances [184, 185, 222]. [4] For cosmological sources, isotropy is likely also a good assumption in many beyond-GR theories; however, isotropy should not be expected to hold in theories with a preferred frame, which are intrinsically anisotropic [236–243]. For simplicity, the rest of this document will treat only the case of an isotropic background, but this does not affect the spirit of the results, which can be easily generalized to the anisotropic case. For predictions of the angular power spectrum of astrophysical GR backgrounds see [244], and for corresponding observational limits that do not assume isotropy see [245].

Assuming both (*vii*) an isotropic background and (*v*) uncorrelated polarizations, on top of (*iii*) stationarity and (*iv*) uncorrelated sky bins, Eq. (10.16) can be written directly in terms of the power spectral density,

$$\left\langle \tilde{h}^*_A(\mathbf{k})\tilde{h}_{A'}(\mathbf{k}')\right\rangle = \frac{1}{8\pi}\delta(f-f')\delta(\hat{n}-\hat{n}')\delta_{AA'}S_A(f). \tag{10.19}$$

If one further assumed (*vi*) equipartition, $S_A(f)$ would be replaced with $S(f)/N$, as explained above. This is the form of the expression most common in recent literature about detection of stochastic gravitational-wave backgrounds (e.g., [222]).

10.2.3 Detection

Because the output of ground-based GW detectors is largely dominated by stochastic instrumental and environmental noise [10, 246], it is not possible to measure the power spectrum of the polarization amplitudes, $S_A(f)$, directly with a single detector at any level of interest. However, this quantity may be inferred by looking instead at the cross-correlation of the output of two or more instruments (see, e.g., [222] for a comprehensive review of data analysis methods).

We assume that each GW detector has a purely linear response to gravitational waves. Therefore, in the Fourier domain, the response of detector I to a plane wave $\tilde{h}_{ab}(\mathbf{k})$ must be expressible as

$$\tilde{h}_I(\mathbf{k}) = \tilde{D}^{ab}_I(\mathbf{k})\,\tilde{h}_{ab}(\mathbf{k}), \tag{10.20}$$

for some tensor $\tilde{D}^{ab}_I(\mathbf{k}) \equiv \tilde{D}^{ab}_I(f,\hat{n})$ representing the detector's frequency- and direction-dependent transfer function. This tensor encodes all relevant information

[4] The same will not necessarily be true for LISA, which will be sensitive to galactic stochastic sources, like the "confusion noise" from white-dwarf binaries [235].

about the detector and the physics of the measurement process [247–254] (for considerations specific to gravity beyond GR, see e.g. [18, 98, 108, 110, 215, 255]). The detector's output $\tilde{h}_I(\mathbf{k})$ (e.g. the calibrated current out of a photodiode) is a gauge-invariant observable. However, the metric perturbation $\tilde{h}_{ab}(\mathbf{k})$ is gauge-dependent; therefore, the detection tensor $\tilde{D}_I^{ab}(\mathbf{k})$ must also depend on the gauge choice, so that the overall gauge dependence on the RHS of Eq. (10.20) exactly cancels.

Assuming a basis of polarization states A that have well defined phase velocities (i.e. they diagonalize the kinetic matrix of the theory), we may use Eq. (10.20) to write the Fourier transform of the signal at detector I explicitly as a sum over polarizations and an integral over sky directions,

$$\tilde{h}_I(f) = \int \sum_A \tilde{F}_I^A(f,\hat{n})\, \tilde{h}_A(f,\hat{n})\, e^{-2\pi i f \hat{n}\cdot\vec{x}_I/v_{\rm p}^A}\, \mathrm{d}\hat{n}\,, \tag{10.21}$$

defining the Fourier-domain response functions as the contraction between the detector and polarization tensors, $\tilde{F}_I^A(f,\hat{n}) \equiv \tilde{D}_I^{ab}(f,\hat{n})\, e_{ab}^A(\hat{n})$, which must also be gauge-dependent.

The time-domain analogue of Eq. (10.20) is given by a convolution,

$$h_I(t,\vec{x}_I) = \int_{-\infty}^{\infty} D_I^{ab}(t)\, h_{ab}(t-\tau,\vec{x}_I)\, \mathrm{d}\tau\,, \tag{10.22}$$

with $\vec{x}_I$ the location of detector I, and $D_I^{ab}(t)$ its impulse response. Since $\tilde{D}_I^{ab}(\mathbf{k})$ is gauge-dependent, the same must be true for $D_I^{ab}(t)$. For an ideal differential arm-length instrument, it is easiest to write down this detector tensor in a synchronous gauge ($h_{0\nu} = 0$ in the detector frame), wherein the end test masses' coordinate locations will not change [220]. In such a gauge, the resulting differential-arm detector tensor is the purely geometric factor

$$D_{ab}(t) = \frac{1}{2}\left(\hat{X}_a\hat{X}_b - \hat{Y}_a\hat{Y}_b\right), \tag{10.23}$$

with $\hat{X}$ and $\hat{Y}$ spacelike unit vectors pointing along the detector arms. For real interferometric detectors, like LIGO and Virgo, Eq. (10.23) is valid *only* in the small-antenna limit (arm length $\ll$ GW wavelength) [250–254].

For any realistic detector, the tensor of Eq. (10.23) will vary in time due to the motion of the instrument with respect to the inertial frame of the wave (e.g. due to Earth's rotation, for ground-based observatories). However, for the cases we are interested in, we can take this variation to be slow with respect to the period of the

waves, so that it can be ignored if Eq. (10.21) is implemented via short-time Fourier transforms. In this ideal "slow-detector" limit, we may then treat the response as time- and frequency-independent to write $\tilde{D}^{ab}_I(\mathbf{k}) = D^{ab}_I(t) \equiv D^{ab}_I$, and so Eq. (10.21) simplifies to

$$\tilde{h}_I(f) = \int \sum_A F^A_I(\hat{n})\, \tilde{h}_A(f,\hat{n})\, e^{-2\pi i f \hat{n}\cdot\vec{x}_I/v^A_{\rm p}}\, \mathrm{d}\hat{n}\,, \tag{10.24}$$

with the frequency-independent *antenna patterns* defined in full analogy to our definition of $\tilde{F}^A_I(f,\hat{n})$ above,

$$F^A_I(\hat{n}) \equiv D^{ab}_I\, e^A_{ab}(\hat{n})\,. \tag{10.25}$$

For details on this simplification, and nuances applicable to anisotropic backgrounds, see Sect. IV in [256].

In the Fourier domain, the cross-correlation between the output of two detectors may then be written in terms of the second moment of the distribution of polarization amplitudes as

$$\begin{aligned}\left\langle \tilde{h}^*_I(f)\tilde{h}_{I'}(f')\right\rangle = \int \mathrm{d}\hat{n}\mathrm{d}\hat{n}' \sum_{AA'} \left\langle \tilde{h}^*_A(\mathbf{k})\tilde{h}_{A'}(\mathbf{k}')\right\rangle \\ \times F^{*A}_I(\hat{n})F^{A'}_{I'}(\hat{n}')\, e^{i(\vec{k}_{A'}\cdot\vec{x}_{I'}-\vec{k}_A\cdot\vec{x}_I)}\,,\end{aligned} \tag{10.26}$$

where, again, assumption (*ii*) of ergodicity is tacitly implied. If we also assume, as we will throughout this paper, that the background is (*iii*) stationary and (*vii*) isotropic, and (*iv*) that sky bins are uncorrelated, we may then use Eq. (10.19) to simplify this to

$$\left\langle \tilde{h}^*_I(f)\tilde{h}_{I'}(f')\right\rangle = \frac{1}{2}\delta(f-f')S_{AA'}(f)\Gamma^{AA'}_{II'}(f), \tag{10.27}$$

where we have defined the *generalized overlap reduction function* for polarizations A, A' and detectors I, I',

$$\Gamma^{AA'}_{II'}(f) \equiv \frac{1}{4\pi}\int \mathrm{d}\hat{n}\, F^{*A}_I(\hat{n})F^{A'}_{I'}(\hat{n})e^{-2\pi i f\hat{n}\cdot\xi^{AA'}_{II'}}, \tag{10.28}$$

in terms of the phase factor $\xi^{AA'}_{II'}(f) \equiv \vec{x}_I/v^A_{\rm p} - \vec{x}_{I'}/v^{A'}_{\rm p}$, which acquires a potential frequency dependence through the phase velocities. If there is one dispersion relation shared by all polarizations (true throughout the rest of this paper), the exponent in Eq. (10.28) can be written as $-2\pi i f\hat{n}\cdot\xi^{AA'}_{II'} = -i\vec{k}\cdot\Delta\vec{x}_{II'}$, in terms of the separation between detectors $\Delta\vec{x}_{II'} \equiv \vec{x}_I - \vec{x}_{I'}$. The overlap reduction functions encode all relevant information pertaining GW polarizations and speed, as well as detector

geometry. The specific definition and normalization chosen here are intended to facilitate generalization of the analysis beyond GR, and are not necessarily standard (see, e.g., Sect. 5.3 of [222] for a review of these functions and their properties).

Because the noise in different instruments will generally be statistically independent [10, 246], by cross-correlating the output of a pair of detectors, one may directly measure the *signal* cross-correlation of Eq. (10.27), and hence infer the polarization power spectra $S_{AA'}(f)$ (as proposed by [179, 180], and studied in multiple works since). In a theory that allows for N independent polarizations, there will be up to $N(N+1)/2$ different $S_{AA'}$ terms (only N if the correlation matrix is diagonal), and at least as many detector pairs ("baselines") will be needed to break all degeneracies between them.

10.2.4 Energy density

Searches for a stochastic gravitational-wave background attempt to measure the Universe's total energy density in gravitational waves as a function of frequency. However, inferring this quantity from direct observables requires theoretical assumptions. Furthermore, the equivalence principle precludes being able to localize energy density in gravitational waves, so this is in fact an *effective* energy density. We elaborate on these important points below and sketch the general procedure to link the effective GW energy density to observables at the detector in (almost) any given theory. Concrete examples of how to apply this are provided in Sec. 10.3.

With an eye to cosmology, the quantity of interest in stochastic searches is usually chosen to be the log-fractional spectrum of the effective GW energy density [180–183, 185],

$$\Omega_{\rm GW}(f) \equiv \frac{1}{\rho_{\rm critical}} \frac{{\rm d}\rho_{\rm GW}}{{\rm d}\ln f}\,, \tag{10.29}$$

with $\rho_{\rm GW}$ the effective GW energy density as a function of frequency, and $\rho_{\rm critical}$ the critical density required to close the universe,

$$\rho_{\rm critical} \equiv \frac{3c^2 H_0^2}{8\pi G}\,, \tag{10.30}$$

where H_0 is the present Hubble parameter [185]. Presenting results of a stochastic background search in terms of this quantity facilitates their cosmological interpretation. More importantly, using an energy density (however parametrized) allows for direct comparison with theoretical models: in order to predict the properties of the GW background, one computes the typical GW power emitted by the system

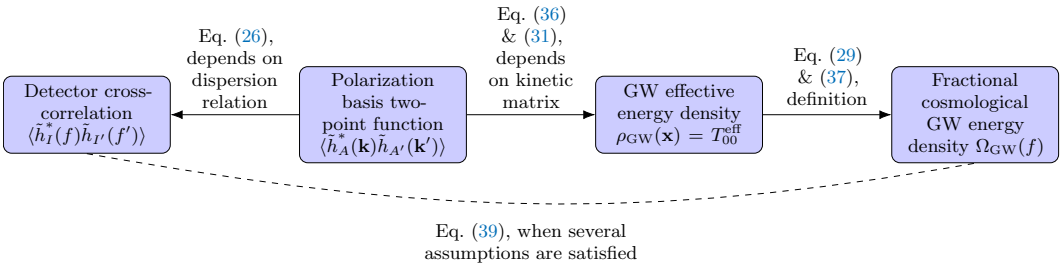

Figure 10.1: *Key quantities appearing in stochastic searches, and how they are related to each other.* The relationships between them are theory-dependent. The primary observable is the detector cross-correlation, but inferences are often stated in terms of the fractional cosmological GW energy density, or a parametrization thereof. Arrows point from more fundamental quantities to derived quantities.

of interest (e.g., compact binaries, cosmic strings, or primordial fluctuations) and then obtains an *energy* spectrum by incoherently adding many such contributions (e.g., using the quadrupole formula with merger rates from population synthesis) [187–210].

However, GW detectors do not measure the effective physical GW energy density, but rather the *amplitude* of the waves at each instrument. In particular, searches for a stochastic background are sensitive to the (incoherent) strain amplitude power, Eq. (10.16). This will remain true for future detection methods, like space missions or pulsar timing. In the case of ground-based observatories, as outlined in Sec. 10.2.3, the stochastic strain amplitudes are probed through the cross-correlation of detector outputs across a network, Eq. (10.26). Thus, whatever the detection method, we will need an object that relates gravitational-wave amplitudes to energies—a mapping that is theory-dependent.

The frequency-domain *effective stress-energy tensor* (ESET) for gravitational waves lets us translate between the more accessible two-point amplitude correlation function, Eq. (10.16), [or the two-detector-output cross-correlation, Eq. (10.26)] and the GW contribution to the energy density, Eq. (10.29). In GR the ESET is given by a simple expression first derived by Isaacson [257, 258] (see Sec. 10.3.1 below), which enables stochastic searches to be parametrized directly in terms of $\Omega_{\rm GW}(f)$ [180–183, 185]. Interestingly, the same relationship has been assumed to hold in most stochastic GW data analysis schemes that allow for departures from GR [107, 112, 213, 214, 216–218], even though the Isaacson formula will not necessarily hold in arbitrary theories [219]. Using the Isaacson formula when inappropriate will

lead to a mistranslation between detector cross-correlations and GW energy densities. This is not only misleading, but can also lead to errors when deriving constraints on theories from observations.

In the context of any specific theory of gravity, the ESET can be derived directly from the action. The ESET is given by a space-time average of the variation of the second-order perturbation of the action with respect to the background (inverse) metric [219],

$$T^{\rm eff}_{ab} = \left\langle\!\!\left\langle -\frac{2}{\sqrt{-g^{(0)}}}\frac{\delta \mathcal{S}^{(2)}_{\rm eff}}{\delta g^{(0)ab}} \right\rangle\!\!\right\rangle, \tag{10.31}$$

where the double angular brackets $\langle\!\langle \cdot \rangle\!\rangle$ indicate an averaging procedure over a spacetime region on the order of several wavelengths (e.g. Brill-Hartle averaging, though other procedures [259] agree when there is a separation of length scales). We briefly summarize the approach here; we refer the interested reader to [219] for more exposition.

The second-order Lagrangian $\mathcal{L}^{(2)}_{\rm eff}$ is obtained from the action $\mathcal{S}[g, \psi]$ after perturbing the metric g_{ab} and other dynamical fields ψ via

$$g_{ab} = g^{(0)}_{ab} + \epsilon h^{(1)}_{ab} + \epsilon^2 h^{(2)}_{ab} + O(\epsilon^3), \tag{10.32}$$

$$\psi = \psi^{(0)} + \epsilon\psi^{(1)} + \epsilon^2\psi^{(2)} + O(\epsilon^3), \tag{10.33}$$

and collecting terms in the action order-by-order in the small parameter ϵ. This gives the expansion

$$\begin{aligned} \mathcal{S}[g, \psi] = \mathcal{S}^{(0)}[g^{(0)}, \psi^{(0)}] & \\ + \epsilon\mathcal{S}^{(1)}[h^{(1)}, \psi^{(1)}; g^{(0)}, \psi^{(0)}] & \\ + \epsilon^2\mathcal{S}^{(2)}[h^{(1,2)}, \psi^{(1,2)}; g^{(0)}, \psi^{(0)}] + O(\epsilon^3), & \end{aligned} \tag{10.34}$$

where $h^{(1,2)}$ means both $h^{(1)}$ and $h^{(2)}$ are present. At order ϵ^0, the action $\mathcal{S}^{(0)}$ generates the ordinary nonlinear background equations of motion for $g^{(0)}$ and $\psi^{(0)}$. At order ϵ^1, the action $\mathcal{S}^{(1)}$ is purely a "tadpole" term which vanishes when $(g^{(0)}, \psi^{(0)})$ are on shell, and therefore does not contribute to any equations of motion. The same is true for the second-order perturbations $(h^{(2)}, \psi^{(2)})$, which appear linearly in $\mathcal{S}^{(2)}$. However, $(h^{(1)}, \psi^{(1)})$ appear quadratically in $\mathcal{S}^{(2)}$: the quadratic action $\mathcal{S}^{(2)}$ then generates the linear equations of motion for $(h^{(1)}, \psi^{(1)})$ when varied with respect to $(h^{(1)}, \psi^{(1)})$; at the same time, the variation with respect to $g^{(0)}$ will be a quadratic functional of $(h^{(1)}, \psi^{(1)})$, and results in the ESET.

From now on we drop the order-counting superscript, letting $h = h^{(1)}$, since we will not encounter $h^{(2)}$. In a local Lorentz frame whose time direction is aligned with the Hubble flow, we can define the position-space effective GW energy density as

$$\rho_{\rm GW} \equiv T^{\rm eff}_{00}[h, h], \tag{10.35}$$

where the double argument $[h, h]$ is just to remind us that $T^{\rm eff}$ is a quadratic functional of h. To use Eq. (10.29), we want $\rho_{\rm GW}$ in momentum space, so we need to make use of a plane-wave expansion like Eq. (10.1). The result will always be a momentum-space integral of the form

$$\rho_{\rm GW}(\mathbf{x}) = \int \widetilde{\mathrm{d}\mathbf{k}}\widetilde{\mathrm{d}\mathbf{k}}' Q^{abcd}(\mathbf{k}, \mathbf{k}') \left\langle \tilde{h}^*_{ab}(-\mathbf{k})\, \tilde{h}_{cd}(\mathbf{k}') \right\rangle e^{i(\mathbf{k}+\mathbf{k}')\cdot\mathbf{x}}, \tag{10.36}$$

where the (gauge-dependent) tensor Q^{abcd} encodes information about the kinetic matrix of the theory in momentum space, and we have used Eq. (10.4) to write $\tilde{h}_{ab}(\mathbf{k}) = \tilde{h}^*_{ab}(-\mathbf{k})$. Notice that here we have replaced the spacetime averaging of Eq. (10.31) with ensemble averaging, based on assumption (*ii*) ergodicity. When the two-point function $\langle \tilde{h}^*_A(\mathbf{k}) \tilde{h}_{A'}(\mathbf{k}') \rangle$ is of the form of Eq. (10.16), the double integral will collapse to a single integral, and the physical energy density will be related to the power spectral density $S_{AA'}(\mathbf{k})$, with some potentially nontrivial frequency dependence arising from Q^{abcd} (we will see several examples below).

When this double integral collapses to a single integral, we can then define the fractional energy density per frequency bin via

$$\rho_{\rm GW} = \int \frac{\mathrm{d}\rho_{\rm GW}}{\mathrm{d}f} \mathrm{d}f = \int \frac{\mathrm{d}\rho_{\rm GW}}{\mathrm{d}\ln f} \frac{\mathrm{d}f}{f}\,. \tag{10.37}$$

With this definition of $\mathrm{d}\rho_{\rm GW}/\mathrm{d}\ln f$, and the relationship between the energy density Eq. (10.36) and a two-point function like Eq. (10.16), it will be possible to relate the power spectral density $S_{AA'}(\mathbf{k})$ to the cosmological fractional energy density $\Omega_{\rm GW}(f)$, Eq. (10.29). The relationships between all these key quantities are illustrated in Fig. 10.1.

Once we have this, we may work directly with $\Omega_{\rm GW}(f)$; in particular, data analysis searches usually assume a power-law model like

$$\Omega_{\rm GW}(f) = \Omega_0 \left(\frac{f}{f_0}\right)^\alpha, \tag{10.38}$$

for some spectral index α, and Ω_0 the characteristic amplitude at some arbitrary reference frequency f_0. This is how LIGO generally parametrizes its searches, e.g. [186]; for a discussion of the validity of this parametrization, see [260].

10.3 Example theories

In this section, we show how different gravitational theories imply different functional relations between the effective fractional energy density spectrum, $\Omega_{\rm GW}(f)$ in Eq. (10.29), the strain cross-power spectrum, $S(f)$ in Eq. (10.16), and, consequently, the cross-correlation between detector outputs, Eq. (10.26) As discussed in Sec. 10.2.3, this last quantity is the relevant observable for ground-based instruments, on which we focus. The relationships between all the key quantities are illustrated in Fig. 10.1. Along the way, we also discuss the expected statistical properties of the polarization amplitudes under each framework, as required purely by the basic structure of the theory (that is, not considering specific source models).

We first demonstrate the procedure by rederiving the standard GR expressions from the Einstein-Hilbert action (Sec. 10.3.1), and then offer a series of beyond-GR examples for which the analogous result is different: we consider the case of Chern-Simons gravity, a theory which is not parity-symmetric (Sec. 10.3.2); this is followed by Brans-Dicke gravity, the prototypical example of a scalar-tensor theory (Sec. 10.3.3); finally, we study Fierz-Pauli gravity (Sec. 10.3.4), in which the graviton is endowed with a mass. The last two examples support nontensorial modes of the metric perturbation (see Sec. 10.2.1).

For all the examples we consider, we find it reasonable to simplify our equations by assuming the stochastic background is (*i*) Gaussian, (*ii*) ergodic, (*iii*) stationary and (*vii*) isotropic, with (*iv*) no correlation between different sky locations. In all cases, then, we find that we can write the cross-correlation between the output of two ideal differential-arm detectors I and I' in the form

$$\left\langle \tilde{h}_I^*(f)\tilde{h}_{I'}(f')\right\rangle = \frac{3H_0^2}{4\pi^2|f|^3}\delta(f-f')\sum_A \Xi_A(f)\,\Omega_A(f)\,\Gamma^A_{II'}(f)\,, \tag{10.39}$$

where the sum is over some polarization basis A that diagonalizes the kinetic matrix of the theory. Here the $\Gamma^A_{II'}(f)$'s are the generalized overlap reduction functions of Eq. (10.28), $\Omega_A(f)$ is the effective fractional energy spectrum in polarization A defined by analogy to Eq. (10.29), and $\Xi_A(f)$ is a model-dependent factor encoding deviations from GR. In Einstein's theory, $\Xi_A(f) = 1$ for tensor polarizations and vanishes otherwise, as we show below.

Many of the results in this section are derived on a flat background and will therefore be erroneous in a cosmological setting. However, because of the vast separation of scales between the gravitational wavelength $\lambda_{\rm GW}$ and the Hubble parameter today H_0,

the error between the flat space results and the cosmologically-correct results will be of fractional order $O(\lambda_{\rm GW} H_0/c)$. This correction has been explicitly computed in GR [261], and while we are not aware of the same computation in beyond-GR theories, it should remain true as long as the theory of gravity respects the separation of scales.

10.3.1 General relativity

The vacuum Einstein field equations can be derived from the Einstein-Hilbert (EH) action,

$$\mathcal{S}_{\rm EH} = \kappa \int {\rm d}\mathbf{x}\sqrt{-g}R\,, \tag{10.40}$$

where $\kappa = c^4/(16\pi G)$, g is the determinant of the metric g_{ab}, and R is the Ricci scalar [220]. We may now expand the metric around some background, $g_{ab} = g^{(0)}_{ab} + \epsilon h_{ab} + \ldots$, as in Eq. (10.32). The source-free linearized equations of motion, on a flat background (so that Riemann vanishes), and in the transverse-traceless gauge ($\nabla_a h^{ab} = 0$ and $h^a{}_a = 0$), take the simple form

$$\Box h_{ab} = 0, \tag{10.41}$$

where $\Box \equiv \nabla^a \nabla_a$ is the d'Alembertian with respect to the background metric. Equation (10.41) leads to the standard geometric optics approximation to GW propagation, from which it follows that GWs show no birefringence and always propagate at the speed of light.

Focus now on the second-order perturbation of Eq. (10.40). On a flat background, the second-order Lagrangian density is given by [219, 262]

$$\begin{aligned}\mathcal{L}^{(2)}_{\rm GR} = \kappa\sqrt{-g}\Bigg[&\frac{1}{2}\left(\nabla_a \bar{h}^{bc}\right)\left(\nabla_b \bar{h}^a{}_c\right) - \frac{1}{4}\left(\nabla_a \bar{h}_{cd}\right)\left(\nabla^a \bar{h}^{cd}\right)\\ &+\frac{1}{8}\left(\nabla_a \bar{h}\right)\left(\nabla^a \bar{h}\right)\Bigg],\end{aligned} \tag{10.42}$$

where all derivatives are taken with respect to the background metric g_{ab}, and $\bar{h}_{ab} \equiv h_{ab} - g_{ab}h^c{}_c/2$ is the trace-reversed metric perturbation. This piece of the Lagrangian density corresponds to $\mathcal{S}^{(2)}$ in Eq. (10.34).

We now apply the transverse-traceless gauge conditions and evaluate the perturbations on-shell (that is, we enforce the first-order equations of motion). Then, varying with respect to the inverse background metric, g^{ab}, as in Eq. (10.31), we obtain, far away

from sources,

$$T_{ab}^{(\mathrm{GR})} = \frac{c^4}{32\pi G}\left\langle\!\!\left\langle \nabla_a h_{cd} \nabla_b h^{cd} \right\rangle\!\!\right\rangle . \tag{10.43}$$

This is the well-known expression for the effective stress-energy carried by a gravitational wave, first derived by Isaacson (and, consequently, known as the *Isaacson formula*) [257, 258].

We may now use Eq. (10.43) to relate $S(f)$ and $\Omega_{\mathrm{GW}}(f)$, as outlined in Sec. 10.2.4. The Isaacson expression implies that, in a local Lorentz frame,

$$\rho_{\mathrm{GW}} = T_{00}^{(\mathrm{GR})} = \frac{c^2}{32\pi G}\left\langle\!\!\left\langle \partial_t h_{ij}\, \partial_t h^{ij} \right\rangle\!\!\right\rangle , \tag{10.44}$$

where we have used the fact that the transverse-traceless metric perturbation will be purely spatial. Plugging in the plane-wave expansion of Eq. (10.1), using the reality condition of Eq. (10.4), and invoking (*ii*) ergodicity, we may rewrite this as

$$\rho_{\mathrm{GW}} = \frac{-c^2}{128\pi^3 G}\int \widetilde{\mathrm{d}\mathbf{k}}\widetilde{\mathrm{d}\mathbf{k}}'\, \omega\omega' \left\langle \tilde{h}^*_{ij}(-\mathbf{k})\, \tilde{h}^{ij}(\mathbf{k}') \right\rangle \times e^{i(\mathbf{k}+\mathbf{k}')\cdot\mathbf{x}}. \tag{10.45}$$

This means that, in GR and in our gauge, ρ_{GW} takes the form of Eq. (10.36) with

$$Q^{abcd}_{\mathrm{GR}} = \frac{-c^2}{128\pi^3 G} g^{ac} g^{bd} \omega\omega' . \tag{10.46}$$

It is convenient at this point to expand the Fourier amplitudes into polarizations. Because GR is parity-symmetric, in this theory all modes are generated and propagate equally, so one is free to choose between linear and circular polarizations; however, working with the former is slightly simpler because the corresponding polarization tensors, Eqs. (10.7a) and (10.7b), are real-valued. Then, summing over A, $A' \in \{+, \times\}$,

$$\rho_{\mathrm{GW}} = \frac{-c^2}{128\pi^3 G}\int \widetilde{\mathrm{d}\mathbf{k}}\widetilde{\mathrm{d}\mathbf{k}}'\, \omega\omega' \left\langle \tilde{h}^*_{A}(-\mathbf{k})\, \tilde{h}_{A'}(\mathbf{k}') \right\rangle \times e^{A}_{\ ij}\, e^{A'ij}\ e^{i(\mathbf{k}+\mathbf{k}')\cdot\mathbf{x}}. \tag{10.47}$$

We now use the fact that the Fourier amplitudes are given by a random process to simplify our expression for ρ_{GW} via Eq. (10.16). Following common practice and for the sake of simplicity, we will assume that the stationary Gaussian background is also (*vii*) isotropic and (*vi*) unpolarized, with equal contributions from the linear

polarizations. Letting the total PSD in tensor polarizations be $S_t \equiv S_+ + S_\times$ with $S_+ = S_\times = S_t/2$, this means

$$S_{AA'}(f,\hat{n}) = \frac{1}{8\pi}\delta_{AA'}S_t(f)\,. \tag{10.48}$$

Again, the assumption (*vi*) of equipartition is justified because GR conserves parity. The correlation of the Fourier polarization amplitudes, Eq. (10.19), then becomes

$$\left\langle \tilde{h}^*_A(\mathbf{k})\tilde{h}_{A'}(\mathbf{k}')\right\rangle = \frac{1}{16\pi}\delta(f-f')\delta(\hat{n}-\hat{n}')\delta_{AA'}S_t(f). \tag{10.49}$$

With this in place, and noting that Eq. (10.14) implies $e_{Aij}e^{Aij} = 4$ when summing over $A = \{+,\times\}$, the effective energy density of Eq. (10.47) simplifies to

$$\rho_{\rm GW} = \frac{\pi c^2}{4G}\int_0^\infty S_t(f) f^2\,\mathrm{d}f\,. \tag{10.50}$$

Comparing with Eq. (10.37), we can immediately read off $\mathrm{d}\rho_{\rm GW}/\mathrm{d}\ln f$, and then, from the definition of $\Omega_{\rm GW}(f)$, Eq. (10.29), we conclude

$$S_t(f) = \frac{3H_0^2}{2\pi^2|f|^3}\Omega_{\rm GW}(f)\,. \tag{10.51}$$

As discussed in Sec. 10.2.3, the actual observable for stochastic-background searches in data from ground-based observatories is the cross-correlation between the outputs of pairs of detectors. For an isotropic background, this is given by Eq. (10.27), which can be written in terms of the fractional energy density by means of Eq. (10.51):

$$\left\langle \tilde{h}^*_I(f)\tilde{h}_{I'}(f')\right\rangle = \frac{3H_0^2}{8\pi^2|f|^3}\delta(f-f')\,\Omega_{\rm GW}(f)\,\Gamma^{\rm t}_{II'}\,, \tag{10.52}$$

where we have defined the total tensor overlap-reduction function as $\Gamma^{\rm t}_{II'} \equiv \Gamma^{++}_{II'}(f) + \Gamma^{\times\times}_{II'}(f)$. This is the desired expression relating the observable strain cross-correlation to the fractional effective-energy density spectrum, that will be predicted by theory. Eq. (10.52) is used in most LIGO and Virgo searches for a stochastic background, via parametrizations like the $\Omega_{\rm GW}(f)$ power-law of Eq. (10.38). Comparing to Eq. (10.39), and recalling $\Omega_{\rm GW} = \Omega_+ + \Omega_\times$ with $\Omega_+ = \Omega_\times = \Omega_{\rm GW}/2$, we see that in GR, $\Xi(f) = 1$ for tensor polarizations, and vanishes otherwise, as expected.

10.3.2 Chern-Simons gravity

Chern-Simons (CS) theory is an extension of GR with motivations ranging from anomaly-cancellation in curved spacetime, low-energy limits of both string theory

and loop quantum gravity, effective field theory of inflation, and more [233, 263–273]. The theory is characterized by the presence of a parity-odd, axion-like scalar field, which couples to curvature through a parity-odd interaction (see [96] for a review). The ESET in this theory was derived in [219], in an asymptotically-flat spacetime and approaching future null infinity ($\mathscr{I}^+$). As noted before, by promoting flat-space results to a cosmological setting, we are making an extremely small error of fractional order $O(\lambda_{\rm GW} H_0/c)$. Below, we provide a sketch of this derivation and show what the result implies for the stochastic background.

As a consequence of its lack of parity symmetry, CS gravity generally predicts birefringent propagation and generation of the metric perturbations, so that one of the circular tensor polarizations is amplified at the expense of the other [274]. Consequently, as is true for any theory lacking parity symmetry, it is not appropriate to assume that the stochastic background is unpolarized [275]. Furthermore, as we will see, the nondynamical version of the theory predicts an expression for the effective GW stress-energy different from the Isaacson formula of Eq. (10.43), and consequently differs from GR via a factor of $\Xi(f) \neq 1$ in Eq. (10.39).

In the absence of matter, CS gravity is given by the Einstein-Hilbert action of Eq. (10.40), plus terms describing the axion-curvature coupling ($\mathcal{S}_{\rm int}$), and dynamics ($\mathcal{S}_\vartheta$) of the scalar field ϑ [96, 233],

$$\mathcal{S}_{\rm CS} = \mathcal{S}_{\rm EH} + \mathcal{S}_{\rm int} + \mathcal{S}_\vartheta \,, \tag{10.53}$$

$$\mathcal{S}_{\rm int} = \frac{\alpha}{4} \int {\rm d}\mathbf{x} \sqrt{-g}\, \vartheta\, {}^*RR \,, \tag{10.54}$$

$$\mathcal{S}_\vartheta = -\frac{\beta}{2} \int {\rm d}\mathbf{x} \sqrt{-g}\, g^{ab} \left(\nabla_a \vartheta\right) \left(\nabla_b \vartheta\right) \,. \tag{10.55}$$

In the above, α is the constant determining the coupling of the CS field to the gravitational sector, while β controls the kinetic energy of the scalar; ${}^*RR$ is the Pontryagin density, which is defined in terms of the Riemann tensor, R_{abcd}, by

$${}^*RR = \frac{1}{2} \epsilon^{abef} R_{abcd} R^{cd}{}_{ef} \,, \tag{10.56}$$

with ϵ_{abcd} the Levi-Civita tensor. This term is parity-odd, and gives CS gravity much of its richness.

Studying the dynamics of the theory, one may show that gravitational waves in CS gravity will present the same tensor (spin-weight 2) propagating degrees of freedom as in GR [233, 234]. On a flat background and in Lorenz gauge ($\nabla_a \bar{h}^{ab} = 0$), metric

perturbations follow the first-order equations of motion [96, 219],

$$\Box\bar{h}_{ab} = -\frac{1}{\kappa}\tilde{T}^{(\vartheta)}_{ab} + \frac{\alpha}{\kappa}\Big[\nabla_c\bar{\vartheta}\nabla_d\Box\bar{h}_{e(a}\epsilon^{cde}{}_{b)} + \nabla_c\nabla_d\bar{\vartheta}\epsilon^{c}{}_{ef(a}\nabla^f\left(\nabla_{b)}\bar{h}^{de} - \nabla^d\bar{h}_{b)}{}^{e}\right)\Big], \tag{10.57}$$

where we split ϑ into a smooth background piece $\bar{\vartheta}$ and a perturbation $\tilde{\vartheta}$, and $\tilde{T}^{(\vartheta)}_{ab}$ is the stress energy sourced quadratically by $\tilde{\vartheta}$,

$$\tilde{T}^{(\vartheta)}_{ab} = \beta\left[\left(\nabla_a\tilde{\vartheta}\right)\left(\nabla_b\tilde{\vartheta}\right) - \frac{1}{2}g_{ab}\left(\nabla^c\tilde{\vartheta}\right)\left(\nabla_c\tilde{\vartheta}\right)\right]. \tag{10.58}$$

Again on a flat background, CS gravity admits an approximately traceless gauge [219], so that $\bar{h}_{ab}$ can be replaced by h_{ab} in these equations, as was done for GR in Eq. (10.41).

In the weak-coupling limit (i.e. $\alpha\nabla\bar{\vartheta} \ll \kappa\lambda_{\rm GW}$, for GW perturbation wavelength $\lambda_{\rm GW} = c/f$), it can be shown that the quadratic Lagrangian density corresponding to $\mathcal{S}^{(2)}$ in Eq. (10.34) can be written as [219]

$$\mathcal{L}^{(2)}_{\rm CS} = \mathcal{L}^{(2)}_{\rm GR} + \Delta\mathcal{L}^{(2)}_{\rm CS}, \tag{10.59}$$

where $\mathcal{L}^{(2)}_{\rm GR}$ is the effective Lagrangian density derived from the Einstein-Hilbert action, Eq. (10.42), and

$$\Delta\mathcal{L}^{(2)}_{\rm CS} \equiv \frac{\alpha}{4}\sqrt{-g}\,\epsilon^{abcd}\left(\nabla_e\nabla_f\bar{\vartheta}\nabla_a\bar{h}_b{}^f\nabla_c\bar{h}_d{}^e + \nabla_a\bar{\vartheta}\nabla^e\bar{h}_b{}^f\nabla_d\nabla_e\bar{h}_{fc}\right). \tag{10.60}$$

From this we may derive the effective GW stress-energy, and relate the energy to the strain cross-correlation, for both the nondynamical and dynamical versions of the theory, as outlined in Sec. 10.2.4.

Nondynamical theory

The nondynamical version of CS gravity is obtained from Eqs. (10.53)–(10.55) by setting $\beta = 0$. This removes the dynamics of the scalar field, fixing it to some *a priori* value. Furthermore, in the *canonical embedding* of this theory [96], we set the field's gradient to be purely timelike in some global frame,

$$\nabla_\alpha\bar{\vartheta} = \mu^{-1}\delta^t{}_\alpha, \tag{10.61}$$

for some constant μ. When expanding h_{ab} as a power series in α in the weak coupling limit, the first-order equations of motion on a flat background, Eq. (10.57), reduce to

a simple wave equation, $\Box\bar{h}_{ab} = 0 + O(\alpha^2)$, as in GR. This implies that the Lorenz gauge is compatible with synchronous gauge (i.e., we can satisfy both $\nabla^\mu h_{\mu\nu} = 0$ and $h_{0\nu} = 0$ in the same frame).

In spite of its name, there is nothing special about the canonical embedding other than its simplicity [96]. Although in the following we assume this particular form for the background scalar field, the qualitative features of our result should be similar in general, possibly with extra terms stemming from any non-zero higher-order derivatives of $\bar{\vartheta}$.

In the canonical embedding of nondynamical CS gravity, it can be shown that the only non-GR contribution to the on-shell ESET comes from the second term in Eq. (10.60), in regions at a great distance from the source [219]. Consequently, we can write:

$$T^{(\mathrm{CS})}_{ab} = T^{(\mathrm{GR})}_{ab} + \Delta T^{(\mathrm{CS})}_{ab}, \tag{10.62}$$

where $T^{(\mathrm{GR})}_{ab}$ is the Isaacson tensor from Eq. (10.43), and $\Delta T^{(\mathrm{CS})}_{ab}$ is the surviving contribution from Eq. (10.60), with components

$$\Delta T^{(\mathrm{CS})}_{\mu\nu} = \frac{\alpha}{2\mu}\left\langle\!\!\left\langle \epsilon_i{}^{jk}\nabla_{(\mu}h^{i\sigma}\nabla_{\nu)}\nabla_k h_{\sigma j}\right\rangle\!\!\right\rangle \tag{10.63}$$

in the global frame. The corresponding non-GR energy density, $\Delta\rho^{(\mathrm{CS})} \equiv \Delta T^{(\mathrm{CS})}_{00}$, over a flat background is

$$\begin{aligned}\Delta\rho^{(\mathrm{CS})} = \frac{-i\alpha}{8\pi^2\mu c^2}\int \widetilde{\mathrm{d}\mathbf{k}}\widetilde{\mathrm{d}\mathbf{k}}'\,\omega\omega'\epsilon^{ijk}k'_k\; e^{i(\mathbf{k}+\mathbf{k}')\cdot\mathbf{x}} \\ \times\left\langle \tilde{h}_i^{*\,\ell}(-\mathbf{k})\,\tilde{h}_{\ell j}(\mathbf{k}')\right\rangle,\end{aligned} \tag{10.64}$$

after expanding over plane-waves in a synchronous gauge, as done in Eq. (10.45), and using the reality condition of Eq. (10.4) to substitute $\tilde{h}_i{}^{\ell}(\mathbf{k}) \to \tilde{h}_i^{*\,\ell}(-\mathbf{k})$. In the notation of Eq. (10.36), $Q^{abcd} = Q^{abcd}_{\mathrm{GR}} + \Delta Q^{abcd}$, where in dCS and in our gauge choice, the components of the correction are given by

$$\Delta Q^{\alpha\beta\gamma\delta} = \frac{-i\alpha}{8\pi^2\mu c^2}g^{\beta\gamma}\epsilon^{\alpha\delta i}k'_i\,\omega\omega'\,. \tag{10.65}$$

We want to expand the perturbation into polarizations, as we did for the GR case in Eq. (10.45). However, it would be inconvenient to do so in terms of the linear plus and cross modes, since these are not actual eigenmodes of the kinetic matrix in CS gravity, and hence their amplitudes will generally be correlated [274]. Instead, we

will work with the right- and left-handed modes of Eq. (10.11), which *do* diagonalize the CS kinetic matrix. Letting $\hat{A} \in \{\mathrm{R}, \mathrm{L}\}$, then

$$\Delta\rho^{(\mathrm{CS})} = \frac{-i\alpha}{8\pi^2\mu c^2}\int \widetilde{\mathrm{d}\mathbf{k}}\widetilde{\mathrm{d}\mathbf{k}}'\ \omega\omega'|\omega'|\epsilon^{ijk}\hat{k}'_k\,(e^{\hat{A}}{}_{i}^{\ell})^*(e^{\hat{A}'}{}_{\ell j}) \times \left\langle \tilde{h}^*_{\hat{A}}(-\mathbf{k})\tilde{h}_{\hat{A}'}(\mathbf{k}')\right\rangle\ e^{i(\mathbf{k}+\mathbf{k}')\cdot\mathbf{x}}. \tag{10.66}$$

Here we have used the fact that, to first order, the GW dispersion relation in canonical nondynamical CS gravity is the same as in GR, so that $k'_i = |\omega'|\hat{k}'_i$.

As a consequence of the birefringence of GWs in CS gravity, it is also no longer reasonable to assume an (*vi*) unpolarized background; rather, we should expect $S_{\mathrm{R}}(f) \neq S_{\mathrm{L}}(f)$. (Although in the canonical embedding there is no amplitude birefringence in GW propagation, wave generation should still be expected to break parity symmetry.). However, we *are* justified in taking the two polarizations to be uncorrelated in this basis, i.e. $S_{\mathrm{RL}}(f) = S_{\mathrm{LR}}(f) = 0$, which is not true in the linear basis.

With the above considerations in mind, we may write the correlation factor in terms of the PSD in each mode as in Eq. (10.19), so that Eq. (10.66) becomes ($\omega' \to -\omega$):

$$\Delta\rho^{(\mathrm{CS})} = \frac{i\alpha\pi^2}{2\mu c^3}\int \mathrm{d}f\,\mathrm{d}\hat{n}|\omega|^3 S_{\hat{A}}(f)\,\delta_{\hat{A}\hat{A}'}\epsilon^{ijk}\hat{k}_k \times (e^{\hat{A}}{}_{i}^{\ell})^*(e^{\hat{A}'}{}_{\ell j})\,. \tag{10.67}$$

With the help of Eqs. (10.15) and (10.10), this simplifies to our final expression for the additional energy density, after integrating over the source direction $\hat{n}$:

$$\Delta\rho^{(\mathrm{CS})} = -\frac{\alpha 8\pi^3}{\mu c^3}\int_0^\infty [S_{\mathrm{R}}(f) - S_{\mathrm{L}}(f)]\,f^3\,\mathrm{d}f. \tag{10.68}$$

Writing the GR contribution also in terms of circular polarizations and adding it to the purely-CS part, it is then straightforward to obtain the total energy density in nondynamical CS gravity:

$$\rho_{\mathrm{GW}} = \frac{\pi c^2}{4G}\int_0^\infty \sum_{\hat{A}} \lambda_{\hat{A}}(f)\,S_{\hat{A}}(f)\,f^2\mathrm{d}f\,, \tag{10.69}$$

where the sum is over circular polarizations, and for convenience we defined

$$\lambda_{\hat{A}}(f) \equiv 1 - 32\pi^2\varepsilon_{\hat{A}}\frac{\alpha G}{\mu c^5}f\,, \tag{10.70}$$

with $\varepsilon_{\mathrm{R/L}} = \pm 1$, as in Eq. (10.10). Because the energy is diagonal in the circular modes, this may also be written as $\rho_{\mathrm{GW}} = \rho_{\mathrm{R}} + \rho_{\mathrm{L}}$, with each term defined as the corresponding summand (pulling the sum up front) in Eq. (10.69).

Using the definition of the fractional energy density spectrum, Eq. (10.29), this means that the strain power in each polarization can be written as

$$S_{\hat{A}}(f) = \frac{3H_0^2}{2\pi^2|f|^3}\lambda_{\hat{A}}^{-1}(f)\,\Omega_{\hat{A}}(f) \tag{10.71}$$

where $\Omega_{\mathrm{R/L}}(f)$ represents the energy density in each polarization, defined in full analogy to Eq. (10.29) such that $\Omega_{\mathrm{GW}} = \Omega_{\mathrm{R}}(f) + \Omega_{\mathrm{L}}(f)$. The observable cross-correlation between the output of two detectors, Eq. (10.27), can then be written as in Eq. (10.39), if we choose the circular tensor polarization basis and let

$$\Xi_{\hat{A}}(f) = \lambda_{\hat{A}}^{-1}(f) \approx 1 + 32\pi^2\varepsilon_{\hat{A}}\frac{\alpha G}{\mu c^5}f\,, \tag{10.72}$$

with the approximation being valid in the weak-coupling limit that we have been working in ($\alpha/\mu \ll \kappa c/f$). As expected, the usual GR expression of Eq. (10.52) is recovered in the limit that the coupling of the scalar field vanishes ($\alpha \to 0$), if we further assume $S_{\mathrm{R}} = S_{\mathrm{L}}$.

Dynamical theory

Perhaps surprisingly, the case of *dynamical* CS gravity is simpler for our purposes. This is because, in the dynamical theory, the functional form of the effective GW stress-energy tensor (about flat spacetime and with $\nabla\bar{\vartheta} \to 0$ far away from sources) is given by the Isaacson formula of Eq. (10.43), as in GR [219]. This notwithstanding, dynamical CS gravity still breaks parity symmetry, featuring birefringent propagation and generation of gravitational waves. Therefore, just as in the nondynamical theory, it would not be justified to take the stochastic background to be unpolarized. Instead, using the circular polarization states, in which the CS kinetic matrix diagonalizes, we find that in the dynamical case

$$\left\langle \tilde{h}_I^*(f)\tilde{h}_{I'}(f')\right\rangle = \frac{3H_0^2}{4\pi^2|f|^3}\delta(f-f')\Omega_{\hat{\mathrm{A}}}(f)\,\Gamma_{II'}^{\hat{\mathrm{A}}}(f)\,, \tag{10.73}$$

with $\hat{A} \in \{L, R\}$, but now allowing $\Omega_L \neq \Omega_R$. Here we also have $\Xi_{\hat{A}}(f) = 1$, as in GR. With at least two detector pairs (e.g. LIGO-Livingston–Virgo, and LIGO-Hanford–Virgo), it should be possible to use this to measure the energy density

in each circular mode. Eq. (10.73) may also be used to parametrize a polarized background in GR, and hence probe polarized cosmological backgrounds like those predicted in [229].

10.3.3 Scalar-tensor theories

Scalar-tensor (ST) theories are defined by the presence of one or more scalar fields that couple to the gravitational sector nonminimally. From a field-theoretic point of view, this family of theories is a natural extension of GR, and, as such, has been extensively studied [101, 276–282]. Scalar-tensor theories are also well-motivated as effective field theories encapsulating the low-energy behavior of quantum gravity completions, like string theory [283–286] or braneworld scenarios [287, 288]. These theories also have important applications to cosmology [289, 290].

The literature contains several formulations of ST theories, with varying degrees of generality and complexity. For simplicity, we will focus on the most basic case, which was introduced by Brans and Dicke in an attempt to make Einstein's theory fully compatible with Mach's principle [101]. Scalar stochastic GW backgrounds have been previously studied in the context of this theory [211]—we revisit some of those results here from the ESET point of view presented in Sec. 10.2.4.

The vacuum action for Brans-Dicke scalar-tensor gravity can be expressed as

$$\mathcal{S}_{\rm ST} = \kappa \int \mathrm{d}\mathbf{x}\sqrt{-g}\left[\phi R - \frac{\omega_{\rm BD}}{\phi}\nabla^a\phi\nabla_a\phi\right], \tag{10.74}$$

for a scalar field ϕ, some constant $\omega_{\rm BD}$, and where, as before, $\kappa = c^4/(16\pi G)$ and R is the Ricci scalar. Matter will follow geodesics of the metric associated with Eq. (10.74); this representation is known as the *Jordan* frame of the theory. Notice that the scalar field has a "scaling symmetry," where if we take $\phi \to C\phi$ for some nonzero real constant C, this constant can be absorbed into κ. If the scalar field asymptotes to a constant ϕ_0 far away from all sources, we can use this scaling symmetry to change the value of ϕ_0 to whatever is most convenient for our calculations, e.g. we can set $\phi_0 = 1$ (note that [211] chooses a different asymptotic value).

Alternatively, it is often useful to recast the ST action in a conformal frame in which the scalar is only minimally coupled to the metric sector. This can be achieved by defining the conformal metric:

$$\underset{\sim}{g}_{ab} \equiv \frac{\phi}{\phi_0} g_{ab}\,. \tag{10.75}$$

In terms of this metric and a redefined scalar field φ, Brans-Dicke theory can be recovered from the action

$$\underset{\sim}{\mathcal{S}}_{\rm ST} = \kappa \int \mathrm{d}\mathbf{x}\sqrt{-\underset{\sim}{g}}\,[\underset{\sim}{R} - 2\underset{\sim}{\nabla}^a\varphi\underset{\sim}{\nabla}_a\varphi], \tag{10.76}$$

where the under-tilded quantities are to be computed using the metric of Eq. (10.75). The new scalar field φ is related to ϕ from Eq. (10.74) by

$$\phi/\phi_0 \equiv e^{-2\alpha_0(\varphi-\varphi_0)}, \tag{10.77}$$

$$\alpha_0 \equiv (3+2\omega_{\rm BD})^{-1/2}, \tag{10.78}$$

where φ_0 is some constant analogous to ϕ_0. Because of its resemblance to the Einstein-Hilbert action of Eq. (10.40), this is known as the *Einstein* representation of the theory. As we will see, Eq. (10.76) is more convenient for theoretical manipulations than Eq. (10.74)—although it should be kept in mind that matter follows geodesics in Eq. (10.74), but not in Eq. (10.76).

As usual, we will perturb the Jordan metric and field to first order by letting $g_{ab} \to g_{ab} + h_{ab}$ and $\phi \to \phi_0 + \delta\phi$, with $h_{ab} \ll g_{ab}$ and $\delta\phi \ll \phi_0$, like in Eq. (10.32). For convenience, we will also define

$$\Phi \equiv -\delta\phi/\phi_0\,. \tag{10.79}$$

Equivalently, we may perturb the Einstein-frame quantities by writing $\underset{\sim}{g}_{ab} \to \underset{\sim}{g}_{ab} + \underset{\sim}{h}_{ab}$ and $\varphi \to \varphi_0 + \delta\varphi$. The two perturbations will be related by the transformation of Eq. (10.77), yielding to first order:

$$\underset{\sim}{h}_{ab} = h_{ab} - \Phi g_{ab}, \tag{10.80a}$$

$$\delta\varphi = \frac{\Phi}{2\alpha_0}, \tag{10.80b}$$

$$\underset{\sim}{g}_{ab} = g_{ab}\,. \tag{10.80c}$$

Studying linearized perturbations in the Jordan frame, it is possible to show that there exists a gauge in which the vacuum linear equations of motion reduce to simple wave equations, $\Box h_{ab} = 0$ and $\Box\Phi = 0$, with the trace of the perturbation satisfying $h = 2\Phi$ [276]. This implies that the metric perturbation may be locally decomposed into spin-weight 2 and spin-weight 0 contributions, in the spirit of Eq. (10.5), such that

$$\tilde{h}_{ab}(\mathbf{k}) = \tilde{h}_+(\mathbf{k})e^+_{ab}(\hat{n}) + \tilde{h}_\times(\mathbf{k})e^\times_{ab}(\hat{n}) + \tilde{\Phi}(\mathbf{k})e^{\rm b}_{ab}(\hat{n}) \tag{10.81}$$

with the polarization tensors as given by Eqs. (10.7a), (10.7b) and (10.7e). It is easy to check, using the linear transformations of Eq. (10.80), that in the Einstein frame this is equivalent to a gauge in which the trace-reversed Einstein metric is given by the same expression as the Jordan metric, i.e.

$$\bar{\underset{\sim}{h}}_{ab} \equiv \underset{\sim}{h}_{ab} - \underset{\sim}{h}\,\eta_{ab}/2 = h_{ab}\,. \tag{10.82}$$

Consequently, $\bar{\underset{\sim}{h}}_{ab}$ is divergenceless ($\nabla^a \bar{\underset{\sim}{h}}_{ab} = 0$), although it is *not* traceless ($\bar{\underset{\sim}{h}} = h = 2\Phi$).

We will now derive an expression for the GW effective stress-energy in the *Einstein* frame, and will then re-express this in terms of the Jordan quantities that are measurable at the detector. The reason for this choice is that, by definition of the Einstein frame, the metric and scalar field separate in the action of Eq. (10.76). This nice feature not only makes *our* computations easier, but also those in the modeling of observational scenarios for the stochastic background—which will generally also offer a prediction of the energy spectrum in the Einstein frame. In any case, there is no difference between the Jordan and Einstein energies to linear order, as given by Eq. (10.80).

From the variation of the second-order perturbation of the Einstein-frame Lagrangian density, Eq. (10.76), with respect to the inverse background metric, $\underset{\sim}{g}^{\mu\nu}$, we can show, as in Sec. 10.2.4, that the effective GW stress-energy tensor will be given by two terms:

$$\underset{\sim}{T}^{(\mathrm{ST})}_{ab} = \underset{\sim}{T}^{(\mathrm{EH})}_{ab} + \Delta\underset{\sim}{T}^{(\mathrm{ST})}_{ab}. \tag{10.83}$$

The first, $\underset{\sim}{T}^{(\mathrm{EH})}_{ab}$, is the contribution from the Einstein-Hilbert part of the action in Eq. (10.76)—this is analogous to $T^{(\mathrm{GR})}_{ab}$ in Eq. (10.43), but is *not identical* to it due to the presence of the scalar and the necessarily different gauge choice with $\bar{\underset{\sim}{h}} = 2\Phi$. In fact, starting from the quadratic Lagrangian density of Eq. (10.42), after evaluating on shell, it may be shown that

$$\underset{\sim}{T}^{(\mathrm{EH})}_{ab} = \frac{\kappa}{2}\left\langle\!\!\left\langle \nabla_a h_{cd} \nabla_b h^{cd} - 2\nabla_a\Phi\nabla_b\Phi \right\rangle\!\!\right\rangle, \tag{10.84}$$

in a synchronous gauge for h_{ab} (Appendix 10.7). Recall that the metric perturbation appearing in this equation can be equivalently taken to be the trace-reversed perturbation in the Einstein frame, or the regular perturbation in the Jordan frame ($h_{ab} = \bar{\underset{\sim}{h}}_{ab}$), because we are working to linear order [Eq. (10.82)].

The second contribution to the stress energy comes from the variation of the second term in Eq. (10.76), and can be shown to be

$$\Delta \underset{\sim}{T}{}^{(\mathrm{ST})}_{ab} = (3 + 2\omega_{\mathrm{BD}})\kappa \left\langle\!\!\left\langle \nabla_a \Phi \nabla_b \Phi \right\rangle\!\!\right\rangle , \tag{10.85}$$

after applying the equations of motion (Appendix 10.7.1). In both Eqs. (10.84) and (10.85), we have simplified the notation by letting $\underset{\sim}{\nabla} \to \nabla$, because these derivatives are taken with respect to the background metric, $\underset{\sim}{g}_{ab} = g_{ab}$ to linear order [Eq. (10.80c)]. Adding together the two contributions, we obtain the total Einstein frame stress energy:

$$\underset{\sim}{T}{}^{(\mathrm{ST})}_{ab} = \frac{\kappa}{2} \left\langle\!\!\left\langle \nabla_a h_{cd} \nabla_b h^{cd} \right\rangle\!\!\right\rangle + 2\kappa \left(1 + \omega_{\mathrm{BD}}\right) \left\langle\!\!\left\langle \nabla_a \Phi \nabla_b \Phi \right\rangle\!\!\right\rangle . \tag{10.86}$$

This agrees with the expression originally found in [276] by a different procedure.

As in previous sections, we may now expand the corresponding effective energy density, $\rho_{\mathrm{GW}} \equiv T_{00}$, into plane-waves to obtain an expression like Eq. (10.36) with $\mathcal{Q}^{abcd} = \mathcal{Q}^{abcd}_{\mathrm{GR}} + \Delta\mathcal{Q}^{abcd}$ and

$$\Delta\mathcal{Q}^{abcd} = -\frac{c^2 \omega \omega'}{128\pi^3 G} \left(\omega_{\mathrm{BD}} + 1\right) g^{ab}\, g^{cd} , \tag{10.87}$$

where we have used the fact that $\Phi = g^{ab} h_{ab}/2$, as implied by Eq. (10.81). The energy density can also be written explicitly in terms of the polarization amplitudes as

$$\rho_{\mathrm{GW}} = -\frac{\kappa}{2c^2} \int \widetilde{\mathrm{d}\mathbf{k}} \widetilde{\mathrm{d}\mathbf{k}}' \, \omega \omega' e^{i(\mathbf{k}+\mathbf{k}')\cdot\mathbf{x}} \times \sum_A \lambda_A \left\langle \tilde{h}^*_A(-\mathbf{k}) \tilde{h}_A(\mathbf{k}) \right\rangle , \tag{10.88}$$

with a sum over the polarizations $A \in \{+, \times, \mathrm{b}\}$. To make the notation more compact, we have also defined the auxiliary factor

$$\lambda_A = \begin{cases} (3 + 2\omega_{\mathrm{BD}}) & \text{if } A = \mathrm{b}, \\ 1 & \text{if } A = +, \times. \end{cases} \tag{10.89}$$

For more details, see Appendix 10.7.2.

We must now make some assumptions about the statistical properties of the Fourier amplitudes. As before, we will assume the simplest case of (*vii*) an isotropic background, with (*v*) uncorrelated polarizations and (*iv*) sky-bins. We can then use the corresponding expression for the correlations, Eq. (10.19), to get:

$$\rho_{\mathrm{GW}} = \frac{\pi c^2}{4G} \sum_A \int_0^\infty \mathrm{d}f \, f^2 \lambda_A S_A(f) . \tag{10.90}$$

From the definition of the fractional energy spectrum, Eq. (10.29), this in turn implies [cf. Eq. (10.37)]

$$S_A(f) = \frac{3H_0^2}{2\pi^2|f|^3}\lambda_A^{-1}\Omega_A(f), \tag{10.91}$$

where Ω_A represents the energy content in polarization A, with $\Omega_{\rm GW} = \sum \Omega_A$ for $A \in \{+, \times, \mathrm{b}\}$, because we took the different modes to be uncorrelated. This is justified because the kinetic matrix of the theory is diagonal for $A \in \{+, \times, \mathrm{b}\}$.

We may use this expression for the power spectral density in each polarization to write the observable cross-correlation between the output of a pair of detectors (I and I'). Using the cross-correlation expression of Eq. (10.27), we find again that we can write this as in Eq. (10.39) with a summation over polarizations $A \in \{+, \times, \mathrm{b}\}$, and the factor

$$\Xi_A(f) = \lambda_A^{-1} = \begin{cases} (3 + 2\omega_{\rm BD})^{-1} & \text{if } A = \mathrm{b}, \\ 1 & \text{if } A = +, \times. \end{cases} \tag{10.92}$$

The GR result of Eq. (10.52) is recovered, as expected, in the limit that $\omega_{\rm BD}$ becomes infinitely large.

10.3.4 Massive gravity

From a field theory perspective, general relativity is nothing but the theory of a nontrivial massless spin-2 particle—the graviton. Therefore, theories of massive gravity, which endow the graviton with a mass, are a natural (and, in some sense, the simplest) extension of Einstein's theory (see [105, 232] for reviews). In its most basic form, linearized massive gravity is given by the Fierz-Pauli (FP) action [291],

$$\mathcal{S}_{\rm FP} = \mathcal{S}^{(2)}_{\rm EH} + \mathcal{S}_m, \tag{10.93}$$

where $\mathcal{S}^{(2)}_{\rm EH}$ is the quadratic piece of the Einstein-Hilbert action of Eq. (10.40), and $\mathcal{S}_m$ is the Fierz-Pauli mass term,

$$\mathcal{S}_m = \frac{1}{2}\kappa \int \mathrm{d}\mathbf{x}\sqrt{-g}\,\mu^2 h_{ab}h_{cd}\,g^{a[b}g^{c]d}, \tag{10.94}$$

for a graviton mass $m = \hbar\mu/c$, and where h_{ab} is a linear perturbation over the background metric g_{ab}, as before. For background diffeomorphism invariance, we explicitly include the $\sqrt{-g}$ term in Eq. (10.94), though the background metric in this action should be thought of as Minkowski (yet potentially in curvilinear coordinates).

Extending this linear theory to a more general, nonlinear one is far from trivial (for reviews see e.g. [105, 232]). Therefore, we will focus only on the linear theory

of Eq. (10.93), and will only comment on the relevance of the linearized analysis for the nonlinear completion at the end of this section. Until then, we will write "massive gravity" to mean Fierz-Pauli theory.

Massive gravity has many interesting features, including the fact that it supports five independent GW polarizations corresponding to the helicity states available to a massive particle: two tensor modes (helicity ±2), two vector modes (helicity ±1), and one scalar mode (helicity 0). Over a flat background, these degrees of freedom propagate following the Klein-Gordon equations of motion describing a massive graviton,

$$\left(\Box - \mu^2\right) h_{ab} = 0\,, \tag{10.95}$$

and are divergenceless and traceless,

$$\nabla_a h^{ab} = 0\,, \tag{10.96a}$$

$$h = g^{ab} h_{ab} = 0\,. \tag{10.96b}$$

These three equations follow from the variation of Eq. (10.93) with respect to the inverse metric perturbation h^{ab} [105, 232] and contain all relevant properties of GWs in this theory. Equation (10.95) immediately gives the dispersion relation $\omega^2 = c^2(|\vec{k}|^2 + \mu^2)$.

Before proceeding, we must discuss the length scales which appear in this calculation. Around a flat background, there are only two length scales of importance: the wavelength of radiation, $\lambda_{\rm GW}$, and the graviton's Compton wavelength, $\lambda_m = h/mc$. Generally speaking, the relevance of corrections to GR due to a nonvanishing graviton mass will depend on the value of the ratio of these two,

$$\alpha(f) \equiv \frac{\lambda_{\rm GW}}{\lambda_m} = \frac{c\mu}{2\pi f}\,, \tag{10.97}$$

or, equivalently, the ratio of the norm of the wave's spatial wavevector to its angular frequency,

$$\beta(f) \equiv \frac{|\vec{k}|c}{2\pi f} = \sqrt{1-\alpha^2}\,. \tag{10.98}$$

This last quantity is just the graviton group velocity in natural units, which is the same as the ratio of the speed of light to the graviton phase velocity ($\beta = v_g/c = c/v_p$). We should expect to recover GR results for vanishing graviton mass, when $\mu \to 0$ and, consequently, $\alpha \to 0$ and $\beta \to 1$. Note that for propagating GW modes, we must have $\alpha < 1$, and consequently, β is real-valued.

When we move to the cosmological setting (or a more general curved background), there is a third length scale at each point: the curvature radius of the background, $L_{\rm BG}$. In order for the Brill-Hartle averaging procedure to be valid, we need a separation of scales, $\lambda_{\rm GW} \ll L_{\rm BG}$, since the B-H average makes errors of order $\lambda_{\rm GW}/L_{\rm BG}$ (this is clearly satisfied when comparing the LIGO/Virgo frequency band with the cosmological curvature radius cH_0^{-1}). Now, in the following, we want to keep the dependence on μ, so we keep terms at the length scale λ_m. This is only compatible with the B-H averaging procedure if we demand the additional separation of scales $\lambda_m \ll L_{\rm BG}$.

We now return to a flat background to develop the results which we later promote to a cosmological background. In a generic frame (that is, without special boosts) with rectangular coordinates, and with the z-axis aligned along the wave's direction of propagation, the equations of motion can be shown to restrict the components of a massive GW to be of the form (Appendix 10.8.1):

$$(h_{\mu\nu}) = \begin{pmatrix} \beta^2 h_{\rm l} & -\beta h_{\rm x} & -\beta h_{\rm y} & -\beta h_{\rm l} \\ -\beta h_{\rm x} & -\frac{1}{2}\alpha^2 h_{\rm l} + h_+ & h_\times & h_{\rm x} \\ -\beta h_{\rm y} & h_\times & -\frac{1}{2}\alpha^2 h_{\rm l} - h_+ & h_{\rm y} \\ -\beta h_{\rm l} & h_{\rm x} & h_{\rm y} & h_{\rm l} \end{pmatrix}, \tag{10.99}$$

for the five linear polarization amplitudes h_A, with $A \in \{+, \times, {\rm x}, {\rm y}, {\rm l}\}$. Here, we have parametrized the single scalar mode allowed by the theory in terms of the longitudinal amplitude (rather than the breathing amplitude, or some linear combination thereof), treating it as the fundamental degree of freedom.[5] It is straightforward to check that the metric of Eq. (10.99) is traceless and divergenceless, as required by the equations of motion (Appendix 10.8.1). The GR case is recovered in the limit of vanishing graviton mass, if we also re-enforce the requirements of transversality and tracelessness by setting $h_{\rm x} = h_{\rm y} = h_{\rm l} = 0$.

We must now determine the functional form of the GW effective stress-energy tensor in FP theory. Varying the effective Lagrangian density from Eq. (10.93) with respect to g^{ab}, as in Eq. (10.31), we may write the effective GW stress-energy tensor in massive gravity as (Appendix 10.8.2)

$$T_{ab}^{\rm (FP)} = T_{ab}^{\rm (EH)} + \Delta T_{ab}^{\rm (FP)}. \tag{10.100}$$

[5] Importantly, note that our definition of the longitudinal mode follows the standard in the GW literature, and does not necessarily agree with the conventions from the massive-gravity theory literature, e.g. Ref. [232] defines the longitudinal tensor as proportional to our $e^{\rm l}_{ab} - e^{\rm b}_{ab}/2$, instead of just $e^{\rm l}_{ab}$.

As in previous examples, $T^{(\mathrm{EH})}_{ab}$ is derived from the Einstein-Hilbert piece of the action, but now evaluated with the new on-shell condition Eq. (10.95), rather than the GR requirement of Eq. (10.41). This gives

$$T^{(\mathrm{EH})}_{ab} = \frac{\kappa}{2}\left\langle\!\!\left\langle \nabla_a h^{cd} \nabla_b h_{cd} \right\rangle\!\!\right\rangle + \kappa\mu^2 \left\langle\!\!\left\langle h_{da} h_b{}^d + \tfrac{1}{4} g_{ab} h^{cd} h_{cd} \right\rangle\!\!\right\rangle . \tag{10.101}$$

The second term in Eq. (10.101) is derived from the Fierz-Pauli mass term of Eq. (10.94) and reduces to

$$\Delta T^{(\mathrm{FP})}_{ab} = -\kappa\mu^2 \left\langle\!\!\left\langle h_{da} h_b{}^d + \tfrac{1}{4} g_{ab} h^{cd} h_{cd} \right\rangle\!\!\right\rangle , \tag{10.102}$$

for on-shell perturbations (Appendix 10.8). This result includes no derivatives of the metric, as is to be expected from Eq. (10.94). Perhaps surprisingly, the mass terms appearing Eq. (10.101) and Eq. (10.102) exactly cancel, resulting in the same functional form for the ESET as in GR,

$$T^{(\mathrm{FP})}_{ab} = \frac{\kappa}{2}\left\langle\!\!\left\langle \nabla_a h^{cd} \nabla_b h_{cd} \right\rangle\!\!\right\rangle . \tag{10.103}$$

This result is in agreement with one derived in [292] based on Noether's theorem on a Minkowski background (though note that Ref. [292] had a slightly different mass term, but this difference cancels out after evaluating the ESET on shell). Despite the fact that the two functionals have the same on-shell expressions, the solutions h_{ab} on which they will be evaluated differ, because they satisfy different linearized equations of motion, Eq. (10.41) vs Eq. (10.95). We caution that, as discussed at the end of this section, the Isaacson expression Eq. (10.103) should not be expected to hold in a nonlinear completion of the theory over arbitrary backgrounds.

Decomposing the metric components into plane-waves, the above expressions imply that the energy density $\rho_{\mathrm{GW}} \equiv T_{00}$, in some frame, may be written as in Eq. (10.36) with $Q^{abcd} = Q^{abcd}_{\mathrm{GR}}$. Breaking up the Fourier amplitudes into polarizations and applying all the usual assumptions (*i*)–(*vii*) about the background, it may then be shown that we can use Eq. (10.19) to write the energy density in terms of the polarization PSDs as (Appendix 10.8.2):

$$\rho_{\mathrm{GW}} = \frac{\pi c^2}{4G} \int_0^\infty \sum_A \lambda_A(f) S_A(f)\, f^2 \mathrm{d}f , \tag{10.104}$$

where the sum is over the five linear polarizations $A \in \{+, \times, \mathrm{x}, \mathrm{y}, \mathrm{l}\}$ as they appear in Eq. (10.99), and we have defined

$$\lambda_A(f) \equiv \begin{cases} 1 & \text{if } A = +, \times, \\ \alpha^2 & \text{if } A = \mathrm{x}, \mathrm{y}, \\ \frac{3}{4}\alpha^4 & \text{if } A = \mathrm{l}\,. \end{cases} \tag{10.105}$$

Clearly, higher powers of α will be strongly suppressed in the limit of small mass we are working in, but leave them in for now nonetheless. Note again that we have assumed that the polarization amplitudes of Eq. (10.99) are statistically independent because they are the fundamental degrees of freedom that diagonalize the kinetic matrix of the theory.

With this expression for ρ_{GW} in hand, the definition of the fractional energy density spectrum, Eq. (10.29), then implies that

$$\Omega_{\mathrm{GW}}(f) = \frac{2\pi^2|f|^3}{3H_0^2} \sum_A \lambda_A(f) S_A(f)\,, \tag{10.106}$$

and, as we have done in previous sections, we may call each summand in this equation $\Omega_A(f)$, with $\Omega_{\mathrm{GW}}(f) = \sum_A \Omega_A(f)$, so that we can write the corresponding polarization spectral density as

$$S_A(f) = \frac{3H_0^2}{2\pi^2|f|^3} \lambda_A^{-1}(f)\, \Omega_A(f)\,. \tag{10.107}$$

We now want to relate the GW energy density to the cross-correlation of the outputs of two interferometric detectors. Instead of Eq. (10.99), we would like to be able to write the GW as a purely spatial metric perturbation ($h_{0\nu} = 0$) in arbitrary frames (i.e. without the need for special boosts). This is so we can have the perturbation be purely spatial in the proper frame of the detector, which would then allow us to use our usual expression for the detector tensor, Eq. (10.23), when computing the output of a measurement.

In GR, the required gauge freedom is afforded by diffeomorphism invariance, which is not directly available to us in massive gravity [105, 232]. However, we may circumvent this restriction by introducing auxiliary fields into the action, designed to reintroduce gauge freedom to the theory (the so-called *Stückelberg trick*). We would then obtain a generalized version of massive gravity that is invariant under infinitesimal coordinate transformations, and which reduces to the usual theory after fixing to a particular gauge (see Appendix 10.8.1).

We refer to the gauge that returns the FP action of Eq. (10.93) as the *unitary* gauge, as opposed to the *synchronous* gauge, in which the metric perturbation can take a purely spatial form without special boosts. In this gauge, a measurement via an interferometric detector in the small-antenna limit can be represented by the double contraction of the metric with the detector tensor of Eq. (10.23), and the metric perturbation can be decomposed as in Eq. (10.6), as explained in Sec. 10.2.3.

Unfortunately, the synchronous polarizations will *not* be statistically independent in the linear basis of Eq. (10.6), which is commonly used in data analysis (e.g. [99]). In fact, the 6 polarization amplitudes in the synchronous gauge, $\underset{\sim}{h}_{\underset{\sim}{A}}$, can be obtained from the 5 in the unitary gauge of Eq. (10.99), h_B, via a (polarization-basis-dependent) transformation matrix, $M_{\underset{\sim}{A}}{}^{B}$, given by

$$(M_{\underset{\sim}{A}}{}^{B}) \equiv \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & \alpha^2 & 0 & 0 \\ 0 & 0 & 0 & \alpha^2 & 0 \\ 0 & 0 & 0 & 0 & -\frac{1}{2}\alpha^2 \\ 0 & 0 & 0 & 0 & \alpha^4 \end{pmatrix}, \tag{10.108}$$

so that $\underset{\sim}{h}_{\underset{\sim}{A}} = M_{\underset{\sim}{A}}{}^{B} h_B$, and where $\underset{\sim}{A} \in \{+, \times, \mathrm{x}, \mathrm{y}, \mathrm{b}, \mathrm{l}\}$ indexes *synchronous* polarization amplitudes $\underset{\sim}{h}_{\underset{\sim}{A}} \in \{\underset{\sim}{h}_{+}, \underset{\sim}{h}_{\times}, \underset{\sim}{h}_{\mathrm{x}}, \underset{\sim}{h}_{\mathrm{y}}, \underset{\sim}{h}_{\mathrm{b}}, \underset{\sim}{h}_{\mathrm{l}}\}$, while $B \in \{+, \times, \mathrm{x}, \mathrm{y}, \mathrm{l}\}$ indexes *unitary* polarization amplitudes $h_B \in \{h_+, h_\times, h_\mathrm{x}, h_\mathrm{y}, h_\mathrm{l}\}$. Had we started with a basis for the unitary metric components different than Eq. (10.99), all our results would still apply after redefining $M_{\underset{\sim}{A}}{}^{B}$ accordingly. We provide an explicit expression for $\underset{\sim}{h}_{ab}$ in terms of the unitary amplitudes in Eq. (10.175) in Appendix 10.8.1. The fact that the 5 unitary amplitudes determine 6 synchronous amplitudes makes it immediately clear that the latter are not statistically independent.

Taking advantage of the synchronous gauge to compute detector responses and taking the unitary polarizations to be uncorrelated, the cross-correlation of two detector outputs may be written directly in terms of the fractional energy spectrum for each unitary polarization via Eq. (10.107),

$$\begin{aligned} \left\langle \tilde{h}^*_I(f) \tilde{h}_{I'}(f') \right\rangle = \frac{3H_0^2}{4\pi^2 |f|^3} \delta(f - f') \sum_B \Omega_B(f) \\ \times \lambda_B^{-1}(f) M_{\underset{\sim}{A}B}(f)\, M_{\underset{\sim}{A}'B}(f)\, \underset{\sim}{\Gamma}{}^{\underset{\sim}{A}\underset{\sim}{A}'}_{II'}(f), \end{aligned} \tag{10.109}$$

with $\lambda_B(f)$ as in Eq. (10.105), $M_{\underset{\sim}{A}B}$ as in Eq. (10.108), and $\underset{\sim}{\Gamma}{}^{\underset{\sim}{A}\underset{\sim}{A}'}_{II'}(f)$ the generalized overlap reduction functions for the synchronous polarizations. These functions

are defined as in Eq. (10.28), with a delay factor corresponding to $v_p = c/\beta$ [cf. Eq. (10.98)], i.e.

$$\xi^{AA'}_{II'} = \frac{\Delta x_{II'}}{c}\sqrt{1-\alpha^2} \approx \frac{\Delta x_{II'}}{c}\left(1 - \frac{1}{2}\alpha^2\right), \tag{10.110}$$

after expanding for small α. The resulting overlap-reduction functions will not be the same (even ignoring differences in normalization) as those used in existing stochastic searches beyond GR [99, 117], because those assume $v_p = c$. However, we should expect that to be a good approximation as long as the extra delay in the time of flight due to the nonvanishing mass, $\delta\xi^{AA'}_{II'} \equiv -\frac{1}{2}\alpha^2\Delta x_{II'}/c$, is small with respect to the timing accuracy of the instruments. For a treatment of overlap-reduction functions without ignoring this correction, see [213].

Regardless of whether we neglect dispersive corrections to the overlap-reduction functions or not, it turns out that, for differential-arm detectors, we have that

$$\Gamma^{\underset{\sim}{A}\underset{\sim}{A}'}_{II'} = \begin{cases} \left(2\delta_{\underset{\sim}{A}\underset{\sim}{A}'} - 1\right)\Gamma^{\mathrm{l}}{}_{II'} & \text{if } \underset{\sim}{A} \text{ or } \underset{\sim}{A}' \text{ in } \{\mathrm{b},\mathrm{l}\}, \\ \delta_{\underset{\sim}{A}\underset{\sim}{A}'}\Gamma^{\underset{\sim}{A}}{}_{II'} & \text{otherwise,} \end{cases} \tag{10.111}$$

as long as differences in the phase velocities of different polarizations are negligible (which is exactly the case for the Fierz-Pauli theory). This relation may be used to put our result of Eq. (10.109) in the form of Eq. (10.39) with

$$\Xi_A(f) \approx \begin{cases} 1 & \text{if } A \text{ in } \{+,\times\}, \\ \alpha^2 & \text{if } A \text{ in } \{\mathrm{x},\mathrm{y}\}, \\ \frac{1}{3}\left(2\alpha^2+1\right)^2 & \text{if } A = \mathrm{l}. \end{cases} \tag{10.112}$$

plus terms of order α^6 and higher. Here, $\Xi_A(f)$ goes smoothly to the GR limit as $\alpha \to 0$ (vanishing graviton mass) for the tensor and vector modes. However, notice that $\Xi_l(f) \to \frac{1}{3}$ (rather than vanishing) in this same limit. This is reminiscent of the vDVZ (van Dam, Veltman, Zakharov) discontinuity [293, 294] (see also [295] for a similar effect, and [232] for more discussion). For interesting details on the derivation of Eqs. (10.109)–(10.112), we refer the reader to Appendix 10.8.3.

Relation to nonlinear massive gravity

There is no problem in thinking of the action Eq. (10.93) as describing a linear spin-2 field h_{ab} on a curved background g_{ab}. However, if we want h_{ab} to represent

metric fluctuations of the gravitational field, the theory must have a nonlinear completion, which is known to have several problems (see e.g. [105, 232] for more discussion). One which we have already mentioned [below Eq. (10.112)] is the vDVZ discontinuity, by which the limit of vanishing graviton mass $\mu \to 0$ does not recover GR (e.g., the scalar degree of freedom does not decouple).

Another major problem is the Boulware-Deser ghost [296], which must be excised order-by-order in the graviton self-interaction. Controlling this ghost degree of freedom to all orders is possible with a specific set of self-interactions, known as de Rham-Gabadadze-Tolley (dRGT) massive gravity [297]. This has been extended to a theory of two interacting metrics by Hassan and Rosen [221], which has dRGT as a careful scaling limit. Bigravity propagates one massive and one massless spin-2 field (7 total degrees of freedom), whereas taking the dRGT limit eliminates the massless mode (leaving only 5 dynamical degrees of freedom, as in the linearized theory).

Indeed when expanded about a Minkowski background to linear order (quadratic in the Lagrangian), dRGT agrees with Fierz-Pauli theory. This might lead one to believe that the preceding FP analysis can be directly lifted to dRGT, or even to bigravity, but this conclusion is unwarranted. The quadratic Lagrangian about nontrivial background-field configurations [298–302] can look rather different from the simple FP Lagrangian.[6]

In fact, the difference from the FP Lagrangian is crucial for the health of such theories, because otherwise the nonlinear theories would also exhibit problematic phenomenology, like the vDVZ discontinuity. However, healthy nonlinear massive gravity theories are protected from vDVZ phenomenology by the Vainshtein screening mechanism [303]. The Vainshtein mechanism leads to a nontrivial, nonlinear field configuration (like a condensate) with a new length scale, the Vainshtein radius. Within this radius, the *effective* couplings for the massive degree of freedom can be very different from what is seen when expanded about the Minkowski background, thus reverting to the phenomenology of general relativity.

In short, the ESET for nonlinear massive gravity on a general background (e.g. one

[6]For a special subclass of "proportional" background configurations in bigravity [299, 302], two linear combinations of the two metrics' perturbations can be combined into the massless and massive eigenstates which diagonalize the kinetic matrix of the quadratic Lagrangian. In this case, the massive mode does have a FP Lagrangian. However, this is likely a special case—as far as we have been able to discern, the transformation to the mass eigenstates has not been performed for a more general background.

exhibiting Vainshtein screening) should be considered an open problem. It seems unlikely that the FP result lifts to the general massive gravity result.

10.4 Conclusion

The detection of a stochastic gravitational-wave background will provide a unique opportunity to study the properties of gravitational waves as they propagate through cosmological distances, and will thus be an invaluable tool to study extensions of general relativity. Properly interpreting the theoretical implications of such a detection will require a detailed understanding of the assumptions that go into the usual searches for a stochastic background, and how the measurement process might be modified in theories beyond general relativity. Towards that goal, in this Chapter we have laid out the formalism underlying searches for stochastic signals in a generic fashion that makes it easily applicable to a large family of theories. We have also surveyed the standard set of assumptions that go into these searches, evaluating their generic applicability, or lack thereof.

First and foremost, we find that most existing treatments of stochastic backgrounds beyond GR fail to consider possible modifications to the effective stress-energy carried by a gravitational wave of a given amplitude and frequency [107, 112, 213, 214, 216–218]. This is important because the *goal* of searches for stochastic backgrounds, within GR or beyond, is precisely to measure the amount of energy that exists in the form of stochastic gravitational waves. Accordingly, data analysis strategies tend to be parametrized directly in terms of an effective energy spectrum, Eq. (10.29). However, this is only possible if one knows the relation between the energy density and the observables at the detector (e.g. the cross-correlation of strain detector outputs)—this will depend on the specific structure of the underlying theory of gravity, and in general need not be the same as in GR. Therefore, parametrizing model-independent searches for backgrounds beyond GR as traditionally done will result in the use of a quantity that should *not* generally be interpreted as the GW energy density, and may thus lead to incorrect comparisons between theory and experiment. Instead, we find it advisable to parametrize theory-agnostic searches using the power spectrum of polarization amplitudes, Eq. (10.16), which have a (mostly) model-independent interpretation. One can always translate amplitudes into effective energies for any specific theory, as sketched in Sec. 10.2.4.

We also reviewed the standard set of simplifying assumptions that the stochastic background is (*i*) Gaussian, (*ii*) ergodic, and (*iii*) stationary, with no correlation

between amplitudes from different (*iv*) sky locations or (*v*) polarizations, and with (*vi*) equipartition of power across polarizations; and also, commonly (although not universally) assumed to be (*vii*) isotropic. While we find that the first four of these premises are generally applicable beyond GR, the same is not true for the rest—this is without considering changes to the potential sources of the background in beyond-GR theories, which may themselves break more of the assumed symmetries. In particular, it is not reasonable to always assume that the usual linear GW polarization amplitudes of Eq. (10.6) will be statistically independent and have well-defined phase velocities, as this will not be true unless the chosen polarization basis diagonalizes the kinetic matrix of the underlying theory of gravity. Similarly, one should be careful in assuming that power will be equipartitioned among polarizations, even for modes with the same spin-weight, as parity-asymmetric theories may predict differences in the generation and propagation of modes with different helicities. Deviations from isotropy should be expected in theories with intrinsically preferred frames.

Finally, we have provided specific examples of beyond-GR theories in which these traditional assumptions break down, and in which the GR expression for the stress-energy of a gravitational wave may receive a correction: Chern-Simons gravity, scalar-tensor theories, and massive gravity. For all these theories, we find that the cross-correlation of the outputs of two ideal differential-arm detectors can be written in terms of the effective GW stress-energy as in Eq. (10.39), with different $\Xi(f)$ factors encoding how each theory departs from GR. This set of examples is not intended to be exhaustive, but merely to show that it is possible to construct viable theories that violate standard assumptions in stochastic searches. This will be important in the interpretation of results like [99, 117] once a stochastic signal is detected.

Acknowledgments

The authors would like to thank Laura Bernard, Tom Callister, Claudia de Rham, Kurt Hinterbichler, Andrew Matas, and Andrew Tolley for useful discussions. M.I. is a member of the LIGO Laboratory. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY–0757058. L.C.S. acknowledges the support of NSF grant PHY–1404569 and the Brinson Foundation. This paper carries LIGO Document Number LIGO–P1700234.

10.5 Appendix: Plane-wave decomposition

Begin with our compact expression for the plane-wave expansion of the metric components, Eq. (10.1):

$$h_{ab}(\mathbf{x}) = \frac{1}{2\pi}\int \tilde{h}_{ab}(\mathbf{k})e^{i\mathbf{k}\cdot\mathbf{x}}\,\widetilde{\mathrm{d}\mathbf{k}}, \tag{10.113}$$

with the integral over the four-wave-vector $\mathbf{k}$ as prescribed by our definition of $\widetilde{\mathrm{d}\mathbf{k}}$ in Eq. (10.2),

$$\widetilde{\mathrm{d}\mathbf{k}} \equiv 2c\,\delta(|\vec{k}|^2 - |\vec{k}_\omega|^2)\,|\vec{k}|^{-1}\mathrm{d}\mathbf{k}\,. \tag{10.114}$$

This definition of the four-dimensional Fourier transform is designed to yield Eq. (10.3), and thus follows the convention of recent stochastic GW background literature (e.g. [183, 185, 222]). This choice, however, differs from the Lorentz-invariant measure most common in field theory (see e.g. Eq. (3.18) in [304] or Eq. (4.4) in [305]),

$$\begin{aligned}\widetilde{\mathrm{d}\mathbf{k}}_{\mathrm{QFT}} &= c\delta(|\vec{k}|^2 - |\vec{k}_\omega|^2)\,\mathrm{d}\mathbf{k}/(2\pi)^3 \\ &= |\vec{k}|\widetilde{\mathrm{d}\mathbf{k}}/2/(2\pi)^3\,.\end{aligned} \tag{10.115}$$

Note that this difference in measures results in a difference in conventions for the Fourier amplitudes. Specifically, this means that $\tilde{h}(k) \propto |\vec{k}|\tilde{h}(k)_{\mathrm{QFT}}$ (the factor of proportionality depends on prefactors outside of the integral).

With the help of Eq. (10.114), Eq. (10.113) can be immediately rewritten as an explicit integral over the four-vector $\mathbf{k}$, transforming each component independently,

$$h_{ab}(\mathbf{x}) = \frac{c}{\pi}\int \tilde{h}_{ab}(\mathbf{k})e^{i\mathbf{k}\cdot\mathbf{x}}\,\delta(|\vec{k}|^2 - |\vec{k}_\omega|^2)\,|\vec{k}|^{-1}\mathrm{d}\mathbf{k}\,. \tag{10.116}$$

Here $\vec{k}_\omega \equiv \vec{k}(\omega)$ encodes the functional dependence of $\vec{k}$ on ω imposed by the specific dispersion relation required by the underlying theory of gravity—in GR, this is just the usual demand that $|\vec{k}_\omega| = \omega/c$). For clarity, we may split the 4-vector $\mathbf{k}$ into frequency and spatial $\vec{k}$-vector,

$$\begin{aligned}h_{ab}(t,\vec{x}) = \frac{1}{\pi}\int_{-\infty}^{\infty}\int_{S^2}\int_0^{\infty} &\tilde{h}_{ab}(\omega,\vec{k})\,e^{i(\vec{k}\cdot\vec{x}-\omega t)} \\ &\times\delta(|\vec{k}|^2 - |\vec{k}_\omega|^2)\,|\vec{k}|\,\mathrm{d}|\vec{k}|\,\mathrm{d}\hat{k}\,\mathrm{d}\omega\,,\end{aligned} \tag{10.117}$$

where we have written the spatial three-integral in polar coordinates such that

$$\mathrm{d}\vec{k} = |\vec{k}|^2\,\mathrm{d}|\vec{k}|\,\mathrm{d}\hat{k}\,, \tag{10.118}$$

with angular domain over the 2-sphere, S^2. In this step, we have also used the fact that $|\vec{k}|$ is non-negative by definition to set its integration limits.

We may now use the fact that, for any continuously differentiable function $g(x)$ with real roots x_i,

$$\delta(g(x)) = \sum_i \frac{\delta(x - x_i)}{|g'(x_i)|}, \tag{10.119}$$

to further simplify the integrand to

$$\begin{aligned} h_{ab}(t, \vec{x}) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{S^2} \int_0^{\infty} \tilde{h}_{ab}(\omega, \vec{k})\, e^{i(\vec{k}\cdot\vec{x} - \omega t)} \\ \times \frac{1}{|\vec{k}_\omega|} \delta(|\vec{k}| - |\vec{k}_\omega|)\, |\vec{k}|\, \mathrm{d}|\vec{k}|\, \mathrm{d}\hat{k}\, \mathrm{d}\omega, \end{aligned} \tag{10.120}$$

where the integration limits have allowed us to ignore the negative root, $|\vec{k}| = -|\vec{k}_\omega|$. It is now straightforward to carry out the integral over the norm $|\vec{k}|$ to obtain:

$$h_{ab}(t, \vec{x}) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{S^2} \tilde{h}_{ab}(\omega, \vec{k}) e^{i(\vec{k}\cdot\vec{x} - \omega t)} \mathrm{d}\hat{k} \mathrm{d}\omega, \tag{10.121}$$

where now $\vec{k}$ is necessarily on shell ($|\vec{k}| = |\vec{k}_\omega|$). Writing this in terms of $f = \omega/2\pi$, $\hat{n} = -\hat{k}$ and $v_{\rm p} \equiv |\vec{k}/\omega|^{-1}$, we immediately recover Eq. (10.3), as promised,

$$h_{ab}(t, \vec{x}) = \int_{-\infty}^{\infty} \int_{\rm sky} \tilde{h}_{ab}(k, \hat{n}) e^{-2\pi i f(t + \hat{n}\cdot\vec{x}/v_{\rm p})} \mathrm{d}\hat{n} \mathrm{d}f, \tag{10.122}$$

thus justifying the second equality in Eq. (10.2),

$$\widetilde{\mathrm{d}\mathbf{k}} = \mathrm{d}\omega\, \mathrm{d}\hat{n}\,. \tag{10.123}$$

10.6 Appendix: Correlation and spectral density

We will reproduce the standard result that assumptions (*iii*) of stationarity and (*iv*) of uncorrelated sky locations allow us to write the cross-correlation of the Fourier amplitudes as in Eq. (10.16),

$$\left\langle \tilde{h}^*_A(\mathbf{k})\, \tilde{h}_{A'}(\mathbf{k}') \right\rangle = \frac{1}{2} \delta(f - f')\, \delta(\hat{n} - \hat{n}')\, S_{AA'}(\mathbf{k}), \tag{10.124}$$

with $S_{AA'}$ the cross-power spectral density of stochastic signals of polarizations A and A'.

The second delta function in Eq. (10.124) is just a direct statement of assumption (*iv*), so focus on the rest of the equation by suppressing the dependence on $\hat{n}$. We are

then left with simple one-dimensional Fourier transforms in the expression for the cross-correlation,

$$\left\langle \tilde{h}_A^*(f)\tilde{h}_{A'}(f')\right\rangle = \left\langle \int h_A(t)e^{2\pi i f t}\mathrm{d}t \times \int h_{A'}(t')e^{-2\pi i f' t'}\mathrm{d}t' \right\rangle . \tag{10.125}$$

Defining $\tau \equiv t' - t$, this can be put in the form:

$$\left\langle \tilde{h}_A^*(f)\tilde{h}_{A'}(f')\right\rangle = \int\int \langle h_A(t)h_{A'}(t+\tau)\rangle \times e^{-2\pi i f'\tau}e^{2\pi i(f-f')t}\mathrm{d}\tau\mathrm{d}t\,. \tag{10.126}$$

Now note that the term in brackets is simply the correlation of $h_A(t)$ and $h_{A'}(t')$, which by assumption of stationarity depends only on the time difference τ, i.e.

$$\langle h_A(t)h_{A'}(t+\tau)\rangle = \langle h_A(0)h_{A'}(\tau)\rangle\ , \tag{10.127}$$

where we have set $t = 0$ for concreteness. We may therefore carry out the integral over t to obtain

$$\left\langle \tilde{h}_A^*(f)\tilde{h}_{A'}(f')\right\rangle = \delta(f-f')\int \langle h_A(0)h_{A'}(\tau)\rangle\, e^{-2\pi i f\tau}\mathrm{d}\tau\,. \tag{10.128}$$

Now, the Wiener-Khinchin theorem [224] guarantees that, if the cross-correlation is continuous, we can always define a function of frequency to give the Fourier transform of the cross-correlation (the integral above); that can be taken as the *definition* of the one-sided cross-power spectral density,

$$S_{XY}(f) \equiv 2\int \langle X(0)Y(\tau)\rangle\, e^{-2\pi i f\tau}\mathrm{d}\tau\,, \tag{10.129}$$

for any two stationary random processes, $X(t)$ and $Y(t)$, and where the prefactor is chosen so that $S(f) \equiv S(|f|)$ is the *one-sided* spectral density. All this means is that we may write

$$\left\langle \tilde{h}_A^*(f)\,\tilde{h}_{A'}(f')\right\rangle = \frac{1}{2}\delta(f-f')S_{AA'}(f), \tag{10.130}$$

or, restoring the $\hat{n}$ dependence,

$$\left\langle \tilde{h}_A^*(\mathbf{k})\,\tilde{h}_{A'}(\mathbf{k}')\right\rangle = \frac{1}{2}\delta(f-f')\delta(\hat{n}-\hat{n}')S_{AA'}(\mathbf{k})\,. \tag{10.131}$$

10.7 Appendix: Scalar-tensor computations

Here we provide details on the computations of the ESET and correlation functions in Brans-Dicke gravity (Sec. 10.3.3). In order to do so, first consider the transformations between the Jordan and Einstein frames. By definition of the Einstein frame, in a generic scalar-tensor theory these can be written as (e.g. Eqs. (34)–(36) in [84])

$$g_{ab} \equiv A^2(\varphi)\underset{\sim}{g}_{ab}\,, \tag{10.132a}$$

$$\phi \equiv A^{-2}(\varphi)\,, \tag{10.132b}$$

for some auxiliary function $A(\varphi)$. We can then use this to define the coupling $\alpha(\varphi)$ as

$$\alpha(\varphi) \equiv \frac{\mathrm{d}\ln A(\varphi)}{\mathrm{d}\varphi}\,. \tag{10.133}$$

To recover the Brans-Dicke theory, we simply expand this coupling to linear order by setting

$$\alpha(\varphi) = \alpha_0 \equiv (2\omega_{\rm BD}+3)^{-1/2}\,, \tag{10.134}$$

so that $\ln A(\varphi) = \alpha_0(\varphi-\varphi_0)$ for some fiducial value φ_0, and Eqs. (10.132) become (Jordan to Einstein)

$$g_{ab} = e^{2\alpha_0(\varphi-\varphi_0)}\underset{\sim}{g}_{ab}\,, \tag{10.135a}$$

$$\phi = e^{-2\alpha_0(\varphi-\varphi_0)}\,. \tag{10.135b}$$

For later convenience, define $\delta\varphi \equiv \varphi-\varphi_0$ and rescale the Jordan field by letting $\phi \to \phi/\phi_0$ for some background value ϕ_0. After doing so, Eqs. (10.135) imply (Einstein to Jordan)

$$\underset{\sim}{g}_{ab} = \frac{\phi}{\phi_0}g_{ab}\,, \tag{10.136a}$$

$$\delta\varphi = -\frac{\ln(\phi/\phi_0)}{2\alpha_0}\,. \tag{10.136b}$$

With the above notation in place, let us perturb the two metrics and scalar fields to first order, and then obtain the relationship between the perturbations in the two frames. Letting $g_{ab} \to g_{ab} + \underset{\sim}{h}_{ab}$ and $\phi \to \phi_0 + \delta\phi$ in the Jordan frame, and $\underset{\sim}{g}_{ab} \to \underset{\sim}{g}_{ab} + \underset{\sim}{h}_{ab}$ and $\varphi \to \varphi_0 + \delta\varphi$ in the Einstein frame, we can then apply the transformations from Eq. (10.136) to write

$$\begin{aligned}\underset{\sim}{g}_{ab} + \underset{\sim}{h}_{ab} &= \phi_0^{-1}\left(\phi_0+\delta\phi\right)\left(g_{ab}+h_{ab}\right)\\ &\approx g_{ab} + \left(h_{ab} + g_{ab}\delta\phi/\phi_0\right).\end{aligned} \tag{10.137}$$

Collecting terms of the same order, this implies that, to first order in the perturbations,

$$g_{ab} = \underset{\sim}{g}_{ab}\,, \tag{10.138a}$$
$$h_{ab} = \underset{\sim}{h}_{ab} + 2\alpha_0\delta\varphi\underset{\sim}{g}_{ab}\,, \tag{10.138b}$$
$$\Phi = 2\alpha_0\delta\varphi\,, \tag{10.138c}$$

where we have defined $\Phi \equiv -\delta\phi/\phi_0$ for convenience. Using this definition to replace the second expression by $\underset{\sim}{h}_{ab} = h_{ab} - \Phi g_{ab}$, it becomes clear these are Eqs. (10.80) provided in the main text.

10.7.1 Effective stress-energy tensor

We wish to compute the effective GW stress energy in the Einstein frame. We will do so by taking advantage of the gauge proposed in [276], in which the trace-reversed Einstein-frame perturbation, $\bar{\underset{\sim}{h}}_{ab}$, satisfies

$$\bar{\underset{\sim}{h}} = 2\Phi\,, \tag{10.139a}$$
$$\nabla^a\bar{\underset{\sim}{h}}_{ab} = 0\,, \tag{10.139b}$$

and follows simple free-wave equations of motion,

$$\Box\bar{\underset{\sim}{h}}_{ab} = 0\,, \tag{10.140a}$$
$$\Box\Phi = 0\,. \tag{10.140b}$$

In this gauge, the Einstein-frame trace-reversed metric perturbation is equal to the regular (non-trace-reversed) perturbation in the Jordan frame: $\bar{\underset{\sim}{h}}_{ab} = h_{ab}$. Thus, $\bar{\underset{\sim}{h}}_{ab}$ may be decomposed into synchronous polarizations as in Eq. (10.81).

To obtain an expression for the GW stress energy in the Einstein frame, we may follow the procedure outlined in Sec. 10.2.4 starting from the action of Eq. (10.76). Perturbing the metric and scalar as described above, and discarding terms higher than second order, we may obtain the quadratic Lagrangian density corresponding to $\mathcal{S}^{(2)}$ in Eq. (10.34),

$$\underset{\sim}{\mathcal{L}}^{(2)} = \underset{\sim}{\mathcal{L}}^{(2)}_{\rm EH} + \kappa\sqrt{-\underset{\sim}{g}}\left[-2\underset{\sim}{g}^{ab}\nabla_a(\delta\varphi)\,\nabla_b(\delta\varphi)\right], \tag{10.141}$$

where $\underset{\sim}{\mathcal{L}}^{(2)}_{\rm EH}$ is the Einstein-Hilbert piece of Eq. (10.42), but in terms of $\underset{\sim}{g}, \underset{\sim}{h}$. The variation of this quantity with respect to $\bar{\underset{\sim}{h}}^{ab}$ and $\delta\varphi$ will lead to the ESET per Eq. (10.31). This will be given by a contribution from the Einstein-Hilbert part of

the action (the Ricci terms above), and another from the rest. We will call those two terms $\underset{\sim}{T}^{(\mathrm{EH})}_{ab}$ and $\Delta\underset{\sim}{T}^{(\mathrm{ST})}_{ab}$ respectively, so that $\underset{\sim}{T}^{(\mathrm{ST})}_{ab} = \underset{\sim}{T}^{(\mathrm{EH})}_{ab} + \Delta\underset{\sim}{T}^{(\mathrm{ST})}_{ab}$.

Focus first on the EH term. This will *not* be the identical to Eq. (10.43) in GR, because there will be an extra contribution from the nonvanishing trace of $\bar{\underset{\sim}{h}}_{ab}$, Eq. (10.139a). To compute it, we may take a shortcut and begin from an expression obtained MacCallum and Taub for the effective EH quadratic Lagrangian contributing to the GW stress-energy far away from the source [262]. The corresponding stress-energy tensor can be written as

$$\begin{aligned}\underset{\sim}{T}^{(\mathrm{EH})}_{ab} = \kappa\Bigg\langle\!\!\Bigg\langle &\frac{1}{2}\nabla_a\bar{\underset{\sim}{h}}^{cd}\nabla_b\bar{\underset{\sim}{h}}_{cd} - \nabla_c\bar{\underset{\sim}{h}}_{da}\nabla^c\bar{\underset{\sim}{h}}_b{}^d - \frac{1}{4}\nabla_a\bar{\underset{\sim}{h}}\nabla_b\bar{\underset{\sim}{h}} \\ &+ \frac{1}{2}\nabla_c\bar{\underset{\sim}{h}}_{ab}\nabla^c\bar{\underset{\sim}{h}} + \underset{\sim}{g}_{ab}\left(\frac{1}{2}\nabla_e\bar{\underset{\sim}{h}}^{fc}\nabla_f\bar{\underset{\sim}{h}}^e{}_c\right. \\ &\left.-\frac{1}{4}\nabla_e\bar{\underset{\sim}{h}}_{cd}\nabla^e\bar{\underset{\sim}{h}}^{cd} + \frac{1}{8}\nabla_e\bar{\underset{\sim}{h}}\nabla^e\bar{\underset{\sim}{h}}\right)\Bigg\rangle\!\!\Bigg\rangle . \end{aligned} \tag{10.142}$$

This expression is valid whenever separation of length-scales allows for a clear definition of the waves over some background. In GR, application of the equations of motion in a transverse-traceless gauge reduces Eq. (10.142) to the Isaacson formula, Eq. (10.43). We proceed similarly here but keeping the trace, using Eqs. (10.140) and (10.139).

First note that the second term in Eq. (10.142) may be re-written by integrating by parts "under the average." This is because the Brill-Hartle average of a total derivative is smaller by a factor of order $O(\lambda_{\mathrm{GW}}/L_{\mathrm{ave}})$ than non-vanishing averages, where L_{ave} is the averaging length scale (see e.g. Sect. IIA in [219]). This then implies that

$$\langle\!\langle -\nabla_a\bar{\underset{\sim}{h}}_{bc}\nabla^a\bar{\underset{\sim}{h}}_d{}^b\rangle\!\rangle = \langle\!\langle\left(\nabla^a\nabla_a\bar{\underset{\sim}{h}}_{bc}\right)\bar{\underset{\sim}{h}}_d{}^b\rangle\!\rangle\left[1 + O\left(\frac{\lambda_{\mathrm{GW}}}{L_{\mathrm{ave}}}\right)\right] . \tag{10.143}$$

Therefore, the second term in Eq. (10.142) vanishes via the equations of motion [Eq. (10.140a)], up to this order. The same logic may be applied to all terms in the second and third lines of Eq. (10.142), which will vanish due to Eq. (10.140a) or Eq. (10.139b). We are then only left with the first and third terms in Eq. (10.142). The first term is just the same quadratic contribution that appears in Eq. (10.43) for GR. Meanwhile, the third term involves the trace of $\bar{\underset{\sim}{h}}_{ab}$, and may thus be written in terms of the scalar field using Eq. (10.139a). The contribution of the Einstein-Hilbert part of the action to the ESET, Eq. (10.142), in ST gravity then reduces to

$$\underset{\sim}{T}^{(\mathrm{EH})}_{ab} = \frac{1}{2}\kappa\langle\!\langle\nabla_a\bar{\underset{\sim}{h}}^{cd}\nabla_b\bar{\underset{\sim}{h}}_{cd}\rangle\!\rangle - \kappa\langle\!\langle\nabla_a\Phi\nabla_b\Phi\rangle\!\rangle . \tag{10.144}$$

Switch now to the contribution from the kinetic term of the scalar field, $\Delta\underset{\sim}{T}^{(\mathrm{ST})}_{ab}$. This will be obtained from the corresponding part of the quadratic Lagrangian of Eq. (10.141), namely $\Delta\underset{\sim}{\mathcal{L}}^{(2)} \equiv -2\sqrt{-\underset{\sim}{g}}\kappa\underset{\sim}{g}^{ab}\nabla_a(\delta\varphi)\nabla_b(\delta\varphi)$. The variation of this quantity may be written as

$$\frac{\delta\Delta\underset{\sim}{\mathcal{L}}^{(2)}}{\delta\underset{\sim}{g}^{ab}} = \kappa\sqrt{-\underset{\sim}{g}}\left[\underset{\sim}{g}_{ab}\underset{\sim}{g}^{cd} - \delta^{c}{}_{(a}\delta^{d}{}_{b)}\right]\nabla_c\delta\varphi\,\nabla_d\delta\varphi\,, \tag{10.145}$$

using the usual fact that $\delta\sqrt{-\underset{\sim}{g}} = -\sqrt{-\underset{\sim}{g}}\,\underset{\sim}{g}_{ab}\delta\underset{\sim}{g}^{ab}/2$, and explicitly symmetrizing the variation of the metric. Therefore, Eq. (10.31) implies that

$$\begin{aligned}\Delta\underset{\sim}{T}^{(\mathrm{ST})}_{ab} &= 2\kappa\left\langle\!\!\left\langle\left(-\underset{\sim}{g}_{ab}\underset{\sim}{g}^{\alpha\beta} + 2\delta^{\alpha}{}_{a}\delta^{\beta}{}_{b}\right)\nabla_\alpha\delta\varphi\,\nabla_\beta\delta\varphi\right\rangle\!\!\right\rangle \\ &= 4\kappa\langle\!\langle\nabla_a\delta\varphi\,\nabla_b\delta\varphi\rangle\!\rangle\,,\end{aligned} \tag{10.146}$$

where one the first term vanished due to Eq. (10.140b), by integration by parts under averaging as before.

We may now write an expression for the total effective stress energy of a scalar-tensor GW in the Einstein frame:

$$\begin{aligned}\underset{\sim}{T}^{(\mathrm{ST})}_{ab} &= \underset{\sim}{T}^{(\mathrm{EH})}_{ab} + \Delta\underset{\sim}{T}^{(\mathrm{ST})}_{ab} \\ &= \kappa\left\langle\!\!\left\langle\frac{1}{2}\nabla_a\bar{\underset{\sim}{h}}^{cd}\nabla_b\bar{\underset{\sim}{h}}_{cd} + \left(\alpha_0^{-2} - 1\right)\nabla_a\Phi\nabla_b\Phi\right\rangle\!\!\right\rangle,\end{aligned} \tag{10.147}$$

where we have used the fact that $\delta\varphi = \Phi/(2\alpha_0)$ to first order, Eq. (10.138c). This may also be written in terms of the Brans-Dicke parameter using the definition of α_0, Eq. (10.134), to obtain our final result presented in Eq. (10.86).

10.7.2 Energy density spectrum

Taking the time-time component of Eq. (10.86) and assuming (*ii*) ergodicity, we immediately obtain an expression for ρ_{GW} in a local Lorentz frame ($\underset{\sim}{g}_{ab} = \eta_{ab}$) from Eq. (10.86),

$$\rho_{\mathrm{GW}} = \frac{\kappa}{2c^2}\left[\left\langle\partial_t\bar{\underset{\sim}{h}}^{ij}\partial_t\bar{\underset{\sim}{h}}_{ij}\right\rangle + 4(\omega_{\mathrm{BD}}+1)\langle\partial_t\Phi\partial_t\Phi\rangle\right], \tag{10.148}$$

or equivalently, because $\Phi = \underset{\sim}{g}^{ab}\underset{\sim}{h}_{ab}/2$ by Eq. (10.81),

$$\begin{aligned}\rho_{\mathrm{GW}} = \frac{\kappa}{2c^2}&\left[g^{ik}g^{jl} + (\omega_{\mathrm{BD}}+1)\,g^{ij}g^{kl}\right] \\ &\times\left\langle\partial_t\bar{\underset{\sim}{h}}_{ij}\partial_t\bar{\underset{\sim}{h}}_{kl}\right\rangle.\end{aligned} \tag{10.149}$$

Expanding out $\bar{\underset{\sim}{h}}_{ij}$ into plane-waves, $\rho_{\rm GW}$ can then be put in the form of Eq. (10.36) with Q^{abcd} as in Eq. (10.87) of the main text.

For convenience, denote each of the two terms in Eq. (10.148) $\rho_{\rm EH}$ and $\rho_{\rm ST}$ respectively, so that $\rho_{\rm GW} = \rho_{\rm EH} + \rho_{\rm ST}$. Making use of all the usual assumptions (*i*)–(*vii*) about the background, we can use Eq. (10.19) to write $\rho_{\rm EH}$ in the same form as Eq. (10.50) in GR,

$$\rho_{\rm EH} = \frac{\pi c^2}{4G} \sum_A \int_0^\infty S_A(f) f^2 \, \mathrm{d}f \,, \tag{10.150}$$

except that now the sum is over $A \in \{+, \times, \mathrm{b}\}$ with $S_{\rm b} = S_\Phi$ by Eq. (10.81). For $\rho_{\rm ST}$, a similar derivation to the one for $\rho_{\rm EH}$ gives the analogous result that

$$\rho_{\rm ST} = \frac{\pi c^2}{4G} (2 + 2\omega_{\rm BD}) \int_0^\infty S_{\rm b}(f) f^2 \, \mathrm{d}f \,. \tag{10.151}$$

Adding both contributions together, we may then write the total energy spectrum compactly as we did in Eqs. (10.89) and (10.90) in the main text.

10.8 Appendix: Massive gravity computations

Here we provide more details for the computation of the GW stress-energy density and correlation functions presented in Sec. 10.3.4. In Appendix 10.8.1 we derive the expressions for the unitary and synchronous metric components, presented respectively in Eqs. (10.99) and (10.108) in the main text. In Appendix 10.8.2, we obtain an expression for the ESET in Fierz-Pauli massive gravity, and one for the energy density $\rho_{\rm GW}$ in terms of the unitary PSDs, making use of statistical assumptions about the background. Finally, in Appendix 10.8.3 we compute an expression for the cross-correlation of the output of two differential-arm detectors in the form of Eq. (10.39). We will make repeated use of the massive Klein-Gordon equation of motion of Eq. (10.95), as well as the fact that the metric perturbation must be divergenceless, Eq. (10.96a), and traceless, Eq. (10.96b). Throughout this appendix, "massive gravity" refers to the Fierz-Pauli theory of Eq. (10.93).

10.8.1 Polarizations

Unitary gauge

We would like to decompose a massive plane GW into a basis of polarization tensors. In GR, diffeomorphism invariance guarantees that we may always find a gauge in

which the perturbation is purely spatial, as in Eq. (10.6). Although this is not possible in FP gravity, we may still write a generic metric perturbation propagating in the z-direction, $h_{\mu\nu}$, as

$$(h_{\mu\nu}) = \begin{pmatrix} h_{00} & h_{01} & h_{02} & h_{03} \\ h_{10} & h_{\rm b} + h_{+} & h_{\times} & h_{\rm x} \\ h_{20} & h_{\times} & h_{\rm b} - h_{+} & h_{\rm y} \\ h_{30} & h_{\rm x} & h_{\rm y} & h_{\rm l} \end{pmatrix}, \tag{10.152}$$

and then apply the constraints from Eqs. (10.96) to cut down the number of degrees of freedom.

First, for a GW with wave-vector k^a, Eq. (10.96a) implies $k^a h_{ab} = 0$. Thus, picking a frame in which the wave travels in the z-direction,

$$(k^\mu) = \left(\omega/c,\, 0,\, 0,\, |\vec{k}|\right), \tag{10.153}$$

lack of divergence, together with symmetry, must mean

$$h_{0\mu} = h_{\mu 0} = -\beta h_{3\mu} = -\beta h_{\mu 3}\,, \tag{10.154}$$

with β as in Eq. (10.98). Eq. (10.154) also implies that

$$h_{00} = -\beta h_{30} = -\beta h_{03} = \beta^2 h_{33}\,. \tag{10.155}$$

However, tracelessness, Eq. (10.96b), also demands

$$h = \eta^{\mu\nu} h_{\mu\nu} = -h_{00} + 2h_{\rm b} + h_{\rm l} = 0\,. \tag{10.156}$$

Therefore, if we choose to get rid of the time-time degree of freedom by writing $h_{00} = 2h_{\rm b} + h_{\rm l}$, Eq. (10.96) requires

$$2h_{\rm b} + h_{\rm l} = \beta^2 h_{\rm l} \implies h_{\rm b} = \frac{1}{2}\left(\beta^2 - 1\right) h_{\rm l}\,, \tag{10.157}$$

so we will only need one scalar polarization. This could have been anticipated from the fact that a symmetric rank-2 tensor in four dimensions can have at most ten independent components, five of which are necessarily constrained by Eq. (10.96), leaving only five degrees of freedom. These correspond to the five possible helicities of a massive spin-2 particle.

Choosing to work in terms of the longitudinal scalar amplitude, our final expression for the metric perturbation in the *unitary gauge* native to FP gravity is Eq. (10.99),

i.e.

$$(h_{\mu\nu}) = \begin{pmatrix} \beta^2 h_{\rm l} & -\beta h_{\rm x} & -\beta h_{\rm y} & -\beta h_{\rm l} \\ -\beta h_{\rm x} & -\frac{1}{2}\alpha^2 h_{\rm l} + h_+ & h_\times & h_{\rm x} \\ -\beta h_{\rm y} & h_\times & -\frac{1}{2}\alpha^2 h_{\rm l} - h_+ & h_{\rm y} \\ -\beta h_{\rm l} & h_{\rm x} & h_{\rm y} & h_{\rm l} \end{pmatrix}, \tag{10.158}$$

where we have used the fact that $(\beta^2 - 1) = \alpha^2$ by Eq. (10.98).

For later convenience, note that the metric perturbation of Eq. (10.158) satisfies

$$h^{ab}(\mathbf{k}) h'_{ab}(\mathbf{k}') = 2 \sum_A C_A(\omega, \omega')\, h_A(\mathbf{k})\, h'_A(\mathbf{k}'), \tag{10.159}$$

with a sum over unitary polarizations A, and for some frequency-dependent normalization coefficients C_A defined by

$$C_A(\omega, \omega') \equiv \begin{cases} 1 & \text{if } A = +, \times, \\ 1 - \beta\beta' & \text{if } A = \mathrm{x}, \mathrm{y}, \\ \frac{3}{2}\,(1 - \beta\beta')^2 - \frac{1}{2}\,(\beta - \beta')^2 & \text{if } A = \mathrm{l}\,. \end{cases} \tag{10.160}$$

The form of these coefficients should not come as a surprise, since they are just terms of the form $e^A{}_{ab} e_A{}^{ab}$, analogous to those in Eq. (10.14), times extra factors arising from the trace and timelike components of Eq. (10.158). Note that C_A acquires its frequency dependence via β and β', Eq. (10.98).

Synchronous gauge

As mentioned in Sec. 10.2.3, it is easiest to compute the influence of a gravitational wave on a LIGO-style detector in the synchronous gauge, because the coordinate locations of the mirrors do not change in this gauge [220]. However, massive gravity lacks the linearized diffeomorphism freedom needed to transform into synchronous gauge. Fortunately, we circumvent the lack of linearized diffeomorphism invariance in massive gravity by using the *Stückelberg trick*: we can add extra auxiliary fields to write the FP action, Eq. (10.93), as a gauge-fixed version of a gauge invariant theory [306]. After adding the Stückelberg fields, ξ_a, we will have the usual freedom to carry out infinitesimal coordinate transformations,

$$\underset{\sim}{h}_{ab} = h_{ab} + \nabla_a \xi_b + \nabla_b \xi_a\,, \tag{10.161}$$

We will want to choose the fields ξ_a such that we can go from the unitary gauge h_{ab} of Eq. (10.158) to a synchronous gauge $\underset{\sim}{h}_{ab}$ in which $\underset{\sim}{h}_{0\nu} = 0$. To do this, pick the same

frame as before, in which **k** is given by Eq. (10.153), and use linearity to consider the transformation of the degrees of freedom in Eq. (10.158), $\{h_+, h_\times, h_{\rm x}, h_{\rm y}, h_{\rm l}\}$, one by one. Below, we will temporarily let $c = 1$ and $k \equiv |\vec{k}|$ for simplicity, but the final result of Eq. (10.175) is insensitive to this. For simplicity, we also let $g_{ab} = \eta_{ab}$. (For more details on the application of this technique to massive gravity, see e.g. Sect. 2.2.2 in [105] or Sect. IV in [232].)

Because the two tensor degrees of freedom, h_+ and $h_\times$, only appear in the spatial part of Eq. (10.158), these modes already satisfy the synchronous gauge condition. Next consider the vector-x amplitude, $h_{\rm x}$: to determine the transformation that would make its contributions to time-like components in Eq. (10.158) vanish, suppose the unitary perturbation is given simply by

$$\left(h_{\mu\nu}\right) = \begin{pmatrix} 0 & -\beta h_{\rm x} & 0 & 0 \\ -\beta h_{\rm x} & 0 & 0 & h_{\rm x} \\ 0 & 0 & 0 & 0 \\ 0 & h_{\rm x} & 0 & 0 \end{pmatrix}, \tag{10.162}$$

and let the single degree of freedom be a simple plane-wave, $h_{\rm x} = A_x \sin(\omega t - kz)$. The goal is to find the form of ξ_μ in Eq. (10.161) that yields $\underset{\sim}{h}_{0\nu} = 0$ in this frame. For instance, for the time-time component, Eq. (10.161) and our requirement that $\underset{\sim}{h}_{00} = 0$ imply

$$\partial_0 \xi_0 = 0 \implies \xi_0 = 0\,. \tag{10.163}$$

In the last step, we integrated over time and used gauge freedom to pick initial conditions in which $\xi_0(\vec{x}) = 0$ for all $\vec{x}$, so that we can ignore the integration constant. Similarly, using this result for ξ_0 and demanding $\underset{\sim}{h}_{01} = 0$, we can also conclude that Eq. (10.161) requires

$$h_{01} + \partial_0 \xi_1 + \partial_1 \xi_0 = 0 \implies \partial_t \xi_1 = \beta h_{\rm x}\,, \tag{10.164}$$

which we can integrate, as we did above, to get

$$\xi_1 = -\frac{\beta}{\omega} A_x \cos(\omega t - kz). \tag{10.165}$$

Since this is the only nonvanishing component of the Stückelberg field relevant to the vector-x amplitude, Eq. (10.161) implies that

$$\underset{\sim}{h}_{13} = \underset{\sim}{h}_{31} = h_{\rm x} + \partial_z \xi_1 = \left(1 - \beta^2\right) h_{\rm x}, \tag{10.166}$$

and $\underset{\sim}{h}_{\mu\nu} = 0$ otherwise, for a unitary metric perturbation whose only non-zero components come from h_{x}, as we supposed above in Eq. (10.162). It can be shown that the same exact argument, applied to h_{y} instead of h_{x}, yields an analogous result,

$$\underset{\sim}{h}_{23} = \underset{\sim}{h}_{32} = h_{\mathrm{y}} + \partial_z \xi_1 = \left(1 - \beta^2\right) h_{\mathrm{y}}, \tag{10.167}$$

if we had started with a unitary metric perturbation whose only non-vanishing degree of freedom was h_{y}.

The case of the longitudinal amplitude, h_{l}, is slightly more complicated, but can be handled in the same way. Suppose the perturbation is given simply by

$$(h_{\mu\nu}) = \begin{pmatrix} \beta^2 h_{\mathrm{l}} & 0 & 0 & -\beta h_{\mathrm{l}} \\ 0 & \frac{1}{2}\left(\beta^2 - 1\right) h_{\mathrm{l}} & 0 & 0 \\ 0 & 0 & \frac{1}{2}\left(\beta^2 - 1\right) h_{\mathrm{l}} & 0 \\ -\beta h_{\mathrm{l}} & 0 & 0 & h_{\mathrm{l}} \end{pmatrix}, \tag{10.168}$$

and let $h_{\mathrm{l}} = A_l \sin(\omega t - kz)$, as we did above for h_{x} (and, implicitly, h_{y}). In this case, the requirement that $\underset{\sim}{h}_{00} = 0$ implies, via Eq. (10.161), that

$$h_{00} + 2\partial_0 \xi_0 = 0 \implies \partial_0 \xi_0 = -\frac{1}{2}\beta^2 h_{\mathrm{l}}, \tag{10.169}$$

and so, integrating over time, we conclude that

$$\xi_0 = \frac{\beta^2}{2\omega} A_l \cos(\omega t - kz), \tag{10.170}$$

where we have neglected integration constants, as before. Now, the result for ξ_0 and the requirement that $\underset{\sim}{h}_{03} = 0$ mean that Eq. (10.161) also implies

$$h_{03} + \partial_0 \xi_3 + \partial_3 \xi_0 = 0 \implies \partial_0 \xi_3 = \left(\beta - \frac{\beta^3}{2}\right) h_{\mathrm{l}}, \tag{10.171}$$

and so, integrating over time,

$$\xi_3 = \frac{\beta}{\omega}\left(\frac{\beta^2}{2} - 1\right) A_l \cos(\omega t - kz)\,. \tag{10.172}$$

Since ξ_0 and ξ_3 are the only nonvanishing components of the Stückelberg field relevant to the longitudinal amplitude, Eq. (10.161) implies that

$$\underset{\sim}{h}_{11} = h_{22} = h_{11} = h_{22} = \frac{1}{2}\left(\beta^2 - 1\right) h_{\mathrm{l}}, \tag{10.173}$$

$$\underset{\sim}{h}_{33} = h_{33} + 2\partial_z \xi_3 = \left(\beta^2 - 1\right)^2 h_{\mathrm{l}}, \tag{10.174}$$

for a unitary metric whose only non-zero components come from $h_{\rm l}$, as we supposed above in Eq. (10.168).

Putting back all degrees of freedom together, we obtain our final expression for the metric perturbation in a synchronous gauge,

$$(\underset{\sim}{h}_{\mu\nu}) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ - \frac{1}{2}\alpha^2 h_{\rm l} & h_\times & \alpha^2 h_{\rm x} \\ 0 & h_\times & -h_+ - \frac{1}{2}\alpha^2 h_{\rm l} & \alpha^2 h_{\rm y} \\ 0 & \alpha^2 h_{\rm x} & \alpha^2 h_{\rm y} & \alpha^4 h_{\rm l} \end{pmatrix}, \tag{10.175}$$

with α as in Eq. (10.97). In the limit of no graviton mass ($\alpha \to 0$), we manifestly recover the transverse-traceless expression familiar from GR without the need for further gauge fixing.

Finally, it will be useful to define a transformation matrix to go from unitary to synchronous polarization amplitudes. The unitary amplitudes are simply the degrees of freedom appearing in Eq. (10.158), while the synchronous ones are just

$$(\underset{\sim}{h}_{\mu\nu}) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \underset{\sim}{h}_{\rm b} + \underset{\sim}{h}_+ & \underset{\sim}{h}_\times & \underset{\sim}{h}_{\rm x} \\ 0 & \underset{\sim}{h}_\times & \underset{\sim}{h}_{\rm b} - \underset{\sim}{h}_+ & \underset{\sim}{h}_{\rm y} \\ 0 & \underset{\sim}{h}_{\rm x} & \underset{\sim}{h}_{\rm y} & \underset{\sim}{h}_{\rm l} \end{pmatrix}, \tag{10.176}$$

in full analogy to Eq. (10.6). Comparing this definition to Eq. (10.175), it can be easily shown that the transformation matrix $M_{\underset{\sim}{A}}{}^B$ satisfying $\underset{\sim}{h}_A = M_{\underset{\sim}{A}}{}^B h_B$ is given by Eq. (10.108).

10.8.2 Effective stress-energy tensor

We wish to obtain an expression for the ESET of GWs in Fierz-Pauli massive gravity, following the procedure outlined in Sec. 10.2.4. To do so, begin with the total FP action of Eq. (10.93), $\mathcal{S}_{\rm FP} = \mathcal{S}_{\rm EH} + \mathcal{S}_m$, with $\mathcal{S}_{\rm EH}$ the Einstein-Hilbert action of Eq. (10.40), and $\mathcal{S}_m$ the contribution from the scalar field given by Eq. (10.94). All computations in this section will be carried out in the *unitary gauge* native to FP gravity, Eq. (10.99), since those polarization amplitudes are the fundamental degrees of freedom that we can take to be uncorrelated in this theory (since they diagonalize its kinetic matrix).

We will consider the two terms in the FP action separately. As in the scalar-tensor case (Appendix 10.7), we may obtain the contribution from the Einstein-Hilbert part

by starting from the MacCallum-Taub expression for the stress energy, Eq. (10.142). Unlike for scalar-tensor, however, we may now ignore all terms showing the trace and let $\bar{h}_{ab} = h_{ab}$, thanks to Eq. (10.96b). With these simplifications, Eq. (10.142) becomes

$$T^{(\mathrm{EH})}_{ab} = \kappa\Big\langle\!\Big\langle \frac{1}{2}\nabla_a h^{cd}\nabla_b h_{cd} - \nabla_c h_{da}\nabla^c h_b{}^d + g_{ab}\left(\frac{1}{2}\nabla_e h^{fc}\nabla_f h^e{}_c - \frac{1}{4}\nabla_e h_{cd}\nabla^e h^{cd}\right)\Big\rangle\!\Big\rangle. \tag{10.177}$$

The first term in this expression yields the Isaacson tensor obtained in GR, Eq. (10.43), except now the sum must include all five polarizations allowed in Eq. (10.99), not just the transverse-traceless ones. The second term may be rewritten via integration by parts "under the average," as discussed around Eq. (10.143), so that it becomes[7]

$$\langle\!\langle -\nabla_c h_{da}\nabla^c h_b{}^d\rangle\!\rangle = \mu^2\langle\!\langle h_{da}h_b{}^d\rangle\!\rangle + (\text{avg. error}), \tag{10.178}$$

after applying the equations of motion, Eq. (10.95). A similar argument shows that the third term vanishes due to Eq. (10.96a), while the fourth and final term takes a similar form as the second one,

$$\langle\!\langle -\nabla_e h_{cd}\nabla^e h^{cd}\rangle\!\rangle = \mu^2\langle\!\langle h^{cd}h_{cd}\rangle\!\rangle + (\text{avg. error})\,. \tag{10.179}$$

Altogether, this means that the contribution to the ESET from the Einstein-Hilbert part of the action is

$$T^{(\mathrm{EH})}_{ab} = \frac{\kappa}{2}\langle\!\langle \nabla_a h^{cd}\nabla_b h_{cd}\rangle\!\rangle + \kappa\mu^2\langle\!\langle h_{da}h_b{}^d + \tfrac{1}{4}g_{ab}h^{cd}h_{cd}\rangle\!\rangle\,. \tag{10.180}$$

Now focus on the contribution from $\mathcal{S}_m$ in Eq. (10.94). This action is already the quadratic action $\mathcal{S}^{(2)}$ needed for Eq. (10.34), namely

$$\mathcal{L}^{(2)}_m = \frac{1}{4}\kappa\mu^2 h^{ab}h^{cd}\sqrt{-g}\,(g_{ab}g_{cd} - g_{ac}g_{bd}), \tag{10.181}$$

where we have explicitly written out the antisymmetrized terms. We have also written h^{ab} with indices up, to match the index position convention used in [219] and thus

[7] The error here arises from the level at which total derivatives average out to over the length L_{ave}. Naturally this length needs to be very large compared to the gravitational wavelength, but its hierarchy with the Compton wavelength is more subtle. To justify keeping the μ^2 terms, we need the averaging error to be small compared to the μ^2 terms.

in deriving Eq. (10.142).[8] The variation of this quantity with respect to the inverse metric can be shown to be

$$\frac{\delta \mathcal{L}_m^{(2)}}{\delta g^{cd}} = \tfrac{1}{2}\kappa\mu^2\sqrt{-g}\left[h_{ca}h_d{}^a + \tfrac{1}{4}g_{cd}h_{ab}h^{ab}\right], \tag{10.182}$$

where we have used the fact that $\delta g_{ab} = -g_{ac}(\delta g^{cd})g_{db}$, $\delta\sqrt{-g} = -\sqrt{-g}\, g_{ab}\delta g^{ab}/2$, and that, on shell, the perturbation is traceless by Eq. (10.96b). The contribution of $\mathcal{S}_m$ to the stress energy may be obtained directly from this variation using Eq. (10.31):

$$\Delta T_{ab}^{(\mathrm{FP})} = -\kappa\mu^2\left\langle\!\!\left\langle h_{da}h_b{}^d + \frac{1}{4}g_{ab}h^{cd}h_{cd}\right\rangle\!\!\right\rangle. \tag{10.183}$$

Adding both contributions computed above, the total ESET in massive gravity, $T_{ab} = T_{ab}^{(\mathrm{EH})} + \Delta T_{ab}^{(\mathrm{FP})}$, is then

$$T_{ab} = \frac{\kappa}{2}\langle\!\langle \nabla_a h^{cd}\nabla_b h_{cd}\rangle\!\rangle, \tag{10.184}$$

as presented in Eqs. (10.100)–(10.103). We further discuss the interpretation of this result in the main text.

We now compute an expression for ρ_{GW} as a function of the PSD of the unitary polarization amplitudes of Eq. (10.158) [Eq. (10.99) in the main text]. Expanding the metric perturbation into plane waves in the local Lorentz frame of the detector (with $g_{ab} = \eta_{ab}$), as in Eq. (10.1), and taking the time-time component of the ESET, we get

$$\begin{aligned}\rho_{\mathrm{GW}} &\equiv \frac{\kappa}{2c^2}\left\langle \partial_t h^{\alpha\beta}\partial_t h_{\alpha\beta}\right\rangle \\ &= \frac{-\kappa}{2c^2}\frac{1}{4\pi^2}\int \left\langle \tilde{h}^*_{\alpha\beta}(-\mathbf{k})\tilde{h}^{\alpha\beta}(\mathbf{k}')\right\rangle e^{i(\mathbf{k}+\mathbf{k}')\cdot\mathbf{x}}\omega\omega'\widetilde{\mathrm{d}\mathbf{k}}\widetilde{\mathrm{d}\mathbf{k}}',\end{aligned} \tag{10.185}$$

assuming (*ii*) ergodicity as usual. The second equality was obtained by proceeding identically as in GR (Sec. 10.3.1). The contraction inside the angular brackets can be rewritten in terms of the unitary polarizations using Eq. (10.159),

$$\begin{aligned}\left\langle \tilde{h}^*_{\alpha\beta}(-\mathbf{k})\tilde{h}^{\alpha\beta}(\mathbf{k}')\right\rangle = \sum C_A(\omega,\omega') \\ \times \left\langle \tilde{h}^*_A(-\mathbf{k})\tilde{h}_A(\mathbf{k}')\right\rangle,\end{aligned} \tag{10.186}$$

[8] This is a somewhat subtle point, since a wrong index position generates implicit dependence on the (inverse) metric. Ultimately it does not matter whether h_{ab} or h^{ab} is treated as the fundamental variable, so long as the same choice is made for all parts of the action when performing the variation with respect to g^{ab}.

for a sum over the degrees of freedom A of Eq. (10.99), and C_A as defined in Eq. (10.160).

Making use of all the usual assumptions (*i*)–(*vii*) about the background, we can then use Eq. (10.19) to write $\rho_{\rm GW}$ as

$$\rho_{\rm GW} = \frac{\pi c^2}{4G}\int_0^\infty \sum_A \lambda_A(f) S_A(f) f^2 {\rm d}f\,, \tag{10.187}$$

for $\lambda_A(f) \equiv C_A(f,f)$, and $S_A(f)$ the PSDs of the unitary polarization amplitudes. Here we have assumed that the polarization amplitudes in the unitary gauge are statistically independent, which is justified because, unlike the synchronous amplitudes, they diagonalize the kinetic matrix of the theory. Note that we recover the GR expression, Eq. (10.50), in the limit of vanishing α, if we also force $S_A(f) = 0$ for nontensorial modes, which is appropriate if these additional degrees of freedom are frozen out [105].

10.8.3 Cross-correlation

Here we derive an expression, in the form of Eq. (10.39), for the cross-correlation of detector outputs as a function of the fractional energy spectrum of massive gravitational waves, Eq. (10.106). Going back to Eq. (10.26), we may write the cross-correlation of the outputs of two detectors as

$$\begin{aligned}\left\langle \tilde{h}_I^*(f)\tilde{h}_{I'}(f')\right\rangle = \int {\rm d}\hat{n}{\rm d}\hat{n}' \left\langle \tilde{\underset{\sim}{h}}{}^*_{\underset{\sim}{A}}(\mathbf{k})\tilde{\underset{\sim}{h}}_{\underset{\sim}{A}'}(\mathbf{k}')\right\rangle \\ \times F_I^{*\underset{\sim}{A}}(\hat{n}) F_{I'}^{\underset{\sim}{A}'}(\hat{n}')\, e^{i(\vec{k}_{\underset{\sim}{A}'}\cdot\vec{x}_{I'} - \vec{k}_{\underset{\sim}{A}}\cdot\vec{x}_I)}\,,\end{aligned} \tag{10.188}$$

where the under-tilded quantities are defined in the *synchronous* gauge of Eq. (10.175). The reason we carry out the expansion in terms of the synchronous amplitudes is that only in the synchronous gauge may we write out the detector response by applying Eq. (10.23). However, we need a relation in terms of the *unitary* degrees of freedom, which diagonalize the kinetic matrix of the theory—we obtain such an expression below.

First, assuming a (*iii*) stationary and (*vii*) isotropic background, with (*iv*) uncorrelated sky bins, we may rewrite the above equation as (see Appendix 10.6)

$$\begin{aligned}\left\langle \tilde{h}_I^*(f)\tilde{h}_{I'}(f')\right\rangle = \delta(f-f')\left\langle \tilde{\underset{\sim}{h}}{}^*_{\underset{\sim}{A}}(f)\tilde{\underset{\sim}{h}}_{\underset{\sim}{A}'}(f)\right\rangle \\ \times \tilde{\Gamma}^{\underset{\sim}{A}\underset{\sim}{A}'}_{II'}(f)\,,\end{aligned} \tag{10.189}$$

where we have pushed all the directional dependence into the generalized overlap reduction functions, $\Gamma^{\tilde{A}\tilde{A}'}_{\tilde{I}\tilde{I}'}(f)$, of Eq. (10.28). Using the transformation of Eq. (10.108), we can now write this directly in terms of the unitary polarization amplitudes,

$$\begin{aligned}\left\langle \tilde{h}^*_I(f)\tilde{h}_{I'}(f')\right\rangle &= \delta(f-f')\left\langle \tilde{h}^*_B(f)\tilde{h}_{B'}(f)\right\rangle \\ &\quad\times M_{\tilde{A}}{}^{B}(f)\, M_{\tilde{A}'}{}^{B'}(f)\,\Gamma^{\tilde{A}\tilde{A}'}_{\tilde{I}\tilde{I}'}(f)\,.\end{aligned} \tag{10.190}$$

Here we have explicitly denoted the frequency dependence in $M_{\tilde{A}}{}^{B}(f)$, which is acquired implicitly via α in Eq. (10.108). Because the unitary polarizations can be taken to be (v) statistically independent, we may rewrite the above equation as a single sum over B,

$$\begin{aligned}\left\langle \tilde{h}^*_I(f)\tilde{h}_{I'}(f')\right\rangle &= \frac{1}{2}\delta(f-f')\sum_B S_B(f) \\ &\quad\times M_{\tilde{A}B}(f)\, M_{\tilde{A}'B}(f)\,\Gamma^{\tilde{A}\tilde{A}'}_{\tilde{I}\tilde{I}'}(f)\,.\end{aligned} \tag{10.191}$$

Using Eq. (10.107), this may be written directly in terms of the fractional energy spectrum for each unitary polarization as in Eq. (10.109).

Without more information about the detectors, Eq. (10.109) would be our final result for massive gravity. However, we may further simplify this for the case of a differential-arm instrument that effects a measurement via the detector tensor of Eq. (10.23). In that case, it may be shown from the definition of the antenna patterns, Eq. (10.25), that $F^{\rm b}_I(\hat{n}) = -F^{\rm l}_I(\hat{n})$ (e.g. [18]). This means that the generalized overlap reduction functions, Eq. (10.28), for the breathing and longitudinal modes will *not* be diagonal. In fact, this is evident from our expression for the $\Gamma^{\tilde{A}\tilde{A}'}_{\tilde{I}\tilde{I}'}$ factors for differential-arm detectors, Eq. (10.111), which follows directly from $F^{\rm b}_I(\hat{n}) = -F^{\rm l}_I(\hat{n})$.

Using Eq. (10.111) and the definitions of $\lambda_B(f)$ and $M_{AB}(f)$, from Eq. (10.105) and Eq. (10.108) respectively, our final result for the cross-correlation of the detector outputs of two differential-arm detectors takes the form of Eq. (10.39) with $\Xi_A(f)$ implicitly defined by

$$\Xi_B(f)\,\Gamma^{B}_{II'} = \lambda_B^{-1}(f) M_{\tilde{A}B}(f) M_{\tilde{A}'B}(f)\Gamma^{\tilde{A}\tilde{A}'}_{\tilde{I}\tilde{I}'}(f). \tag{10.192}$$

This reduces to the main result of Eq. (10.112), to quadratic order in α.

Chapter 11

MEASURING THE SPEED OF CONTINUOUS GRAVITATIONAL WAVES

11.1 Introduction

General relativity (GR) predicts that, far away from sources and other regions of strong spacetime curvature, gravitational waves (GWs) propagate along null geodesics of the background metric (see, e.g., Sec. 35.3 in [220]). This implies that, in this geometric-optics limit, GWs must travel at the speed of light ($c_g = c$). However, this fact is a consequence of the specific structure of the equations of motion imposed by GR and, hence, need not hold for arbitrary theories of gravity [83]. Specific examples of modifications of GR that predict $c_g \neq c$ are massive gravity theories or, more generally, theories in which spacetime is endowed with prior geometries (e.g. bimetric theories) or extra dimensions (see, e.g., [84, 105]).

Until recently, direct model-independent constraints on the speed of GWs were hard to come by. In fact, before the observations by LIGO and Virgo, the only serious contender came from energetic cosmic rays: if it was the case that $c_g < c$, these particles should have been slowed down on their way to Earth due to gravitational Cherenkov radiation, not reaching the observed energies (the idea goes back to at least the 1980s [307, 308]). Recent constraints placed in this manner implied $c - c_g < 2 \times 10^{-15} c$ ($c - c_g < 2 \times 10^{-19} c$), assuming the cosmic rays are galactic (extragalactic) in origin [309].

Other constraints were derived under the framework of linearized massive gravity, but were not direct observations of the GW speed. This includes so-called "static" constraints obtained from limits on a Yukawa coupling in the Newtonian potential of objects in the solar system [310, 311], as well as "dynamic" constraints from the energy loss in pulsar binaries [292]. More recently, constraints on the graviton mass were also derived from the lack of superradiance in supermassive black holes [312], but these are also highly model-dependent. In all these cases, an implied bound on the GW speed can be derived by assuming a specific form for the GW dispersion relation.

The first GW detection, GW150914, placed stringent constraints on the mass of the

graviton from the observed lack of dispersion in the signal waveform. However, this was still not a direct measurement of the GW speed, as it could identify only deformations in the shape of the expected signal, not changes in its overall time of arrival. As long as the wave train traveled undisturbed (i.e., "undispersed"), this test would not have been able to detect $c_g \neq c$.

LIGO and Virgo observations did eventually provide the first clear measurement of the GW speed. The first such result was obtained in [313] through a Bayesian analysis of the relative time of arrival of three binary-black-hole signals (GW150914, GW151226, GW170104) at different detectors. The bound obtained was $0.55c < c_g < 1.42c$ at 90% credibility. This was quickly and dramatically surpassed by the joint observation of GW and electromagnetic (EM) radiation from the binary-neutron-star merger, GW170817. Contingent on some assumptions regarding the relative emission time of the GW and EM signals, as well as the properties of the intergalactic medium, GW170817 allowed us to place a bound of $-3 \times 10^{-15}c \leq c_g - c \leq 7 \times 10^{-16}c$ [123]—a fantastic improvement over previous constraints. As a caveat, note that these bounds may be circumvented if delays between the GW and EM signals are allowed to be greater than 10 s, or if the EM radiation can precede the gravitational emission [314–316].

Here we explore another possibility for directly measuring the speed of gravity that will become available once we detect continuous monochromatic signals (CWs), for example, from rapidly rotating neutron stars. This technique relies on sensitivity to the Doppler modulation (in this context, also known as "Rømer modulation") of the expected signal caused by Earth's motion with respect to the source. This is similar to the strategy used by Ole Rømer in 1676 to produce the first quantitative estimate of the speed of light itself [317]. Although we do not expect it to be competitive with the bound place by GW170817, this measurement would offer an independent alternative, not limited by the same systematics.

11.2 Background

11.2.1 Signal

The projection of a continuous (quasimonochromatic) GW onto a detector can be written as (see Chapters 4, 7 and 8 for more exposition):

$$h(t) = F_+(t)a_+ \cos\phi(t) + F_\times(t)a_\times \sin\phi(t)\,, \tag{11.1}$$

where the F_+ and $F_\times$ factors respectively give the instrument's response to the plus (+) and cross (×) GW polarizations and $\phi(t)$ is the phase of the signal. The $a_{+/\times}$ factors give the amplitude of each polarization and depend both on the intrinsic emitted power and the radiation angular pattern. For the canonical triaxial-ellipsoid source (e.g. a spinning neutron star with a small asymmetry in its moment of inertia) these are given by

$$a_+ = h_0 \frac{1}{2}\left(1 + \cos^2 \iota\right), \tag{11.2}$$

$$a_\times = h_0 \cos \iota, \tag{11.3}$$

in terms of the source inclination, ι, and intrinsic amplitude h_0. For the purpose of analyzing the data, it is convenient to factorize this signal as [57, 67]:

$$h(t) = \Lambda(t) e^{i\phi(t)} + \Lambda^*(t) e^{-i\phi(t)}, \tag{11.4}$$

having defined the "heterodyned signal",

$$\Lambda(t) \equiv \frac{1}{2}\left[a_+ F_+(t) - i a_\times F_\times(t)\right], \tag{11.5}$$

with an asterisk (*) indicating complex conjugation.

The phase of the signal can be approximated via a Taylor expansion on the source frequency, ν,[1] and by using the fact that for this kind of quadrupolar source we expect $\phi(t) = 2\phi_{\rm rot}(t)$, where $\phi_{\rm rot}$ is the angular phase of the rotating source. That way, we can write:

$$\phi(t) = 4\pi\left(\nu t_b + \frac{1}{2}\dot{\nu} t_b^2 + \frac{1}{6}\ddot{\nu} t_b^3\right) + \phi_0, \tag{11.6}$$

where ϕ_0 is a fiducial phase offset and where the "barycentric time" t_b is itself a function of detector time t that gives the time by a clock inertial with respect to the source. In searches targeted at known pulsars, the rotation frequency and its derivatives are known thanks to exquisitely-precise pulsar timing techniques [55, 56].

11.2.2 Barycentering

The time of arrival of GW phase-fronts at a detector on Earth is recorded in a local frame that is not an inertial frame with respect to the source. Therefore, the measured signal from a persistent GW is modulated in frequency by Doppler and relativistic effects. We can implement this through a time-of-arrival function [56] that, assuming our detector lies on Earth's surface, can be expressed as [56]

$$t_b(t) = t + \Delta_E(t) + \Delta_S(t) + \Delta_B(t) + \Delta_R(t) \tag{11.7}$$

[1] In other chapters, we have denoted the source frequency by "$f_{\rm rot}$" instead of "ν".

where Δ_E is the "Einstein delay", Δ_S the "Shapiro delay", $Delta_B(t)$ the "binary delay" and Δ_R the "Rømer delay". Eq. (11.7) effectively transforms detector time, t, into time at the solar-system barycenter (SSB), t_b, which is itself assumed to be inertial with respect to the source [56].

The Einstein delay accounts for the relativistic time dilation between the local Lorentz frame (LLF) of the SSB and that of a detector on Earth's surface, amounting to a secular lag of $\sim 0.5\,\mathrm{s\,yr^{-1}}$ between the two clocks [318].The Shapiro delay results from the bending of the waves due to the spacetime curvature around the SSB and, as such, depends on the location of all the major bodies in the solar system at the time of arrival [319]. The binary term encodes all potential effects due to the source being in a binary system [320]. None of these three effects depend explicitly on the speed of gravity or have sufficient impact on the measurement to be relevant for our purposes. For simplicity, we shall assume our target is isolated.

The only term left, the Rømer delay, is the key to our study. This is simply the classical time it takes the signal to travel between the detector and the SSB [56, 318]:

$$\Delta_R(t) = \frac{\vec{r} \cdot \hat{n}}{c_g}, \tag{11.8}$$

where $\vec{r}$ is the location vector of the detector with respect to the SSB, $\hat{n}$ is a unit vector pointing from the SSB towards the source and c_g is the GW speed. This term results in a modulation driven by Earth's yearly motion around the SSB and, to a lesser extent, by its daily rotation around its spin axis. Notice that here, unlike in Δ_E and Δ_S, the dynamical speed of a gravitational perturbations appears explicitly in the denominator. Thus, a change in GW speed should be directly measurable through the Rømer delay.

11.2.3 Searches

We will focus on targeted CW analyses, in which the intrinsic phase evolution of the signal can be assumed to be known, for example, thanks to EM observations (see the Background section in Chapter 8). With knowledge of the phase evolution, one can simplify the problem by digitally heterodyning (i.e. multiplying by a complex exponential with phase evolution matching the expected signal) the time-domain GW data so that the signal they contain becomes [57, 66]:

$$h'(t) \equiv h(t)e^{-i\phi(t)} = \Lambda(t) + \Lambda^*(t)e^{-i2\phi(t)}, \tag{11.9}$$

with $\Lambda(t)$ as in Eq. (11.5).

A series of low-pass filters can then be applied to remove the second term in Eq. (11.9), which enables the down-sampling of the data by averaging over minute-long time bins. As a result, $\Lambda(t)$ is the only contribution from the original signal left in our binned data, B, which will now look like

$$B(t_k) = \Lambda(t_k) + n(t_k), \tag{11.10}$$

where $n(t_k)$ is the heterodyned, filtered and downsampled noise in bin k, which carries no information about the GW signal. From Eq. (11.5), note that $\Lambda(t)$ evolves only slowly, with a period of a one and half a sidereal day, and that all high-frequency components intrinsic to the GW signal have been removed by the data reduction process, i.e. $\phi(t)$ does not appear in Eq. (11.10).

At this step, Bayesian inference, or any other regression technique (e.g. χ^2 minimization, as used below) may be used to retrieve the signal from the noise. Heterodyning makes this process much more efficient by reducing the size of the data to be analyzed, but it relies on precise knowledge of the phase evolution, including the barycentering delays.

11.3 Method

Our goal is to measure the GW speed, c_g, starting from data that has been preprocessed as described in Sec. 11.2.3 under the assumption that $c_g = c$. To do this, consider that if $c_g \neq c$ the cancellation of the phase in Eq. (11.9) will be imperfect. In particular, there will be a remainder in given by the c_g-dependent terms in the Rømer delay of Eq. (11.8).

Letting $\Delta\phi \equiv \phi_{c_g} - \phi_c$, for ϕ_{c_g} and ϕ_c respectively the true and assumed phase evolutions, from Eq. (11.12) and Eq. (11.7) we will have:

$$\begin{aligned}\Delta\phi(t) = 4\pi \Bigg[&(\vec{r}\cdot\hat{n})\left(\frac{1}{c_g} - \frac{1}{c}\right)\left(\nu + \dot{\nu}t + \frac{1}{2}\ddot{\nu}t^2\right) + (\vec{r}\cdot\hat{n})^2\left(\frac{1}{c_g^2} - \frac{1}{c^2}\right)\frac{1}{2}\left(\dot{\nu} + \ddot{\nu}t\right) \\ &+ (\vec{r}\cdot\hat{n})^3\left(\frac{1}{c_g^3} - \frac{1}{c^3}\right)\frac{1}{6}\ddot{\nu}\Bigg]\end{aligned} \tag{11.11}$$

or, defining $\delta \equiv c/c_g$ and $\Delta_{R'} \equiv t + \vec{r}(t)\cdot\hat{n}/c$,

$$\begin{aligned}\Delta\phi(t) = 4\pi \Big[&\Delta_{R'}\left(\delta - 1\right)\left(\nu + \dot{\nu}t + \tfrac{1}{2}\ddot{\nu}t^2\right) + \tfrac{1}{2}\Delta_{R'}^2\left(\delta^2 - 1\right)\left(\dot{\nu} + \ddot{\nu}t\right) \\ &+ \tfrac{1}{6}\Delta_{R'}^3\left(\delta^3 - 1\right)\ddot{\nu}\Big]\end{aligned} \tag{11.12}$$

After low-pass filtering (Sec. 11.2.3), the signal in the data will take the form:

$$h'(t) = \Lambda(t)e^{i\Delta\phi(t)}, \tag{11.13}$$

instead of simply $\Lambda(t)$ as in Eq. (11.10).

Eq. (11.13) implies that $c_g \neq c$ causes the expected signal to retain some fast phase evolution (encoded in $\Delta\phi$), instead of only the daily modulation due to the antenna patterns in Eq. (11.5). The frequency of the leftover phase modulation will be given by the intrinsic GW frequency, $\sim 2\nu$, with additional components at Earth's orbital frequency ($\omega_{\rm orb} \approx 2 \times 10^{-7}$ Hz) due to Δ_R. For example, for signals from the Crab pulsar (PSR J0534+2200) we expect $2\nu = 60$ Hz, which is much faster than the frequency characteristic of the antenna patterns, $\omega_{\rm sid} = 2\pi/(\text{sidereal day}) \approx 7 \times 10^{-5}$ Hz. (See Background section in Chapters 7 and 8 for in-detail discussions of the antenna pattern functions.)

Remarkably, this allows us to use the data that has already been heterodyned and reduced assuming $c_g = c$ to detect whether GW speed deviates from the speed of light, without requiring to go back to raw LIGO data (which would have been computationally costly).

Computing the Rømer delay

In order to model and fit for the phase evolution of the signal in Eq. (11.13), it is necessary to know the Rømer delay for every time bin, and for any value of c_g. To compute Δ_R to sufficient precision, the location of the detector and the source with respect to the SSB must be well known. This requires consulting solar system ephemerides, which can be obtained through the HORIZONS system developed by the Solar System Dynamics Group of the Jet Propulsion Laboratory [321].[2] Note that ephemerides come in different versions and one must make sure to use the same one for the $c_g \neq c$ analysis as in the original data reduction assuming $c_g = c$. Otherwise, ephemerides errors might be mistaken for deviations in the GW speed.

From the solar-system ephemeris, one obtains the location of the geocenter, $\vec{r}_\oplus$, with respect to the SSB as a function of time. To compute Δ_R in Eq. (11.8), we need the location of the detector on the surface, $\vec{r}$, which can be computed by adding the

[2] Solar-system ephemerides from JPL are included as part of the LIGO Algorithm Library and can be accesed through the LALBarycenter routine, which has the added advantage of being pre-corrected to be a function of signal arrival GPS time [162].

vector from the geocenter to the observatory $\vec{s}$:

$$\vec{r} = \vec{r}_\oplus + \vec{s} \tag{11.14}$$

Clearly, $\vec{s}$ depends on the latitude and longitude of the site and varies with time as Earth rotates in its axis over a sidereal day. On top of this, we must also factor in the precession and nutation of Earth's axis.

The other component of the Rømer delay is the source-direction vector, $\hat{n}$. In celestial coordinates (ICRS[3]), this is given by [56]

$$\hat{n} = \hat{n}_0 + \vec{\mu}_\perp t - \left(\frac{1}{2}|\vec{\mu}_\perp|^2 \hat{n}_0 + \mu_\parallel \vec{\mu}_\perp\right) t^2, \tag{11.15}$$

where the time t is measured with respect to the ICRS standard, and

$$\hat{n}_0 = \begin{pmatrix} \cos\alpha_\star \cos\delta_\star \\ \sin\alpha_\star \cos\delta_\star \\ \sin\delta_\star \end{pmatrix}, \tag{11.16}$$

with $\alpha_\star$ and $\delta_\star$ the right ascension and declination of the source in ICRS. The additional terms encode corrections due to the proper motion of the source in the plane of the sky ($\vec{\mu}_\perp$) and along the line of sight ($\vec{\mu}_\parallel$). These extra corrections tend to be small: for instance, in the case of the Crab pulsar $\mu_\alpha = 13$ mas yr^{-1} and $\mu_\delta = 7$ mas yr^{-1} [152]. We will neglect these small terms here.

Putting all these pieces together, we can compute Eq. (11.8) as a function of time. This will show significant variation over a sidereal day and a sidereal year due to Earth's motion on its spin and orbit, as shown in Fig. 11.1). With this information, it is straightforward to compute the phase evolution given in Eq. (11.12) and shown in Fig. 11.2 for various values of $\delta = c/c_g$. This, in turn, is used to compute the $c_g \neq c$ signal template of Eq. (11.13), which is needed for searches and injections. The real and imaginary parts of such templates are represented in Fig. 11.3, where we have simulated a signal from the Crab pulsar as seen by the LIGO Hanford (H1) detector. As expected from the discussion above, changes in GW speed are manifested only in the phase evolution of the signal, and not the amplitude (Fig. 11.4).

11.4 Analysis

In order to quantify how well we can detect variations in GW speed, we fabricate several instantiations of data with noise corresponding to the sensitivity of the H1

[3]International Celestial Reference System [322].

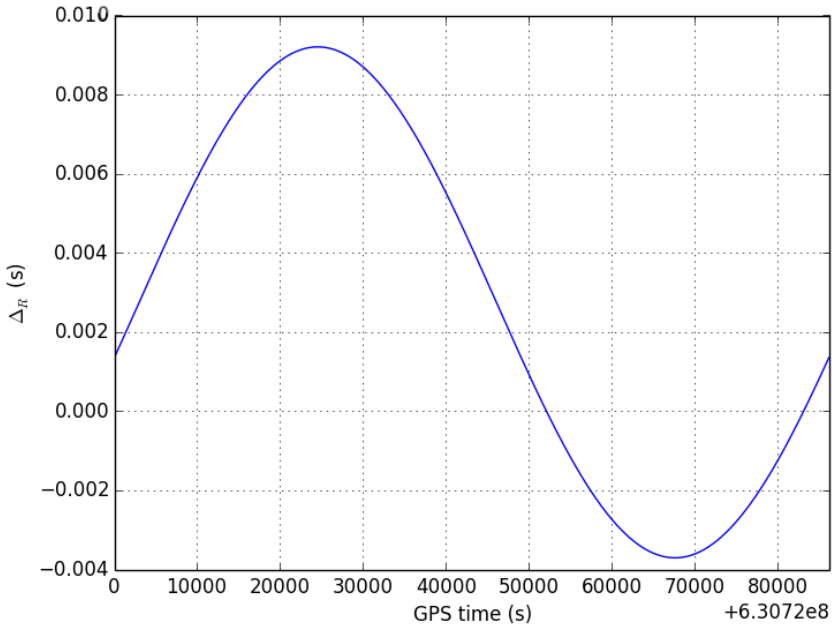

Figure 11.1: Rømer delay between the Crab pulsar and the LIGO Hanford detector over 1 sidereal day.

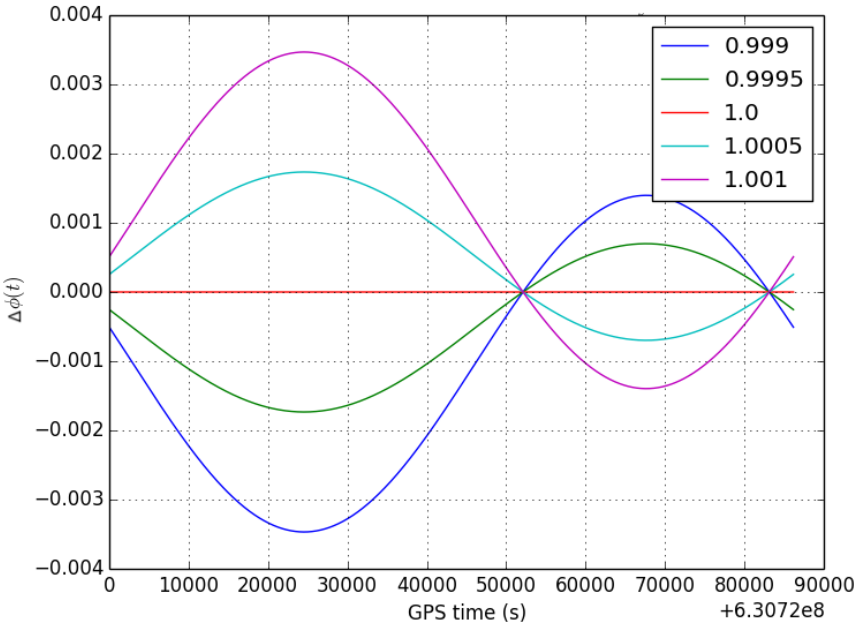

Figure 11.2: Phase evolution of signal (radians) over one sidereal day for different δ after heterodyning, Eq. (11.12). When $\delta = 1$, i.e. $c_g = c$, the phase is constant (red).

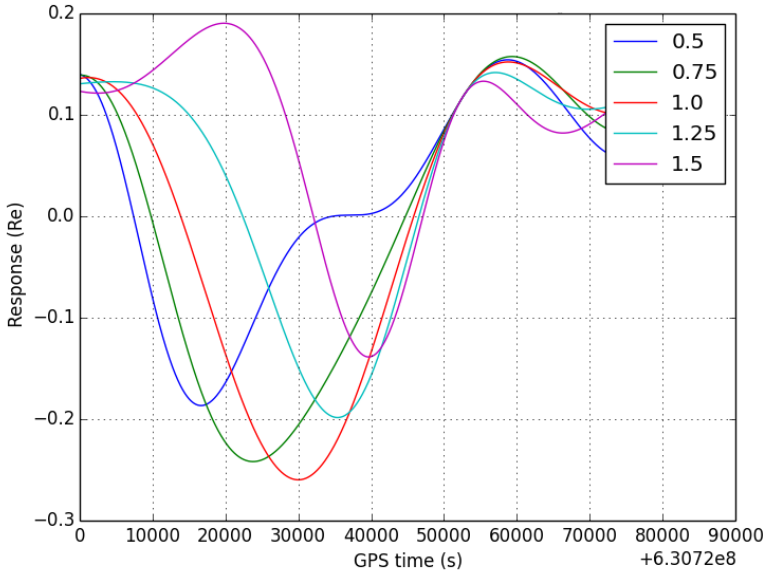

(a) Real

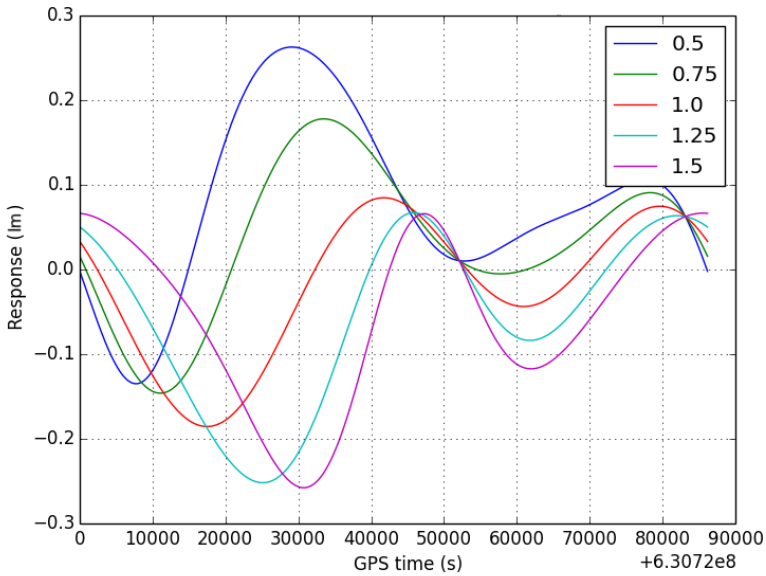

(b) Imaginary

Figure 11.3: Heterodyned signal from the Crab pulsar as seen by the LIGO Hanford detector, Eq. (11.13), real (top) and imaginary (bottom) parts. Color corresponds to different values of $\delta = c/c_g$.

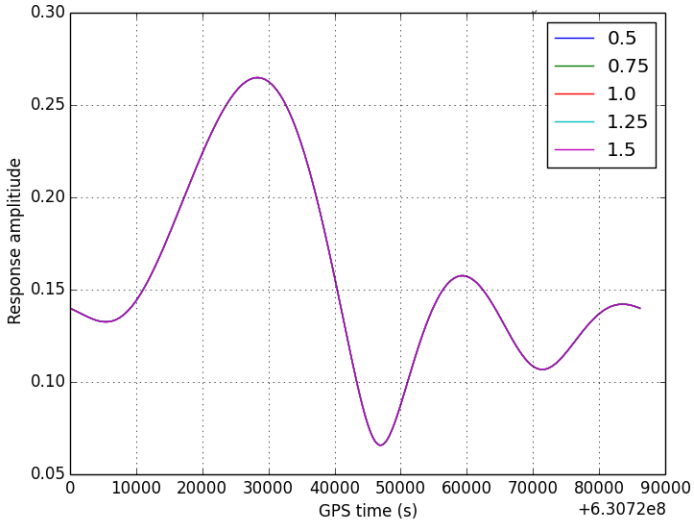

Figure 11.4: The amplitude modulation of the signal is unaffected by changes in GW speed, by the fact that curves are identical irrespective of the value of $\delta = c/c_g$.

detector during Initial LIGO's sixth science run (S6) [323]. Into these, we inject simulated signals from the Crab pulsar corresponding to different GW speeds. We then attempt to recover the correct value of $\delta = c/c_g$ for different strengths of injected signal. We use both noise extracted from LIGO data and random Gaussian noise of comparable standard deviation. The former are obtained by the reheterodyning procedure outlined in Chapter 7 so that the search is blind to astrophysical signals. We focus on the case $\delta > 1$, i.e. $c_g < c$.

The search is performed by use of a matched filter: for each injection, we compute the cross-correlation between the injected data and a template assuming different values of the speed parameter, δ_{rec} (Fig. 11.5a). Of course, in the absence of noise, the match is perfect and we are able to recover the injected value of the speed parameter (δ_{inj}) flawlessly, no matter the amplitude of the injected signal (Fig. 11.5b). When noise is present, the quality of the match depends on the signal-to-noise ratio (SNR).

The results of the analysis for different injection strengths are shown in Fig. 11.6. For a strong-enough signal, we are able to recover all values of δ_{inj} quite well. However, when the signal is of the order of the noise (standard deviation $\sim 10^{-23}$), the peak-finding method fails and returns the values at the boundaries of the search range (horizontal lines in Fig. 11.6). Note that LIGO data does not seem to behave significantly different from Gaussian noise, as expected (see Appendix in Chapter 7).

The goodness-of-fit in the recovered value of δ for a given injection strength is

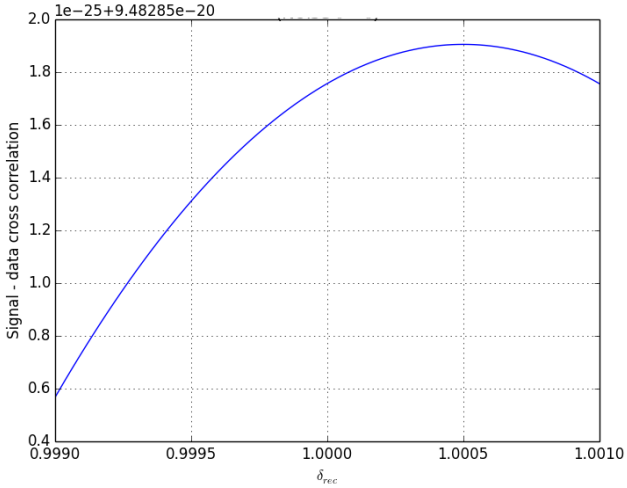

(a) $\delta_{inj} = 1.0005$

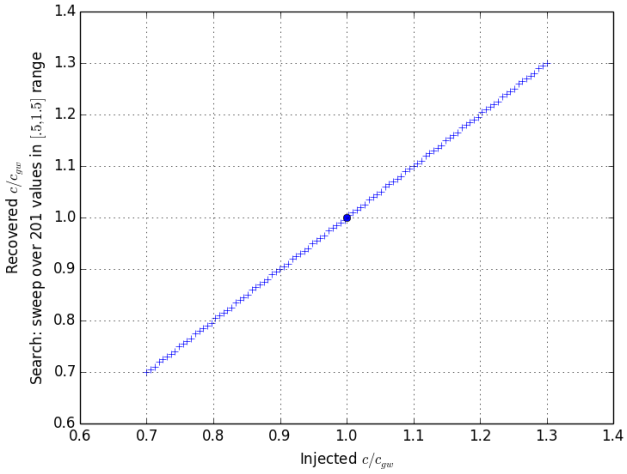

(b) $\delta_{inj} = 0.7, 1.3$

Figure 11.5: The cross-correlation between a template and a signal peaks when $\delta_{inj} = \delta_{rec}$ (left). Finding this peak we can recover δ_{inj} (right). Note that these plots are both tests performed in absence of noise.

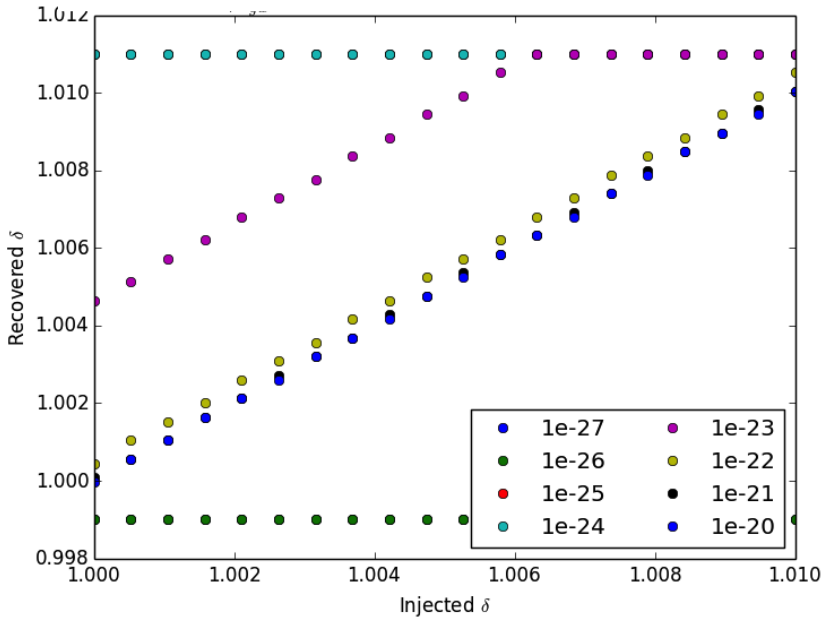

(a) Gaussian noise

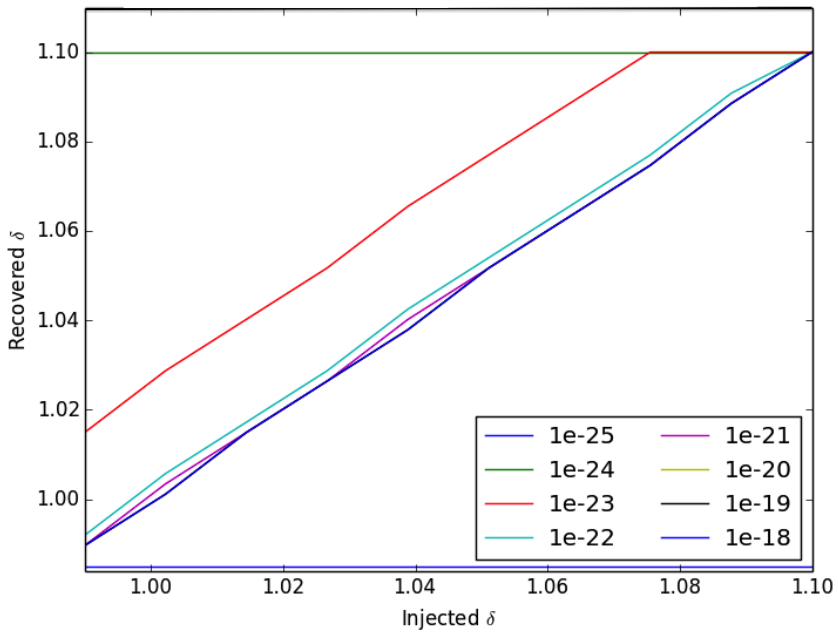

(b) Crab-heterodyned LIGO H1 S6 noise

Figure 11.6: Recovered vs. injected δ for different signal strengths in Gaussian (left) and LIGO (right) noise. A slope of 1 indicates good detection. When the signal is too faint, the peak-finding makes the detection curve be a horizontal line at the maximum or minimum value of the δ_{rec} range. Note that the standard deviation of the noise is of the order 10^{-23} in both cases.

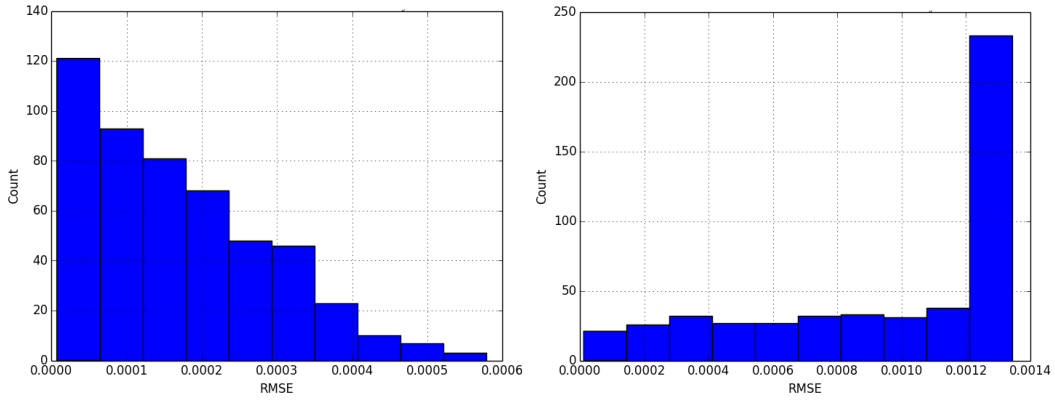

(a) $h_0 = 10^{-22}$ (b) $h_0 = 10^{-23}$

Figure 11.7: Histograms of RMSE for recovery of signals with $\delta = 1.0, 1.01$ at strong (left) and weak (right) injection strengths. For a strong signal we are able to recover $\delta \pm 0.0005$, while for the weak case there are clear systematic errors.

evaluated through the root-mean-square error, i.e.

$$\text{RMSE} = \sqrt{\sum (\delta_{inj} - \delta_{rec})^2 / N} \tag{11.17}$$

where N is the number of different injected signals (different δ_{inj}'s) for a given h_0. Repeating the same search multiple times but with varying noise, we obtain the uncertainty on our measurement of δ for certain SNR. For a strong signal, we are able to recover δ up to ± 0.0005 (Fig. 11.7a). However, when the signal is of the order of the noise, it is clear that the peak-finding procedure introduces systematic errors that push δ_{rec} towards the boundary of the search range.

11.5 Conclusions

We explored the possibility of measuring the speed of gravity using continuous gravitational waves. The key to the strategy is the Rømer modulation caused by Earth's relative motion with respect to the source, which depends inversely on the GW speed. In the context of searches targeted at specific sources, like pulsars, we showed how signal templates would have to be modified to account for $c_g \neq c$. We also demonstrated that we can carry out a measurement of the speed using data that has been manipulated in standard ways assuming $c_g = c$, making the computations more computationally efficient. Finally, we implemented a simple frequentist technique to recover the GW speed from simulated signals. From injections on Gaussian and blinded LIGO-S6 data we determined that, for a very loud signal, we should be able

to measure c/c_g down to levels of $\sim 10^{-3}$ with a statistical error of 5×10^{-4}. This would not be competitive with existing constraints obtained from GW170817, but would provide an independent measurement with different systematics. In any case, further investigation is required to corroborate these results and gauge our ability to make this measurement with weaker signals. Although we made some preliminary progress towards a full Bayesian implementation of this measurement in [324], a full-fledged analysis is work in progress.

Acknowledgments

M.I. would like to give special thanks to Alan Weinstein for his extensive support and guidance, as well as Carver Mead for useful discussions. M.I. would also like to acknowledge Matthew Pitkin, without whose data this work would not have been possible. Thanks also to Kenneth Libbrecht and the Caltech SFP office.

Chapter 12

TOWARDS CONSTRAINING GENERIC GRAVITATIONAL-WAVE DISPERSION RELATIONS

12.1 Introduction

In general relativity (GR), gravitational waves (GWs) travel along null geodesics ($|\mathbf{p}|^2 = 0$, for the graviton four-momentum $\mathbf{p}$) and are, thus, non-dispersive. However, this prediction is a specific consequence of GR's equations of motion and need not hold in generic theories. In fact, modifications to the GW dispersion are equivalent to a change in the wave equation for metric perturbations, which in turn derive from the action of the theory and can take many forms, in principle.

The simplest modification to the GW dispersion relation is that of a massive graviton: $|\mathbf{p}|^2 = -m_{\rm g}^2$, in a local Lorentz frame (LLF) with metric signature +2 (see [232] for an excellent review on massive gravity, including a discussion beyond the linear regime). Other popular dispersion relations incorporate more complex behaviors through additional momentum-dependent terms, which can parametrize desirable features like violations of Lorentz invariance ([236–243, 325]). LIGO and Virgo have placed bounds on the mass of the graviton starting with the detection of GW150914 [82] and on Lorentz invariance with GW170104 [4], the latter based on the phenomenological model of [325]. More generally, the dispersion could incorporate terms that depend not only on the magnitude of the wave's momentum (wavevector) but also its direction and spatiotemporal location, with functional relations of arbitrary complexity [326].

Here, we study the effect of an arbitrary dispersion relation on GW phenomenology. In Sec. 12.2, we provide some background information on post-Newtonian (PN) and inspiral-merger-ringdown (IMR) waveforms. In Sec. 12.3, we derive the observable effect of such a dispersion on the GW waveform following the heuristic methods of [325, 327], with the main result being Eq. (12.43). In Sec. 12.4, we apply this expression to specific forms of $f(\mathbf{p})$ (massive graviton, simplest anisotropic term).

12.2 Background

In this section, we review the basics of waveforms from compact-binary coalescences (CBCs) in GR as a background to the modifications to be examined in following

sections. The summary of post-Newtonian (PN) waveforms follows Chapter 4.1 in [328].

12.2.1 Post-Newtonian waveforms

Consider two compact objects, with masses m_1 and m_2 and no spins, orbiting each other and slowly inspiraling in as the system loses energy through gravitational radiation. In the Newtonian approximation (0PN), the dynamics of an inspiral can be treated, like the traditional two-body problem, by means of the equivalent one-body problem in the center of mass frame. Using Kepler's laws and standard expressions for quadrupolar radiation in GR, the GW waveform for a non-cosmological, non-precessing source can be written:

$$h_+(t) = A(\tau)\ \frac{1}{2}(1+\cos^2\iota)\ \cos\Phi(\tau), \tag{12.1}$$

$$h_\times(t) = A(\tau)\ \cos\iota\ \sin\Phi(\tau), \tag{12.2}$$

for the plus (+) and cross (×) polarizations, defined in a frame aligned with the orbital angular momentum (polarization angle $\psi = 0$), and with time parameter τ. The terms that depend on the inclination angle, ι, encode the quadrupolar symmetry of the angular emission pattern—in particular, the fact that this is the $l = m = 2$ mode of the metric perturbation.

Meanwhile, the amplitude factor depends on the masses of the binary components and can be written as:

$$A(\tau) = \frac{\mathcal{R}}{r}\left(\frac{5\mathcal{R}}{c\tau}\right)^{\frac{1}{4}}, \tag{12.3}$$

where r is the distance to the source, and we have defined $\mathcal{R} \equiv G\mathcal{M}/c^2$ as the lengthscale corresponding to the *chirp mass* $\mathcal{M}$, itself given by

$$\mathcal{M}^5 = \mu^3 M^2 = \frac{(m_1 m_2)^3}{m_1 + m_2}, \tag{12.4}$$

in terms of the reduced mass $\mu = m_1 m_2/(m_1 + m_2)$ and the total mass $M = m_1 + m_2$.

The phase evolution is approximated by $\Phi(\tau) = 2\pi f(\tau)\,\tau$, with

$$f(\tau) = \frac{1}{\pi}\left(\frac{5}{256}\frac{1}{\tau}\right)^{\frac{3}{8}}\left(\frac{\mathcal{R}}{c}\right)^{-\frac{5}{8}}, \tag{12.5}$$

The time parameter, $\tau = t_c - t$, is the time to coalescence. The time of coalescence t_c can be formally defined as the value of t for which our approximation to the frequency

diverges. Physically, t_c marks the moment when the quasi-circular approximation breaks down and the inspiralling objects merge. Also note that, if t refers to the time at the source, then at the detector the waveform should be parametrized by the *r*etarded time $t_{\rm ret} = t - r/c$.

To generalize the above discussion to far-away sources, consider that all quantities measured at the detector will be cosmologically redshifted with respect to the local frame of the source. Detector-frame frequencies and times can be obtained from the source-frame ones via redshift factors of $(1+z)$. In particular, using geometrized units ($G = c = 1$), all quantities with dimensions of [mass]^n will be multiplied by a factor of $(1+z)^n$ in going from the emitter to the detector frames. That way, the comoving distance r becomes the luminosity distance $d_L = r(1+z)$, the emitted frequency f_e becomes the detected frequency $f_d = f_e(1+z)^{-1}$ and the source-frame chirp mass $\mathcal{M}_e$ becomes the detector-frame chirp mass $\mathcal{M}_d = \mathcal{M}(1+z)$. This last quantity is the chirp mass that would be inferred from Eq. (12.3) were redshift not taken into account.

Finally, it is of interest to compute the Fourier transform (FT) of the above signal. Expanding $h(t)$ around the time $\tilde{t}$ and setting $\tilde{f} \equiv f(\tilde{t})$, the stationary phase approximation (SPA) allows us to write the FT, $\tilde{h}(\tilde{f})$, in the form:

$$\tilde{h}(\tilde{f}) = \frac{A(\tilde{t})}{\sqrt{\dot{f}(\tilde{t})}} e^{i\Psi(\tilde{f})}, \tag{12.6}$$

where, for an inspiral-only signal, the amplitude and phase functions in the emission frame are respectively:

$$A(\tilde{f}) = \frac{4}{5}\frac{\mathcal{M}_{\rm e}}{a_0 \Sigma(\chi_{\rm e})} \left(\pi \mathcal{M} \tilde{f}_{\rm e}\right)^{\frac{2}{3}}, \tag{12.7}$$

$$\Psi(\tilde{f}) = 2\pi \int_{f_c}^{\tilde{f}} (t - t_c) {\rm d}f + 2\pi \tilde{f} t_c + \Psi_0. \tag{12.8}$$

Here a_0 and $\Sigma(\chi_{\rm e})$ are defined by the cosmological metric explored below, while the offset Ψ_0 is related to the waveform phase at coalescence Φ_c, by $\Psi_0 = -\Phi_c - \pi/4$. Frequencies and masses subscripted by "e" are measured at the emitter, while unsubscripted ones are the redshifted values seen by the detector.

12.2.2 Phenomenological waveforms

While the above two expressions are only valid in the adiabatically-Newtonian regime, we can make use of inspiral-merger-ringdown (IMR) waveforms that cover all stages

of the binary evolution [329]. Phenomenological IMR waveforms tuned to numerical relativity (NR) can be parametrized in an SPA-inspired fashion:

$$\tilde{h}(f) = A_{\rm eff}(f) e^{-i\Psi_{\rm eff}(f)}, \tag{12.9}$$

$$A_{\rm eff}(f) \equiv C \begin{cases} \left(f/f_{\rm merg}\right)^{-7/6} \,, & f < f_{\rm merg} \\ \left(f/f_{\rm merg}\right)^{-2/3} \,, & f_{\rm merg} \le f < f_{\rm ring} \\ w\mathcal{L}(f, f_{\rm ring}, \sigma) \,, & f_{\rm ring} \le f < f_{\rm cut} \end{cases}, \tag{12.10}$$

$$\Psi_{\rm eff}(f) = 2\pi f t_{\rm c} + \Psi_0 + \sum_{k=0}^{7} \psi_k f^{(k-5)/3}. \tag{12.11}$$

In these expressions, the notation has been simplified by letting $\tilde{f} \to f$. Furthermore, $f_{\rm merg}$ and $f_{\rm ring}$ represent the characteristic frequencies marking the beginning of the merger and ringdown stages respectively, while $f_{\rm cut}$ corresponds to the point where the waveform is terminated. Also, C is an overall normalization that can be obtained by matching the inspiral section of the waveform to the its PN approximation; in the case of an optimally-oriented binary (i.e. aligned with the detector), this is:

$$C = \frac{M^{\frac{5}{6}} f_{\rm merg}^{-\frac{7}{6}}}{\pi^{\frac{2}{3}} r} \left(\frac{5\eta}{24}\right)^{\frac{1}{2}}, \tag{12.12}$$

for the *asymetric mass ratio* $\eta = m_1 m_2 / M^2$. The last line of Eq. (12.10) is a Lorentzian $\mathcal{L}(f, f_{\rm ring})$ weighted by a constant w such that:

$$w \equiv \frac{\pi\sigma}{2} \left(\frac{f_{\rm ring}}{f_{\rm merg}}\right)^{-2/3}, \tag{12.13}$$

$$\mathcal{L}(f, f_{\rm ring}, \sigma) \equiv \frac{1}{2\pi} \frac{\sigma}{(f - f_{\rm ring})^2 + \sigma^2/4}. \tag{12.14}$$

Finally, and most importantly, the ψ_k coefficients in Eq. (12.11) are phenomenological phase parameters obtained by matching to NR waveforms.

Note that more recent parametrization incorporate better description of the ringdown by supplementing the Lorentzian with an exponential decay [33, 34]. The ability to incorporate precession due to component spins has also been added [35].

12.3 Phenomenology

We wish to compute the waveform of a gravitational wave (GW) following a generic dispersion relation of the form:

$$|\mathbf{p}|^2 = g_{\mu\nu} p^\mu p^\nu = -f(\mathbf{p}) \,, \tag{12.15}$$

where $f(\mathbf{p})$ is, for the moment, an arbitrary function of the four-momentum $\mathbf{p}$, and so that GR is recovered for $f(\mathbf{p}) = 0$.

A modified dispersion relation, like Eq. (12.15), will manifest itself as a frequency-dependent phase shift in the observed gravitational waveform.

12.3.1 Redshift and dispersion

We would like to consider cosmological sources, so it is necessary to consider a background Friedmann-Lemaitre-Robertson-Walker (FLRW) metric:

$$\mathrm{d}s^2 = -c^2\mathrm{d}t^2 + a^2(t)\left[\mathrm{d}\chi^2 + \Sigma^2(\chi)\left(\mathrm{d}\theta^2 + \sin^2\theta\,\mathrm{d}\phi^2\right)\right] . \tag{12.16}$$

This metric presents a Killing tensor $K_{\mu\nu} = a^2(t)\left(g_{\mu\nu} + U_\mu U_\nu\right)$, with $\mathbf{U} \equiv (1,0,0,0)$ the four-velocity of a comoving observer. Therefore, the quantity

$$K^2 \equiv K_{\mu\nu}p^\mu p^\nu = a^2(t)\left[g_{\mu\nu}p^\mu p^\nu + \left(g_{\mu\nu}U^\mu p^\nu\right)^2\right] \tag{12.17}$$

is conserved along the trajectory of a particle with 4-momentum $\mathbf{p}$.

For the case of a graviton of (local) energy $E = p^0 c$, emitted radially at $(t = t_\mathrm{e}, \chi = \chi_\mathrm{e})$ and received at $(t = t_\mathrm{a}, \chi = 0)$, we have that

$$\mathbf{p} = \left(p^0, p^\chi, 0, 0\right) = \left(\frac{E}{c}, p^\chi, 0, 0\right) , \tag{12.18}$$

so Eq. 12.17 becomes:

$$K^2 = a^2(t)\left[-c^2\left(p^0\right)^2 + a^2(t)\,(p^\chi)^2 + c^2\left(p^0\right)^2\right] = \left[a^2(t)\,p^\chi\right]^2 = \left(p_\chi\right)^2 . \tag{12.19}$$

In other words, $p_\chi = a^2(t)\,p^\chi$ is conserved (but not p^χ). Taking advantage of this, we can write:

$$p_\chi = a^2(t)\,p^\chi = p_\chi(t_\mathrm{e}) = a^2(t_\mathrm{e})\,p^\chi(t_\mathrm{e}) \tag{12.20}$$

Using this result, Eq. (12.15) implies:

$$g_{00}\left(p^0\right)^2 + g_{11}\,(p^\chi)^2 = -f(\mathbf{p}) \implies -E^2c^{-2} + a^2(t)\,(p^\chi)^2 = -f(\mathbf{p}) , \tag{12.21}$$

which yields the following three equivalent expressions:

$$E = c\sqrt{f(\mathbf{p}) + a^2(t)\,(p^\chi)^2}, \tag{12.22}$$

$$p^{\chi} = -a^{-1}(t)\sqrt{E^2c^{-2} - f(\mathbf{p})}, \tag{12.23}$$

$$p_{\chi} = -a(t)\sqrt{E^2c^{-2} - f(\mathbf{p})}. \tag{12.24}$$

Note these are all implicit relations between $\vec{p}$ and E because, in its most generic form, $f(\mathbf{p})$ may contain terms dependent on both $\vec{p}$ and E.

12.3.2 Delay

Consider a wave packet emitted radially towards the observer from a source at $\chi = \chi_{\rm e}$. The group velocity is given by Eq. (12.18):

$$v_g \equiv \frac{\mathrm{d}\chi}{\mathrm{d}t} = \frac{\mathrm{d}\chi}{\mathrm{d}\lambda}\frac{\mathrm{d}\lambda}{\mathrm{d}t} = p^{\chi}\frac{c}{p^t} = c^2\frac{p^{\chi}}{E} = c^2\frac{k^{\chi}}{\omega}, \tag{12.25}$$

where χ marks the packet's location, λ is an affine parameter and $\mathbf{p}$, $\mathbf{k}$ are respectively the four momentum and wave vector associated with the wave. (Note that the factor of c in the third equality appears simply because $p^0 = \mathrm{d}x^0/\mathrm{d}\lambda$ and $x^0 = ct$.) Because of the expanding background, ω and k^{χ} are not constant; to make the time dependence explicit, use Eq. (12.20) to write Eq. (12.25) as:

$$v_g = \frac{\mathrm{d}\chi}{\mathrm{d}t} = \frac{c^2}{a^2(t)}\frac{k_{\chi}}{\omega(t;k_{\chi})}, \tag{12.26}$$

and recall $\omega(t;k_{\chi})$ is given by means of the result for the energy derived above, Eq. (12.21), with $E = \hbar\omega$ as usual.

We recover GR by setting $f(\mathbf{p}) = 0$, which means:

$$-\omega^2c^{-2} + a^2(t)(k^{\chi})^2 = 0 \implies -\omega^2c^{-2} + a^{-2}(t)(k_{\chi})^2 = 0 \implies k_{\chi} = -a(t)\omega c^{-1}, \tag{12.27}$$

where I picked the negative root because the wave propagates toward the origin; thus, for GR, Eq. (12.26) becomes the familiar result for the cosmological redshift:

$$\left.\frac{\mathrm{d}\chi}{\mathrm{d}t}\right|_{\rm GR} = \frac{c}{a(t)}\,. \tag{12.28}$$

Consider now a wave packet of energy E' emitted radially from $\chi_{\rm e}$ at time $t'_{\rm e}$ and second one with energy E emitted from the same location at time $t_{\rm e} = t'_{\rm e} + \Delta t_{\rm e}$. We can integrate Eq. (12.26) for each wave packet and equate both expressions. For the first one we have simply:

$$\chi_{\rm e} = c^2k'_{\chi}\int_{t'_{\rm e}}^{t'_{\rm a}}\mathrm{d}t\; a^{-2}(t)\left[\omega(t;k'_{\chi})\right]^{-1}, \tag{12.29}$$

while for the second one:

$$\begin{aligned}\chi_{\rm e} &= c^2 k_\chi \int_{t_{\rm e}}^{t_{\rm a}} {\rm d}t\; a^{-2}(t)\omega^{-1}(t;k_\chi) = c^2 k_\chi \int_{t'_{\rm e}+\Delta t_{\rm e}}^{t'_{\rm a}+\Delta t_{\rm a}} {\rm d}t\; a^{-2}(t)\omega^{-1}(t;k_\chi) \\ &= c^2 k_\chi \left[\int_{t'_{\rm e}}^{t'_{\rm a}} {\rm d}t\; a^{-2}(t)\omega^{-1}(t;k_\chi) - \int_{t'_{\rm e}}^{t'_{\rm e}+\Delta t_{\rm e}} {\rm d}t\; a^{-2}(t)\omega^{-1}(t;k_\chi) \right. \\ &\qquad \left. + \int_{t'_{\rm a}}^{t'_{\rm a}+\Delta t_{\rm e}} {\rm d}t\; a^{-2}(t)\omega^{-1}(t;k_\chi)\right] \end{aligned} \tag{12.30}$$

Assuming $a(t)$ does not change significantly in the short times $\Delta t_{\rm e}$ and $\Delta t_{\rm a}$, this may be approximated as:

$$\begin{aligned}\chi_{\rm e} \approx c^2 k_\chi \left[\int_{t'_{\rm e}}^{t'_{\rm a}} {\rm d}t\; a^{-2}(t)\omega^{-1}(t;k_\chi) - \Delta t_{\rm e}\; a^{-2}(t'_{\rm e})\omega^{-1}(t'_{\rm e};k_\chi) \right. \\ \left. + \Delta t_{\rm a}\; a^{-2}(t'_{\rm a})\omega^{-1}(t'_{\rm a};k_\chi)\right], \end{aligned} \tag{12.31}$$

Equating these two expressions for $\chi_{\rm e}$, we obtain the key result

$$\begin{aligned}\Delta t_{\rm a} = \frac{a^2(t'_{\rm a})\omega(t'_{\rm a};k_\chi)}{k_\chi} \int_{t'_{\rm e}}^{t'_{\rm a}} {\rm d}t\; a^{-2}(t) \left[\frac{k'_\chi}{\omega(t;k'_\chi)} - \frac{k_\chi}{\omega(t;k_\chi)}\right] \\ + \Delta t_{\rm e}\; \frac{a^2(t'_{\rm a})}{a^2(t'_{\rm e})} \frac{\omega(t'_{\rm a};k_\chi)}{\omega(t'_{\rm e};k_\chi)} \end{aligned} \tag{12.32}$$

Now, in GR, $f(\mathbf{p}) = 0$ means $-\omega^2 c^{-2} + a^{-2}(t)k_\chi^2 = 0$, so we can write:

$$\omega(t)|_{\rm GR} = c k_\chi a^{-1}(t), \tag{12.33}$$

$$k_\chi\big|_{\rm GR} = a(t_{\rm a})\omega(t_{\rm a})c^{-1} = \omega(t_{\rm a})c^{-1} \tag{12.34}$$

where, in the last equality, we have used the fact that the present scale factor is defined to be unity: $a(t_{\rm a}) = a_0 = 1$. The GR expressions can be used to get a general prescription for cases in which the departure from GR is small ($f(\mathbf{k})/|\mathbf{k}|^2 << 1$). In that case, we can write ω and k_χ as expansions around the GR value:

$$\omega(t;k_\chi) = \frac{c k_\chi}{a(t)}\left[1 + \delta_\omega(t;k_\chi)\right] = \frac{c k_\chi}{a(t)}\left[1 + \delta_\omega(t;\omega_{\rm a})\right], \tag{12.35}$$

$$k_\chi = a(t_{\rm a})\frac{\omega(t_{\rm a})}{c}\left(1 + \delta_{k_\chi}\right) = \frac{\omega_{\rm a}}{c}\left(1 + \delta_{k_\chi}\right), \tag{12.36}$$

for dimensionless perturbations δ_ω and δ_{k_χ} If we then apply the redshift formula $1 + z = a_{\rm a}/a_{\rm e}$, the time delay between the two wave–packets simplifies to (suppressing

the functional dependence of δ_ω):

$$\begin{aligned}\Delta t_\mathrm{a} &= \Delta t_\mathrm{e}\,(1+z)^2\,\frac{a(t'_\mathrm{e})}{a(t'_\mathrm{a})} + c\int\frac{\mathrm{d}t}{a^2(t)}\left[\frac{a(t)}{c(1+\delta'_\omega)} - \frac{a(t)}{c(1+\delta_\omega)}\right] \\ &= \Delta t_\mathrm{e}\,(1+z) + \int\frac{\mathrm{d}t}{a(t)}\left[(1-\delta'_\omega) - (1-\delta_\omega)\right], \end{aligned} \tag{12.37}$$

and we obtain the final expression for the delay in the time of arrival of the two wave packets:

$$\Delta t_\mathrm{a} = \Delta t_\mathrm{e}\,(1+z) + \int\frac{\mathrm{d}t}{a(t)}\left[\delta_\omega(t;\omega_\mathrm{a}) - \delta_\omega(t;\omega'_\mathrm{a})\right]. \tag{12.38}$$

Note that I have written $\delta_\omega(t;\omega_\mathrm{a})$ rather than $\delta_\omega(t;k_\chi)$ because k_χ is related to ω_a by means of Eq. (12.36). In general, one obtains $\delta_\omega(t;\omega_\mathrm{a})$ by plugging in Eq. (12.36) into the expression for $\delta_\omega(t;k_\chi)$; however, because we are only interested in first–order deviations to the phase, it suffices to use the GR relation $k_\chi = \omega_\mathrm{a}/c$. In other words, I define: $\delta_\omega(t;\omega_\mathrm{a}) \equiv \delta_\omega(t;k_\chi = \omega_\mathrm{a}/c)$

Formally, this is valid because we can characterize the modification to the dispersion by some small parameter ϵ, the precise form of which depends on the particular theory and that must exist if the deviation is to be perturbative. If this is the case, then we can write:

$$\begin{aligned}\omega &\approx \omega_\mathrm{GR}\left[1+\delta_\omega(t;k_\chi)\right] \\ &= \omega_\mathrm{GR}\left[1+\epsilon f(t;k_\chi)\right] \\ &\approx \omega_\mathrm{GR}\left[1+\epsilon f(t;k_{\chi,\mathrm{GR}}[1+\epsilon g(\omega_\mathrm{a})])\right] \\ &\approx \omega_\mathrm{GR}\left[1+\epsilon f(t;k_{\chi,\mathrm{GR}})\,(1+\epsilon h(t;\omega_\mathrm{a}))\right] \\ &\approx \omega_\mathrm{GR}\left[1+\epsilon f(t;k_{\chi,\mathrm{GR}})\right] \\ &= \omega_\mathrm{GR}\left[1+\delta_\omega(t;k_{\chi,\mathrm{GR}})\right] \end{aligned} \tag{12.39}$$

keeping terms to first order in ϵ and with the functions $f(t;k_\chi) = \partial_\epsilon\omega|_{\epsilon=0}$, $g(\omega_\mathrm{a}) = \partial_\epsilon k_\chi\big|_{\epsilon=0}$ and $h(t;\omega_\mathrm{a}) = \partial_\epsilon f(t;k_\chi)\big|_{\epsilon=0}$ derived from the appropriate series expansions. Thus, we are indeed justified in writing Eq. (12.38) with $\delta_\omega(t;\omega_\mathrm{a}) \equiv \delta_\omega(t;k_\chi = \omega_\mathrm{a}/c)$.

12.3.3 Phase shift

We can translate the time delay of Eq. (12.38) into a phase shift by means of Eq. (12.8). Modifying the notation slightly (exchanging $f \leftrightarrow \tilde{f}$), the Fourier phase is given by

$$\begin{aligned}\Psi(f_{\rm a}) &= 2\pi \int_{f_{\rm c,a}}^{f_{\rm a}} (t_{\rm a} - t_{\rm c,a}) {\rm d}\tilde{f}_{\rm a} + 2\pi f_{\rm a} t_{\rm c,a} + \Psi_0 \\ &= 2\pi \int_{f_{\rm c,e}}^{f_{\rm e}} (t_{\rm a} - t_{\rm c,a})(1+z)^{-1} {\rm d}\tilde{f}_{\rm e} + 2\pi f_{\rm a} t_{\rm c,a} + \Psi_0 \,, \end{aligned} \tag{12.40}$$

where, in the last line, I have written out the integral in terms of the emitted frequencies. Identifying $\Delta t_{\rm a} = t_{\rm a} - t_{\rm c,a}$, be can now use Eq. (12.38) to write:

$$\begin{aligned}\Psi(f_{\rm a}) &= \Psi_0 + 2\pi f_{\rm a} t_{\rm c} + 2\pi \int_{f_{\rm c,e}}^{f_{\rm e}} \Delta t_{\rm e} {\rm d}\tilde{f}_{\rm e} + \\ &+ \frac{2\pi}{1+z} \int_{f_{\rm c,e}}^{f_{\rm e}} \int_{t_{\rm e}}^{t_{\rm a}} \frac{1}{a(t)} \left[\delta_f(t;\tilde{f}_{\rm a}) - \delta_f(t; f_{\rm a,c})\right] {\rm d}t {\rm d}\tilde{f}_{\rm e}, \\ &= \Psi_0 + 2\pi f_{\rm a} t_{\rm c} + 2\pi \int_{f_{\rm c,e}}^{f_{\rm e}} \Delta t_{\rm e} {\rm d}\tilde{f}_{\rm e} + \\ &+ 2\pi \int_{f_{\rm c,a}}^{f_{\rm a}} \int_{t_{\rm e}}^{t_{\rm a}} \frac{1}{a(t)} \left[\delta_f(t;\tilde{f}_{\rm a}) - \delta_f(t; f_{\rm a,c})\right] {\rm d}t {\rm d}\tilde{f}_{\rm a}, \end{aligned} \tag{12.41}$$

where I have defined $\delta_f(t; f_{\rm a}) \equiv \delta_\omega(t; 2\pi f_{\rm a})$. (Note that, if we instead defined δ_f by $f = f_{\rm GR}(1+\delta_f)$, then $\delta_f = \delta_\omega$.) Consequently, we can again simplify our notation ($f_{\rm a} \to f$) to write:

$$\Psi(f) = \Psi_{\rm GR}(f) + \Delta\Psi(f), \tag{12.42}$$

with the deviation from GR given by

$$\Delta\Psi(f) = \int_{f_{\rm c}}^{f} \int_{t_{\rm e}}^{t_{\rm a}} \frac{2\pi}{a(t)} \left[\delta_f(t;\tilde{f}) - \delta_f(t; f_{\rm c})\right] {\rm d}t {\rm d}\tilde{f}. \tag{12.43}$$

This phase shift can now be used to obtain the generic waveform implied by $f(\mathbf{p})$, starting from the corresponding GR waveform. Note that this modification can be applied to the whole waveform, even trough the merger and ringdown stages, like in Eq. (12.9).

12.4 Examples

In this section, we derive the phase shift corresponding to example dispersion relations: a massive graviton and the simplest possible anisotropic correction.

12.4.1 Massive graviton

In a theory in which the graviton is endowed with a mass $m_{\rm g}$ (e.g. [291]), the linearized GW dispersion relation over the FLRW metric becomes

$$-\frac{\omega^2}{c^2}+a^{-2}(t)k_\chi^2+\frac{m_{\rm g}^2c^2}{\hbar^2}=0\,, \tag{12.44}$$

and, accordingly, the angular frequency can be written as:

$$\omega=\frac{ck_\chi}{a(t)}\sqrt{1+\frac{a^2(t)}{k_\chi^2}\frac{m_{\rm g}^2c^2}{\hbar^2}}\approx\frac{ck_\chi}{a(t)}\left(1+\frac{1}{2}\frac{a^2(t)}{k_\chi^2}\frac{m_{\rm g}^2c^2}{\hbar^2}\right), \tag{12.45}$$

where I have assumed that the graviton wavelength far exceeds the GW wavelength ($\lambda_{\rm g}\gg\lambda$) and kept only terms to first order in $m_{\rm g}^2$ to obtain the second relation. Using this, we can now write $\omega=\omega_{\rm GR}(1+\delta_\omega)$ as in Eq. (12.35), with

$$\delta_\omega(t;\omega_{\rm a})=\frac{1}{2}\frac{a^2(t)}{k_\chi^2}\frac{m^2c^2}{\hbar^2}\approx\frac{a^2(t)}{2}\frac{m^2c^4}{\hbar^2\omega_{\rm a}^2}, \tag{12.46}$$

where I have replaced k_χ with its GR value, i.e. $k_\chi\to w_{\rm a}/c$, to keep terms only of first order in $m_{\rm g}^2$ (as justified above). Plugging this into Eq. (12.43), we obtain the desired phase shift:

$$\begin{aligned}\Delta\psi(f)&=\int_{f_{\rm c}}^{f}\int_{t_{\rm e}}^{t_{\rm a}}\frac{2\pi}{a(t)}\frac{a^2(t)}{2}\frac{m^2c^4}{\hbar^2}\frac{1}{4\pi^2}\left(\frac{1}{\tilde f^2}-\frac{1}{f_{\rm c}^2}\right){\rm d}t{\rm d}\tilde f\\&=\frac{m_{\rm g}^2c^4}{4\pi\hbar^2}\int_{t_{\rm e}}^{t_{\rm a}}a(t){\rm d}t\int_{f_{\rm c}}^{f}\left(\frac{1}{\tilde f^2}-\frac{1}{f_{\rm c}^2}\right){\rm d}\tilde f\\&=-\frac{m_{\rm g}^2c^3}{4\pi\hbar^2}\frac{D}{1+z}\left(\frac{1}{f}+\frac{f}{f_{\rm c}}-\frac{2}{f_{\rm c}}\right),\end{aligned} \tag{12.47}$$

where I have defined D as in [327], namely:

$$D\equiv(1+z)c\int_{t_{\rm e}}^{t_{\rm a}}a(t){\rm d}t. \tag{12.48}$$

One could choose to stop here. However, inspection of Eq. (12.11) reveals that terms linear in f are degenerate with $t_{\rm c}$, while those independent of f can be absorbed by the fiducial phase Ψ_0. Ignoring those degenerate terms, our final expression of the massive graviton phase shift is then:

$$\Delta\psi(f)=-\frac{m_{\rm g}^2c^3}{4\pi\hbar^2}\frac{D}{1+z}\frac{1}{f}=-\frac{\pi D}{1+z}\frac{\lambda}{\lambda_{\rm g}^2}, \tag{12.49}$$

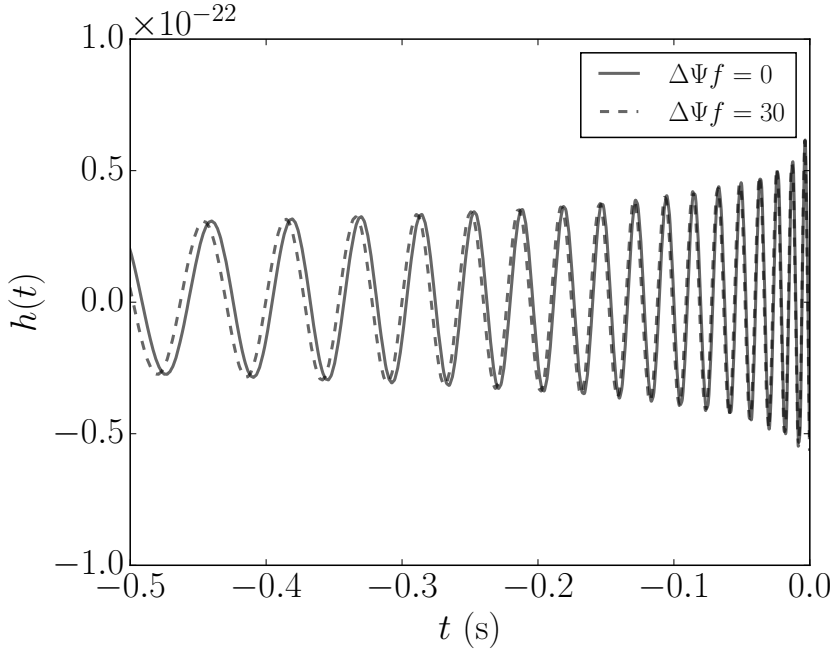

Figure 12.1: Dispersed (dashed) and undispersed (solid) example waveforms.

for $\lambda_g = h/(m_g c)$ and $\lambda = c/f$. The effect of this extra phase shift is represented in Fig. 12.1 for a very *large* dispersion of $\Delta\Psi = 30\,\mathrm{Hz}\, f^{-1}$.

This is equivalent to the result from [327]. Indeed, after restoring factors of c, Eqs. (3.8c) and (3.9) in [327] imply:

$$\Delta\Psi_{\mathrm{Will}} = -\beta u^{-1} = -\frac{\pi^2 D\mathcal{M}}{\lambda_g^2(1+z)}\frac{c}{\pi\mathcal{M}f} = -\frac{\pi D}{1+z}\frac{\lambda}{\lambda_g^2}\,, \tag{12.50}$$

where I have used $\beta = \pi^2 D\mathcal{M}/\lambda_g^2(1+z)$ and $u = \pi\mathcal{M}f/c$, as defined in [327].

12.4.2 First anisotropic term

The simplest anisotropic modification to the dispersion relation can be obtained by defining adding a term of the form $\vec{v}\cdot\hat{k}$, for an arbitrary spatial vector $\vec{v}$:

$$-\frac{\omega^2}{c^2} + a^{-2}(t)k_\chi^2 + v^\chi = 0\,, \tag{12.51}$$

where $v^\chi = v^i k_i/(k^j k_j) = -v^i n_i$ for $\hat{n} = -\hat{k}$ giving the location of the source in the sky. Since, v^χ is constant by construction, this extra term behaves in the same way as that for the massive graviton and is thus degenerate with it when looking at a single source at a time. Explicitly, the phase shift for this sort of anisotropic dispersion will

be

$$\Delta\psi(f) = -\frac{c^2 v^\chi}{4\pi}\frac{D}{1+z}\frac{1}{f}, \tag{12.52}$$

in full analogy to Eq. (12.49). As before, this result hinges on the assumption that the extra term in the dispersion relation (in this case, v^χ) can be treated as a perturbation over the GR solution.

The degeneracy between this anisotropic effect and a massive graviton can be broken by simultaneously analyzing multiple events coming from different locations in the sky. In order to do this, it is necessary to coherently study all detected events in a joint Bayesian-inference analysis—this is computationally challenging with current techniques, but feasible if the signals are short.

12.5 Conclusion

We have generalized previous work in [325, 327] to provide a recipe for the observable waveform-dephasing effect corresponding to a generic dispersion, Eq. (12.15), with corrections $f(\mathbf{p})$ depending arbitrarily on the graviton four-momentum $\mathbf{p}$. The steps are outlined in Sec. 12.3.2, with the main result being Eq. (12.38) and Eq. (12.43). We demonstrated the application of the recipe with two simple examples (massive graviton and a simple anisotropic dispersion). In the future, this recipe could be used to implement data analysis techniques to constrain arbitrary dispersion relations, as has been done for a massive graviton [82] and a Lorentz-violating correction [4]. The implementation of this test for very complex corrections (high number of terms in $f(\mathbf{p})$) might require new computational techniques.

As an important caveat, not that the derivations provided here, as those in [325, 327], relied on heuristic arguments considering the group velocities of two gravitons in a thought experiment. It would not be surprising if a proper analysis, solving the wave equation corresponding to an arbitrary dispersion relation for the right boundary conditions, yielded different results—in fact, early work seems to indicate that this might be the case [330].

Chapter 13

DIRECTED SEARCHES FOR GRAVITATIONAL WAVES FROM ULTRALIGHT BOSONS

M. Isi, L. Sun, R. Brito, and A. Melatos, "Directed searches for gravitational waves from ultralight bosons," (2018), M.I. conceived the project and led the writing of the manuscript, arXiv:1810.03812 [gr-qc] .

13.1 Introduction

After decades of dedicated effort, the Advanced Laser Interferometer Gravitational-wave Observatory (aLIGO) [331] and Advanced Virgo [332] detectors have inaugurated gravitational-wave (GW) astronomy with the observation of several compact-binary coalescences (CBCs) [1–6, 178]. Arguably, one of the most exciting prospects in this new era of astronomy is to use gravitational waves to learn about fundamental physics. Common examples of this are attempts to probe the nature of gravity by testing general relativity [82, 333], or to probe the nature of nuclear matter through the neutron-star equation of state [6, 334, 335]. Another exciting possibility is that of using gravitational waves to learn about particle physics. In particular, it may be possible to search for new ultralight bosons with gravitational-wave detectors, a prospect that has recently garnered much attention [336–343]. In this paper, we explore the potential of achieving this using directed searches for continuous gravitational waves with ground-based detectors.

There are strong theoretical reasons to believe in the existence of new weakly-interacting, ultralight scalar (spin 0) or vector (spin 1) particles. The prime example of this is the axion, a (pseudo-)scalar particle originally proposed to explain the strong constraints on the existence of charge-parity (CP) violating terms in the strong nuclear force sector [344–346]. This quantum-chromodynamics (QCD) axion is among the best-motivated extensions of the standard model, but there are others. For example, string theory predicts the existence of a variety of axion-like particles (potentially including the QCD axion) with masses populating each decade between 10^{-33} eV and 10^{-10} eV as a result of the compactification of extra spatial dimensions [347]. Similarly, a hidden sector of light vector particles also naturally arises in

compactifications of string theory [348]. Besides motivations from particle and high-energy physics, these bosons are also popular dark matter candidates (see e.g. [349–351]).

Because of their weak couplings to the standard model and their vanishingly-small mass, all these proposed new particles would be extremely hard to detect by conventional means. In particular, all existing constraints on the existence of the QCD axion rely on its expected coupling to standard model particles, a property that is heavily model-dependent [352]. Such observations loosely constrain the mass of the QCD axion to be $\lesssim 10^{-3}$ eV [353], but values of order 10^{-10} eV or lower are favored by theory [336]. For other kinds of conjectured ultralight bosons, whose potential interactions with the standard model are very weak or inexistent, constraints and detection can only provided through their gravitational coupling.

Given the substantial challenge in detecting ultralight bosons, there has been considerable excitement about the proposal to look for these particles by taking advantage of the universal character of gravitational couplings. The idea hinges on the phenomenon of superradiance [354–358], by which macroscopic clouds of these bosons should form around rapidly spinning black holes (BHs) and, in turn, produce a varied set of observational signatures [336–343, 359–363]. Indeed, constraints based on BH spin measurements of x-ray binaries have already been put, and exclude roughly the mass interval $[10^{-12}, 10^{-11}]$ eV for non-interacting massive scalar fields [340, 360] and $[10^{-13}, 10^{-11}]$ eV for non-interacting massive vector fields [343, 360]. Constraints derived from observations of x-ray binaries should however be interpreted with caution, since there is large uncertainty about the age and history of these systems, as well as caveats about the systematics affecting their spin measurements [364, 365].

A much more clean observational signature is the emission of potentially-detectable gravitational waves at a frequency of roughly twice the boson mass [339–343]. As it turns out, this means that clouds formed around stellar mass BHs should emit signals within the most sensitive band of ground-based detectors, probing boson masses in the theoretically interesting mass-range of the order of 10^{-15} eV through 10^{-11} eV [339, 342].

In this paper, we explore the prospect for the direct detection of continuous gravitational waves emitted by boson clouds. In particular, we focus on searches in data from present and future ground-based detectors directed at known BHs. As the main observational scenario, we consider the followup of remnants from compact-binary

coalescences detected through gravitational waves, but also examine BH candidates known from electromagnetic observations. Besides treating the data analysis, we study in detail the morphology of boson signals and use numerical calculations, combined with the latest analytic results, to estimate their amplitude and other relevant features. This allows us to more accurately predict the potential signals that may be expected from clouds around a given BH. In the case of scalar clouds, we use those estimates to obtain detection horizons for second-generation and proposed third-generation instruments. Along the way, we review the theoretical basics of BH-boson superradiance in a language that we hope facilitates future work by gravitational-wave analysts interested in the topic.

In Sec. 13.2, we review the theory of boson clouds around BHs and the gravitational-wave emission mechanism. In Sec. 13.3 we discuss at length the specific morphology of the expected signals, as seen by ground-based detectors. In Sec. 13.4, we introduce hidden Markov model tracking as an ideal strategy to search for these signals, validate the method and estimate the sensitivity through Monte-Carlo simulations; we also discuss implications for the followup of compact-binary mergers and x-ray binaries as potential sources. Finally, we provide a summary and conclusions in Sec. 13.5.

13.2 The boson cloud

The physics of boson fields around BHs has been extensively studied in different limits using both analytic and numerical methods [336, 342, 347, 357, 358, 366–371]. We summarize those results below, in a way that is best suited for the data-analysis framework introduced later in the paper. We explain how a macroscopic boson cloud spontaneously arises around fast-spinning BHs and proceeds to emit large amounts of gravitational radiation, providing some essential mathematical detail.

Consider first a Kerr BH of mass M and angular momentum J. The characteristic length associated with the BH mass will be $r_g \equiv GM/c^2$, or half the Schwarzschild radius. The other characteristic length, given by the BH spin, is the usual Kerr parameter, $a \equiv J/(Mc)$, from which we can in turn define the dimensionless spin, $\chi \equiv ac^2/(GM)$. In terms of these quantities, the radius of the hole's outer horizon, in Boyer-Lindquist coordinates, is

$$r_+ = r_g \left(1 + \sqrt{1-\chi^2}\right) \equiv r_g\, \bar{r}_+ , \tag{13.1}$$

where $\bar{r}_+$ is defined here to be dimensionless. At this location, the BH will then have

a frame-dragging angular speed (with respect to infinity) of

$$\Omega_{\rm BH} = \frac{1}{2}\frac{c}{r_g}\frac{\chi}{1+\sqrt{1-\chi^2}} \equiv \frac{c}{r_g}\,\overline{\Omega}_{\rm BH}\,, \tag{13.2}$$

where $\overline{\Omega}_{\rm BH}$ is defined here to be dimensionless. (See, e.g., [372] for a recent review of the Kerr metric.)

Now, imagine that, beyond the usual particles in the standard model, there exists an ultralight boson of mass

$$m_b \equiv \mu/c^2\,, \tag{13.3}$$

where μ is the boson's rest energy. The corresponding length and time scales are given by the Compton wavelength, $\lambda_\mu \equiv 2\pi\lambda\!\!\!^{-}_\mu \equiv h/(m_b c)$, and angular frequency, $\omega_\mu = c/\lambda\!\!\!^{-}_\mu = \mu/\hbar$. As soon as the BH is born,[1] (perhaps, as a result of stellar collapse or a binary coalescence) the usual quantum fluctuations in the boson field will cause pairs of particles to spontaneously appear in the hole's vicinity, causing a number of them to unavoidably fall in. What happens next will depend on the properties of the BH and the infalling excitations: under most circumstances, the particle will simply disappear behind the horizon never to return; however, for the right sets of parameters, the excitation in the boson field will scatter off the BH with a boost in amplitude, effectively increasing the number of particles (occupation number) around the BH [354–356, 358, 373–375].

From the second law of BH thermodynamics [356] (or more generic kinematic arguments [347, 358, 375]), we may expect the boson-wave amplification to occur when the following *superradiance condition* is satisfied:

$$\omega_\mu/m < \Omega_{\rm BH}\,, \tag{13.4}$$

where m is the (magnetic) quantum number corresponding to the projection of the particle's *total* angular momentum along the BH spin direction. This amplification extracts energy from the BH just as in the classical Penrose process [354], which is itself another manifestation of BH superradiance (see [358] for a review). Because of the field's nonzero mass, a scattered boson will generally tend to be bound to the BH, attracted by its gravitational pull. Consequently, scattered particles may remain confined in that region, facilitating successive scatterings and the associated compounded amplification of the field. This process is similar to the "BH bomb"

[1] Meaning, as soon as it is sufficiently close to the ideal Kerr metric.

devised by Press and Teukolsky [355], with the mirror replaced by the boson mass [367–369, 376–378].

We may anticipate that the boson confinement will be maximal when the field and BH have comparable characteristic lengthscales, i.e. $\lambda_\mu \sim r_g$. If this is the case, then the field amplitude will grow at an exponential rate in a (quasi-)bound state around the BH (Sec. 13.2.1). As the field grows, it draws energy and angular momentum from the BH until the condition of Eq. (13.4) is no longer satisfied. Because the field is bosonic, there are no limits to the occupation number of any given energy level: the number of particles in a superradiant state will grow exponentially to form a macroscopic "cloud", with a mass of up to order ∼10% of the BH's [379, 380]. This cloud will slowly fade away, as its energy is radiated away in the form of gravitational waves over very long timescales compared to the superradiant rate (Sec. 13.2.2).

13.2.1 Black-hole and boson interactions

Energy levels

The qualitative picture laid out above is backed up by analytic and semi-analytic calculations of boson fields over a Kerr background [336, 342, 347, 357, 358, 366–369]. In the nonrelativistic ($\alpha/j \ll 1$) regime implied by Eq. (13.4), the influence of the BH is effectively reduced to a simple inverse-radius gravitational potential. This potential causes the bosons to present quasi-bound energy eigenstates essentially identical to those in the hydrogen atom, but with gravity replacing electromagnetism as the relevant interaction. More carefully solving the Schrödinger equation over a Kerr background, one indeed finds that, in this regime, the system has hydrogenic energy levels [366, 368],

$$E_{\bar{n}} \approx \mu \left(1 - \frac{1}{2}\frac{\alpha^2}{\bar{n}^2} + \dots\right), \tag{13.5}$$

for $\bar{n} = n + l + 1$ the principal quantum number, n the radial quantum number, and l the *orbital* azimuthal quantum number.[2] As usual, we have $|j - s| \leq l \leq j + s$ and $-j \leq m \leq j$, where j and s are respectively the total and spin angular-momentum quantum numbers. All superradiant levels are hydrogenic with a spectrum well-described by Eq. (13.5) [369]. The quantity α in Eq. (13.5) plays exactly the same

[2] Although the spin parameter does not appear at leading order in α [369], the fact that the BH is spinning does affect the angular part of the boson eigenfunctions: these have to be described using spin-weighted *spheroidal*, rather than spherical, harmonics [342] (see Sec. 13.3). Higher-order corrections to the energy eigenvalues, including corrections due to the BH spin, can be found in [361].

role as the fine-structure constant in the hydrogen atom, and takes the value of the ratio of the two relevant lengthscales (or, equivalently, timescales):

$$\alpha \equiv \frac{r_g}{\lambda\!\!\!^{-}_{\mu}} = \frac{GM}{c}\frac{m_b}{\hbar} = \frac{GM}{c^3}\omega_\mu\,, \tag{13.6}$$

where $\lambda\!\!\!^{-}_{\mu} \equiv \lambda_\mu/(2\pi)$. Importantly, Eq. (13.4) implicitly constrains α as a function of the BH spin. If we want superradiance to take place, then Eq. (13.4) and Eq. (13.2) demand:

$$\alpha < \frac{1}{2}m\chi\left(1+\sqrt{1-\chi^2}\right)^{-1} < \frac{m}{2}\,, \tag{13.7}$$

where the second inequality is obtained by noting $0 \leq \chi < 1$. Because $m \leq j$, this condition justifies working in the nonrelativistic, $\alpha < j$, limit in the first place [336].

Cloud growth

Unlike the hydrogen atom, however, the BH-boson system is non-Hermitian due to the ingoing boundary condition for waves at the horizon. This means that the occupation number of the different energy eigenstates need not be constant—in fact, they will most certainly *not* be so for the superradiant states we are interested in. For small α/m, the occupation number[3] of a given quantum state will grow exponentially at a rate that may be analytically approximated as [343, 366, 368]:

$$\Gamma_{jlmn} \approx 2\alpha^{2j+2l+5}\bar{r}_+\left(m\Omega_{\rm BH}-\omega_\mu\right)C_{jlmn}\,, \tag{13.8}$$

with C_{jlmn} a dimensionless factor, and $\bar{r}_+$ the dimensionless radius defined in Eq. (13.1). In the case of a scalar boson, the orbital angular momentum is necessarily the total angular momentum ($j = l$), and it can be shown that [366, 368]:

$$\begin{aligned} C^{(\rm scalar)}_{jlmn} = {} & \frac{2^{4l+2}(2l+n+1)!}{(l+n+1)^{2l+4}n!}\left[\frac{l!}{(2l)!(2l+1)!}\right]^2 \\ & \times\prod_{k=1}^{l}\left[k^2\left(1-\chi^2\right)+\frac{4r_+^2}{c^2}\left(m\Omega_{\rm BH}-\omega_\mu\right)^2\right]. \end{aligned} \tag{13.9}$$

Expressions for vector C_{jlmn} up to leading order in α can be found in Appendix A of [343]. The validity of this approximation in the regime of interest has been confirmed numerically for scalars [369] and, more recently, vectors [360, 371, 381–383].

Irrespective of boson spin, there are three key features of the occupation growth rate, Γ in Eq. (13.8), that can be distilled from the above results:

[3] The rate of change of the occupation number is *twice* that of the field amplitude, which is itself given by the imaginary part of the wavefunction frequency: $\Gamma = 2\,\mathrm{Im}(\omega)$.

1. the sign of Γ depends solely on $\left(m\Omega_{\rm BH} - \omega_\mu\right)$, implying that indeed energy levels satisfying Eq. (13.4) will grow exponentially, while others will be depleted;

2. Γ is a high power of α, growing with the sum $j + l$;

3. for a fixed angular momentum (j, l) and α, Γ decreases rapidly with n.

Because we are working in the small-α limit, these three facts mean that, for a given system (i.e. a given α and χ), *the fastest growth will occur for the fundamental radial harmonic of the level with the* smallest *possible total angular-momentum, j, that still supports a magnetic number, $|m| \leq j$, sufficiently large to satisfy Eq.* (13.4). In other words, if the boson has spin-weight $s = 0, 1$, the level with the fastest superradiant growth in a given system will have angular quantum numbers $\{j,\, l,\, m\}$ given by

$$j = l + s = m = \text{ceil}(\alpha/\overline{\Omega}_{\rm BH})\,, \tag{13.10}$$

where "ceil" stands for the operation of rounding up to the closest integer. In particular, the fastest-possible level over *all* values of α and χ will then be

$$j = l + s = m = 1\,,\; n = 0\,. \tag{13.11}$$

Given this, it follows from Eq. (13.8) that vector clouds will tend to grow significantly faster than scalar ones ($\Gamma^{(\rm v)}/\Gamma^{(\rm s)} \sim \alpha^{-2}$).

Final state

As the particle number grows, the energy and angular momentum required to populate the boson energy levels are extracted from the BH.[4] Consequently, the BH quickly loses mass and spin until Eq. (13.4) is asymptotically saturated and the growth rate, Eq. (13.8), vanishes. As implied by Eq. (13.7), the spin of the BH at the end of this process will then be

$$\chi_f = \frac{4\alpha_f m}{4\alpha_f^2 + m^2}\,, \tag{13.12}$$

where α_f is given by Eq. (13.6) for the *final* BH mass. If no other processes (like accretion) take place in the relevant timescale, the final mass of the cloud (M_c) will simply be given by the difference between the initial (M_i) and final (M_f) BH masses.

[4] As this happens, Eq. (13.4) guarantees that the BH area increases, satisfying the second law of thermodynamics [356, 375].

If only one level is populated, then it may be shown that this will be approximately [342]

$$M_c = M_i - M_f \approx M_i \frac{\alpha_i \chi_i}{m}, \tag{13.13}$$

with the last equality being valid for $\alpha \lesssim 0.1$.

A more exact value for this quantity may be obtained by numerically solving a set of difference equations, e.g. Eqs. (17)–(21) in [357], assuming a quasi-adiabatic evolution. If superradiant growth is the dominant factor, so that we can ignore other processes like GW emission and accretion, these are just

$$\dot{M} = -\Gamma_{jlmn} M_c\,, \tag{13.14a}$$

$$\dot{M}_c = -\dot{M}\,, \tag{13.14b}$$

$$\dot{J} = -m\Gamma_{jlmn}\omega_{\bar{n}}^{-1} c^2 M_c \tag{13.14c}$$

$$\dot{J}_c = -\dot{J}\,, \tag{13.14d}$$

where dots indicate time derivatives, $(M,\ J)$ and $(M_c,\ J_c)$ are the instantaneous mass and angular momentum for the BH and boson cloud respectively. In the following sections, we will use the exact value for M_c computed this way to characterize the signal. In any case, it may be shown that the boson cloud may reach a size of at most $M_c \approx 0.1 \times M_i$ [379, 380].

The time it takes a single-level cloud to grow to its full size is simply the time it takes the BH to reach the spin of Eq. (13.12). In the absence of significant interaction with the environment (e.g. through accretion or strong gravitational-wave emission), this is inversely linked to the "instability timescale" implied by of Eq. (13.8), namely

$$\tau_{\rm inst} \equiv 1/\Gamma_{jlmn}\,. \tag{13.15}$$

This corresponds to an e-folding in the occupation number of level (j, l, m, n). In the nonrelativistic limit ($\alpha \ll 1$), $\tau_{\rm inst}$ can be approximated for the dominant scalar level ($l = m = 1, n = 0$) by [342]

$$\tau_{\rm inst}^{(\rm s)} \approx 27\,{\rm days} \left(\frac{M}{10\,M_\odot}\right)\left(\frac{0.1}{\alpha}\right)^9 \frac{1}{\chi_i}, \tag{13.16}$$

which is generally slower than the timescale for the dominant vector level ($j = m = 1$, $n = 0$), approximated by [343]

$$\tau_{\rm inst}^{(\rm v)} \approx 2\,{\rm minutes} \left(\frac{M}{10\,M_\odot}\right)\left(\frac{0.1}{\alpha}\right)^7 \frac{1}{\chi_i}, \tag{13.17}$$

in the same small-α limit.

As it turns out, gravitational interactions between levels might prevent the simultaneous population of more than one state [336]. Whether this is true or not, the hierarchy of involved timescales also suggests that we need only consider single-level clouds. For instance, for most parameters, the $l = m = 2$ scalar level has an instability timescale larger or comparable to the depletion timescale of the fastest-growing level $l = m = 1$, so we should expect the latter to be unoccupied by the time the former reaches any significant size. The same holds for vectors, as discussed in [343]. This means that we need only consider a single level at a time: as soon as the BH is born, the fastest-growing level will be quickly populated; because the field has integer spin, there is no limit to how many particles can occupy it, and the growth will continue until the BH reaches the spin of Eq. (13.12); the next level, with second-largest Γ, will begin to grow only after the first one is depleted (through gravitational wave emission, as discussed below, or any other reason).

All of the results in this section were obtained from perturbative analyses that consider the nonrelativistic behavior of a boson field over a static Kerr background. These do not take into account effects like back-reaction of the field onto the background metric, gravitational wave emission by the boson condensate, or interaction between energy levels. However, the validity of Eqs. (13.5)–(13.8) is confirmed by full numerical relativity simulations for the case of vector fields [370, 371, 379, 384].

13.2.2 Gravitational-wave emission

Once the boson cloud has reached a macroscopic size, it will emit a significant amount of gravitational radiation. There are three main mechanisms by which this may happen: (i) emission due to annihilation of bosons into gravitons; (ii) boson transitions between energy levels, analogous to electron jumps in the hydrogen atom; and (iii) abrupt collapse of the cloud due to particle self-interactions ("bosenova").

Due to the high occupation numbers involved, the first two processes can be described purely classically, with GW emission stemming from a time-varying quadrupole (and higher-multipoles, to a lesser degree) in the cloud's stress energy. Transitions only become important if more than one level is occupied with comparable numbers. Therefore, it could take over thousands of years after the birth of the BH for such a signal to become detectable [339, 340], making transitions interesting for very old BHs only. Unfortunately, the typical duration of transition signals would be of order years or shorter [339], which makes their observation from old potential

sources highly unlikely. Meanwhile, bosenovae are only relevant in the presence of large boson self-interactions [385, 386]; in particular, they are not expected to occur for the QCD axion [339]. In any case, the typical duration of bosenova signals would be of the order of milliseconds and the numerical simulations required to produce their waveforms are still in their infancy [385, 386]—this makes bosenovae a possibly relevant for unmodeled-burst analyses (e.g., [113, 387, 388]), rather than the continuous-wave searches we are concerned with.

Given the above considerations, we restrict ourselves to annihilation signals, which are the best understood and most relevant for ground-based detectors. To understand this, consider a BH that has been maximally spun down due to the growth of the boson cloud surrounding it. We should expect this cloud to be composed of a vast number of particles in a coherent state corresponding to the fastest-growing energy level, as determined by Eq. (13.8). Indeed, we may think of the cloud as a macroscopic object with particle density given by the norm of the boson wavefunction, and rotating with angular frequency $\omega_{\bar{n}} = m\Omega_{\rm BH}$, as implied by saturation of Eq. (13.4). Treating this object purely classically, we may then expect the cloud to radiate gravitational waves at twice its rotational frequency [48].

A detailed description of the cloud's gravitational-wave emission can be obtained by using the Teukolsky formalism to solve the linearized Einstein equations for the cloud's stress energy, given by the wavefunction of the relevant quantum state [342]. As a result, one indeed finds that the cloud emits gravitational waves with angular frequency

$$\tilde{\omega} = 2\omega_{\bar{n}}\,, \tag{13.18}$$

and that the emission pattern is described by a set of spin-weighted spheroidal harmonics corresponding to the spin of the final BH, with azimuthal numbers $\tilde{l} \geq 2j$, and magnetic quantum number *fixed* to $\tilde{m} = 2m$ (see Sec. 13.3.1). Both vectors and scalars emit GWs with the same angular pattern (for a given j and m), with the fastest-growing level radiating mostly in the $\tilde{l} = \tilde{m} = 2$ mode.

The gravitational power radiated in each angular mode $\tilde{l}$ may be written as

$$\dot{E}_{\rm GW}(\tilde{l}, \tilde{m}, \tilde{\omega}) = \frac{1}{4\pi}\frac{c^5}{G}\left(\frac{c}{r_g\tilde{\omega}}\right)^2\left(\frac{M_c}{M_f}\right)^2 \mathcal{A}^2_{\tilde{l}\tilde{m}}(\alpha, \chi_i)\,, \tag{13.19}$$

where M_f and $r_g = GM_f/c^2$ are respectively the mass and lengthscale of the *final* BH, and $\mathcal{A}_{\tilde{l}\tilde{m}}(\alpha, \chi_i)$ is a dimensionless factor. For scalars, we compute the $\mathcal{A}_{\tilde{l}\tilde{m}}$'s numerically using BH perturbation theory as in [342],[5] but this is not currently

[5]Note that $\mathcal{A} = |Z| \times (M^2/M_c)$, for $|Z|$ as defined in [342].

feasible for vectors. Regardless of boson spin, as long as Eq. (13.7) is satisfied, the emitted power will be a steep function of α [337, 339, 342]. In fact, for small α, the power emitted in the dominant angular mode ($\tilde{l} = 2$) by the fastest-growing scalar level ($l = m = 1$, $n = 0$) may be roughly approximated as [357]

$$\dot{E}_{\rm GW}^{(\rm s)} \approx 7 \times 10^{41}\,{\rm erg/s} \left(\frac{\alpha}{0.1}\right)^{16} \chi_i^2\,, \tag{13.20}$$

and for the fastest-growing vector level ($j = m = 1$, $l = n = 0$) as [343]

$$\dot{E}_{\rm GW}^{(\rm v)} \approx 2 \times 10^{49}\,{\rm erg/s} \left(\frac{\alpha}{0.1}\right)^{12} \chi_i^2\,. \tag{13.21}$$

The difference in the α dependence in these two expressions arises from the fact that the fastest vector level has no orbital angular momentum ($l = 0$), and so lies closer to the BH yielding a more compact cloud.

The energy in the gravitational radiation is drawn from the cloud itself, which slowly fades away as its component particles annihilate into gravitons [336, 339]. As a result, the signal will be almost monochromatic, with a slowly decreasing amplitude and slowly increasing frequency. The evolution of both these quantities is tied to the timescale implied by Eq. (13.19),

$$\tau_{\rm GW} \equiv M_c c^2 / \dot{E}_{\rm GW}\,. \tag{13.22}$$

This "gravitational-wave timescale" is just the time it takes for half of the rest-energy of the cloud to be radiated away, and can be thought of as the typical duration of the signal. Using the approximations of Eq. (13.13) and Eq. (13.20), we get for the dominant scalar level in the nonrelativistic limit, a signal duration of

$$\tau_{\rm GW}^{(\rm s)} \approx 6.5 \times 10^4\,{\rm yr} \left(\frac{M}{10\,M_\odot}\right) \left(\frac{0.1}{\alpha}\right)^{15} \frac{1}{\chi_i}. \tag{13.23}$$

Similarly, using Eq. (13.21), for the dominant vector level we get

$$\tau_{\rm GW}^{(\rm v)} \approx 1\,{\rm day} \left(\frac{M}{10\,M_\odot}\right) \left(\frac{0.1}{\alpha}\right)^{11} \frac{1}{\chi_i}. \tag{13.24}$$

Clearly, the vector processes tend to take place at a much faster pace than scalar ones, as expected from the higher radiated power. In both cases, however, the duration of the signal is significantly shorter than the time it takes for the cloud to grow, as given by Eq. (13.15). This an important, general feature that justifies the separate treatment of the early growth and late emission stages in the first place.

13.3 The signal

Having reviewed the physics of boson clouds around BHs in Sec. 13.2, we will now focus on the properties of the gravitational signal produced by one of these systems, as seen by differential-armlength detectors on the ground. Given that we only expect one quantum state to be significantly populated at any given time, we will restrict our discussion to signals from single-level clouds. In any case, the signal from a multilevel cloud can be produced trivially by the addition of several single-level waveforms described below, assuming negligible interaction between levels. We describe the signal morphology in Sec. 13.3.1 and elaborate on the most salient features in Sec. 13.3.2.

13.3.1 Waveform

Consider a cloud made up of bosons in a single quantum state that has just stopped growing, after drawing enough energy and angular momentum from its host BH to saturate Eq. (13.4). In that case, as anticipated in Sec. 13.2.2, we expect the cloud to emit a continuous gravitational signal with a small spin-up and amplitude depending on the properties of both BH and boson. The strain signal, h_I, seen by a given differential-armlength detector, I, can be written in the usual form as a sum over polarizations,

$$h^I(t) = F^I_+(t)\, a_+ \cos\phi(t) + F^I_\times(t)\, a_\times \sin\phi(t), \tag{13.25}$$

where the F^I_p's are the antenna-response functions of detector I to signals of plus (+) and cross (×) polarizations (see, e.g., Appendix B in [168] for explicit expressions). These depend implicitly on the relative location and orientation of the detector and the source, usually parametrized by its right ascension ($\alpha_\star$), declination ($\delta_\star$), and polarization angle (ψ). This last parameter determines how the frame in which the polarizations are defined is oriented in the plane of the sky; for our purposes, this will be the angle between the spin of the BH and the projection of the celestial North onto the plane normal to the line of sight.

The polarization amplitudes, a_+ and $a_\times$, are made up of contributions from several angular multipoles, indexed by the wave azimuthal number, $\tilde{l} \geq 2j$, and with fixed magnetic number, $\tilde{m} = 2m$,

$$a_{+/\times} = -\sum_{\tilde{l}\geq 2l} h_0^{(\tilde{l})} \left[{}_{-2}S_{\tilde{l}\tilde{m}\tilde{\omega}} \pm {}_{-2}S_{\tilde{l}-\tilde{m}-\tilde{\omega}} \right], \tag{13.26}$$

with the plus (minus) sign on the right-hand side corresponding to the + (×) polarization [342]. Because the boson cloud does not emit GWs isotropically, these amplitudes depend on the orientation of the source relative to the detector.

The angular dependence is encoded in Eq. (13.26) via the *spin-weighted spheroidal harmonics* [389], which are analogous to the usual spin-weighted spherical harmonics but account for the nonsphericity of the space around the Kerr BH. As such, these are functions of BH spin, signal frequency and orientation with respect to the source:

$$ {}_sS_{\tilde{l}\tilde{m}\tilde{\omega}} \equiv {}_sS_{\tilde{l}\tilde{m}}(a\tilde{\omega}/c, \cos\iota), \tag{13.27}$$

where for us the spin weight is always $s = -2$, as needed to describe GWs. The inclination ι is defined as the angle between the BH spin and the line of sight. We compute these eigenfunctions numerically using Leaver's method [390, 391].

The characteristic amplitude of each mode can be written as

$$h_0^{(\tilde{l})} = \frac{c^4}{G}\frac{M_c}{M_{BH}^2}\frac{1}{2\pi^2 f^2 r}\mathcal{A}_{\tilde{l}\tilde{m}}(\alpha, \chi_i), \tag{13.28}$$

where the dimensionless factor $\mathcal{A}_{\tilde{l}\tilde{m}}$ encodes the relative amount of energy that the source deposits in each mode, as in Eq. (13.19). Assuming the BH has been fully spun down by the cloud, this is only a function of the initial BH spin and the fine-structure constant α, and can be computed numerically from BH perturbation theory following [342]. In the nonrelativistic limit ($\alpha \ll 1$), this can be approximated by [337, 340, 357]

$$h_0^{(\mathrm{s})} \approx 8\times 10^{-28}\left(\frac{M}{10M_\odot}\right)\left(\frac{\alpha}{0.1}\right)^7\left(\frac{\mathrm{Mpc}}{r}\right)\left(\frac{\chi-\chi_f}{0.1}\right) \tag{13.29}$$

for the dominant scalar mode ($l = m = 1, n = 0, \tilde{l} = \tilde{m} = 2$), and by [343]

$$h_0^{(\mathrm{v})} \approx 4\times 10^{-24}\left(\frac{M}{10M_\odot}\right)\left(\frac{\alpha}{0.1}\right)^5\left(\frac{\mathrm{Mpc}}{r}\right)\left(\frac{\chi-\chi_f}{0.1}\right) \tag{13.30}$$

for the dominant vector mode ($j = m = 1, n = 0, \tilde{l} = \tilde{m} = 2$), corresponding to the approximations in Eq. (13.20) and Eq. (13.21) respectively. Equations Eq. (13.29) and Eq. (13.30) include an explicit spin-dependent correction factor to account for the dependence of the mass cloud on α, with χ_f itself a function of α defined in Eq. (13.12). Although not fully accurate, these expressions will be useful when studying the scalings of the expected signal amplitude—especially in the case of vectors, for which the $\mathcal{A}$ factors have yet to be computed numerically.

In both the scalar and vector cases, the amplitude of the signal will decrease as a function of time, starting from the peak values given by Eq. (13.28). This is due to the progressive dissipation of the boson cloud sourcing the GW signal. This weakening occurs over a timescale of the order of the signal duration, Eq. (13.22).

In a frame inertial with respect to the source, the phase evolution of the signal corresponds to a simple monotone that may potentially evolve slowly in frequency. In the frame of the detector, extra timing corrections are needed, so that, in terms of the time t measured at Earth, the phase evolution can be written as

$$\phi(t) = 2\pi \sum_{j=0}^{N} \frac{\partial_t^{(j)} f}{(j+1)!} [t - t_0 + \delta t(t)]^{(j+1)} + \phi_0 \,, \tag{13.31}$$

where $\partial_t^{(j)} f$ is the j^{th} time derivative of f, the GW frequency measured at fiducial time t_0, and ϕ_0 is a phase offset. The timing corrections $\delta t(t)$ account for delays due to the relative motion of the source and detector, general- and special-relativistic effects, as well as potential corrections due to the presence of a companion if the source is part of a binary (e.g. [75, 320, 392]).

As implied by Eq. (13.18), the source-frame frequency will be given by $f = \tilde{\omega}/(2\pi) = \omega_{\bar{n}}/\pi$. Noting that $\omega_{\bar{n}} \approx \omega_\mu$ by Eq. (13.5), this may be approximated as a function of BH mass and fine-structure constant as

$$f \approx \frac{\alpha}{\pi r_g} \approx 645\,\text{Hz} \left(\frac{10 M_\odot}{M}\right) \left(\frac{\alpha}{0.1}\right). \tag{13.32}$$

This means that stellar-mass BHs should support boson clouds that emit gravitational waves at frequencies within the sensitive band of existing and planned ground-based detectors.

Once the superradiance instability has shut down, the GW signal will expect a (slight) positive change in frequency (a "spinup"). This is expected on purely classical grounds, as is typical of any gravitationally-bound system (as in the characteristic "chirp" of compact-binary coalescences). The value of $\partial_t^{(1)} f \equiv \dot{f}$ can be computed from the rate of change in the cloud's binding energy [343]. For scalars, Eq. (13.20) implies a signal frequency derivative (see Appendix 13.6)

$$\dot{f}^{(\text{s})} \approx 3 \times 10^{-14}\,\text{Hz/s} \left(\frac{10 M_\odot}{M}\right)^2 \left(\frac{\alpha}{0.1}\right)^{19} \chi_i^2 \,, \tag{13.33}$$

while, for vectors Eq. (13.21) implies [343]

$$\dot{f}^{(\text{v})} \approx 1 \times 10^{-6}\,\text{Hz/s} \left(\frac{10 M_\odot}{M}\right)^2 \left(\frac{\alpha}{0.1}\right)^{15} \chi_i^2 \,. \tag{13.34}$$

The frequency drift is faster for vectors, as corresponds to quicker cloud dissipation. Besides this spinup, we will also want to allow for higher-order derivatives of the frequency in Eq. (13.31) to incorporate potential perturbations caused by the astrophysical environment, presence of a companion, level interactions, or theoretical uncertainty.

13.3.2 Projected properties

We are interested in making statements about the presence of ultralight bosons based on searches for GW signals from known BHs. This means that we need to know what strain frequencies and amplitudes we may expect from our target BH, without knowing the true mass of the boson (if it exists). In principle, a particular BH could support clouds for a range of boson masses, which we will parametrize implicitly via the fine-structure constant, α of Eq. (13.6). Although a particular single-level cloud is expected to emit a GW quasimonotone, an unknown α means that we could expect signals at a variety of frequencies. This is clear from Eq. (13.20), in the small-α limit, or more generally from the fact that the $\mathcal{A}_{\tilde{l}\tilde{m}}$ factors in Eq. (13.28) will be nonzero for a range of α's. In other words, a given BH could "resonate" with different bosons, allowing us to probe a (narrow) range of particle masses; hence there is a *band* of signal frequencies to be potentially expected from a given BH.

In this section, we study in detail the properties of continuous signals that can be expected from clouds around a given BH. Although so far we have kept the discussion general, we now focus on scalars to provide concrete examples. We use numerical techniques to compute the power emitted by different systems, obtaining estimates that should be more reliable than previously published projections. In particular, for each value of α and initial BH parameters, we numerically solve the differential equations governing the evolution of the cloud to obtain the final BH parameters, as in [357]. We then use the numerical results of [342] for the $\mathcal{A}_{\tilde{l}\tilde{m}}$ factors to compute the radiated amplitude by means of Eq. (13.28). Unfortunately, at the moment, it is not possible to do this for vectors, since the corresponding perturbative calculations are significantly more difficult and have yet to be carried out.

To summarize the key points from the discussion below: any given BH can allow us to probe a narrow range of boson masses set by its mass and, to a lesser extent, its spin (Figs. 13.1–13.3); heavier BHs will "resonate" with lighter bosons and produce louder signals at lower frequencies (Fig. 13.4); and signals from heavier BHs (lighter bosons) will both grow and vanish more slowly, resulting in smaller

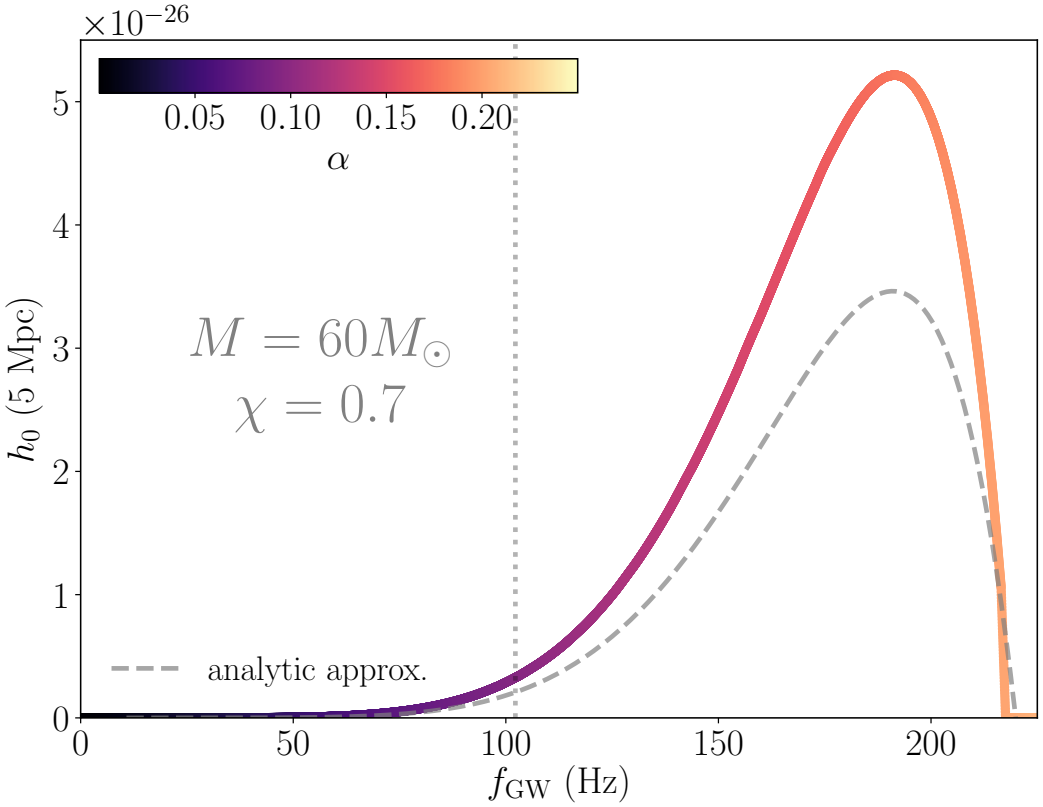

Figure 13.1: *Strain amplitude vs frequency for example black hole (scalar cloud).* The x-axis shows different frequencies at which we might expect a GW signal from a scalar ($l = m = 1$, $n = 0$) cloud around a BH with initial mass of $M_i = 60\ M_\odot$ and initial spin $\chi_i = 0.70$. The colored curve shows the corresponding characteristic amplitude, h_0 in Eq. (13.28) assuming $r = 5$ Mpc, parametrized by the fine-structure constant, α in Eq. (13.6), as indicated by the colorbar. For reference, the other curve shows the small-α approximation of Eq. (13.29), including the spin correction responsible for the amplitude turnover. Points to the left of the vertical dotted line have $\tau_{\rm inst} > 1$ yr.

frequency derivatives (Figs. 13.6–13.8). We expect similar conclusions to hold for vectors, except that the overall radiated power will be stronger and the timescales shorter, cf. Eq. (13.8) and Eq. (13.21).

Amplitude and frequency

Begin with the example of a BH with initial mass $M = 60\ M_\odot$ and dimensionless spin $\chi = 0.70$, parameters consistent with the remnant from LIGO's first detection [1]. Assume that the BH is then maximally spun down by the presence of a *scalar* cloud, and consider the amplitude of GWs emitted immediately after. We will assume the

cloud is dominated by the fastest-growing energy level ($l = m = 1$, $n = 0$), and restrict ourselves to the dominant GW mode ($\tilde{l} = \tilde{m} = 2$). For concreteness, place the source 5 Mpc away and consider the amplitude of the signal as seen from Earth for different values of α. Our BH could potentially support clouds emitting GWs at different frequencies and with different amplitudes, depending on the true value of the boson mass.

The above fact is clearly illustrated in Fig. 13.1, which shows the characteristic amplitude of waves produced by clouds around our example BH for different initial values of α. In the best case scenario for this BH, if there existed a boson with $\mu \approx 4 \times 10^{-13}\,\mathrm{eV}$ such that $\alpha \approx 0.179$, then we would observe a signal with characteristic strain amplitude $h_0 \approx 5.2 \times 10^{-26}(5\,\mathrm{Mpc}/r)$ at $f \approx 191\,\mathrm{Hz}$, corresponding to the peak in Fig. 13.1. For this value of the boson mass, at the end of the superradiant process the BH will have reached a final spin of $\chi = 0.62$, having lost 1.7% of its mass to the cloud.

Fig. 13.1 also implicitly defines the range of frequencies of interest for searches directed at this BH to be, say, within 150 Hz and 200 Hz. This range could be broader or narrower, depending on the sensitivity of the search and how long one waits from the birth of the BH to make an observation. For instance, points to the left of the dotted gray line correspond to boson masses for which the signal would take longer than 1 year to reach the amplitude shown (more on timescales in Sec. 13.3.2 below). According to Eq. (13.7), the maximum value of α for this source is ~0.2, for which the amplitude vanishes. Note that, because μ depends linearly on α, this means that any given BH will allow us to probe a very narrow range of boson masses.

While the overall shape of the curve in Fig. 13.1 will be generally the same for all BHs, the location and width of the peak will be a strong function of the initial BH mass and spin. This is represented in Fig. 13.2, in which we have fixed the BH distance and spin to the values above, but allowed its mass to vary. Color in this figure represents the strain amplitude of GWs emitted at a given frequency (y-axis) as a function of initial BH mass (x-axis), while the gray dashed line marks the peak amplitude for a given BH mass. Although α is not explicitly shown, it should be clear from Fig. 13.1 that moving vertically towards higher frequencies and amplitudes corresponds to increasing α. In fact, a vertical cut of Fig. 13.2 at $M_i = 60\,M_\odot$ would yield Fig. 13.1.

Fig. 13.2 shows that heavier BHs can support clouds that emit GWs at lower frequencies but greater amplitudes. This was expected from the discussion in Sec.

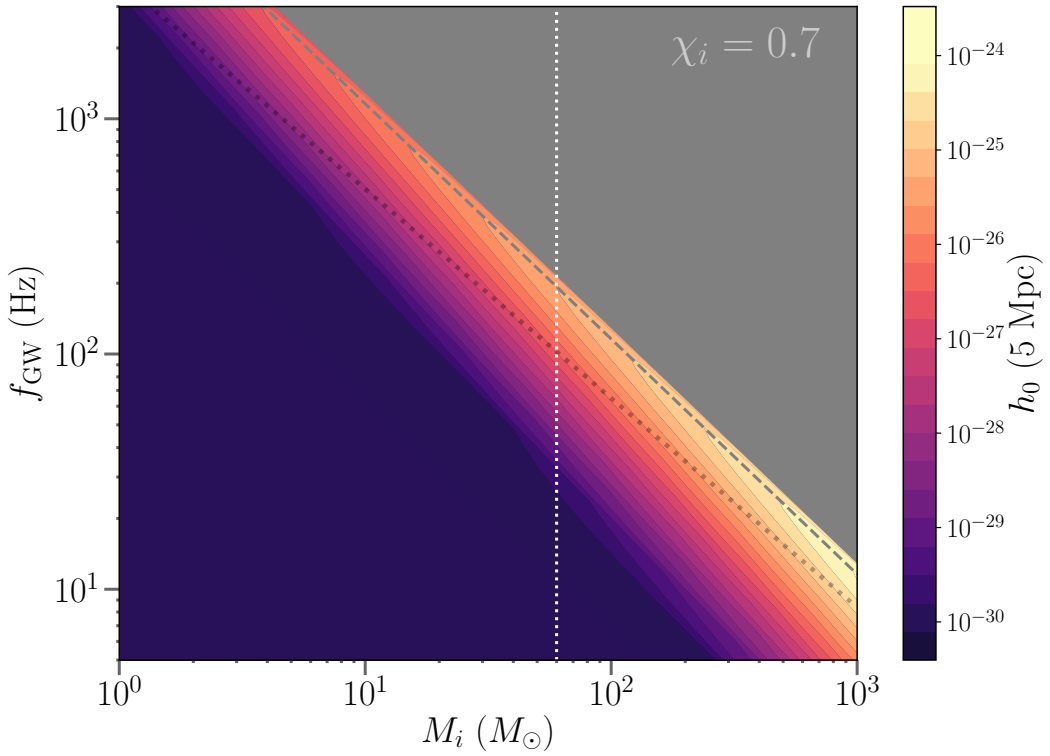

Figure 13.2: *Strain amplitude at different GW frequencies vs initial black-hole mass for fixed initial spin (scalar cloud).* Color shows the characteristic strain amplitude, Eq. (13.28), from a scalar cloud ($l = m = 1$, $n = 0$) that would be emitted at different frequencies (y-axis), i.e. for different α's (not shown), vs BH mass (x-axis). The gray dashed line marks the peak amplitude. The source is assumed to lie at $r = 5$ Mpc, with initial spin $\chi_i = 0.70$. For ease of display, we set an arbitrary lower cutoff of $h_0 \geq 10^{-30}$: amplitudes for points in the bottom purple region vanish asymptotically for lower f, while amplitudes for points in the top gray region vanish identically since superradiance cannot occur for those parameters. Points below the dotted gray line have $\tau_{\rm inst} > 1$ yr. A vertical cross-section at $M = 60\ M_\odot$ (vertical dotted line) yields Fig. 13.1.

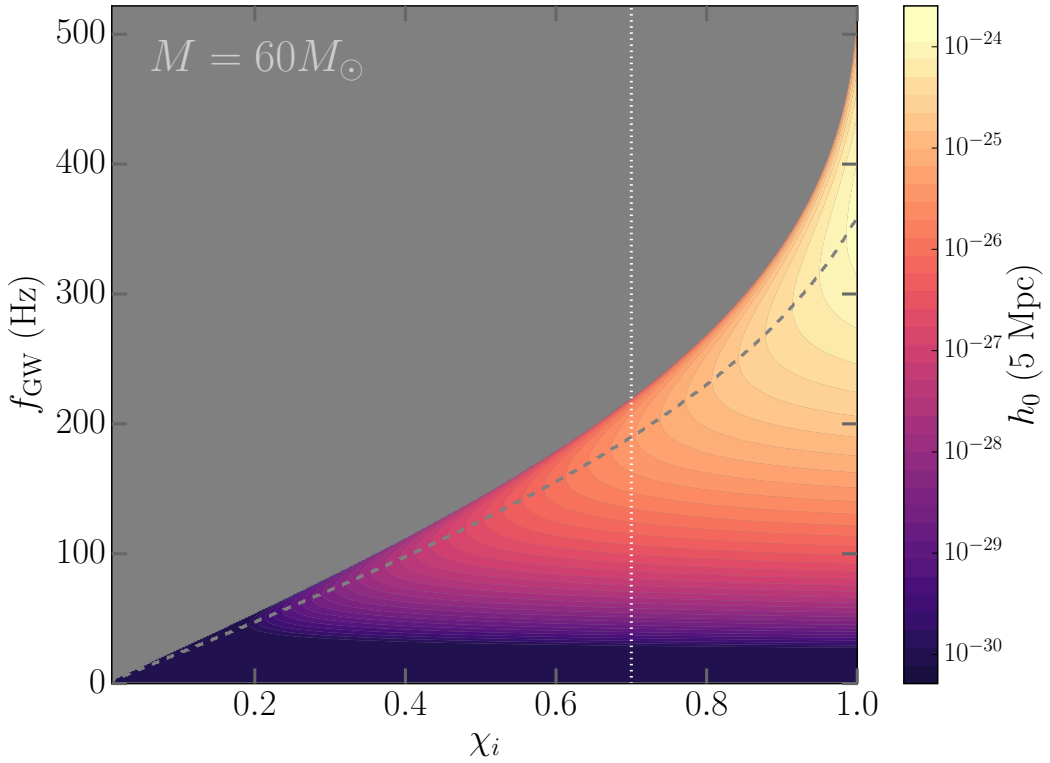

Figure 13.3: *Strain amplitude at different GW frequencies vs BH initial spin for fixed mass (scalar cloud).* Color shows the characteristic strain amplitude, Eq. (13.28), from a scalar cloud ($l = m = 1$, $n = 0$) that would be emitted at different frequencies (y-axis), i.e. for different α's (not shown), vs BH initial spin (x-axis). The gray dashed line marks the peak amplitude. The source is assumed to lie at $r = 5$ Mpc, with mass $M = 60\ M_\odot$. For ease of display, we set an arbitrary lower cutoff of $h_0 \geq 10^{-30}$: as can be inferred from Fig. 13.1, amplitudes for points in the bottom purple region vanish asymptotically for lower f, while amplitudes for points in the top gray region vanish identically since superradiance cannot occur for those parameters. A vertical cross-section at $\chi_i = 0.70$ (vertical dotted line) yields Fig. 13.1.

13.2.2: (i) heavier BHs are also larger, and so must be the boson cloud surrounding it, thus yielding lower GW frequencies; and (ii) heavier BHs result in a heavier cloud, as dictated by Eq. (13.13), which will in turn radiate more strongly, per Eq. (13.19). Because the overall radiated power decreases with BH mass, this also means that the band of detectable frequencies is narrower for lighter BHs, which is also visible in Fig. 13.2. The fact that the peak frequency (dashed line) decreases linearly with BH mass was already anticipated in Eq. (13.32). This can be understood from the observations that (i) $f \sim \omega/\pi \sim 2/\lambda\!\!\!^{-}_{\mu}$ and that (ii) $\lambda\!\!\!^{-}_{\mu} \sim r_g$ for the boson and BH sizes to match and maximize superradiance. As one moves vertically up the plot (increasing f or, equivalently, α), the emitted power vanishes abruptly at a point defined by the saturation of Eq. (13.7); in this case, because $\chi_i = 0.70$ and $m = 1$, this corresponds to $\alpha = 0.2$. Finally, as in Fig. 13.1, points below the dotted gray line would take longer than 1 year from the birth of the BH to reach the displayed amplitudes.

The properties of the GW emission will also vary with the initial spin of the BH. This is illustrated in Fig. 13.3, in which we have fixed the BH distance and mass ($M_i = 60\ M_\odot$), but allowed its initial spin to vary. As in Fig. 13.2, color represents the characteristic strain amplitude emitted at a given frequency (y-axis) for different values of the initial spin (x-axis); again, the dashed line traces the peak amplitude. It is no surprise to find that BHs with greater initial spins yield louder GWs: the faster the BH is spinning before the superradiant process kicks off, the longer the cloud may grow without saturating Eq. (13.4) and, consequently, the more mass it will extract before its growth stalls. The fact that higher-spin BHs result in heavier clouds is reflected in Eq. (13.13), and the corresponding dependence of h_0 on χ_i can be glimpsed from Eq. (13.20), which is valid for $\alpha \ll 1$. Similarly, for a fixed initial BH mass, a faster initial spin results in a lighter final BH [see Eq. (13.13) again], and so yields higher GW emission frequencies per Eq. (13.32). As in Fig. 13.2, the upper frequency cutoff is given by Eq. (13.12).

As suggested by Fig. 13.2 and Fig. 13.3, the characteristic GW amplitude emitted by a boson cloud as a function of frequency may show interesting structure as the initial BH mass and spin are varied. However, in many situations, it suffices to know the expected amplitude of the *peak* emission from a given system. This information is summarized in Fig. 13.4, which displays the characteristic amplitude and frequency for the optimal cloud as a function of BH mass, and for different values of the initial spin. The curves in the bottom panel can be understood as constant-spin cuts of the

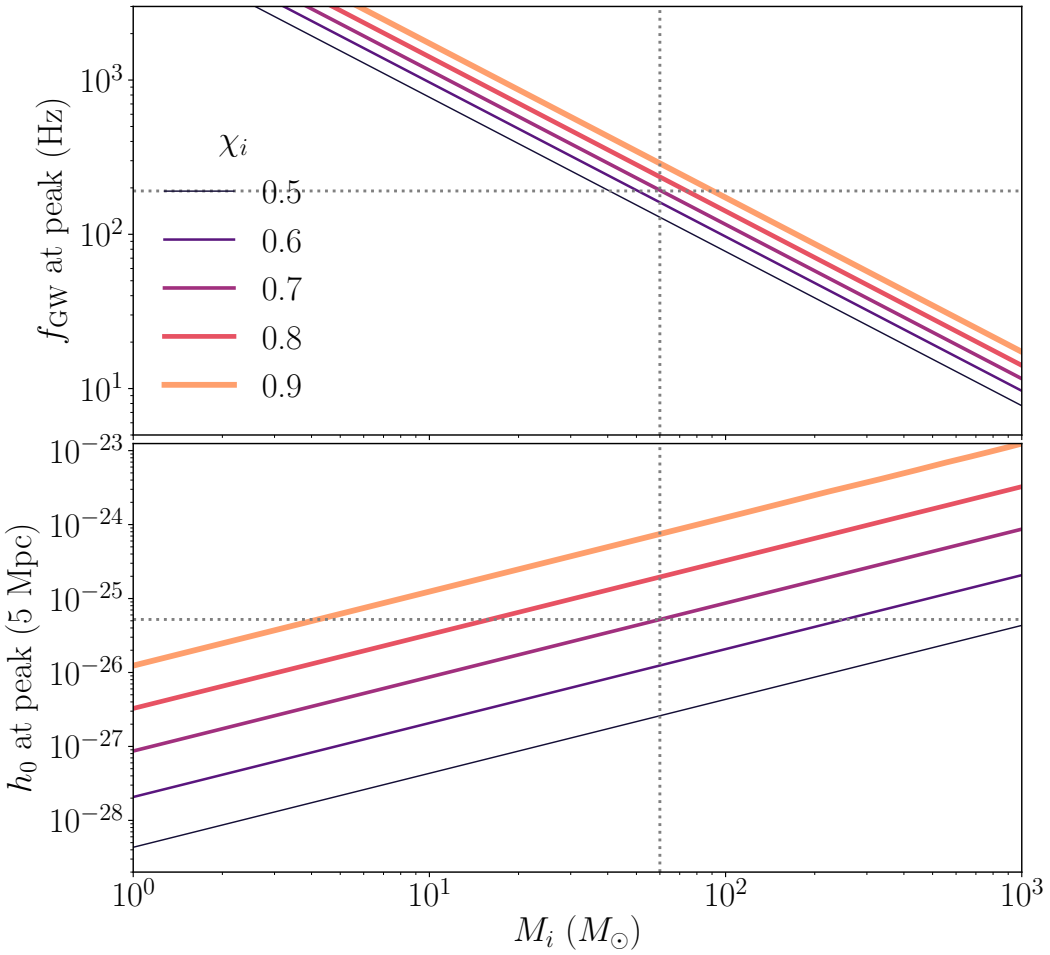

Figure 13.4: *Optimal strain frequency and amplitude vs initial BH mass for different initial spins (scalar cloud).* Frequency (top) and characteristic amplitude (bottom) of the strain produced by the best-possible cloud (best-possible α) as a function of initial BH mass. Different curves correspond to different initial spins, showing that higher spins result in stronger emission. We assume that the source is situated at $r = 5$ Mpc, and that the scalar cloud is dominated by the fastest level ($l = m = 1$, $n = 0$). The intersection of the $\chi_i = 0.70$ line with a vertical cut at $M = 60\ M_\odot$ (dotted vertical line) give the amplitude and frequency of the peak in Fig. 13.1 (dotted horizontal lines).

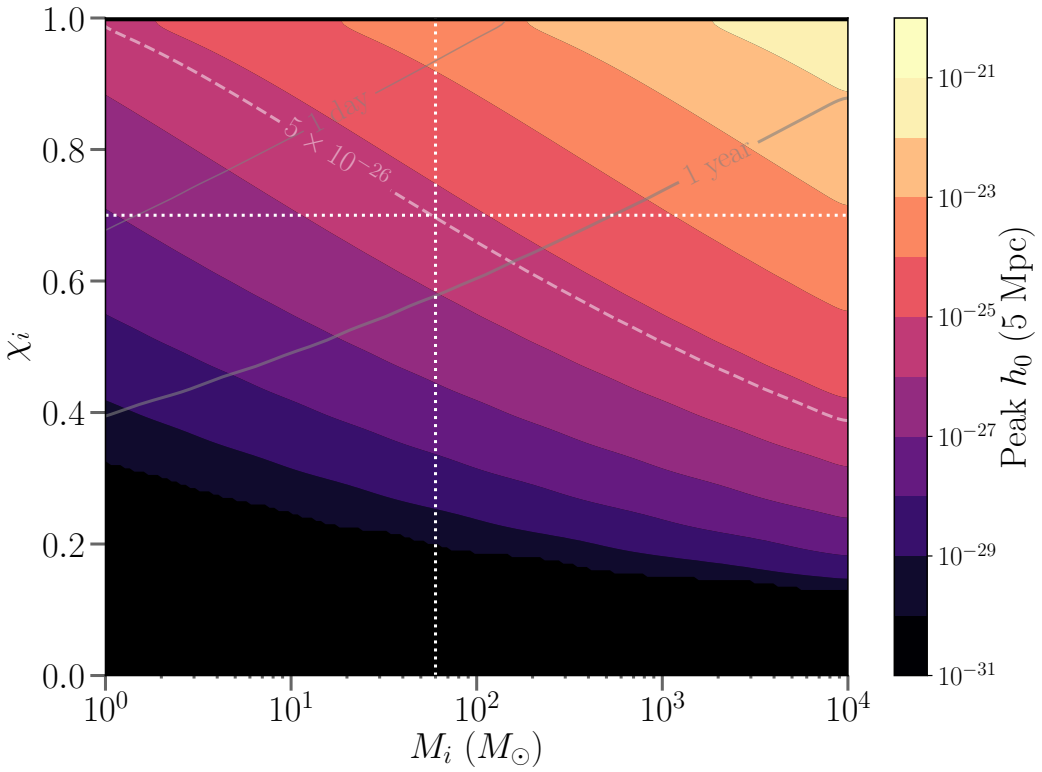

Figure 13.5: *Optimal strain amplitude vs initial BH parameters (scalar cloud).* Color gives the characteristic strain amplitude emitted by the best-possible cloud matched to a BH with the indicated initial mass (x-axis) and spin (y-axis). Horizontal cuts yield the curves shown in the bottom panel of Fig. 13.4. The intersection of the dotted white lines ($M_i = 60\ M_\odot$, $\chi_i = 0.70$) corresponds to the peak of Fig. 13.1. Gray lines mark $\tau_{\rm inst} = 1$ day and $\tau_{\rm inst} = 1$ yr for reference.

full mass-spin plane shown in Fig. 13.5. As for the other colormaps, the dotted white lines in that plot mark the values of our example BH ($M_i = 60\ M_\odot$, $\chi_i = 0.70$), which can at best yield an amplitude of $h_0 = 5.2 \times 10^{-26}(5\ \text{Mpc}/r)$ (peak of Fig. 13.1). Gray lines mark representative values of the instability timescale of Eq. (13.15) (see Sec. 13.3.2). Fig. 13.4 and Fig. 13.5 once again reflect the fact that greater strains are obtained for heavier BHs with larger initial spins. Some representative values are shown in Table 13.1, where the bold row corresponds to the intersection of the dotted lines in Fig. 13.2.

Timescales

The figures discussed so far provide important information about the expected strain as a function of frequency when searching for signals from a given BH, but it is

Table 13.1: Parameters of optimal scalar cloud for representative BHs. A "k" next to a value stands for "$\times 10^3$". The bold row corresponds to the intersection of the dotted lines in Fig. 13.2.

M_i	χ_i	μ	α_i	f	h_0	$\tau_{\rm inst}$	$\tau_{\rm GW}$
$M_\odot$		10^{-13} eV		Hz	5Mpc/r	day	yr
3	0.90	122	0.273	5.8k	4×10^{-26}	0.1	2
10	0.90	36	0.273	1.7k	1×10^{-25}	0.3	6
60	**0.70**	**4.0**	**0.179**	**191**	$\mathbf{5\times10^{-26}}$	**39**	**8k**
60	0.90	6.0	0.273	290	7×10^{-25}	2	38
200	0.85	1.6	0.243	77	1×10^{-24}	12	511
300	0.95	1.4	0.311	66	8×10^{-24}	4	40

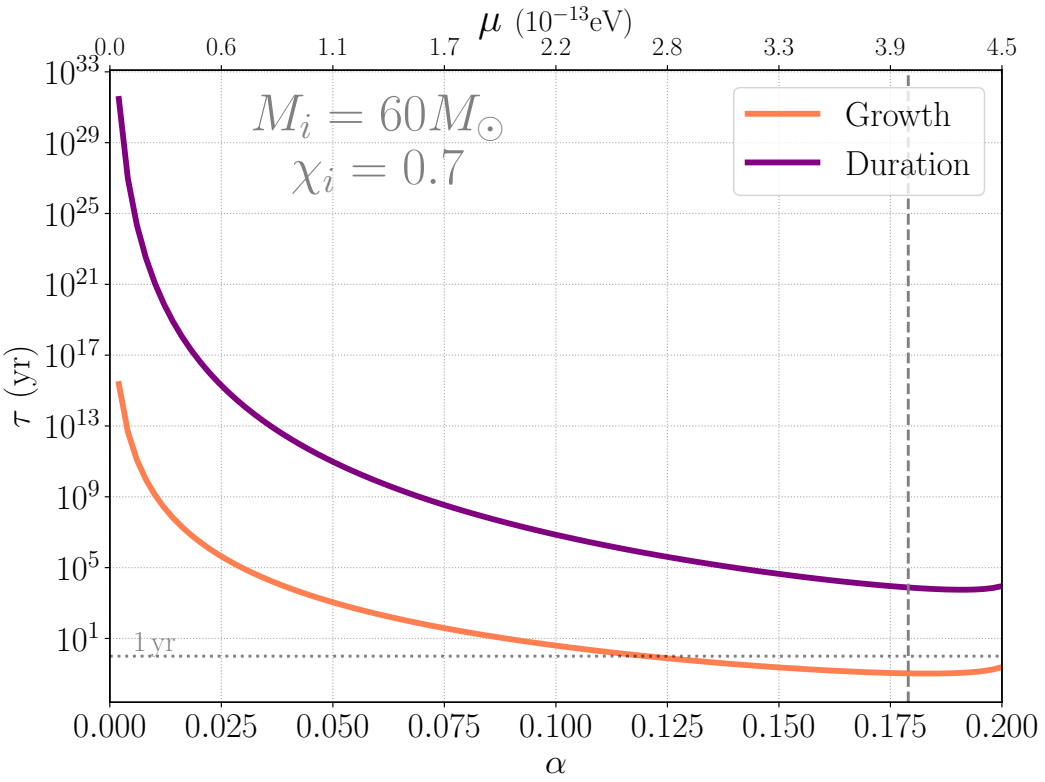

Figure 13.6: *Signal growth and duration timescales for example BH (scalar cloud).* Curves show the signal duration (purple, top), Eq. (13.15), and growth (orange, bottom), Eq. (13.22), timescales for a scalar cloud ($l = m = 1$, $n = 0$) as a function of the fine-structure constant α from Eq. (13.6). The BH is assumed to have an initial mass of $M = 60\ M_\odot$ and spin of $\chi = 0.70$. The vertical dashed line marks the value α that yields peak emission for such BH, for which $\tau_{\rm inst} = 39$ days and $\tau_{\rm GW} = 7.5\times10^3$ yr. Note that values of $\alpha > 0.2$ preclude superradiance given this spin, cf. Eq. (13.7).

important to also consider the timescales introduced in Sec. 13.2. There are two timescales associated with the gravitational signal: the time it takes to reach its peak amplitude and its duration thereafter.

The signal-growth timescale depends strongly on α, as is illustrated in Fig. 13.6 (orange curve) for our example BH ($M = 60\ M_\odot$, $\chi = 0.70$). This high sensitivity on α means that, when analyzing real data, it will be important to only consider values of the boson mass that could have yielded a detectable signal given the age of the BH being targeted. In particular, *strain upper-limits can only be meaningfully translated into boson-mass constraint if the BH is sufficiently old to support a cloud that would emit gravitational waves of such amplitude*. If a search is carried out before such time, one should instead look for a weaker and still growing signal. This will be especially important for young BHs.

For instance, for the BH in Fig. 13.6 this means that, to constrain the presence of the best-matching boson (vertical dashed line: $\alpha \approx 0.179$, i.e. $\mu \approx 4 \times 10^{-13}$ eV), one must wait *at least* 1 month from the moment the BH is born before looking for a GW signal in the data. During that first month, the cloud is still growing and the signal has not reached its peak ($h_0 \approx 5.2 \times 10^{-26}$, according to Fig. 13.1), meaning it might be too weak and unstable for detection. Thus, absence of a detectable signal during that initial period would *not* be evidence against the existence of the boson. The same is true for any other value of α, but the peak strains will be weaker (cf. Fig. 13.1) and the times required to reach them possibly much longer (if $\alpha < 0.179$, for our example with $\chi_i = 0.70$).

The second relevant timescale, the signal duration, is also strongly dependent on α. This is also illustrated in Fig. 13.6 (purple curve) for our example BH ($M_i = 60\ M_\odot$, $\chi_i = 0.70$). For the boson that best matches this BH ($\alpha \approx 0.179$), the characteristic duration of the signal is $\sim 7.5 \times 10^3$ yr. Similarly to the situation described above, absence of a detectable signal long after this would *not* constitute evidence against the existence of such a boson. This is because, if one waits too long, the fastest energy level will have been depleted, and one should instead look for signals corresponding to the next level, cf. Eq. (13.11). This feature is especially important when targeting old BHs.

Both timescales are a function of BH mass, as reflected in Fig. 13.7 and Fig. 13.8, and scale inversely with BH spin. The color in Fig. 13.7 and Fig. 13.8 corresponds to the growth and duration timescales respectively, both assuming a spin of $\chi = 0.70$. The horizontal dashed line marks $\alpha = 0.179$, the value of the fine-structure constant

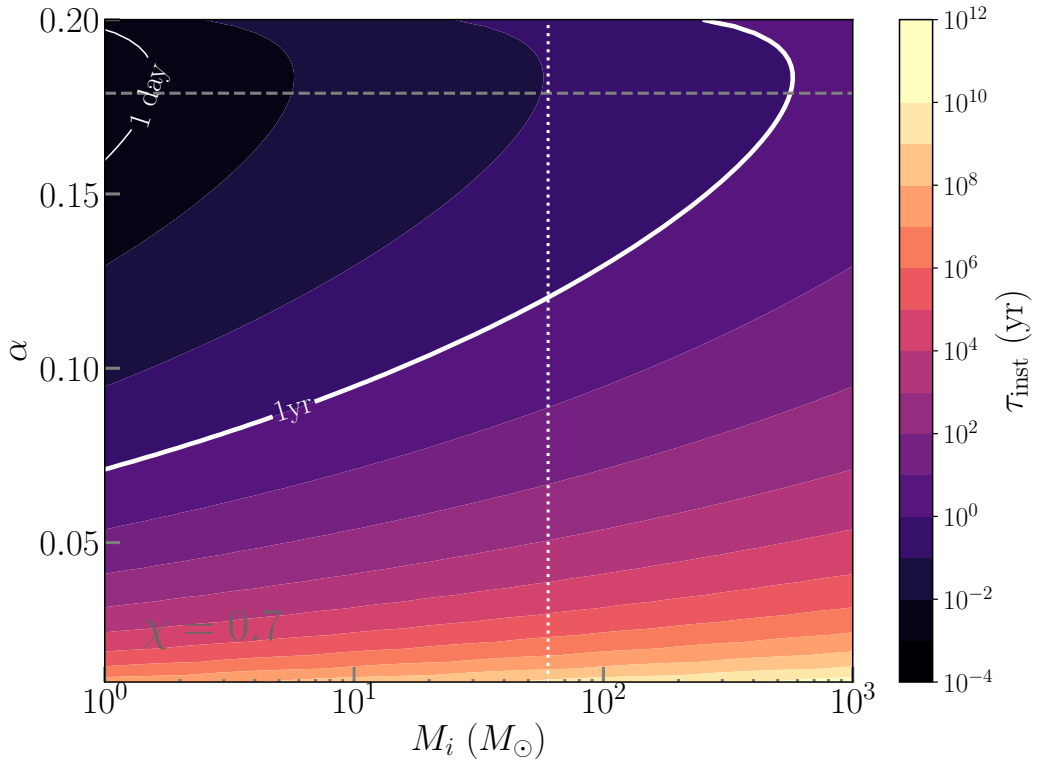

Figure 13.7: *Superradiant-instability timescale (scalar cloud).* Color shows the characteristic growth time, Eq. (13.15), for a scalar cloud ($l = m = 1$, $n = 0$) as a function of BH mass (x-axis) and fine-structure constant (y-axis). The BH is assumed to have an initial spin of $\chi = 0.70$. The highlighted contours correspond to $\tau_{\rm inst} = 1$ day (top, thin) and $\tau_{\rm inst} = 1$ yr (bottom, thick). The vertical dotted line gives the instability timescales for $M = 60\ M_\odot$. The horizontal dashed line corresponds to the value of α that yields optimal GW emission for a BH with initial spin $\chi = 0.70$. Note that values of $\alpha > 0.2$ preclude superradiance given this spin, cf. Eq. (13.7).

that yields peak emission for a BH with that spin. Meanwhile, a vertical cut along the dotted lines ($M = 60\ M_\odot$) would produce the orange and purple curves in Fig. 13.6, respectively for Fig. 13.7 and Fig. 13.8. Although both timescales vary widely for different α's, for any given system (M, χ, α), $\tau_{\rm inst}$ is always orders of magnitude shorter than $\tau_{\rm GW}$ which allows the treatment of the cloud growth and signal emission as two different regimes, as explained in Sec. 13.2.

13.4 Directed searches

There are multiple observational signatures of BH superradiance that could be used to probe the boson-mass space [336–343, 360–363, 393]. Among all these, we will focus on the prospect for direct detection of the continuous gravitational waves

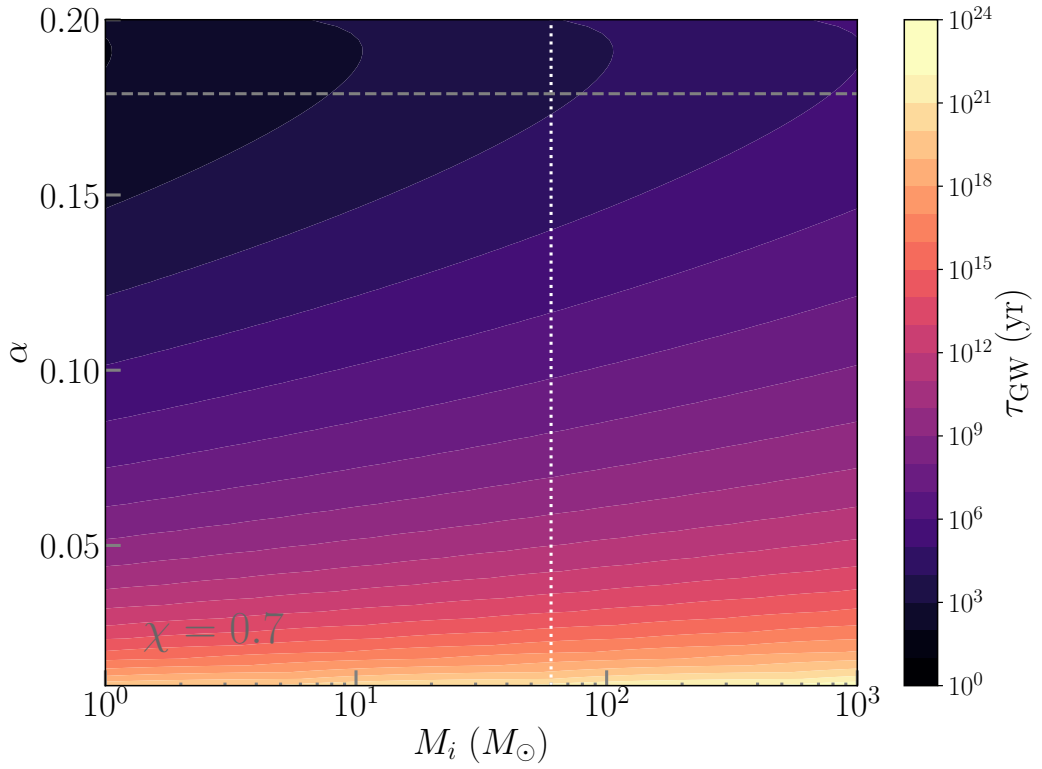

Figure 13.8: *Signal duration timescale (scalar cloud).* Color shows the characteristic signal duration, Eq. (13.22), for a scalar cloud ($l = m = 1$, $n = 0$) as a function of BH mass (x-axis) and fine-structure constant (y-axis). The BH is assumed to have an initial spin of $\chi = 0.70$. The vertical dotted line gives the instability timescales for $M = 60\ M_\odot$. The horizontal dashed line corresponds to the value of α that yields optimal GW emission for a BH with initial spin $\chi = 0.70$. Note that values of $\alpha > 0.2$ preclude superradiance given this spin, cf. Eq. (13.7).

expected from these sources (Sec. 13.3). In particular, we will restrict ourselves to searches *directed* at specific well-localized targets, rather than searches covering the whole sky (see, e.g., [75] for a review of continuous-wave searches). This means that we are interested in studying known (potential) BHs that could have the right mass, spin and age to possibly harbor a radiating boson cloud. In order to apply existing search strategies, we would also like the cloud to be stable enough to make sure the signal lasts sufficiently long and evolves slowly enough to be considered "persistent" (we sharpen these criteria below).

Because the properties of the central BH can be measured *a priori*, directed searches can potentially make unambiguous statements about the existence of ultralight bosons without relying on BH population models, which carry much uncertainty. If a signal were found from a given target, detailed measurements of its morphology (see Sec.

13.3.1) would provide invaluable information about the mass and dynamics of the new particle. On the other hand, if a signal were *not* found, knowledge of the BH parameters could allow us to place stringent constraints on the existence of bosons in the corresponding mass range. Furthermore, having a specific sky location allows us to probe deeper in the noise and explore a greater range of parameters to farther distances. This comes at the price, of course, of the restriction to BHs that are already known, which may limit the use of the method in practice if no suitable BHs are discovered to target.

In the following, we introduce hidden Markov model (HMM) tracking as an ideal method to carry out directed searches for these signals (Sec. 13.4.1). We evaluate its sensitivity with Monte-Carlo simulations and use the results to estimate the scalar-cloud detection horizons for future detectors (Sec. 13.4.2). This discussion is agnostic as to the origin of the target BH, assuming only a known location and reasonably constrained intrinsic parameters. The conclusions are, therefore, generally applicable to any known stellar-mass BH, but we devote special attention to remnants from compact-binary mergers and holes in x-ray binaries (Sec. 13.4.3). As we discuss below, vector signals present unique data-analysis and theoretical challenges, so we focus mainly (though not uniquely) on scalars.

13.4.1 Search method

Hidden Markov model tracking is an efficient search strategy for detecting quasi-monochromatic gravitational waves [394, 395]. It was developed with rapidly-spinning neutron stars in mind and has been applied in searches directed at several targets [6, 396–398]. This strategy is well suited to searches for gravitational waves from boson clouds because its computational efficiency enables the coverage of a wide range of signal parameters, and because it does not rely on restrictive waveform models. Furthermore, it allows small deviation from restrictive waveform models, unlike other coherent or semi-coherent search methods that rely on Taylor-series-based matched filters and are, thus, more model-restricted (e.g., [399, 400]). This makes it ideal to search for signals over a broad frequency band (cf. Fig. 13.1), even when the location of the source is only loosely known and when there is potential uncertainty in the signal morphology.

Algorithm overview

The goal of HMM tracking is to find the most likely path that a putative signal takes in the time-frequency plane, contingent on the observed noisy data [394, 395]. To do so, it divides the f-t plane into pixels, assuming the signal is monochromatic over a period $T_{\rm drift}$ and splitting the frequency axis into bins of width $\Delta f = 1/(2T_{\rm drift})$. The signal power in each bin is then estimated coherently using the $\mathcal{F}$-statistic [74, 399], a frequency-domain estimator that accounts for the motion of the Earth and is widely used in continuous-wave searches [75]. At each time step i, this statistic is computed for each discrete frequency bin j by coherently integrating over the time interval $(t_i, t_i + T_{\rm drift})$. Henceforth "$f_{0j}$" denotes the central value of the signal frequency, f_0, in the jth bin.[6] If the total observation time is $T_{\rm obs}$, then the values of $\mathcal{F}(f_0)$ for the $N_T = T_{\rm obs}/T_{\rm drift}$ blocks of duration $T_{\rm drift}$ are combined incoherently as described in [394, 395].

Based on this information, the HMM algorithm computes the likelihood of different signal paths, assuming the signal can only transition between adjacent frequency bins from one time step to the next. For application to boson signals, we assume the transition probability $A_{f_{0j}f_{0k}}$ between frequency bins f_{0j} and f_{0k} to be

$$A_{f_{0j+1}f_{0j}} = A_{f_{0j}f_{0j}} = \frac{1}{2}, \tag{13.35}$$

and to vanish otherwise (see [394, 395] for details). The choice of Eq. (13.35) amounts to favoring signals with a positive frequency derivative, in agreement with the signal model of Sec. 13.3. We choose a uniform prior $\Pi_{f_{0j}} = N_Q^{-1}$ on f_0 over the frequency band being searched, where N_Q is the total number of frequency bins. The result of the HMM tracking algorithm is summarized by a figure of merit representing the significance of the optimal path relative to all others (see [395, 396]). This quantity can then be treated as a regular (frequentist) detection statistic, and its background can be computed over several noise-only instantiations data to assign detection significances.

Frequency-derivative tolerance

Although, in principle, this method would be able to handle signals with arbitrary frequency evolutions, allowing for large frequency drifts ($\dot{f}$) comes with a significant

[6] Here we follow the HMM-tracking literature by using "f_0" to denote the estimator for the (unknown) frequency of the signal, rather than "f" for frequency in general [394, 395].

reduction in sensitivity. In order to allow for a maximum frequency derivative $\max(\dot{f})$ we must choose $T_{\mathrm{drift}} \leq \Delta f / |\max(\dot{f})|$, so as to guarantee that

$$\left| \int_{t}^{t+T_{\mathrm{drift}}} \mathrm{d}t' \dot{f}(t') \right| \leq \Delta f \,, \tag{13.36}$$

for $0 < t < T_{\mathrm{obs}}$ and where the frequency resolution is set to $\Delta f = (2T_{\mathrm{drift}})^{-1}$, as mentioned above. Therefore, tracking a signal with higher $\dot{f}$ requires reducing the coherent-integration time over which the $\mathcal{F}$ statistic is computed, which in turn diminishes the sensitivity of the search [395].

The implementation of the $\mathcal{F}$-statistic-based HMM used by the LIGO and Virgo Collaborations can currently track quasimonochromatic signals with derivatives of at most $\dot{f} \sim 10^{-8}\,\mathrm{Hz\,s^{-1}}$ [395]. This is more than enough to accommodate the majority of scalar signals [Eq. (13.33)], but puts most (although not all) of the vector parameter space out of reach [Eq. (13.34)]. Detecting shorter-lived signals may be possible by extending the current method to track not only the signal frequency, but also its first time derivative with a two-dimensional HMM as in [395]. Such a strategy would also naturally handle noticeable decays in signal amplitude over the observation run. The adaptation of the methods in [395] to boson signals will be subject of future work.[7]

Computing cost

The computing time for one central processing unit (CPU) over a total observing time T_{obs} in a frequency band from f_{min} to f_{max} is given by [395]

$$\mathcal{T} = 2\kappa\beta N_{\mathrm{ifo}} T_{\mathrm{drift}} T_{\mathrm{obs}} T_{\mathrm{SFT}}^{-1} N_{\mathrm{sky}} (f_{\mathrm{max}} - f_{\mathrm{min}}) \,, \tag{13.37}$$

where T_{SFT} is the length of the short Fourier transforms (SFTs) used to compute the $\mathcal{F}$-statistic [74], N_{ifo} is the number of interferometers, N_{sky} is the number of sky locations, β is the percentage of time that the interferometers collect data ("duty cycle"), and κ is the time to compute the $\mathcal{F}$-statistic per template per SFT. The value of κ depends on T_{SFT} and the CPU architecture; we adopt the recent estimate that $\kappa = 4 \times 10^{-8}\,\mathrm{s}$ for $T_{\mathrm{SFT}} = 1800\,\mathrm{s}$.

[7] The HMM tracking based on 1-s short Fourier transforms described in [398] can be used to search for long-duration transient signals with timescales of order $\sim 10^2$–10^4 s. The timescale of a vector signal is much longer than that, hence longer short Fourier transforms are needed. However, the Doppler modulation due to the motion of the Earth with respect to the solar system barycenter is not negligible when the length of short Fourier transforms is longer than $\sim$ 10–100 s [400].

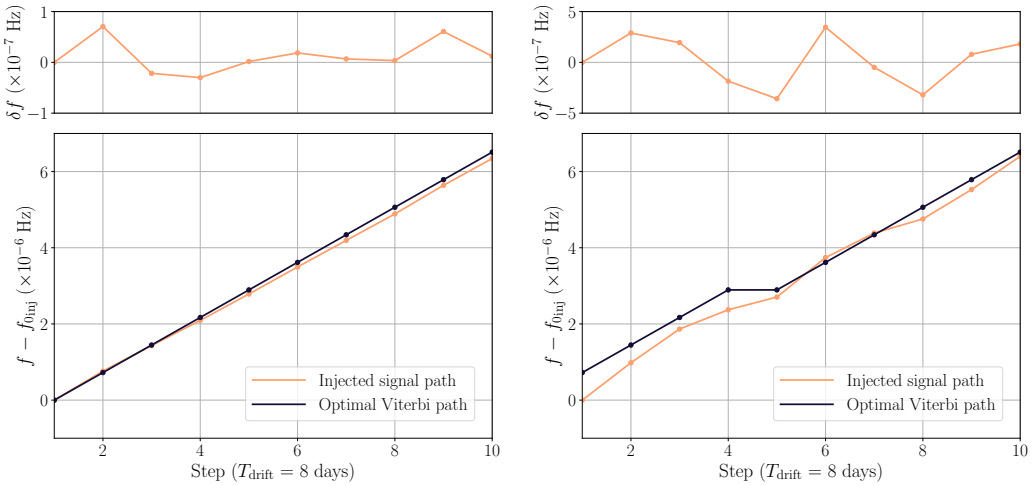

(a) Less variability

(b) More variability

Figure 13.9: *HMM sample tracking paths*. Injected $f_0(t)$ (light curves) and optimal Viterbi paths (dark curves) for the injected signals with (a) weaker random walk $|\delta f| \leq 0.1\Delta f$ and (b) stronger random walk $|\delta f| \leq 0.5\Delta f$. The top panels show the random walk δf added to the injected signals at each step, which is too small to be seen by eye in the bottom panel of (a). The horizontal axis is in units of HMM steps with each step spanning for $T_{\mathrm{drift}} = 8\,\mathrm{d}$. Good matches are obtained in both (a) and (b) with $\varepsilon_{f_0} = 0.16\Delta f$ and $0.50\Delta f$, respectively. Injection parameters are in Table 13.2 and the injected signal strain is $h_0 = 5 \times 10^{-26}$.

We normally divide the full frequency band into multiple 1-Hz subbands to allow parallelized computing. For example, if we have 10^2 cores running in parallel, a search for $T_{\mathrm{obs}} = 80$ days over frequency band spanning 100–200 Hz in two detectors and with a fixed sky location ($N_{\mathrm{sky}} = 1$) takes about 7 min to complete. This estimation is consistent with the real cost of our simulations below.

13.4.2 Sensitivity estimates

We would like to study the sensitivity of ground-based detectors to continuous GW signals from boson clouds around known BHs. For this purpose, we simulate signals consistent with the morphology described in Sec. 13.3.1 and study how well they can be recovered using the HMM tracking (Sec. 13.4.2). We then translate expected strain sensitivities into detection horizons for boson signals with current and future ground-based detectors (Sec. 13.4.2). Finally, we explore the impact of uncertainties in the source's sky location (Sec. 13.4.2).

Table 13.2: Injection parameters.

Parameter	Symbol	Value
Initial Frequency	$f_{0\mathrm{inj}}$	201.2 Hz
First derivative of $f_{0\mathrm{inj}}$	$\dot{f}_{\mathrm{inj}}$	1×10^{-12} Hz s^{-1}
Right ascension	$\alpha_\star$	$23^{\mathrm{h}}23^{\mathrm{m}}26.0^{\mathrm{s}}$
Declination	$\delta_\star$	$58°48'0.0''$
Inclination	$\cos\iota$	$[-1,\ 1]$
Polarization	ψ	$[0,\ 2\pi]$
Initial phase	Φ_0	$[0,\ 2\pi]$
Gaussian noise ASD	$S_h^{1/2}(f)$	4×10^{-24} Hz$^{-1/2}$

Table 13.3: Search parameters.

Parameter	Symbol	Value
Search frequency band	f	201–202 Hz
Coherent time	T_{drift}	8 d
Bin size	Δf	7.23×10^{-7} Hz
Total observing time	T_{obs}	80 d
Number of steps	N_T	10

Strain sensitivity

To study our sensitivity to waves from boson clouds, we inject synthetic signals with parameters consistent with the morphology described in Sec. 13.3.1 into simulated Gaussian noise corresponding to two aLIGO detectors at design sensitivity [401]. The signal frequency and frequency derivative were chosen to be roughly in agreement with an optimal scalar cloud around the example BH discussed in Sec. 13.3.2—that is, $M_i = 60\ M_\odot$ and $\chi_i = 0.70$, consistent with the GW150914 remnant [31]. On top of an overall positive frequency drift, we let the signal frequency jump randomly at eight-day steps to simulate small uncertainties in the signal model, i.e. we add a random frequency fluctuation δf at each time step. We vary the magnitude of δf, as well as the intrinsic amplitude h_0, for different injections. We choose source inclinations randomly such that $\cos\iota$ is uniformly distributed over the range $[-1, 1]$, while we pick polarization angles ψ and initial phases Φ_0 uniformly over $[0, 2\pi]$. Unless otherwise stated, we fix the sky location to the values in Table 13.2. (See Sec. 13.3.1 for definitions of all the signal parameters.) The HMM tracking is conducted with the settings shown in Table 13.3, directed at the true sky location of the injection. We choose $T_{\mathrm{obs}} = 8$ d assuming $\dot{f}^{(s)} \lesssim 10^{-12}$ Hz/s, in agreement with the scenario considered above.

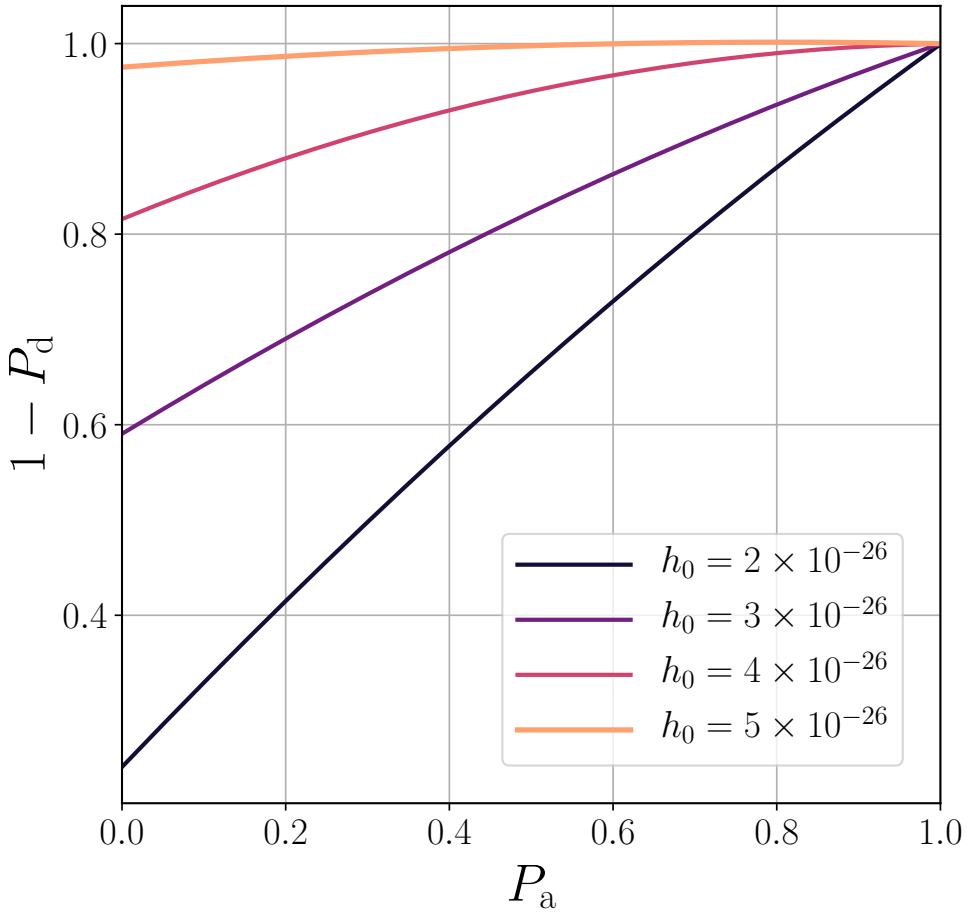

Figure 13.10: Receiver operator characteristic (ROC) curves for the injections with parameters in Table 13.2. The four curves (from top to bottom) correspond to the four representative wave strains $h_0/10^{-26} = 5$, 4, 3, and 2. The horizontal and vertical axes indicate the false alarm probability P_{a} and detection probability $1 - P_{\mathrm{d}}$, respectively. Each curve is based on 200 realizations with randomly chosen polarization and inclination angles and initial phase.

To demonstrate that HMM can accurately reconstruct boson signals, Fig. 13.9 presents two tracking examples for injected signals with $h_0 = 5 \times 10^{-26}$ and parameters in Table 13.2. The frequency random walks are such that $|\delta f| \leq 0.1\Delta f$ and $|\delta f| \leq 0.5\Delta f$ for panels (a) and (b), respectively. The optimal HMM paths (dark curves) match the injected path $f_0(t)$ (light curves) closely: the root-mean-square error (RMSE) between the optimal HMM paths and the actual signals are $\varepsilon = 1.16 \times 10^{-7}$ Hz $= 0.16\Delta f$ for panel (a) and 3.65×10^{-7} Hz $= 0.50\Delta f$ for panel (b). These small discrepancies are mostly due to the frequency discretization carried out by the HMM algorithm.

We next quantify the efficiency of HMM tracking at detecting signals of different

amplitudes. For concreteness, we assume a small uncertainty in the signal model by setting $|\delta f| \leq 0.1\Delta f$, as in Fig. 13.9a. Figure 13.10 shows the receiver operator characteristic (ROC) curves for injected signals with four values of h_0, ranging from 2×10^{-26} to 5×10^{-26}. For each signal amplitude, these curves show the detection probability ($1 - P_{\rm d}$, where $P_{\rm d}$ is the false-dismissal probability) as a function of required false-alarm probability threshold ($P_{\rm a}$). For instance, if we demand a false alarm probability $P_{\rm a} = 1\%$, we can expect to detect a signal with $h_0 = 4 \times 10^{-26}$ (5×10^{-26}) with 84% (98%) of the time.

The detection threshold in continuous-wave searches is traditionally defined to be 95% false-dismissal rate at 1% false-alarm probability [75, 396]. In our case, for an observation time of $T_{\rm obs} = 80$ days with two aLIGO design detectors, this corresponds to a strain amplitude of $h_0^{95\%} = 4.7 \times 10^{-26}$ for unknown inclination. Based on this, we will consider boson signals "detectable" if they reach an amplitude of $h_0^{95\%}$ or higher for the observation conditions.

From the empirical result that $h_0^{95\%} = 4.7 \times 10^{-26}$ obtained for the simulations above, it is straightforward to estimate how the sensitivity of the search would scale for different detector networks and observation times. The sensitivity scaling will be given by [395]

$$h_0^{95\%}(f) \propto N_{\rm ifo}^{-1/2} S_h(f)^{1/2} \left(T_{\rm drift} T_{\rm obs}\right)^{-1/4} , \tag{13.38}$$

assuming a network of $N_{\rm ifo}$ detectors with power-spectral density (PSD) $S_h(f)$ at the signal frequency.[8] Using this, Fig. 13.11 presents projected 95%-confidence strain upper limits, $h_0^{95\%}$, for different detectors as a function of GW frequency (assuming there is no detection). We show results for aLIGO design sensitivity (gray), as well as proposed third-generation detectors: LIGO Voyager (yellow), LIGO Cosmic Explorer (purple) and Einstein Telescope in the "D" configuration (red) [402–405]. All curves in Fig. 13.11 were produced assuming $N_{\rm ifo} = 1$ and $T_{\rm obs} = 1$ year, but it is straightforward to rescale them for different configurations using Eq. (13.38). Finally, note that these curves were obtained by effectively marginalizing over source orientation: they represent the value of $h_0^{95\%}$ marginalized over the distributions of $\cos\iota$ and ψ in the ranges shown in Table 13.2. To obtain the values of $h_0^{95\%}$ corresponding to *optimal* source orientation, one should divide the curves of Fig. 13.11 by a factor of ∼2.8 [396].

[8] This requires that all detectors have comparable sensitivities given by $S_h(f)$; were this not the case, the $N_{\rm ifo}^{-1/2} S_h(f)^{1/2}$ factor would have to be replaced by the effective PSD $S_{\rm eff}(f)^{1/2}$, which is given by the harmonic mean of the PSDs for each detector.

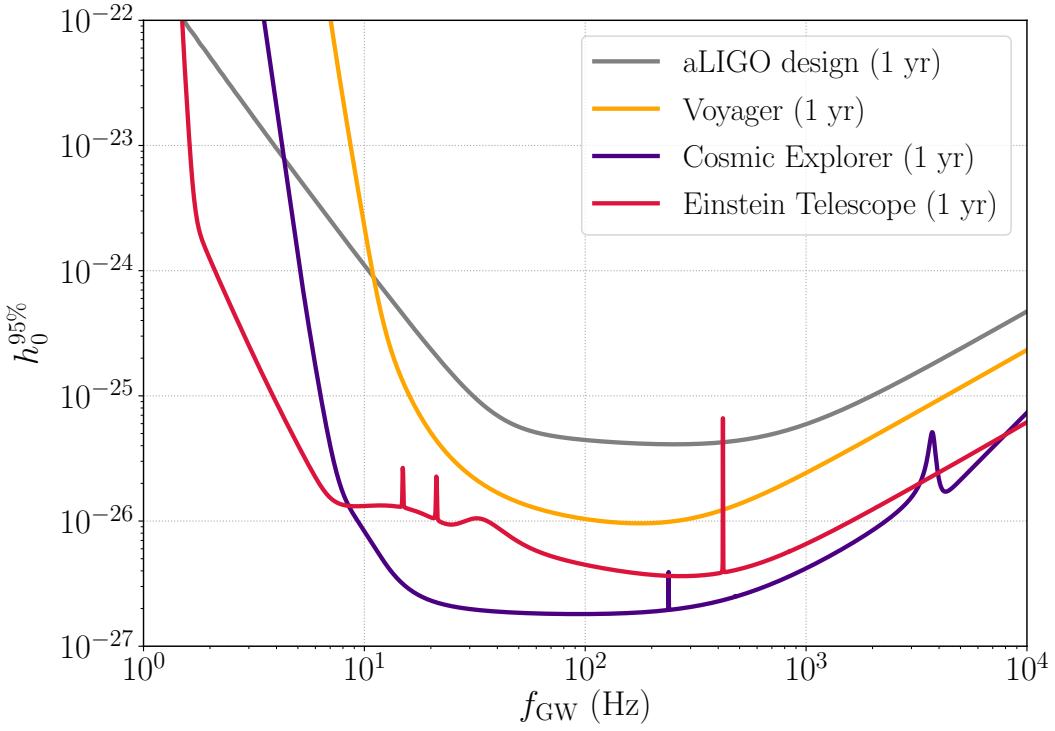

Figure 13.11: *Sensitivity vs GW frequency for different detectors.* Value of $h_0^{95\%}$ marginalized over source orientation for design aLIGO (gray), LIGO Voyager (yellow), LIGO Cosmic Explorer (purple) and the Einstein Telescope (red). All curves assume one year of continuous observation by a single detector of the indicated type.

The sky position of the source with respect to the detector does not impact the search sensitivity significantly because the variation due to the antenna pattern is averaged out when the integration time is much longer than a day. We verify this by injecting signals with $h_0 = 5 \times 10^{-26}$ at different sky positions, and with all other parameters as in Table 13.2. As before, the HMM tracking is conducted with the settings shown in Table 13.3, directed at the true sky location of the injection. The detection efficiencies for each sky location are listed in Table 13.4. As before, each row is based on 200 realizations with randomly chosen $\cos\iota$, ψ, and Φ_0. The standard deviation of detection efficiencies at these eight sky positions is only 0.02.

Detection horizons for scalar clouds

It is useful to translate the projected strain sensitivities of Fig. 13.11 into detection horizons for boson signals from BHs with different parameters. The detection horizon is the farthest distance up to which we should expect to be able to detect an

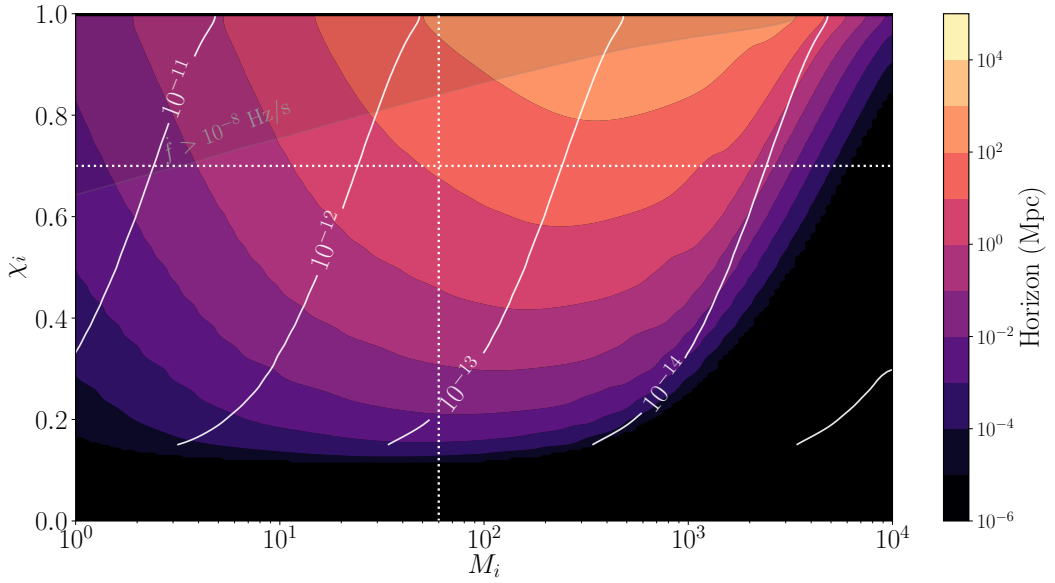

(a) aLIGO design

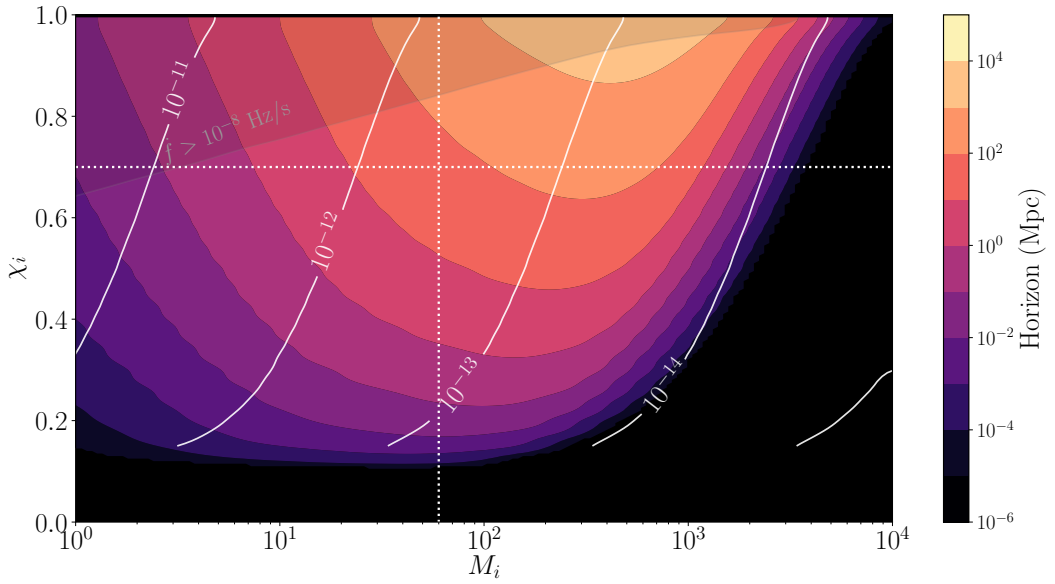

(b) LIGO Voyager

Figure 13.12: *Detection horizons for scalar clouds (aLIGO and LIGO Voyager).* Maximum-detectable luminosity distances (color) for optimal scalar clouds ($l = m = 1$, $n = 0$, $\tilde{l} = \tilde{m} = 2$) around BHs with the indicated initial mass (x-axis) and spin (y-axis) for different detectors. White contour lines indicate the values of the corresponding boson rest-energy μ/eV. The shaded region (top-left) marks parameters that would yield signals evolving prohibitively fast ($\dot{f}_{\rm det} > 10^{-8}$) for existing search methods, based on Eq. (13.33). The dotted white lines highlight the mass and spin (60 $M_\odot$, 0.70) of the GW150914-like example discussed repeatedly in the main text. Values correspond to HMM tracking for one year of continuous observation by a single detector, accounting for signal redshifts and variability in maximum $T_{\rm drift}$ allowed by the expected signal, cf. Eq. (13.36).

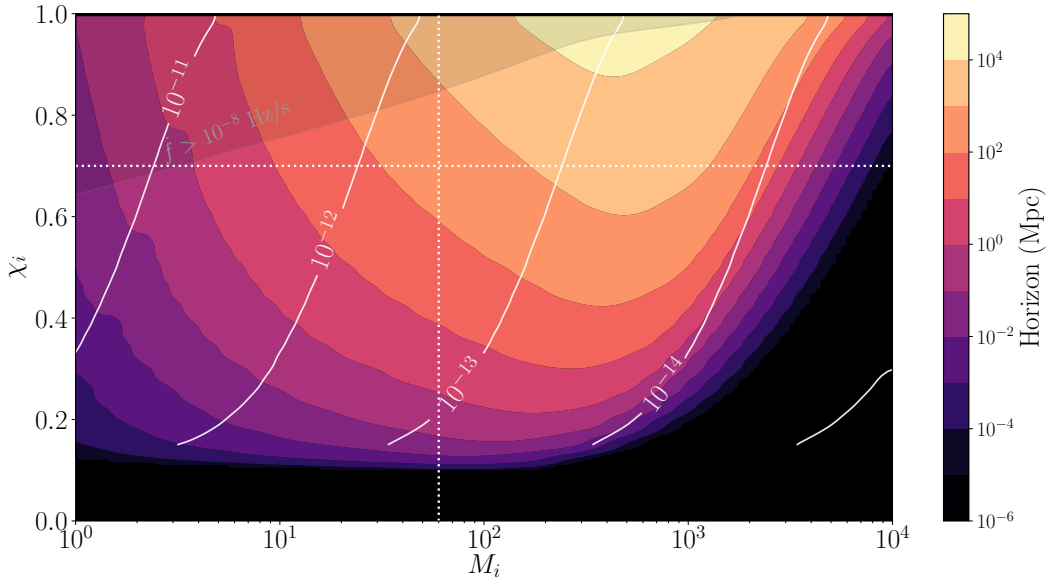

(c) LIGO Cosmic Explorer

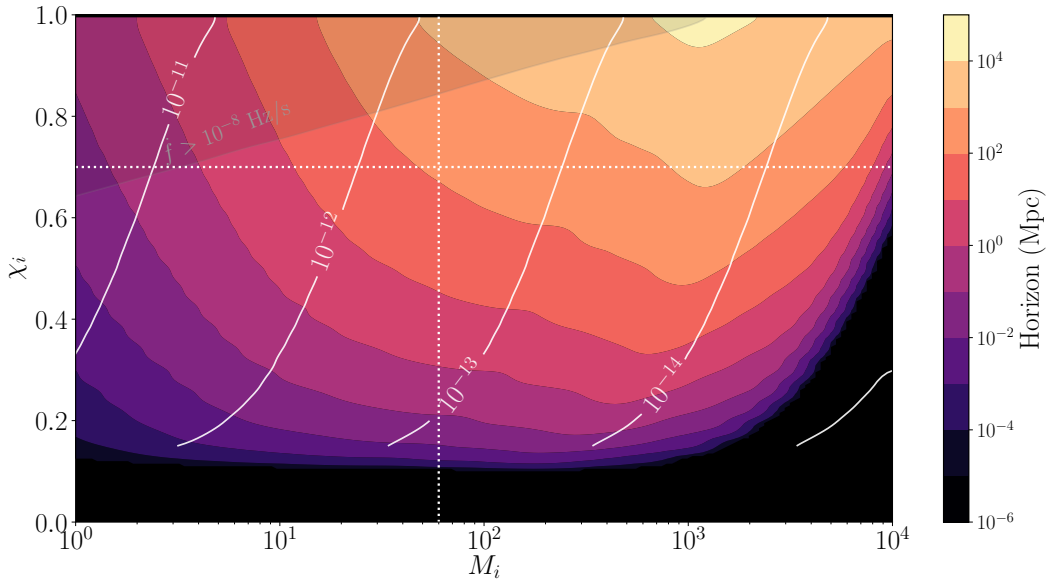

(d) Einstein Telescope

Figure 13.12: *Detection horizons for scalar clouds (Cosmic Explorer and Einstein Telescope).* Maximum-detectable luminosity distances (color) for optimal scalar clouds ($l = m = 1$, $n = 0$, $\tilde{l} = \tilde{m} = 2$) around BHs with the indicated initial mass (x-axis) and spin (y-axis) for different detectors. White contour lines indicate the values of the corresponding boson rest-energy μ/eV. The shaded region (top-left) marks parameters that would yield signals evolving prohibitively fast ($\dot{f}_{\rm det} > 10^{-8}$) for existing search methods, based on Eq. (13.33). The dotted white lines highlight the mass and spin (60 $M_\odot$, 0.70) of the GW150914-like example discussed repeatedly in the main text. Values correspond to HMM tracking for one year of continuous observation by a single detector, accounting for signal redshifts and variability in maximum $T_{\rm drift}$ allowed by the expected signal, cf. Eq. (13.36).

Table 13.4: Detection efficiency vs sky location ($h_0 = 5 \times 10^{-26}$)

Right ascension	Declination	Detection efficiency
23h 23m 26.0s	58°48′0.0″	0.98
23h 23m 26.0s	−59°35′0.0″	0.97
23h 23m 26.0s	00°02′0.0″	0.92
23h 23m 26.0s	88°48′0.0″	0.98
23h 23m 26.0s	−89°18′0.0″	0.98
05h 23m 26.0s	58°48′0.0″	0.98
11h 23m 26.0s	58°48′0.0″	0.98
17h 23m 26.0s	58°48′0.0″	0.99

optimal boson signal—namely, a signal from a boson cloud that perfectly matches its host BH (to maximize *intrinsic* strain) and is optimally oriented with respect to the detector (to maximize *measured* strain). Consequently, horizons are a measure of how well we can do in the best-case scenario and are, thus, not generally representative of most detections (see e.g. [406] for an overview of distance measures in GW astronomy)—yet, they are a straightforward proxy for the reach of our instruments to this type of source. We compute this quantity for *scalar* clouds based on the results from Sec. 13.3.2; we defer computation of horizons for vector clouds until better numerical estimates of their intrinsic amplitudes become available and analysis methods suitable for higher frequency derivatives are developed.

Figure 13.12 shows the horizon luminosity distance (color) for scalar signals as a function of initial BH mass M_i and spin χ_i, for (a) design Advanced LIGO, (b) LIGO Voyager, (c) LIGO Cosmic Explorer, and (d) the Einstein Telescope. White contours indicate the values of the boson rest-energy μ (eV) that we would be able to probe with a BH of that mass and spin. The shaded region marks values for which we expect the signal to evolve too rapidly for current data analysis techniques to handle, based on Eq. (13.33) and Eq. (13.36). Note that this varies slightly among plots due to minor differences in redshift. In all cases we assume one year of uninterrupted observation by a single detector.

The horizon plots were obtained by finding the luminosity distance at which an optimal cloud for the given BH parameters would become barely detectable, i.e. $h_0(d_L) = h_0^{95\%}(f_{\rm det})$ for detector-frame signal frequency $f_{\rm det}$. In order to obtain the relevant value of $h_0^{95\%}$, we rescale the curves of Fig. 13.11 using Eq. (13.38) to account for variations in $T_{\rm drift}$. This is needed because the expected $\dot{f}_{\rm det}$ varies widely over the parameter space, affecting the maximum-allowed coherence time, cf. Eq.

(13.36).[9] We also take into account the fact that both frequencies and frequency derivatives get redshifted as the signal makes its way to Earth,[10] i.e. $f_{\rm det} = f_{\rm src}(1+z)^{-1}$ and $\dot{f}_{\rm det} = \dot{f}_{\rm src}(1+z)^{-2}$ for a BH at redshift z and source-frame frequency $f_{\rm src}$. Finally, we rescale the curves in Fig. 13.11 to obtain values corresponding to optimal source orientation, as explained above. All these different factors modulate the intrinsic strain inferred from Fig. 13.5 to yield Fig. 13.12.

For all detectors, the horizon generally increases with initial BH mass and spin, as expected from Fig. 13.5. Furthermore, higher masses and spins are expected to yield smaller $\dot{f}$'s, which enables longer coherent times (longer $T_{\rm drift}$'s) and, thus, slightly higher sensitivity, cf. Eq. (13.38). Yet, this tendency is offset by the fact that heavier systems yield lower frequencies (Fig. 13.4), causing the horizon to quickly drop as signals reach the lower end of the detector's sensitive band (cf. Fig. 13.11). Moreover, signals from clouds around heavier BHs can more easily get redshifted out of the band. At the other end of the spectrum, the instruments we consider tend to be more sensitive at higher frequencies, but these correspond to lower masses and, thus, lower radiated power (for a given χ_i). On the other hand, increasing the BH spin yields both higher GW amplitudes and, to an extent, frequencies. Unfortunately, however, lower masses and higher spins also result in high $\dot{f}$'s that make much of that part of parameter space inaccessible to current methods (shaded regions).

All this means, roughly, that the farthest horizons will be obtained for BHs with masses in the range $10^2 \lesssim M_i/M_\odot \lesssim 10^3$ and spin as high as possible, corresponding to boson masses within $10^{-14} \lesssim \mu/{\rm eV} \lesssim 10^{-12}$ (depending on χ_i). Even outside this range, these horizons are significantly more distant than the sources at which these searches are generally directed, which tend to lie within the Milky Way (see e.g. [75])

As a concrete example, consider again a GW150914-like remnant with $M_i = 60\ M_\odot$ and $\chi_i = 0.70$. As we saw back in Sec. 13.3.2, this BH would be best matched by a scalar boson with $\mu = 4\times10^{-13}$ eV. A scalar of that mass would yield an optimal cloud ($l = m = 1$, $n = 0$) that radiates gravitational waves ($\tilde{l} = \tilde{m} = 2$) at a source-frame frequency of 191 Hz with characteristic amplitude $h_0 = 5.2 \times 10^{-26}(5\,{\rm Mpc}/d_L)$, corresponding to the peak in Fig. 13.1. From the intersection of the dotted lines in Fig. 13.12a, we see that such a signal would be detectable, at most, up to 11 Mpc

[9] In an actual analysis, we might want to set a $T_{\rm drift}$ shorter than that implied by Eq. (13.33) in order to allow for theoretical uncertainty in the predicted value of the frequency derivative.

[10] Assuming standard ΛCDM cosmology with present parameters: $\Omega_m = 0.308$, $\Omega_\Lambda = 0.692$, $\Omega_k = 0.0$, $h = 0.678$ [407].

Table 13.5: Scalar-cloud horizons (Mpc) for representative BHs (boson and signal parameters shown in Table 13.1).

M_i ($M_\odot$)	χ_i	aLIGO	Voy	CE	ET
3	0.90	0.2	0.4	2	2
10	0.90	3	6	35	24
60	**0.70**	**11**	**49**	$\mathbf{3\times10^2}$	$\mathbf{1\times10^2}$
60	0.90	58	2×10^2	1×10^3	7×10^2
200	0.85	2×10^2	6×10^2	5×10^3	1×10^3
300	0.95	5×10^2	2×10^3	2×10^4	4×10^3

away with one aLIGO detector at design sensitivity observing continuously for 1 yr—or, equivalently, ∼20 Mpc for three such detectors. This agrees with previous estimates in [340].

Prospects are even better for third-generation detectors, with farther horizons over most of the parameter space. However, as is evident from the dark regions in Fig. 13.12, this is not true for higher masses because all instruments we consider sacrifice sensitivity at lower frequencies in favor of higher ones (Fig. 13.11); as a result, a second-generation network would be preferable in that region of parameter space. This notwithstanding, third-generation detectors would offer significant improvements for mostly any target with $M_i \lesssim 10^3 M_\odot$. In particular, we find that LIGO Cosmic Explorer could reach ranges of over 10^4 Mpc ($z \gtrsim 1.4$) for fast-spinning BHs over a range of masses (Fig. 13.12c). The Einstein Telescope could also reach such distances, but for a more limited choice of parameters, and would have shorter reach for most of the sources we consider; on the other hand, this instrument *would* outperform all others at higher masses (lower frequencies). Some representative values are presented in Table 13.5 to ease comparison between instruments.

Effect of sky-location uncertainty

We would like to understand the effect of uncertainty in the source sky location on the HMM tracking, mainly motivated by the prospect of following up compact-binary mergers. In order to find a continuous signal coming from some area in the sky, we would have to analyze the gravitational-wave data with the HMM multiple times assuming slightly different sky locations to tile the patch where the source is thought to lie. The number of iterations (the number of "templates") needed, $N_{\rm sky}$, is

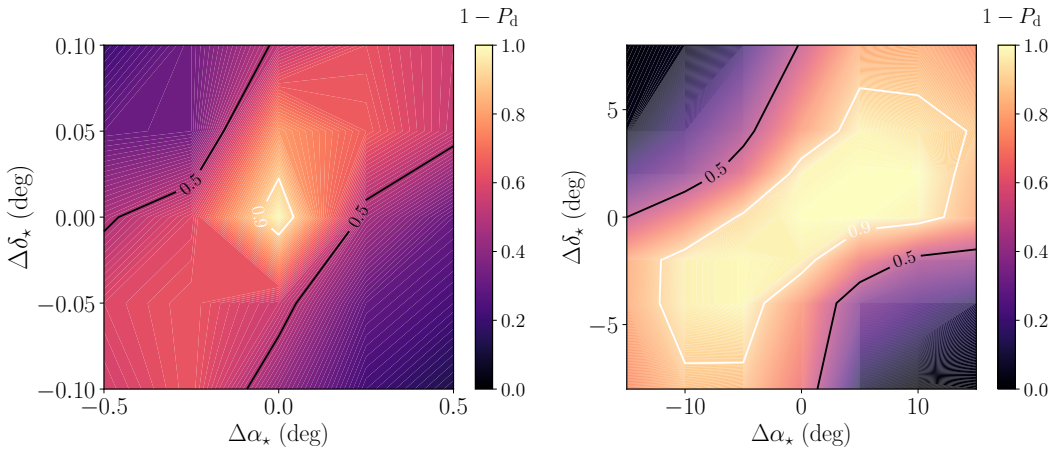

Figure 13.13: *Sky resolution*. Color shows detection efficiency ($1 - P_d$, for P_d the false-dismissal probability) as a function of offsets in right ascension (x-axis) and declination (y-axis) with respect to the true location for injections with $h_0 = 5\times10^{-26}$ (left) and $h_0 = 2 \times 10^{-25}$ (right). All other injection parameters are as in Table 13.2 and search settings are shown in Table 13.3. The left (right) plot was interpolated from a square grid with 5 (7) sky locations on each side.

determined both by the size of the target area and by the sky resolution of the analysis, which is in turn tied to the frequency resolution of the search and the amplitude of the signal.

To estimate N_{sky}, we run two sets of simulations by injecting signals with $h_0 = 5\times10^{-26}$ and $h_0 = 2\times10^{-25}$ into fabricated Gaussian noise for two aLIGO detectors at design sensitivity. All other parameters, including the sky location, are as listed in Table 13.2. HMM searches are then conducted using the settings shown in Table 13.3, but for a grid of sky locations in the neighborhood of the injected signal. In other words, for each injection amplitude ($h_0 = 5 \times 10^{-26}$ or $h_0 = 2 \times 10^{-25}$) and for each sky location assumed by the search, we inject a signal with random orientation and phase, but location fixed to the value in Table 13.2; we repeat this 200 times to obtain detection probabilities for each of the search locations.

The results of this study are summarized in Fig. 13.13 for both the soft (left panel) and loud (right panel) injections. Color in these figures encodes the detection efficiency ($1 - P_d$) for searches assuming a sky location indicated by their offset in right ascension ($\Delta\alpha_\star$) and declination ($\Delta\delta_\star$) with respect to the true location ($\Delta\alpha_\star = \Delta\delta_\star = 0$). The white (black) contours mark points at which signals were detected 90% (50%) of the time at 1% false-alarm probability. Notice that, for the weaker signal, the 90%-contour encloses an area of ~0.01 deg^2 around the true

location, while for the stronger signal this is roughly four orders of magnitude larger. We may take the size of the 90% contours as indicative of the spacing of the sky grid needed to capture a signal.

In an actual search, we need $N_{\rm sky} \sim 10^3$ sky templates per $\deg^2$ to detect a weak signal near the detection limit. The sky resolution generally agrees with other coherent or semi-coherent CW search methods [408, 409]. Here we discuss the search feasibility given the required number of sky templates. As a representative example, consider that the existing three-detector network (Advanced LIGO and Advanced Virgo) was able to localize the binary-neutron start merger GW170817 to a sky region spanning $\sim 30\,\deg^2$ with 90% credibility [6]. Based on Fig. 13.13, we would then need $N_{\rm sky} \sim 10^4$ sky templates to obtain 90% detection efficiency of a signal with $h_0 = 5 \times 10^{-26}$ lying somewhere inside the GW170817 90%-credible region; by contrast, $N_{\rm sky} \sim 3$ would suffice for a signal with $h_0 = 2 \times 10^{-25}$. Based on estimates for standard computing architectures and algorithm settings, finding a signal at 200 Hz with $h_0 = 5 \times 10^{-26}$ in a $\sim 30\,\deg^2$-region ($T_{\rm obs} = 80\,\mathrm{d}$, $N_{\rm ifo} = 2$) would take 5 days of computing on 1k CPUs, which is feasible but not cheap. Note that $N_{\rm sky}$ scales as f^2 [408–411], so more templates will be required at higher frequencies. Because computing cost scales directly with $N_{\rm sky}$, the burden will be vastly reduced once more gravitational detectors join the network and the sky locations reach the projected $O(1\,\deg^2)$ [412].

13.4.3 Potential sources

The discussion thus far has been largely unconcerned with the kind of BH being targeted. In this section, we flesh out the implications of the above conclusions for two types of promising sources: remnants from compact-binary mergers (Sec. 13.4.3) and BHs in x-ray binaries (Sec. 13.4.3).

Merger remnants

As pointed out before [340, 343], nearby CBC remnants would be ideal targets for searches for gravitational signals from ultralight bosons. Because we witness their birth first hand, the age of remnant BHs is perfectly known and their mass and spin well constrained. This would enable accurate estimation of the continuous-signal amplitude that should be expected for any given α (Sec. 13.3.2), allowing us to potentially place interesting constraints on the existence of matching bosons.

Furthermore, the location and orientation inferred from the initial chirp would allow us to take advantage of existing infrastructure for directed searches for continuous waves in LIGO and Virgo data (Secs. 13.4.1 and 13.4.2).

Ideally, we would follow up any and all mergers, as soon as a reasonable time has passed for the cloud to form (Sec. 13.3.2). In practice, however, we may be limited by the uncertainty in the sky location. Signals detected with only two instruments will be too loosely localized to allow for followup (e.g., LIGO's first detection was localized to sky region of $\sim 260\,\mathrm{deg}^2$ [1]). Fortunately, as we saw in Sec. 13.4.2, the localization provided by a three-detector network would already be manageable with existing computational resources (e.g. $\sim 28\,\mathrm{deg}^2$ for the binary neutron star [6]). This will be further improved when more, and more sensitive, detectors join the effort: a network including LIGO India [133, 134] and KAGRA [135, 136], on top of the three existing detectors operating at design sensitivity, is expected to routinely locate events to an area of order $\sim 1\,\mathrm{deg}^2$ [412]. Regardless of the number of instruments, events with an electromagnetic counterpart (e.g., mergers involving a neutron star) will always be sufficiently well localized.

Extracting information about bosons from one of these observations would also require good knowledge of the remnant distance and orientation. This is required to translate strain (h_0) constraints into limits on radiated-power ($\dot{E}_{\rm GW}$), which can then be turned into statements about the existence of a boson with a given mass [cf. Eqs. (13.19) and (13.28)]. In particular, if the source were too far away, the signal from the hypothetical cloud would be undetectable at Earth, rendering constraints on its amplitude moot. Thus, if the distance is not determined by other means (e.g. association with a host galaxy), the implications for bosons will be contingent on the uncertainty in the luminosity distance inferred from the CBC observation.

In the case of a scalar cloud, for most remnant masses and spins, the source would have to be relatively close for the signal to be detectable by ground-based detectors (Fig. 13.12). For a second-generation network at design sensitivity, the horizon would lie below $100\,\mathrm{Mpc} \times \sqrt{N_{\rm ifo}}$ for most signal parameters (Fig. 13.12a). Given that we have yet to observe a BH merger that close [1–6, 178], this projection is not too auspicious. Yet, note that the horizon can reach close to $10^3\,\mathrm{Mpc} \times \sqrt{N_{\rm ifo}}$ in some regions of parameter space—although taking advantage of this with existing algorithms would require a population CBCs yielding remnants with $M \gtrsim 100\,M_\odot$ ($\dot{f} < 10^{-8}$ Hz/s). Estimates of rates from BH population models were provided in [340] based on the nonrelativistic approximation to the amplitude Eq. (13.29).

As we saw in Sec. 13.4.2, prospects are better for next-generation detectors, especially Cosmic Explorer (Fig. 13.12c). Even then, good reach to remnants with $M \sim \mathcal{O}(10M_\odot)$ would require spins roughly $\gtrsim 0.85$, possibly less depending on the mass. Although we have not yet observed any such events [1–6, 178], numerical-relativity simulations routinely produce remnants with such spins [413–416]. Note that the horizons for boson signals are always significantly closer than those for compact-binary coalescences [402].

The vector case is slightly different. Detection horizons are in principle considerably farther for vector clouds due to the intrinsically higher radiated power (see Sec. 13.2.2), making most remnant masses and spins accessible. However, more radiated power also means shorter cloud lifetimes and, consequently, faster rates of change for the signal amplitude and frequency [cf. Eq. (13.34)]. For much of the parameter space, the expected signal would then evolve too rapidly for existing continuous-wave algorithms to handle (see discussion in Sec. 13.4.2). Therefore, the more powerful and quickly-evolving vector signals would currently not be detectable, effectively reducing our horizon to such sources. Detection rates for vectors taking this into account were estimated in [343] by using the nonrelativistic approximation of Eq. (13.30).

Even for scalar signals, the restriction to small frequency derivatives is quite detrimental, preventing us from accessing higher boson masses (lower BH masses). Because the estimates of Eq. (13.33) and Eq. (13.34) are only approximate, there is still sense in searching for signals with $\dot{f} > 10^{-8}\,\mathrm{Hz\,s^{-1}}$ in the shaded regions of Fig. 13.12—although a negative result would be harder to interpret as evidence against the existence of a boson in that mass range. As suggested above, this is strong motivation to adapt analysis techniques that can handle quickly evolving continuous signals to make them suitable for boson searches—this is work in progress.

Finally, note that we expect to infer the remnant mass and spin from the CBC signal with enough precision to obtain a reasonably accurate prediction of the cloud GW amplitude for any given boson mass. For instance, the mass and dimensionless spin of the GW150914 remnant were each measured with one-sided relative errors of under 10% at 90% credibility, which is sufficiently narrow to make a followup search possible (Sec. 13.4.2). The characteristic magnitude of such errors is expected to be significantly reduced for detections at higher signal-to-noise ratio, which should be commonplace once the current network achieves design sensitivity and for next-generation detectors (see, e.g., [412, 417, 418]).

In any case, a simplistic way to deal with parameter uncertainty would be to compute the optimal strain for a cloud around a BH corresponding to the upper bounds of the mass and spin credible intervals. A value computed that way would itself be an upper limit on the boson strain, because this quantity scales directly with mass and spin (Fig. 13.5). Alternatively, a rigorous statistical analysis would take in the full-dimensional posterior probability density on the BH parameters (intrinsic and extrinsic) and marginalize over all parameters to obtain a posterior on the expected boson strain as a function of α. The development of this more sophisticated strategy is work in progress.

X-ray binaries

Another type of potentially interesting targets is known BHs in x-ray binaries (see e.g. [419, 420] for reviews treating such systems). The relevance of x-ray binaries to this research program has been pointed out since the outset (e.g. [336, 338, 339, 343]). They have the advantage of being much closer and better-located in the sky compared to the CBC remnants, with good measurements of their mass and, in some cases, spin [421]. In fact, some limits on the boson-mass space have already been placed contingent on these measurements, roughly excluding the mass interval $10^{-12} \lesssim \mu/\mathrm{eV} \lesssim 10^{-11}$ for scalars [340] and $10^{-13} \lesssim \mu/\mathrm{eV} \lesssim 10^{-11}$ for vectors [343]. Unfortunately, there is large uncertainty about the age and history of these systems, as well as important caveats about the systematics affecting their spin measurements [364, 365]. Furthermore, the effect of the active astrophysical environments surrounding these BHs is only understood at the order-of-magnitude level [339, 343, 361]. For all these reasons, boson constraints derived from existing observations of x-ray binaries should be interpreted with caution.

There are also data-analysis challenges intrinsic to signals coming from sources in a binary system: the Doppler modulation due to the motion of the source within the binary causes the signal power to spread over multiple frequencies. The signal power must then be collected from "orbital sidebands" that span a frequency band B given by [394]

$$B \approx 4\pi \frac{f_0 a_0}{cP}\,, \tag{13.39}$$

where a_0 is the binary's projected semimajor axis, P its orbital period, and c is the speed of light. Frequency-domain matched filters, like those presented in [394, 422], can be applied to sum up the distributed signal power using (imperfect) knowledge of

Table 13.6: Cygnus X-1 parameters.

Parameter	Symbol	Value	Ref.
Mass ($M_\odot$)	M	14.8 ± 1.0	[420, 423]
Spin	χ	≥ 0.95	[420]
Right ascension	$\alpha_\star$	$19^{\rm h}58^{\rm m}22^{\rm s}$	[424]
Declination	$\delta_\star$	$34^\circ 12' 0.6''$	[424]
Inclination (deg)	ι	27.1 ± 0.8	[423]
Distance (kpc)	r	$1.86^{+0.12}_{-0.11}$	[420]
Orbital period (days)	P	5.6	[423, 425]
Proj. semimajor axis ($R_\odot$)	a_0	$31.6^{+4.1}_{-3.1}$	[425]

the orbital parameters. Those methods would generally demand $B \lesssim 0.5\,\text{Hz}$ in order to achieve the required sensitivity (see Sec. 13.4.2). Besides, it becomes prohibitively expensive to detect a weak signal from a binary if the orbital parameters (e.g., P, a_0, and time of passage through the orbit's ascending node T_P) are poorly measured. More details can be found in Sec. III B of Ref. [422].

As an example, consider the nearby Cygnus X-1 binary, which has been proposed as an interesting target for boson searches [338]. If we take source parameters in Table 13.6 and assume $f_0 \sim 1\,\text{kHz}$ given $M \approx 15 M_\odot$, then the power of a signal from Cygnus X-1 would span a frequency band $B \approx 2\,\text{Hz}$. In addition, the uncertainty in a_0 is non-negligible and the knowledge of T_P is limited. Unfortunately, for most of the interesting x-ray systems, the orbital parameters are not well measured electromagnetically, and the sidebands would be broader than ~1 Hz. This means that, for now, the strain upper limits that can be placed on such sources using the existing methods will be too loose to be of interest.

13.5 Conclusion

Black-hole superradiance could be the key that allows gravitational-wave detectors to uncover evidence of new ultralight bosons beyond the standard model, thus extending the reach bringing particles physics within the reach of gravitational-wave science. In this paper, we explored the prospect for achieving this exciting goal by looking for the continuous gravitational signals expected from scalar and vector clouds, using searches directed at known BHs.

We began by reviewing the physics of boson clouds (Sec. 13.2) and examined in detail the properties of continuous signals from clouds around a known BH (Sec. 13.3). In doing so, we hoped to provide a bridge between the theory and data-analysis

literatures. We then used numerical techniques, combined with the latest analytic results, to compute the features of gravitational waves emitted by scalar clouds around BHs with different initial parameters (Figs. 13.1–13.8).

We put forward the use of hidden Markov model (HMM) tracking [6, 394–396] as an ideal method to carry out directed searches for boson signals (Sec. 13.4.1). This strategy is well suited to searches for gravitational waves from boson clouds because its computational efficiency enables the coverage of a wide range of signal parameters, and because it does not rely on restrictive waveform models. This makes it ideal to search for signals over a broad frequency band (cf. Fig. 13.1), even when the location of the source is only loosely known and when there is potential uncertainty in the signal morphology. We demonstrated this through a series of Monte-Carlo simulations (Sec. 13.4.2).

From our simulations, we obtained an empirical estimate of the sensitivity of directed searches to boson signals in data from future ground-based detectors: aLIGO design, LIGO Voyager, LIGO Cosmic Explorer and the Einstein telescope (Fig. 13.11). For scalar clouds, we translated the expected strain sensitivities into detection horizons for those four detectors (Fig. 13.12), assuming one year of observation by a single detector. We found that, for a second-generation network at design sensitivity, the horizon would lie below $100\,\mathrm{Mpc} \times \sqrt{N_{\mathrm{ifo}}}$ for most signal parameters; prospects are better for next-generation detectors, especially Cosmic Explorer, for which horizons could reach up to $\sim 10^5$ Mpc. Generally speaking, these horizons lie much farther than the sources at which continuous-wave searches are generally directed [75], but significantly closer than horizons for compact-binary coalescences [402]. Some representative values are shown in Table 13.1.

In computing signal amplitudes from scalar clouds, we numerically solved the evolution equations governing cloud growth and made use of numerical estimates from BH perturbation theory to obtain the radiated power [342]. Furthermore, to estimate horizons, we incorporated the effect of redshifts on the signal frequency and frequency derivative. We also took into account that the settings of the search algorithm should be varied across parameter space for optimal performance. This allowed us to obtain sensitivity estimates that should be more reliable than previously published projections.

Finally, we discussed implications for the followup of remnants from compact-binary coalescences (Sec. 13.4.3), as well as BHs in x-ray binaries (Sec. 13.4.3). We explored the impact of uncertainties in the source's sky location and showed that

HMM tracking will be able to efficiently cover the localization credible-regions obtained from CBC signals with a network of at least three detectors. We also discussed the challenges intrinsic to vector signals, which make their analysis difficult in spite of their higher radiated power. We emphasized the strong motivation to extend existing search techniques to handle signals with higher frequency derivatives, so as to bring a significant portion of the scalar and vector signal space into reach. The implementation of such techniques, as well as development of statistical strategies to rigorously handle uncertainty in BH parameters, is work in progress.

Acknowledgments

M.I. and R.B. would like to thank Asimina Arvanitaki, Masha Baryakhtar, William East and Robert Lasenby for organizing the meeting "Searching for New Particles with Black Hole Superradiance" held at the Perimeter Institute for Theoretical Physics in May 2018. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Economic Development & Innovation. M.I. and L.S. are members of the LIGO Laboratory. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY–0757058. L.S. was supported by an Australian Research Training Program Stipend Scholarship and the Albert Shimmins Fund at earlier stages of this project. The research was also supported by Australian Research Council (ARC) Discovery Project DP110103347 and the ARC Centre of Excellence for Gravitational Wave Discovery CE17010004. This paper carries LIGO Document Number LIGO–P1800270.

13.6 Appendix: Frequency drift

The gravitational self-energy of the cloud affects the boson's eigenfrequencies and, consequently, the gravitational wave frequency. As the cloud dissipates due to gravitational emission, this causes an increase in the emitted signal frequency, similar to what happens in a compact-binary coalescence. Therefore, we may treat the system adiabatically to obtain the frequency drift from the radiated power. This computation was presented in [343] for vectors, and we reproduce it here for scalars.

The gravitational self-energy of a bound state per particle is given by

$$U_c = -G\frac{m_b}{M_c}\int \frac{\rho(r,\theta,\phi)\, m(r)}{r}\mathrm{d}^3\mathbf{x}, \tag{13.40}$$

where m_b is the mass of the boson field, M_c the overall mass in the cloud, $\rho(r, \theta, \phi)$ is its density and $m(r)$ is the the mass of cloud enclosed in the radius r, namely

$$m(r) = \int_0^r \rho(r, \theta, \phi)\, \mathrm{d}^3\mathbf{x}\,. \tag{13.41}$$

The rate of change of the GW frequency can then estimated by [343]

$$\dot{f} \simeq \frac{1}{2\pi\hbar} \times 2\dot{U}_c\,. \tag{13.42}$$

As the cloud dissipates, the total mass of the system decreases, causing the binding energy to increase ($U_c < 0$) and the GW frequency to increase.

The dominant scalar field mode can be approximated by the $\ell = m = 1$ hydrogen wave function, with a density ρ given by Eq.(11) in Ref. [357]. After some algebra one finds that, at leading order in α,

$$\dot{f} \simeq \frac{93}{1024}\frac{c\,\alpha^3}{\pi G}\frac{\dot{E}_{\rm GW}}{M^2} = \frac{93}{1024}\frac{c^3\alpha^3}{\pi G}\frac{M_c}{M^2}\frac{1}{\tau_{\rm GW}}\,, \tag{13.43}$$

where we have used $\dot{M}_c = -\dot{E}_{\rm GW}c^{-2}$ and $\dot{E}_{\rm GW} = M_c c^2/\tau_{\rm GW}$. Using Eq. (13.13) to approximate $M_c \sim \alpha M_i \chi_i$ (valid in the limit $\alpha \ll 1$ for $m = 1$) and Eq. (13.23) to write $\tau_{\rm GW} \sim GM^2/(0.025 M_c \alpha^{14} c^3)$, we get

$$\dot{f} \simeq 3 \times 10^{-14}\,\mathrm{Hz/s}\left(\frac{10M_\odot}{M}\right)^2\left(\frac{\alpha}{0.1}\right)^{19}\chi_i^2\,. \tag{13.44}$$

Chapter 14

CONCLUSION

In this thesis, I have explored several avenues for learning about fundamental physics from gravitational-wave (GW) observations. Because little can be done in the absence of signals, I also developed data-analysis methods to improve LIGO and Virgo's chances of detecting both transient and persistent signals. Below, we reiterate the key results for each chapter and discuss future related work.

We began in Chapter 3 by demonstrating the potential of measures of signal coherence for enhancing confidence in detections of compact-binary coalescences. Our approach is extremely successful at identifying instrumental artifacts, and can thus vastly increase the inferred significance of weak signals. Further work is needed to gauge the gains afforded by Bayesian coherence in relation to the newest statistical analyses implemented by LIGO and Virgo during the second observation run. In the future, we will extend our methods to handle an arbitrary number of detectors and will work with other members of the LIGO and Virgo Collaborations to incorporate these techniques into the staple searches (this joint project is ongoing).

In Chapter 4, we explored how to improve our ability to detect persistent GWs from known pulsars. We introduced a robust technique to empirically compute false-alarm probabilities for searches targeted at such sources, without assuming the instrumental noise is Gaussian. This work filled an important gap in the field, which generally lacked robust strategies for determining the significance of continuous-wave detection candidates without idealizing the noise. Our method has already been applied to follow up several detection candidates in actual analyses of LIGO data and will continue to be used in future runs.

In Chapter 5, we began our study of the basic properties of GWs. We showed that it is possible to measure the local geometric structure of GW polarizations by using transient signals seen by multiple detectors. We argued that we can (partially) do this without much worrying about the detailed phase evolution of each polarization. However, this comes with the limitation of only being able to distinguish between "extreme" scenarios, in which the signal is purely tensor, purely vector or purely scalar. The implementation of techniques to address the case of mixed polarizations (e.g. tensor plus some small nontensor component) is currently work in progress.

Although that case is most interesting to theorists developing extensions to general relativity, we will not be able to take full advantage of those methods until more detectors join the global network.

In Chapter 6, we presented the first direct observational statements about the local geometry of GW polarizations. These were obtained from GW170814 and GW170817 using the methods of Chapter 5. We concluded that GWs are more likely to come in the polarizations predicted by general relativity, as opposed to either of the two fully-nontensorial alternatives. We evaluated this via Bayes factors, but also presented ancillary evidence, like reconstructed sky locations.

Persistent GWs would be a far better resource for probing GW polarizations than transients. We first demonstrated this in the context of searches targeted at known pulsars in the study of Chapter 7. Those methods, that were tested on initial LIGO data, grew into the full-fledged Bayesian analysis presented in Chapter 8. This multi-stage Bayesian infrastructure can detect signals of any polarization (consistent with relativity, or otherwise) and, in the presence of a signal, can also assess agreement with general relativity and place upper-limits on the power of nontensorial components. The method was successfully applied to data from Advanced LIGO's first observation run to obtain direct constraints on nontensorial strain amplitudes for 200 pulsars, as discussed in Chapter 9. I am currently working on expanding the implementation to handle signals at the first, not just the second, harmonic of the pulsar rotation frequency. The analysis of data from Advanced LIGO and Virgo's second observation run is in progress.

The statistical approach I developed for the case of continuous waves can be straightforwardly applied applied to other signal morphologies. For instance, I have contributed to the implementation of this strategy for the case of stochastic GW backgrounds [10]. An analysis of such a background would be sensitive to alternative polarizations produced at the source or generated during the propagation of the signal on its way to Earth, while relying on very few assumptions. In Chapter 10, we explored the impact of these assumptions on the scope of the analysis, with examples in the context of specific theories. Our conclusions are important in interpreting the results of analyses of data from present and future detectors, like [117].

It is easy to imagine that, due to effects beyond general relativity, GWs could travel in vacuum at speeds other than the speed of light, potentially also showing dispersion or birefringence. Some of these effects are already under study by LIGO, most notably the possibility of a massive graviton. In Chapter 11 we explored how to

measure GW speed using continuous waves from known pulsars. In Chapter 12, we briefly examined the possibility of detecting complex modes of GW propagation, like anisotropic dispersion relations. This is a fertile avenue of research, and much more work is required in this direction.

Finally, the study of gravity itself is not the only kind of fundamental science that can be done with GWs. In Chapter 13, we explored the exotic, yet compelling, possibility of discovering new particle physics with GW detectors. Specifically, we focused on the prospect of directed searches to detect continuous GWs from ultralight-boson condensates around spinning black-holes. We show that existing data-analysis techniques may be applied for this purpose, and show that future-generation detectors will be able to set interesting constraints on the existence of such bosons.

The topics covered in this thesis represent only a sliver of what can be potentially learned about fundamental physics with gravitational waves. Besides the detailed analysis of the properties of exotic sources and other topics not covered here, the possibilities for doing fundamental science in general will soon increase dramatically, as we transition into the era of *precision* GW astronomy. In the upcoming years, existing GW detectors will achieve their design sensitivities and will be joined by new observatories in India and Japan. These ground-based instruments will subsequently be complemented in space by Laser Interferometer Space Antenna (LISA), a joint ESA & NASA mission which will have access to lower sections of the GW spectrum. Meanwhile, still lower GW frequencies will be probed by pulsar timing arrays. All this, plus continued partnerships with electromagnetic and neutrino observatories, means that we will soon be inundated with an unprecedented panoply of data surely holding many fundamental clues about nature. Our exploration of the universe through gravitational waves has just begun.

BIBLIOGRAPHY

[1] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **116**, 061102 (2016), arXiv:1602.03837 [gr-qc] .

[2] B. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **116**, 241103 (2016), arXiv:1606.04855 .

[3] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. **X6**, 041015 (2016), arXiv:1606.04856 [gr-qc] .

[4] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **118**, 221101 (2017), arXiv:1706.01812 .

[5] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **119**, 141101 (2017), arXiv:1709.09660 .

[6] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **119**, 161101 (2017), arXiv:1710.05832 .

[7] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **120**, 091101 (2018), arXiv:1710.05837 .

[8] J. Aasi *et al.* (LIGO Scientific Collaboration), Class. Quantum Grav. **32**, 074001 (2015), arXiv:1411.4547 .

[9] F. Acernese *et al.* (Virgo Collaboration), Class. Quantum Grav. **32**, 024001 (2015), arXiv:1408.3978 .

[10] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Class. Quant. Grav. **33**, 134001 (2016), arXiv:1602.03844 .

[11] J. Veitch and A. Vecchio, Phys. Rev. D **81**, 062003 (2010), arXiv:0911.3820 .

[12] K. Cannon *et al.*, Astrophys. J. **748**, 136 (2012), arXiv:1107.2665 .

[13] K. Cannon, C. Hanna, and D. Keppel, Phys. Rev. **D88**, 024025 (2013), arXiv:1209.0718 .

[14] C. Messick *et al.*, Phys. Rev. **95**, 042001 (2017), arXiv:1604.04324 .

[15] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, Astrophys. J. **849**, 118 (2017), arXiv:1705.01513 .

[16] S. A. Usman *et al.*, Class. Quant. Grav. **33**, 215004 (2016), arXiv:1508.02357 .

[17] T. Dal Canton *et al.*, Phys. Rev. **D90**, 082004 (2014), arXiv:1405.6731 .

[18] M. Isi, M. Pitkin, and A. J. Weinstein, Phys. Rev. D **96**, 042001 (2017), arXiv:1703.07530 .

[19] J. Skilling, Bayesian Anal. **1**, 833 (2006).

[20] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin, W. Del Pozzo, F. Feroz, J. Gair, C.-J. Haster, V. Kalogera, T. Littenberg, I. Mandel, R. O'Shaughnessy, M. Pitkin, C. Rodriguez, C. Röver, T. Sidery, R. Smith, M. Van Der Sluys, A. Vecchio, W. Vousden, and L. Wade, Phys. Rev. D **91**, 042003 (2015), arXiv:1409.7215 .

[21] T. G. F. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, M. Agathos, J. Veitch, K. Grover, T. Sidery, R. Sturani, and A. Vecchio, Phys. Rev. D **85**, 082003 (2012), arXiv:1110.0530 .

[22] D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial*, 2nd ed. (Oxford University Press, New York, 2006) p. 264.

[23] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms David J.C. MacKay* (Cambridge University Press, Cambridge, 2003).

[24] R. Lynch, S. Vitale, R. Essick, E. Katsavounidis, and F. Robinet, Phys. Rev. **D95**, 104046 (2017), arXiv:1511.05955 .

[25] N. J. Cornish and T. B. Littenberg, Class. Quantum Grav. **32**, 135012 (2015), arXiv:1410.3835 .

[26] J. Powell, M. Szczepanczyk, and I. S. Heng, Phys. Rev. **D96**, 123013 (2017), arXiv:1709.00955 .

[27] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, Phys. Rev. D **89**, 064023 (2014), arXiv:1311.5738 .

[28] M. Pitkin, C. Gill, J. Veitch, E. Macdonald, and G. Woan, *Gravitational waves. Numerical relativity - data analysis. Proceedings, 9th Edoardo Amaldi Conference, Amaldi 9, and meeting, NRDA 2011, Cardiff, UK, July 10-15, 2011*, J. Phys. Conf. Ser. **363**, 012041 (2012), arXiv:1203.2856 [astro-ph.HE] .

[29] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Astrophys. J. **839**, 12 (2017), arXiv:1701.07709 .

[30] A. Nitz, I. Harry, D. Brown, *et al.*, "ligo-cbc/pycbc: Latest release," (2017).

[31] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **116**, 241102 (2016), arXiv:1602.03840 [gr-qc] .

[32] R. Smith, S. E. Field, K. Blackburn, C.-J. Haster, M. Pürrer, V. Raymond, and P. Schmidt, Phys. Rev. D **94**, 044031 (2016), arXiv:1604.08253 .

[33] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, Phys. Rev. D **93**, 044006 (2016), arXiv:1508.07250 .

[34] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, Phys. Rev. D **93**, 044007 (2016), arXiv:1508.07253 .

[35] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Phys. Rev. Lett. **113**, 151101 (2014), arXiv:1308.3271 .

[36] T. B. Littenberg and N. J. Cornish, Phys. Rev. **D91**, 084034 (2015), arXiv:1410.3852 .

[37] C. Biwer, D. Barker, J. C. Batch, *et al.*, Phys. Rev. **D95**, 062002 (2017), arXiv:1612.07864 .

[38] J. B. Kanner, T. B. Littenberg, N. Cornish, M. Millhouse, E. Xhakaj, F. Salemi, M. Drago, G. Vedovato, and S. Klimenko, Phys. Rev. **D93**, 022002 (2016), arXiv:1509.06423 .

[39] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, Springer Series in Statistics (Springer New York, New York, NY, 2009).

[40] J. D. E. Creighton and W. G. Anderson, *Gravitational-Wave Physics and Astronomy* (Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany, 2011) p. 188.

[41] B. P. Abbott et al., (LIGO Scientific Collaboration, and Virgo Collaboration), Phys. Rev. Lett. **119**, 161101 (2017), arXiv:1710.05832 .

[42] B. P. Abbott et al., (LIGO Scientific Collaboration, and Virgo Collaboration), Astrophys. J. **851**, L35 (2017), arXiv:1711.05578 .

[43] B. P. Abbott, (The LIGO Scientific Collaboration, and The Virgo Collaboration), Phys. Rev. D **94**, 102002 (2016), arXiv:1606.09619 .

[44] J. Aasi, (The LIGO Scientific Collaboration, and The Virgo Collaboration), Astrophys. J. **813**, 39 (2015), arXiv:1412.5942 .

[45] B. P. Abbott, (The LIGO Scientific Collaboration, and The Virgo Collaboration), Phys. Rev. D **94**, 042002 (2016), arXiv:1605.03233 .

[46] J. Aasi, (The LIGO Scientific Collaboration, and The Virgo Collaboration), Phys. Rev. D **90**, 062010 (2014), arXiv:1405.7904 [gr-qc] .

[47] J. Aasi, (The LIGO Scientific Collaboration, and The Virgo Collaboration), Phys. Rev. D **93**, 042007 (2016), arXiv:1510.03621 [astro-ph.IM] .

[48] K. S. Thorne, in *Three hundred years Gravit.*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, 1987) Chap. 9, pp. 330 – 458.

[49] J. Aasi et al., (The LIGO Scientific Collaboration, and The Virgo Collaboration), Astrophys. J. **785**, 119 (2014), arXiv:1309.4027 .

[50] K. Riles, Mod. Phys. Lett. **A32**, 1730035 (2017), arXiv:1712.05897 [gr-qc] .

[51] M. Pitkin, M. Isi, J. Veitch, and G. Woan, "A nested sampling code for targeted searches for continuous gravitational waves from pulsars," (2017), arXiv:1705.08978 .

[52] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).

[53] B. P. Abbott et al., (LIGO Scientific Collaboration, and Virgo Collaboration), Phys. Rev. **D96**, 122004 (2017), arXiv:1707.02669 .

[54] D. I. Jones and N. Andersson, Mon. Not. R. Astron. Soc. **331**, 203 (2002).

[55] G. Hobbs, R. Edwards, and R. Manchester, Mon. Not. R. Astron. Soc. **369**, 655 (2006), arXiv:0603381 [astro-ph] .

[56] R. T. Edwards, G. B. Hobbs, and R. N. Manchester, Mon. Not. R. Astron. Soc. **372**, 1549 (2006), arXiv:0607664 [astro-ph] .

[57] R. Dupuis and G. Woan, Phys. Rev. D **72**, 102002 (2005), arXiv:0508096 [gr-qc] .

[58] P. Astone, A. Colla, S. D'Antonio, S. Frasca, and C. Palomba, Phys rev. D. **90**, 042002 (2014), arXiv:1407.8333 .

[59] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **116**, 131103 (2016), arXiv:1602.03838 [gr-qc] .

[60] D. V. Martynov *et al.*, Phys. Rev. **D93**, 112004 (2016), [Addendum: Phys. Rev.D97,no.5,059901(2018)], arXiv:1604.00439 [astro-ph.IM] .

[61] B. P. Abbott et al., (LIGO Scientific Collaboration, and Virgo Collaboration), Class. Quant. Grav. **35**, 065010 (2018), arXiv:1710.02185 [gr-qc] .

[62] P. Covas *et al.* (LSC Instrument Authors), Phys. Rev. **D97**, 082002 (2018), arXiv:1801.07204 [astro-ph.IM] .

[63] P. Leaci, P. Astone, M. A. Papa, and S. Frasca, *Gravitational waves. Proceedings, 8th Edoardo Amaldi Conference, Amaldi 8, New York, USA, June 22-26, 2009*, J. Phys. Conf. Ser. **228**, 012006 (2010).

[64] G. D. Meadors, K. Kawabe, and K. Riles, Class. Quant. Grav. **31**, 105014 (2014), arXiv:1311.6835 [astro-ph.IM] .

[65] J. C. Driggers *et al.* (LSC Instrument Authors), (2018), arXiv:1806.00532 [astro-ph.IM] .

[66] T. M. Niebauer, A. Rüdiger, R. Schilling, L. Schnupp, W. Winkler, and K. Danzmann, Phys. Rev. D **47**, 3106 (1993).

[67] M. Pitkin, *Searches for continuous and transient gravitational waves from known neutron stars and their astrophysical implications*, Ph.D. thesis, University of Glasgow (2006).

[68] P. Jaranowski and A. Królak, Class. Quant. Grav. **27**, 194015 (2010), arXiv:1004.0324 [gr-qc] .

[69] P. Astone, S. D'Antonio, S. Frasca, and C. Palomba, Class. Quantum Gravity **27**, 194016 (2010).

[70] P. Astone, A. Colla, S. D'Antonio, S. Frasca, and C. Palomba, Journal of Physics Conference Series **363**, 012038 (2012), arXiv:1203.6733 [astro-ph.IM] .

[71] P. Astone, A. Colla, S. D'Antonio, S. Frasca, C. Palomba, and R. Serafinelli, Phys. Rev. D **89**, 062008 (2014), arXiv:1403.1484 .

[72] S. Mastrogiovanni, P. Astone, S. D'Antonio, S. Frasca, G. Intini, P. Leaci, A. Miller, C. Palomba, O. J. Piccinni, and A. Singhal, Classical and Quantum Gravity **34**, 135007 (2017), arXiv:1703.03493 [gr-qc] .

[73] P. Astone, A. Colla, S. D'Antonio, S. Frasca, C. Palomba, and R. Serafinelli, Phys. Rev. D **89**, 062008 (2014), arXiv:1403.1484 [astro-ph.IM] .

[74] R. Prix, *The $\mathcal{F}$-statistic and its implementation in ComputeFstatistic v2*, Tech. Rep. LIGO-T0900149 (LIGO Laboratory, 2011).

[75] K. Riles, Mod. Phys. Lett. A **32**, 1730035 (2017).

[76] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. D **93**, 122003 (2016), arXiv:1602.03839 [gr-qc] .

[77] J. Abadie *et al.* (LIGO Collaboration, Virgo Collaboration), Phys.Rev. **D85**, 082002 (2012), arXiv:1111.7314 [gr-qc] .

[78] C. Biwer *et al.*, Phys. Rev. **D95**, 062002 (2017), arXiv:1612.07864 [astro-ph.IM] .

[79] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **120**, 031104 (2018), arXiv:1709.09203 [gr-qc] .

[80] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. D **96**, 122006 (2017), arXiv:1710.02327 [gr-qc] .

[81] J. D. Hunter, Computing In Science & Engineering **9**, 90 (2007).

[82] B. P. Abbott et al., (The LIGO Scientific Collaboration, and The Virgo Collaboration), Phys. Rev. Lett. **116**, 221101 (2016).

[83] C. M. Will, *Theory and experiment in gravitational physics*, revised ed. (Cambridge University Press, Cambridge, 1993).

[84] C. M. Will, Living Rev. Relativ. **17** (2014), 10.12942/lrr-2014-4.

[85] K. Chatziioannou, N. Yunes, and N. Cornish, Phys. Rev. D **86**, 022004 (2012).

[86] J. M. Weisberg, D. J. Nice, and J. H. Taylor, Astrophys. J. **722**, 1030 (2010), arXiv:1011.0718 .

[87] P. C. C. Freire, N. Wex, G. Esposito-Farèse, J. P. W. Verbiest, M. Bailes, B. A. Jacoby, M. Kramer, I. H. Stairs, J. Antoniadis, and G. H. Janssen, Mon. Not. R. Astron. Soc. **423**, 3328 (2012).

[88] I. H. Stairs, Living Rev. Relativ. **6**, 5 (2003).

[89] N. Wex, "Testing Relativistic Gravity with Radio Pulsars," (2014), arXiv:1402.5594 .

[90] Z. Berezhiani, D. Comelli, F. Nesti, and L. Pilo, Phys. Rev. Lett. **99**, 131101 (2007), arXiv:0703264 [hep-th] .

[91] S. Hassan, A. Schmidt-May, and M. von Strauss, J. High Energy Phys. **2013**, 86 (2013), arXiv:1208.1515 .

[92] K. Max, M. Platscher, and J. Smirnov, Phys. Rev. Lett. **119**, 111101 (2017), arXiv:1703.07785 .

[93] P. Brax, A.-C. Davis, and J. Noller, Phys. Rev. D **96**, 023518 (2017), arXiv:1703.08016 .

[94] B. Pontecorvo, Sov. Phys. JETP **6**, 429 (1957).

[95] B. Pontecorvo, Sov. Phys. JETP **26**, 984 (1968).

[96] S. Alexander and N. Yunes, Phys. Rep. **480**, 1 (2009), arXiv:0907.2562 .

[97] D. M. Eardley, D. L. Lee, A. P. Lightman, R. V. Wagoner, and C. M. Will, Phys. Rev. Lett. **30**, 884 (1973).

[98] D. Eardley, D. Lee, and A. Lightman, Phys. Rev. D **8**, 3308 (1973).

[99] T. Callister, A. S. Biscoveanu, N. Christensen, M. Isi, A. Matas, O. Minazzoli, T. Regimbau, M. Sakellariadou, J. Tasson, and E. Thrane, Phys. Rev. X **7**, 041058 (2017), arXiv:1704.08373 .

[100] K. S. Thorne, D. L. Lee, and A. P. Lightman, Phys. Rev. D **7**, 3563 (1973).

[101] C. Brans and R. H. Dicke, Phys. Rev. **124**, 925 (1961).

[102] D. Andriot and G. L. Gómez, J. Cosmol. Astropart. Phys. **2017**, 048 (2017), arXiv:1704.07392 .

[103] A. P. Lightman and D. L. Lee, Phys. Rev. D **8**, 3293 (1973).

[104] N. Rosen, Ann. Phys. (N. Y). **84**, 455 (1974).

[105] C. de Rham, Living Rev. Relativ. **17** (2014), 10.12942/lrr-2014-7, arXiv:1401.4173 .

[106] C. Mead, "Gravitational Waves in G4v," (2015), arXiv:1503.04866 .

[107] A. Nishizawa, A. Taruya, K. Hayama, S. Kawamura, and M.-a. Sakagami, Phys. Rev. D **79**, 082002 (2009).

[108] A. Błaut, Phys. Rev. D **85**, 043005 (2012).

[109] M. Isi, A. J. Weinstein, C. Mead, and M. Pitkin, Phys. Rev. D **91**, 082002 (2015).

[110] E. Poisson and C. M. Will, *Gravity: Newtonian, Post-Newtonian, Relativistic* (Cambridge University Press, Cambridge, 2014).

[111] K. J. Lee, F. A. Jenet, and R. H. Price, Astrophys. J. **685**, 1304 (2008).

[112] S. J. Chamberlin and X. Siemens, Phys. Rev. D **85**, 082001 (2012), arXiv:1111.5661 .

[113] N. J. Cornish and T. B. Littenberg, Class. Quantum Gravity **32**, 135012 (2014), arXiv:1410.3835 .

[114] T. B. Littenberg, J. B. Kanner, N. J. Cornish, and M. Millhouse, Phys. Rev. D **94**, 044050 (2016), arXiv:1511.08752 .

[115] C. M. Will, Living Rev. Rel. **17**, 4 (2014), arXiv:1403.7377 [gr-qc] .

[116] M. Isi and A. J. Weinstein, *Probing gravitational wave polarizations with signals from compact binary coalescences*, Tech. Rep. LIGO-P1700276 (LIGO Laboratory, 2017) arXiv:1710.03794 [gr-qc] .

[117] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **120**, 201102 (2018), arXiv:1802.10194 [gr-qc] .

[118] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), "Tests of General Relativity with GW170817," (2018), arXiv:1811.00364 [gr-qc] .

[119] M. M. Kasliwal and S. Nissanke, Astrophys. J. **789**, L5 (2014), arXiv:1309.1554 [astro-ph.HE] .

[120] L. P. Singer, L. R. Price, B. Farr, A. L. Urban, C. Pankow, *et al.*, Astrophys.J. **795**, 105 (2014), arXiv:1404.5623 [astro-ph.HE] .

[121] C. P. L. Berry *et al.*, Astrophys. J. **804**, 114 (2015), arXiv:1411.6934 [astro-ph.HE] .

[122] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, *8th Gravitational Wave Data Analysis Workshop (GWDAW 2003) Milwaukee, Wisconsin, December 17-20, 2003*, Class. Quant. Grav. **21**, S1809 (2004), arXiv:gr-qc/0412119 [gr-qc] .

[123] B. P. Abbott et al., (The LIGO Scientific Collaboration, the Virgo Collaboration, Fermi Gamma-ray Burst Monitor, and INTEGRAL, Astrophys. J. **848**, L13 (2017), arXiv:1710.05834 .

[124] B. P. Abbott *et al.* (GROND, SALT Group, OzGrav, DFN, INTEGRAL, Virgo, Insight-Hxmt, MAXI Team, Fermi-LAT, J-GEM, RATIR, IceCube, CAASTRO, LWA, ePESSTO, GRAWITA, RIMAS, SKA South Africa/MeerKAT, H.E.S.S., 1M2H Team, IKI-GW Follow-up, Fermi GBM, Pi of Sky, DWF (Deeper Wider Faster Program), Dark Energy Survey, MASTER, AstroSat Cadmium Zinc Telluride Imager Team, Swift, Pierre Auger, ASKAP, VINROUGE, JAGWAR, Chandra Team at McGill University, TTU-NRAO, GROWTH, AGILE Team, MWA, ATCA, AST3, TOROS, Pan-STARRS, NuSTAR, ATLAS Telescopes, BOOTES, CaltechNRAO, LIGO Scientific Collaboration, High Time Resolution Universe Survey, Nordic Optical Telescope, Las Cumbres Observatory Group, TZAC Consortium, LOFAR, IPN, DLT40, Texas Tech University, HAWC, ANTARES, KU, Dark Energy Camera GW-EM, CALET, Euro VLBI Team, ALMA), Astrophys. J. **848**, L12 (2017), arXiv:1710.05833 [astro-ph.HE] .

[125] T. Baker, E. Bellini, P. G. Ferreira, M. Lagos, J. Noller, and I. Sawicki, Phys. Rev. Lett. **119**, 251301 (2017), arXiv:1710.06394 [astro-ph.CO] .

[126] P. Creminelli and F. Vernizzi, Phys. Rev. Lett. **119**, 251302 (2017), arXiv:1710.05877 [astro-ph.CO] .

[127] J. Sakstein and B. Jain, Phys. Rev. Lett. **119**, 251303 (2017), arXiv:1710.05893 [astro-ph.CO] .

[128] J. M. Ezquiaga and M. Zumalacárregui, Phys. Rev. Lett. **119**, 251304 (2017), arXiv:1710.05901 [astro-ph.CO] .

[129] D. Langlois, R. Saito, D. Yamauchi, and K. Noui, Phys. Rev. **D97**, 061501 (2018), arXiv:1711.07403 [gr-qc] .

[130] A. Dima and F. Vernizzi, Phys. Rev. **D97**, 101302 (2018), arXiv:1712.04731 [gr-qc] .

[131] C. de Rham and S. Melville, (2018), arXiv:1806.09417 [hep-th] .

[132] S. Boran, S. Desai, E. O. Kahya, and R. P. Woodard, Phys. Rev. **D97**, 041501 (2018), arXiv:1710.06168 [astro-ph.HE] .

[133] B. Iyer *et al.*, *LIGO-India, Proposal of the Consortium for Indian Initiative in Gravitational-wave Observations (IndIGO)*, Tech. Rep. LIGO-M1100296 (LIGO Laboratory, 2011).

[134] C. S. Unnikrishnan, Int. J. Mod. Phys. **D22**, 1341010 (2013), arXiv:1510.06059 [physics.ins-det] .

[135] K. Somiya (KAGRA), *Gravitational waves. Numerical relativity - data analysis. Proceedings, 9th Edoardo Amaldi Conference, Amaldi 9, and meeting, NRDA 2011, Cardiff, UK, July 10-15, 2011*, Class. Quant. Grav. **29**, 124007 (2012), arXiv:1111.7185 [gr-qc] .

[136] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto (KAGRA), Phys. Rev. **D88**, 043007 (2013), arXiv:1306.6747 [gr-qc] .

[137] S. G. Turyshev, Annu. Rev. Nucl. Part. Sci. **58**, 30 (2008), arXiv:0806.1731 .

[138] J. P. Ostriker and J. E. Gunn, Astrophys. J. **157**, 1395 (1969).

[139] I. Shklovskii, Sov. Astron. **13**, 562 (1970).

[140] G. M. Harry, Class. Quantum Gravity **27**, 084006 (2010).

[141] A. J. Weinstein, *Gravitational waves. Numerical relativity - data analysis. Proceedings, 9th Edoardo Amaldi Conference, Amaldi 9, and meeting, NRDA 2011, Cardiff, UK, July 10-15, 2011*, Class. Quant. Grav. **29**, 124012 (2012), arXiv:1112.1057 [gr-qc] .

[142] W. Del Pozzo, J. Veitch, and A. Vecchio, Phys. Rev. **D83**, 082002 (2011), arXiv:1101.1391 [gr-qc] .

[143] M. Agathos, W. Del Pozzo, T. G. F. Li, C. Van Den Broeck, J. Veitch, and S. Vitale, Phys. Rev. D **89**, 082001 (2014).

[144] N. Yunes, R. O'Shaughnessy, B. J. Owen, and S. Alexander, Phys. Rev. D **82**, 064017 (2010).

[145] P. Hořava, A. Mohd, C. M. Melby-Thompson, and P. Shawhan, Gen. Relativ. Gravit. **46**, 1720 (2014), arXiv:1405.0249 .

[146] M. Zimmermann and E. Szedenits, Phys. Rev. D **20**, 351 (1979).

[147] W. Althouse, L. Jones, and A. Lazzarini, *Determination of Global and Local Coordinate Axes for the LIGO Sites (LIGO-T980044)*, Tech. Rep. (LIGO, 2001).

[148] B. Allen, *Gravitational Wave Detector Sites*, Tech. Rep. (University of Wisconsin–Milwaukee, 1996) arXiv:9607075 [gr-qc] .

[149] C. Ng and R. W. Romani, Astrophys. J. **673**, 411 (2008).

[150] B. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Astrophys. J. **713**, 671 (2010), arXiv:0909.3583 .

[151] A. Olive et al. and Particle Data Group, Chinese Phys. C **38**, 090001 (2014).

[152] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, Astron. J. **129**, 1993 (2005), arXiv:0412641 [astro-ph] .

[153] The LIGO Scientific Collaboration, Reports Prog. Phys. **72**, 076901 (2009), arXiv:0711.3041 .

[154] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein, and B. Stephens, in *Proc. LISA Symp. X* (IOP Publishing, Gainesville, Florida, 2014) arXiv:1410.4839 .

[155] I. M. Chakravarty, J. D. Roy, and R. G. Laha, *Handbook of methods of applied statistics Volume I* (John Wiley and Sons, New York, 1967) pp. 392–394.

[156] T. W. Anderson and D. A. Darling, Ann. Math. Stat. **23**, 193 (1952).

[157] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, Phys. Rev. D **58**, 084020 (1998), arXiv:9804044 [gr-qc] .

[158] R. Bondarescu, S. A. Teukolsky, and I. Wasserman, Phys. Rev. D **79**, 104003 (2009), arXiv:0809.3448 .

[159] D. I. Jones, Mon. Not. R. Astron. Soc. **453**, 53 (2015), arXiv:1501.05832 .

[160] J. D. Scargle, Astrophys. J. **504**, 405 (1998), arXiv:astro-ph/9711233 [astro-ph] .

[161] X. Fan, Y. Chen, and C. Messenger, Phys. Rev. **D94**, 084029 (2016), arXiv:1607.06735 [gr-qc] .

[162] `https://wiki.ligo.org/DASWG/LALSuite`.

[163] E. Jaynes, IEEE Trans. Syst. Sci. Cybern. **4**, 227 (1968).

[164] D. Foreman-Mackey, J. Open Source Softw. **1**, 1 (2016).

[165] C. R. Keeton, Mon. Not. R. Astron. Soc. **414**, 1418 (2011), arXiv:1102.0996 .

[166] E. Berti, E. Barausse, V. Cardoso, L. Gualtieri, P. Pani, U. Sperhake, L. C. Stein, N. Wex, K. Yagi, T. Baker, C. P. Burgess, F. S. Coelho, D. Doneva, A. D. Felice, P. G. Ferreira, P. C. C. Freire, J. Healy, C. Herdeiro, M. Horbatsch, B. Kleihaus, A. Klein, K. Kokkotas, J. Kunz, P. Laguna, R. N. Lang, T. G. F. Li, T. Littenberg, A. Matas, S. Mirshekari, H. Okawa, E. Radu, R. O'Shaughnessy, B. S. Sathyaprakash, C. V. D. Broeck, H. A. Winther, H. Witek, M. E. Aghili, J. Alsing, B. Bolen, L. Bombelli, S. Caudill, L. Chen, J. C. Degollado, R. Fujita, C. Gao, D. Gerosa, S. Kamali, H. O. Silva, J. G. Rosa, L. Sadeghian,

M. Sampaio, H. Sotani, and M. Zilhao, Class. Quantum Gravity **32**, 243001 (2015), arXiv:1501.07274 .

[167] B. P. Abbott and (The LIGO Scientific Collaboration), Phys. Rev. D **95**, 062003 (2017), arXiv:1602.03845 .

[168] W. Anderson, P. Brady, J. Creighton, and É. Flanagan, Phys. Rev. D **63**, 042003 (2001).

[169] C. Ng and R. W. Romani, Astrophys. J. **601**, 479 (2004), arXiv:0310155 .

[170] R. D. Ferdman, I. H. Stairs, M. Kramer, R. P. Breton, M. A. McLaughlin, P. C. C. Freire, A. Possenti, B. W. Stappers, V. M. Kaspi, R. N. Manchester, and A. G. Lyne, Astrophys. J. **767**, 85 (2013), arXiv:1302.2914 .

[171] B. J. Rickett, W. A. Coles, C. F. Nava, M. A. McLaughlin, S. M. Ransom, F. Camilo, R. D. Ferdman, P. C. C. Freire, M. Kramer, A. G. Lyne, and I. H. Stairs, Astrophys. J. **787**, 161 (2014), arXiv:arXiv:1404.1120v1 .

[172] W. W. Zhu, I. H. Stairs, P. B. Demorest, D. J. Nice, J. A. Ellis, S. M. Ransom, Z. Arzoumanian, K. Crowter, T. Dolch, R. D. Ferdman, E. Fonseca, M. E. Gonzalez, G. Jones, M. L. Jones, M. T. Lam, L. Levin, M. A. McLaughlin, T. Pennucci, K. Stovall, and J. Swiggum, Astrophys. J. **809**, 41 (2015), arXiv:arXiv:1504.00662v1 .

[173] R. E. Kass and A. E. Raftery, J. Am. Stat. Assoc. **90**, 773 (1995).

[174] Jeffreys Harold, *Theory of probability*, 3rd ed. (Clarendon Press, Oxford, 1998).

[175] C. P. Robert, N. Chopin, and J. Rousseau, Stat. Sci. **24**, 191 (2009), arXiv:arXiv:0804.3173v7 .

[176] B. P. Abbott et al., (The LIGO Scientific Collaboration, and The Virgo Collaboration), (2017), arXiv:1707.02669 .

[177] B. P. Abbott et al., (The LIGO Scientific Collaboration, and The Virgo Collaboration), Phys. Rev. D **96**, 062002 (2017), arXiv:1707.02667 .

[178] B. P. Abbott et al., (The LIGO Scientific Collaboration, and The Virgo Collaboration), Astrophys. J. Lett. **851** (2017), 10.3847/2041-8213/aa9f0c.

[179] L. P. Grishchuk, JETP Lett. **23**, 293 (1976).

[180] P. F. Michelson, Mon. Not. R. Astron. Soc. **227**, 933 (1987).

[181] N. Christensen, *On measuring the stochastic gravitational radiation background with laser interferometric antennas*, thesis, Massachusets Institute of Technology (1990).

[182] N. Christensen, Phys. Rev. D **46**, 5250 (1992).

[183] E. E. Flanagan, Phys. Rev. D **48**, 2389 (1993), arXiv:astro-ph/9305029 .

[184] B. Allen, in *Relativistic gravitation and gravitational radiation. Proceedings, School of Physics, Les Houches, France, September 26-October 6, 1995* (1996) pp. 373–417, arXiv:gr-qc/9604033 .

[185] B. Allen and J. Romano, Phys. Rev. D **59**, 102001 (1999).

[186] B. P. Abbott et al., (The LIGO Scientific Collaboration, and The Virgo Collaboration), Phys. Rev. Lett. **118**, 121101 (2017), arXiv:1612.02029 .

[187] A. A. Starobinsky, ZhETF Pisma Redaktsiiu **30**, 719 (1979).

[188] R. Easther, J. T. Giblin, Jr., and E. A. Lim, Phys. Rev. Lett. **99**, 221301 (2007), arXiv:astro-ph/0612294 .

[189] N. Barnaby, E. Pajer, and M. Peloso, Phys. Rev. **D85**, 023525 (2012), arXiv:1110.3327 .

[190] J. L. Cook and L. Sorbo, Phys. Rev. **D85**, 023534 (2012), [Erratum: Phys. Rev.D86,069901(2012)], arXiv:1109.0022 .

[191] M. S. Turner, Phys. Rev. **D55**, R435 (1997), arXiv:astro-ph/9607066 .

[192] R. Easther and E. A. Lim, JCAP **0604**, 010 (2006), arXiv:astro-ph/0601617 .

[193] M. Kamionkowski, A. Kosowsky, and M. S. Turner, Phys. Rev. **D49**, 2837 (1994), arXiv:astro-ph/9310044 .

[194] M. Gasperini and G. Veneziano, Astropart. Phys. **1**, 317 (1993), arXiv:hep-th/9211021 .

[195] M. Gasperini and G. Veneziano, Mod. Phys. Lett. **A8**, 3701 (1993), arXiv:hep-th/9309023 .

[196] M. Gasperini, JCAP **1612**, 010 (2016), arXiv:1606.07889 .

[197] C. Caprini and D. G. Figueroa, (2018), arXiv:1801.04268 .

[198] X.-J. Zhu, E. J. Howell, D. G. Blair, and Z.-H. Zhu, Mon. Not. Roy. Astron. Soc. **431**, 882 (2013), arXiv:1209.0595 .

[199] S. Marassi, R. Schneider, G. Corvino, V. Ferrari, and S. Portegies Zwart, Phys. Rev. **D84**, 124037 (2011), arXiv:1111.6125 .

[200] P. D. Lasky, M. F. Bennett, and A. Melatos, Phys. Rev. **D87**, 063004 (2013), arXiv:1302.6033 .

[201] P. A. Rosado, Phys. Rev. **D86**, 104007 (2012), arXiv:1206.1330 .

[202] X.-J. Zhu, X.-L. Fan, and Z.-H. Zhu, Astrophys. J. **729**, 59 (2011), arXiv:1102.2786 .

[203] S. Marassi, R. Schneider, and V. Ferrari, Mon. Not. Roy. Astron. Soc. **398**, 293 (2009), arXiv:0906.0461 .

[204] A. Buonanno, G. Sigl, G. G. Raffelt, H.-T. Janka, and E. Muller, Phys. Rev. **D72**, 084001 (2005), arXiv:astro-ph/0412277 .

[205] P. Sandick, K. A. Olive, F. Daigne, and E. Vangioni, Phys. Rev. **D73**, 104024 (2006), arXiv:astro-ph/0603544 .

[206] T. W. B. Kibble, J. Phys. **A9**, 1387 (1976).

[207] T. Damour and A. Vilenkin, Phys. Rev. **D71**, 063510 (2005), arXiv:hep-th/0410222 .

[208] X. Siemens, V. Mandic, and J. Creighton, Phys. Rev. Lett. **98**, 111101 (2007), arXiv:astro-ph/0610920 .

[209] S. Sarangi and S. H. H. Tye, Phys. Lett. **B536**, 185 (2002), arXiv:hep-th/0204074 .

[210] M. Maggiore, Phys. Rep. **331**, 283 (2000).

[211] M. Maggiore and A. Nicolis, Phys. Rev. D **62**, 024004 (2000).

[212] M. Gasperini and C. Ungarelli, Phys. Rev. D **64**, 064009 (2001).

[213] A. Nishizawa and K. Hayama, Phys. Rev. D **88**, 064005 (2013), arXiv:1307.1281 .

[214] A. Nishizawa, A. Taruya, and S. Kawamura, Phys. Rev. D **81**, 104043 (2010), arXiv:0911.0525 .

[215] J. R. Gair, J. D. Romano, and S. R. Taylor, Phys. Rev. D **92**, 102003 (2015), arXiv:1506.08668 .

[216] N. J. Cornish, L. O'Beirne, S. R. Taylor, and N. Yunes, (2017), arXiv:1712.07132 [gr-qc] .

[217] L. O'Beirne and N. J. Cornish, (2018), arXiv:1804.03146 [gr-qc] .

[218] N. Yunes and X. Siemens, Living Rev. Relativ. **16**, 9 (2013), arXiv:1304.3473 .

[219] L. C. Stein and N. Yunes, Phys. Rev. D **83**, 064038 (2011), arXiv:1012.3144 .

[220] C. M. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (W. H. Freeman, San Francisco, 1973).

[221] S. F. Hassan and R. A. Rosen, JHEP **02**, 126 (2012), arXiv:1109.3515 .

[222] J. D. Romano and N. J. Cornish, Living Rev. Relativ. **20**, 2 (2017), arXiv:1608.06889 .

[223] J. Maldacena, J. High Energy Phys. **2003**, 013 (2003), arXiv:astro-ph/0210603 .

[224] P. Jaranowski and A. Królak, *Analysis of Gravitational-Wave Data* (Cambridge University Press, Cambridge, 2009).

[225] L. P. Grishchuk, A. G. Doroshkevich, and V. M. Yudin, Zh. Eksp. Teor. Fiz. **69**, 1857 (1975).

[226] M. Lachieze-Rey and J.-P. Luminet, Phys. Rept. **254**, 135 (1995), arXiv:gr-qc/9605010 .

[227] N. J. Cornish, D. N. Spergel, and G. D. Starkman, *Topology of the universe. Proceedings, Conference, Cleveland, USA, October 17-19, 1997*, Class. Quant. Grav. **15**, 2657 (1998), arXiv:astro-ph/9801212 .

[228] G. P. Smith, M. Jauzac, J. Veitch, W. M. Farr, R. Massey, and J. Richard, Mon. Not. Roy. Astron. Soc. **475**, 3823 (2018), arXiv:1707.03412 [astro-ph.HE] .

[229] S. H. S. Alexander, M. E. Peskin, and M. M. Sheikh-Jabbari, Phys. Rev. Lett. **96**, 081301 (2006), arXiv:hep-th/0403069 .

[230] N. Seto and A. Taruya, Phys. Rev. Lett. **99**, 121101 (2007), arXiv:0707.0535 .

[231] N. Seto and A. Taruya, Phys. Rev. D **77**, 103001 (2008), arXiv:0801.4185 .

[232] K. Hinterbichler, Rev. Mod. Phys. **84**, 671 (2012), arXiv:1105.3735 .

[233] R. Jackiw and S.-Y. Pi, Phys. Rev. D **68**, 104012 (2003), arXiv:gr-qc/0308071 .

[234] C. F. Sopuerta and N. Yunes, Phys. Rev. D **80**, 064006 (2009).

[235] A. J. Ruiter, K. Belczynski, M. Benacquista, S. L. Larson, and G. Williams, Astrophys. J. **717**, 1006 (2010), arXiv:0705.3272 .

[236] V. A. Kostelecky and S. Samuel, Phys. Rev. **D39**, 683 (1989).

[237] V. A. Kostelecky and R. Potting, Nucl. Phys. **B359**, 545 (1991).

[238] V. A. Kostelecky and R. Potting, Phys. Rev. **D51**, 3923 (1995), arXiv:hep-ph/9501341 .

[239] V. A. Kostelecky, Phys. Rev. **D69**, 105009 (2004), arXiv:hep-th/0312310 .

[240] V. A. Kostelecký and M. Mewes, Phys. Lett. B **757**, 510 (2016), arXiv:1602.04782 .

[241] T. Jacobson and D. Mattingly, Phys. Rev. **D64**, 024028 (2001), arXiv:gr-qc/0007031 [gr-qc] .

[242] C. Eling, T. Jacobson, and D. Mattingly, in *Deserfest: A celebration of the life and works of Stanley Deser. Proceedings, Meeting, Ann Arbor, USA, April 3-5, 2004* (2004) pp. 163–179, arXiv:gr-qc/0410001 [gr-qc] .

[243] T. Jacobson, *Proceedings, Workshop on From quantum to emergent gravity: Theory and phenomenology (QG-Ph): Trieste, Italy, June 11-15, 2007*, PoS **QG-PH**, 020 (2007), arXiv:0801.1547 [gr-qc] .

[244] G. Cusin, I. Dvorkin, C. Pitrou, and J.-P. Uzan, Phys. Rev. Lett. **120**, 231101 (2018), arXiv:1803.03236 .

[245] B. P. Abbott et al., (The LIGO Scientific Collaboration, and The Virgo Collaboration), Phys. Rev. Lett. **118**, 121102 (2017), arXiv:1612.02030 .

[246] P. B. Covas et al. (LSC Instrument Authors), Phys. Rev. D **97**, 082002 (2018), arXiv:1801.07204 .

[247] R. L. Forward, Phys. Rev. D **17**, 379 (1978).

[248] P. Linsay, P. Saulson, R. Weiss, and S. E. Whitcomb, *A Study of a Long Baseline Gravitational Wave Antenna System*, Tech. Rep. LIGO-T830001 (National Science Foundation, 1983).

[249] F. B. Estabrook, Gen. Relativ. Gravit. **17**, 719 (1985).

[250] R. Schilling, *1st International LISA Symposium on Gravitational Waves Oxfordshire, England, July 9-12, 1996*, Class. Quant. Grav. **14**, 1513 (1997).

[251] M. Rakhmanov, *Response of LIGO to Gravitational Waves at High Frequencies and in the Vicinity of the FSR (37.5 kHz)*, Tech. Rep. LIGO-T060237 (LIGO Laboratory, 2005).

[252] M. Rakhmanov, J. D. Romano, and J. T. Whelan, *Proceedings, 12th Workshop on Gravitational wave data analysis (GWDAW-12): Cambridge, USA, December 13-16, 2007*, Class. Quant. Grav. **25**, 184017 (2008), arXiv:0808.3805 .

[253] M. Rakhmanov, Class. Quant. Grav. **26**, 155010 (2009), arXiv:1407.5376 .

[254] R. Essick, S. Vitale, and M. Evans, Phys. Rev. **D96**, 084004 (2017), arXiv:1708.06843 .

[255] M. E. Tobar, T. Suzuki, and K. Kuroda, Phys. Rev. **D59**, 102002 (1999).

[256] B. Allen and A. C. Ottewill, Phys. Rev. D **56**, 545 (1997), arXiv:gr-qc/9607068 .

[257] R. A. Isaacson, Phys. Rev. **166**, 1263 (1968).

[258] R. A. Isaacson, Phys. Rev. **166**, 1272 (1968).

[259] R. M. Zalaletdinov, General Relativity and Gravitation **28**, 953 (1996).

[260] T. Callister, L. Sammut, S. Qiu, I. Mandel, and E. Thrane, Phys. Rev. **X6**, 031018 (2016), arXiv:1604.02513 .

[261] B. Bonga and J. S. Hazboun, Phys. Rev. **D96**, 064018 (2017), arXiv:1708.05621 [gr-qc] .

[262] M. A. H. MacCallum and A. H. Taub, Commun. Math. Phys. **30**, 153 (1973).

[263] R. Delbourgo and A. Salam, Phys. Lett. **40B**, 381 (1972).

[264] T. Eguchi and P. G. O. Freund, Phys. Rev. Lett. **37**, 1251 (1976).

[265] L. Alvarez-Gaume and E. Witten, Nucl. Phys. **B234**, 269 (1984).

[266] M. B. Green and J. H. Schwarz, Phys. Lett. **B149**, 117 (1984).

[267] A. Ashtekar, A. P. Balachandran, and S. Jo, Int. J. Mod. Phys. **A4**, 1493 (1989).

[268] B. A. Campbell, M. J. Duncan, N. Kaloper, and K. A. Olive, Nucl. Phys. **B351**, 778 (1991).

[269] B. A. Campbell, N. Kaloper, R. Madden, and K. A. Olive, Nucl. Phys. **B399**, 137 (1993), arXiv:hep-th/9301129 .

[270] S. Weinberg, Phys. Rev. **D77**, 123541 (2008), arXiv:0804.4291 [hep-th] .

[271] V. Taveras and N. Yunes, Phys. Rev. **D78**, 064070 (2008), arXiv:0807.2652 [gr-qc] .

[272] S. Alexander, L. S. Finn, and N. Yunes, Phys. Rev. D **78**, 066005 (2008), arXiv:0712.2542 .

[273] S. Mercuri and V. Taveras, Phys. Rev. **D80**, 104007 (2009), arXiv:0903.4407 [gr-qc] .

[274] S. Alexander and J. Martin, Phys. Rev. D **71**, 063526 (2005).

[275] C. R. Contaldi, J. Magueijo, and L. Smolin, Phys. Rev. Lett. **101**, 141101 (2008), arXiv:0806.3082 .

[276] D. L. Lee, Phys. Rev. **D10**, 2374 (1974).

[277] T. Damour and G. Esposito-Farèse, Phys. Rev. D **46**, 4128 (1992).

[278] T. Damour and G. Esposito-Farèse, Phys. Rev. D **54**, 1474 (1996), arXiv:gr-qc/9602056 .

[279] T. Chiba, T. Harada, and K.-i. Nakao, Progress of Theoretical Physics Supplement **128**, 335 (1997).

[280] Y. Fujii and K. Maeda, *The scalar-tensor theory of gravitation* (Cambridge University Press, 2007).

[281] T. P. Sotiriou and V. Faraoni, Rev. Mod. Phys. **82**, 451 (2010), arXiv:0805.1726 [gr-qc] .

[282] T. P. Sotiriou, *Proceedings of the 7th Aegean Summer School : Beyond Einstein's theory of gravity. Modifications of Einstein's Theory of Gravity at Large Distances.: Paros, Greece, September 23-28, 2013*, Lect. Notes Phys. **892**, 3 (2015), arXiv:1404.2955 [gr-qc] .

[283] T. Damour and A. M. Polyakov, Nucl. Phys. **B423**, 532 (1994), arXiv:hep-th/9401069 .

[284] T. Damour, F. Piazza, and G. Veneziano, Phys. Rev. **D66**, 046007 (2002), arXiv:hep-th/0205111 .

[285] T. Damour, F. Piazza, and G. Veneziano, Phys. Rev. Lett. **89**, 081601 (2002), arXiv:gr-qc/0204094 .

[286] T. R. Taylor and G. Veneziano, Phys. Lett. **B213**, 450 (1988).

[287] L. Randall and R. Sundrum, Phys. Rev. Lett. **83**, 4690 (1999), arXiv:hep-th/9906064 .

[288] L. Randall and R. Sundrum, Phys. Rev. Lett. **83**, 3370 (1999), arXiv:hep-ph/9905221 .

[289] V. Faraoni, *Cosmology in scalar tensor gravity* (2004).

[290] T. Clifton, P. G. Ferreira, A. Padilla, and C. Skordis, Phys. Rept. **513**, 1 (2012), arXiv:1106.2476 .

[291] M. Fierz and W. Pauli, Proc. R. Soc. A Math. Phys. Eng. Sci. **173**, 211 (1939).

[292] L. S. Finn and P. J. Sutton, Phys. Rev. **D65**, 044022 (2002), arXiv:gr-qc/0109049 [gr-qc] .

[293] H. van Dam and M. J. G. Veltman, Nucl. Phys. **B22**, 397 (1970).

[294] V. I. Zakharov, JETP Lett. **12**, 312 (1970), [Pisma Zh. Eksp. Teor. Fiz.12,447(1970)].

[295] P. Van Nieuwenhuizen, Phys. Rev. **D7**, 2300 (1973).

[296] D. G. Boulware and S. Deser, Phys. Rev. **D6**, 3368 (1972).

[297] C. de Rham, G. Gabadadze, and A. J. Tolley, Phys. Rev. Lett. **106**, 231101 (2011), arXiv:1011.1232 [hep-th] .

[298] L. Alberte, Int. J. Mod. Phys. **D21**, 1250058 (2012), arXiv:1110.3818 [hep-th] .

[299] S. F. Hassan, A. Schmidt-May, and M. von Strauss, JHEP **05**, 086 (2013), arXiv:1208.1515 [hep-th] .

[300] L. Bernard, C. Deffayet, and M. von Strauss, Phys. Rev. **D91**, 104013 (2015), arXiv:1410.8302 [hep-th] .

[301] L. Bernard, C. Deffayet, and M. von Strauss, JCAP **1506**, 038 (2015), arXiv:1504.04382 [hep-th] .

[302] L. Bernard, C. Deffayet, A. Schmidt-May, and M. von Strauss, Phys. Rev. **D93**, 084020 (2016), arXiv:1512.03620 [hep-th] .

[303] A. I. Vainshtein, Phys. Lett. **39B**, 393 (1972).

[304] M. Srednicki, *Quantum Field Theory* (Cambridge University Press, Cambridge, 2007).

[305] L. H. Ryder, *Quantum Field Theory*, 2nd ed. (Cambridge University Press, Cambridge, 2007).

[306] E. C. G. Stueckelberg, Helv. Phys. Acta **30**, 209 (1957).

[307] C. M. Caves, Annals Phys. **125**, 35 (1980).

[308] K. S. Thorne, in *Gravitational Radiation; Proc. Adv. Study Institute, Les Houches, Haute-Savoie, Fr. June 2-21, 1982 (A84-35026 16-90).*, edited by N. Deruelle and T. Piran (North-Holland Publishing Company, Amsterdam, 1983) pp. 1–57.

[309] G. D. Moore and A. E. Nelson, JHEP **09**, 023 (2001), arXiv:hep-ph/0106220 [hep-ph] .

[310] C. Talmadge, J. P. Berthias, R. W. Hellings, and E. M. Standish, Phys. Rev. Lett. **61**, 1159 (1988).

[311] S. R. Choudhury, G. C. Joshi, S. Mahajan, and B. H. J. McKellar, Astropart. Phys. **21**, 559 (2004), arXiv:hep-ph/0204161 [hep-ph] .

[312] R. Brito, V. Cardoso, and P. Pani, Phys. Rev. **D88**, 064006 (2013), arXiv:1309.0818 [gr-qc] .

[313] N. Cornish, D. Blas, and G. Nardini, Phys. Rev. Lett. **119**, 161102 (2017), arXiv:1707.06101 [gr-qc] .

[314] R. Ciolfi and D. M. Siegel, Astrophys. J. **798**, L36 (2015), arXiv:1411.2015 [astro-ph.HE] .

[315] L. Rezzolla and P. Kumar, Astrophys. J. **802**, 95 (2015), arXiv:1410.8560 [astro-ph.HE] .

[316] D. Tsang, J. S. Read, T. Hinderer, A. L. Piro, and R. Bondarescu, Phys. Rev. Lett. **108**, 011102 (2012), arXiv:1110.0467 [astro-ph.HE] .

[317] Philosophical Transactions (1665-1678) **12**, 893 (1677).

[318] D. C. Backer and R. W. Hellings, Ann. Rev. Astron. Astrophys. **24**, 537 (1986).

[319] I. I. Shapiro, Phys. Rev. Lett. **13**, 789 (1964).

[320] R. Blandford and S. A. Teukolsky, The Astrophysical Journal **205**, 580 (1976).

[321] J. D. Giorgini, D. K. Yeomans, A. B. Chamberlin, P. W. Chodas, R. A. Jacobson, M. S. Keesey, J. H. Lieske, S. J. Ostro, E. M. Standish, and R. N. Wimberly, in *AAS/Division for Planetary Sciences Meeting Abstracts #28*, Bulletin of the American Astronomical Society, Vol. 28 (1996) p. 1158.

[322] E. F. Arias, P. Charlot, M. Feissel, and J.-F. Lestrade, Astron. Astrophys. **303**, 604 (1995).

[323] J. Aasi *et al.* (VIRGO, LIGO Scientific), Class. Quant. Grav. **32**, 115012 (2015), arXiv:1410.7764 [gr-qc] .

[324] J. Mattinson, M. Isi, and A. J. Weinstein, *Using Continuous GWs from Known Pulsars to Measure Gravitational Wave Speed*, Tech. Rep. LIGO-T1600262 (LIGO Laboratory, (2016)).

[325] S. Mirshekari, N. Yunes, and C. M. Will, Phys. Rev. **D85**, 024041 (2012), arXiv:1110.2720 [gr-qc] .

[326] R. Tso, M. Isi, Y. Chen, and L. Stein, in *Proceedings, 7th Meeting on CPT and Lorentz Symmetry (CPT 16): Bloomington, Indiana, USA, June 20-24, 2016* (2017) pp. 205–208, arXiv:1608.01284 [gr-qc] .

[327] C. M. Will, Phys. Rev. **D57**, 2061 (1998), arXiv:gr-qc/9709011 [gr-qc] .

[328] M. Maggiore, *Gravitational Waves* (Oxford University Press (OUP), 2007).

[329] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, A. M. Sintes, J. T. Whelan, B. Brügmann, P. Diener, N. Dorband, J. Gonzalez, M. Hannam, S. Husa, D. Pollney, L. Rezzolla, L. Santamaría, U. Sperhake, and J. Thornburg, Physical Review D **77** (2008), 10.1103/physrevd.77.104017.

[330] N. Johnson-McDaniel, private communication (2018).

[331] J. Aasi *et al.* (The LIGO Scientific Collaboration), Class. Quant. Grav. **32**, 074001 (2015), arXiv:1411.4547 [gr-qc] .

[332] F. Acernese *et al.* (Virgo), Class. Quant. Grav. **32**, 024001 (2015), arXiv:1408.3978 [gr-qc] .

[333] N. Yunes, K. Yagi, and F. Pretorius, Phys. Rev. **D94**, 084002 (2016), arXiv:1603.08955 [gr-qc] .

[334] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), (2018), arXiv:1805.11579 [gr-qc] .

[335] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), (2018), arXiv:1805.11581 [gr-qc] .

[336] A. Arvanitaki and S. Dubovsky, Physical Review D **83**, 044026 (2011).

[337] H. Yoshino and H. Kodama, Progress of Theoretical and Experimental Physics **2014**, 43E02 (2014), 1312.2326 .

[338] H. Yoshino and H. Kodama, Progress of Theoretical and Experimental Physics **2015** (2015), 10.1093/ptep/ptv067.

[339] A. Arvanitaki, M. Baryakhtar, and X. Huang, Physical Review D **91**, 084011 (2015).

[340] A. Arvanitaki, M. Baryakhtar, S. Dimopoulos, S. Dubovsky, and R. Lasenby, Physical Review D **95**, 043001 (2017).

[341] R. Brito, S. Ghosh, E. Barausse, E. Berti, V. Cardoso, I. Dvorkin, A. Klein, and P. Pani, Physical Review Letters **119**, 131101 (2017).

[342] R. Brito, S. Ghosh, E. Barausse, E. Berti, V. Cardoso, I. Dvorkin, A. Klein, and P. Pani, Physical Review D **96**, 064050 (2017).

[343] M. Baryakhtar, R. Lasenby, and M. Teo, Phys. Rev. D **96**, 035019 (2017), arXiv:1704.05081 .

[344] R. D. Peccei and H. R. Quinn, Physical Review Letters **38**, 1440 (1977).

[345] R. D. Peccei and H. R. Quinn, Physical Review D **16**, 1791 (1977).

[346] S. Weinberg, Physical Review Letters **40**, 223 (1978).

[347] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, Physical Review D **81**, 123530 (2010).

[348] M. Goodsell, J. Jaeckel, J. Redondo, and A. Ringwald, JHEP **11**, 027 (2009), arXiv:0909.0515 [hep-ph] .

[349] J. Jaeckel and A. Ringwald, Ann. Rev. Nucl. Part. Sci. **60**, 405 (2010), arXiv:1002.0329 [hep-ph] .

[350] R. Essig *et al.*, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013* (2013) arXiv:1311.0029 [hep-ph] .

[351] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten, Phys. Rev. **D95**, 043541 (2017), arXiv:1610.08297 [astro-ph.CO] .

[352] J. E. Kim and G. Carosi, Rev. Mod. Phys. **82**, 557 (2010), arXiv:0807.3125 [hep-ph] .

[353] G. G. Raffelt, Lect. Notes Phys. **741**, 51 (2006), arXiv:0611350 [hep-ph] .

[354] R. Penrose, Nuovo Cimento Rivista Serie **1** (1969).

[355] W. H. Press and S. A. Teukolsky, Nature **238**, 211 (1972).

[356] J. D. Bekenstein, Phys. Rev. D **7**, 949 (1973).

[357] R. Brito, V. Cardoso, and P. Pani, Class. Quant. Grav. **32**, 134001 (2015), arXiv:1411.0686 [gr-qc] .

[358] R. Brito, V. Cardoso, and P. Pani, *Superradiance*, Lecture Notes in Physics, Vol. 906 (Springer International Publishing, Cham, 2015).

[359] A. D. Plascencia and A. Urbano, JCAP **1804**, 059 (2018), arXiv:1711.08298 [gr-qc] .

[360] V. Cardoso, Ó. J. Dias, G. S. Hartnett, M. Middleton, P. Pani, and J. E. Santos, J. Cosmol. Astropart. Phys. **2018**, 043 (2018), arXiv:1801.01420 .

[361] D. Baumann, H. S. Chia, and R. A. Porto, (2018), arXiv:1804.03208 [gr-qc] .

[362] O. A. Hannuksela, R. Brito, E. Berti, and T. G. F. Li, (2018), arXiv:1804.09659 [astro-ph.HE] .

[363] J. Zhang and H. Yang, (2018), arXiv:1808.02905 [gr-qc] .

[364] C. S. Reynolds, Space Sci. Rev. **183**, 277 (2014), arXiv:1302.3260 [astro-ph.HE] .

[365] J. E. McClintock, R. Narayan, and J. F. Steiner, Space Sci. Rev. **183**, 295 (2014), arXiv:1303.1583 [astro-ph.HE] .

[366] I. M. Ternov, V. R. Khalilov, G. A. Chizhov, and A. B. Gaina, Sov. Phys. J. **21**, 1200 (1978), [Izv. Vuz. Fiz.21N9,109(1978)].

[367] T. J. Zouros and D. M. Eardley, Ann. Phys. (N. Y). **118**, 139 (1979).

[368] S. Detweiler, Phys. Rev. D **22**, 2323 (1980).

[369] S. R. Dolan, Phys. Rev. D **76**, 084001 (2007), arXiv:0705.2880 .

[370] H. Witek, V. Cardoso, A. Ishibashi, and U. Sperhake, Phys. Rev. D **87**, 043513 (2013), arXiv:1212.0551 .

[371] W. E. East, Phys. Rev. D **96** (2017), 10.1103/PhysRevD.96.024004, arXiv:1705.01544 .

[372] S. A. Teukolsky, Class. Quantum Gravity **32**, 124006 (2014), arXiv:1410.2130 .

[373] Y. B. Zel'dovich, Pis'ma Zh. Eksp. Teor. Fiz. **14**, 270 [JETP Lett. **14**, 180 (1971)] (1971).

[374] Y. B. Zel'dovich, Zh. Eksp. Teor. Fiz **62**, 2076 [Sov.Phys. JETP **35**, 1085 (1972)] (1972).

[375] J. D. Bekenstein and M. Schiffer, Phys. Rev. **D58**, 064014 (1998), arXiv:gr-qc/9803033 [gr-qc] .

[376] T. Damour, N. Deruelle, and R. Ruffini, Lett. Nuovo Cim. **15**, 257 (1976).

[377] H. Furuhashi and Y. Nambu, Prog. Theor. Phys. **112**, 983 (2004), arXiv:gr-qc/0402037 [gr-qc] .

[378] M. J. Strafuss and G. Khanna, Phys. Rev. **D71**, 024034 (2005), arXiv:gr-qc/0412023 [gr-qc] .

[379] W. E. East and F. Pretorius, Phys. Rev. Lett. **119**, 041101 (2017), arXiv:1704.04791 [gr-qc] .

[380] C. A. R. Herdeiro and E. Radu, Phys. Rev. Lett. **119**, 261101 (2017), arXiv:1706.06597 [gr-qc] .

[381] P. Pani, V. Cardoso, L. Gualtieri, E. Berti, and A. Ishibashi, Phys. Rev. **D86**, 104017 (2012).

[382] V. P. Frolov, P. Krtou, D. Kubizak, and J. E. Santos, Phys. Rev. Lett. **120**, 231103 (2018).

[383] S. R. Dolan, "Instability of the Proca field on Kerr spacetime," (2018), arXiv:1806.01604 [gr-qc] .

[384] W. E. East, (2018), arXiv:1807.00043 [gr-qc] .

[385] H. Yoshino and H. Kodama, Prog. Theor. Phys. **128**, 153 (2012).

[386] H. Yoshino and H. Kodama, Class. Quant. Grav. **32**, 214001 (2015).

[387] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, Class. Quant. Grav. **25**, 114029 (2008).

[388] S. Klimenko *et al.*, Phys. Rev. **D93**, 042004 (2016).

[389] S. A. Teukolsky, Astrophys. J. **185**, 635 (1973).

[390] E. W. Leaver, Proc. R. Soc. A Math. Phys. Eng. Sci. **402**, 285 (1985).

[391] E. Berti, V. Cardoso, and M. Casals, Phys. Rev. D **73**, 024013 (2006), arXiv:0511111 [gr-qc] .

[392] J. H. Taylor and J. M. Weisberg, Astrophys. J. **345**, 434 (1989).

[393] S. D'Antonio *et al.*, (2018), arXiv:1809.07202 [gr-qc] .

[394] S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, Physical Review D **93**, 123009 (2016).

[395] L. Sun, A. Melatos, S. Suvorova, W. Moran, and R. J. Evans, Phys. Rev. D **97**, 043013 (2018).

[396] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. D **95**, 122003 (2017).

[397] B. P. Abbott *et al.*, "Search for gravitational waves from a long-lived remnant of the binary neutron star merger GW170817," (2018), in prep.

[398] L. Sun, A. Melatos, *et al.*, "Application of hidden Markov model tracking to the search for long-duration transient gravitational waves from the remnant of the binary neutron star merger GW170817," (2018), in prep.

[399] P. Jaranowski, A. Królak, and B. F. Schutz, Physical Review D **58**, 063001 (1998).

[400] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, Physical Review D **77**, 082001 (2008).

[401] D. H. Shoemaker *et al.*, LIGO Report No. T0900288 (2009).

[402] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Class. Quant. Grav. **34**, 044001 (2017), arXiv:1607.08697 [astro-ph.IM] .

[403] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Instrument Science White Paper*, Tech. Rep. LIGO-T1600119 (LIGO Laboratory, 2016).

[404] E. T. S. Team, *Einstein gravitational wave Telescope conceptual design study*, Tech. Rep. ET-0106C-10 (European Gravitational Observatory, 2011).

[405] B. Sathyaprakash *et al.*, *Gravitational waves. Numerical relativity - data analysis. Proceedings, 9th Edoardo Amaldi Conference, Amaldi 9, and meeting, NRDA 2011, Cardiff, UK, July 10-15, 2011*, Class. Quant. Grav. **29**, 124013 (2012), [Erratum: Class. Quant. Grav.30,079501(2013)], arXiv:1206.0331 [gr-qc] .

[406] H.-Y. Chen, D. E. Holz, J. Miller, M. Evans, S. Vitale, and J. Creighton, "Distance measures in gravitational-wave astrophysics and cosmology," (2017), arXiv:1709.08079 [astro-ph.CO] .

[407] P. A. R. Ade *et al.* (Planck), Astron. Astrophys. **594**, A13 (2016), arXiv:1502.01589 [astro-ph.CO] .

[408] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Physical Review D **57**, 2101 (1998).

[409] P. R. Brady and T. Creighton, Physical Review D **61**, 082001 (2000).

[410] K. Wette and R. Prix, Phys. Rev. **D88**, 123005 (2013), arXiv:1310.5587 [gr-qc] .

[411] K. Wette, Phys. Rev. **D92**, 082003 (2015), arXiv:1508.02372 [gr-qc] .

[412] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Living Rev. Relativ. **21**, 3 (2018).

[413] A. H. Mroue *et al.*, Phys. Rev. Lett. **111**, 241104 (2013), arXiv:1304.6077 [gr-qc] .

[414] G. Lovelace *et al.*, Class. Quant. Grav. **32**, 065007 (2015), arXiv:1411.7297 [gr-qc] .

[415] M. A. Scheel, M. Giesler, D. A. Hemberger, G. Lovelace, K. Kuper, M. Boyle, B. Szilagyi, and L. E. Kidder, Class. Quant. Grav. **32**, 105009 (2015), arXiv:1412.1803 [gr-qc] .

[416] T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, M. A. Scheel, and B. Szilagyi, Class. Quant. Grav. **33**, 165001 (2016), arXiv:1512.06800 [gr-qc] .

[417] S. M. Gaebel and J. Veitch, Class. Quant. Grav. **34**, 174003 (2017), arXiv:1703.08988 [astro-ph.IM] .

[418] S. Vitale and M. Evans, Phys. Rev. **D95**, 064052 (2017), arXiv:1610.06917 [gr-qc] .

[419] R. A. Remillard and J. E. McClintock, Ann. Rev. Astron. Astrophys. **44**, 49 (2006), arXiv:astro-ph/0606352 [astro-ph] .

[420] M. Middleton, in *Astrophysics of Black Holes: From Fundamental Aspects to Latest Developments*, Astrophysics and Space Science Library, Vol. 440, edited by C. Bambi (2016) p. 99.

[421] M. C. Miller and J. M. Miller, Phys. Rept. **548**, 1 (2014), arXiv:1408.4145 [astro-ph.HE] .

[422] S. Suvorova, P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans, Physical Review D **96**, 102006 (2017).

[423] J. A. Orosz, J. E. McClintock, J. P. Aufdenberg, R. A. Remillard, M. J. Reid, R. Narayan, and L. Gou, The Astrophysical Journal **742**, 84 (2011).

[424] F. van Leeuwen, Astron. Astrophys. **474**, 653 (2007), arXiv:0708.1752 [astro-ph] .

[425] L. Iorio, Astrophysics and Space Science **315**, 335 (2008), arXiv:0707.3525 [gr-qc] .

[426] K. S. Thorne and R. D. Blandford, *Modern Classical Physics: Optics, Fluids, Plasmas, Elasticity, Relativity, and Statistical Physics* (Princeton University Press, Princeton, 2017).

Appendix A

MEASURING A GRAVITATIONAL WAVE

A.1 Frames and gauges

The detector response for LIGO is usually computed in one of two coordinate systems: the transverse-traceless (TT) gauge and the local-Lorentz frame (LLF) of the beamsplitter. In reality, in both cases we use a frame (in the special-relativistic sense) centered on and comoving with the beamsplitter (corner station), with the x- and y-axes along the detector arms. The difference is in the gauge for the metric—that is, equivalently, in the specific way the coordinates are laid down. Given this, a more specific name for what is usually just denoted "LLF" would be "the Fermi-normal coordinate system in the detector frame," cf. Eq. (13.73) in [220].

Note there's a third, more correct option: the proper frame of the beam-splitter, which is analogous to the LLF, except one does not assume the masses are freely falling, i.e. we perturb over the background Earth's spacetime curvature, rather than the Minkowski metric. To lowest order, this just modifies the g_{00} term to be something like $-(1-2\Phi)$, where Φ is the Newtonian gravitational potential. Something similar could be done for the TT gauge. In any case, the spatial terms in the proper frame of the beam splitter will be the same as in nonaccelerated LLF or TT, depending on the gauge. See Box 37.1 in [220] for more exposition.

In the TT gauge, to linear order in the perturbation, the metric takes the form of $g_{ab} = \eta_{ab} + h_{ab}$ (see, e.g., Eq. (27.80) in [426]). Using geodesic deviation, it may then be shown that the coordinate locations of the masses do not change with time as the wave goes by, i.e. they remain at fixed coordinates L, measured from the beam-splitter. However, the proper distances between the masses and the beamsplitter do change. This is manifested in the propagation of light, which gets red/blue shifted as it propagates up and down the tubes (see Exercise 27.19 in [426]). The phase-shift measurable by the detector can be obtained by solving for the phase in the cavity, using the fact that k_x and k_y are conserved (but note the same is not true for k^x and k^y!), and that the coordinate locations of the masses do not change. The result is given in Eq. (27.89) of [426].

In the LLF, the metric is just the flat metric, $g_{ab} = \eta_{ab}$, up to corrections of order

$h \times (L/\lambda)^2$——very small for LIGO and Virgo (see, e.g., Eq. (27.93) in [426]). These corrections also control the effect of the GW perturbation on the laser light in the cavities, which is consequently very small (unlike in TT). While the effect on the light is of order $\ll h$, the relative motion of the masses due to geodesic deviation is of order h. We can get the phase shift measured by the instrument by treating light as usual in a Newtonian setting (i.e., no gravitational redshift, or any other complicated effects) with masses moving with an effective acceleration given by h through the equation of geodesic deviation. The result is given in Eq. (27.96) of [426].

In summary: in the TT gauge, the masses do not move, but light gets blue and redshifted, leading to a change in accumulated phase; in the LLF (Fermi-normal) gauge, there's no effect on the light, but the masses move and so change the roundtrip time and accumulated phase; as expected, both observables are the same, no matter which gauge is used to compute them.

Importantly, also note that the LLF gauge is valid only in the small-antenna limit. Otherwise, the local semi-Newtonian viewpoint breaks down and, mathematically, the expansion in powers of L/λ fails to converge. Meanwhile, the TT gauge is usable always (as long as there's separation of scales, so we can separate the wave from the background in the first place).

From the above analysis it is also clear why our instruments are designed to measure separations between different points in space and not, say, difference in the ticking rates of clocks: locally, effects of the GW on the time (non-synchronous) components of the metric are vanishingly small. Consequently, this is why we analyze the geodesic deviation and focus on the electric components of the Riemann tensor, R_{0i0j}.

A.2 Detector output

In the LLF gauge (when it is valid), the effect of the GW on the test masses can be approximated via the equation of geodesic deviation: we compute the electric components of the Riemann tensor, R_{0i0j}, at the origin (beam-splitter) to get the acceleration of the test masses with respect to the origin

$$\frac{\mathrm{d}^2 x^i}{\mathrm{d}\tau^2} = -R_{0i0j} x^j = \frac{1}{2} \frac{\mathrm{d}^2 A_{ij}}{\mathrm{d}\tau^2} x^j \,, \tag{A.1}$$

after defining A_{ij} as (negative) twice the double time-integral of R_{0i0j}. Letting $x^i = x^i(0) + \Delta x^i(t)$, and integrating the above equation twice, then:

$$\Delta x^i = \frac{1}{2} A_{ij} x^j(0), \tag{A.2}$$

For two arms, this implies

$$\begin{aligned} \Delta L &= X - Y \\ &= (L + \hat{X}_i \Delta X^i) - (L + \hat{Y}_i \Delta Y^i) \\ &= \hat{X}_i \Delta X^i - \hat{Y}_i \Delta Y^i \\ &= \frac{1}{2} L \left(\hat{X}^i \hat{X}^j - \hat{Y}^i \hat{Y}^j \right) A_{ij} \end{aligned} \tag{A.3}$$

Let

$$A = \frac{\Delta L}{L} = \frac{1}{2} \left(\hat{X}^i \hat{X}^j - \hat{Y}^i \hat{Y}^j \right) A_{ij} \tag{A.4}$$

for the two arms X and Y. It can be shown, e.g. from Eq. (27.96) in [426], that detectors measure a phase shift in the laser light such that $\Delta\phi \propto LA$. For more details see Eqs. (9.50) and (27.89) in [426], and the text above Eq. (13.95) in [110].

Appendix B

TABULATED RESULTS: FIRST SEARCH FOR NONTENSORIAL GRAVITATIONAL WAVES FROM KNOWN PULSARS

Here we present tabulated results from the analysis discussed in Chapter 9.

Table B.1: Results obtained using orientation information (see Chapter 9 for details).

	f_{GW} (Hz)	$\log \mathcal{B}_N^{GR}$	$\log \mathcal{B}_N^{GR+s}$	$\log \mathcal{B}_N^{GR+v}$	$\log \mathcal{B}_N^{GR+sv}$	$\log \mathcal{B}_N^{s}$	$\log \mathcal{B}_N^{v}$	$\log \mathcal{B}_N^{sv}$	$\log O_N^S$	$\log O_I^S$	$h_s^{95\%}$	$h_v^{95\%}$ $\times 10^{-26}$	$h_t^{95\%}$
J0205+6449	30.41	-0.19	-0.07	-0.32	-0.26	0.06	-0.09	-0.10	-0.12	-0.53	14.79	7.99	3.70
J0534+2200	59.33	-0.19	-0.40	-0.76	-0.88	-0.25	-0.55	-0.74	-0.47	-0.74	1.74	1.69	2.99
J0835-4510	22.37	-0.06	-0.15	-0.29	-0.37	-0.07	-0.21	-0.26	-0.19	-0.67	16.66	13.78	15.27
J1709-4429	19.51	-0.07	-0.02	-0.10	-0.07	0.01	-0.05	-0.04	-0.05	-0.60	61.33	36.31	34.34
J1952+3252	50.59	-0.24	-0.43	-0.81	-0.98	-0.23	-0.58	-0.77	-0.50	-0.80	2.00	1.53	1.49
J2229+6114	38.72	-0.24	-0.40	-0.62	-0.77	-0.19	-0.45	-0.61	-0.43	-0.81	3.36	2.70	2.92

Table B.2: Results obtained without using orientation information (see Chapter 9 for details).

	f_{GW} (Hz)	$\log \mathcal{B}_N^{s}$	$\log \mathcal{B}_N^{v}$	$\log \mathcal{B}_N^{t}$	$\log \mathcal{B}_N^{sv}$	$\log \mathcal{B}_N^{st}$	$\log \mathcal{B}_N^{vt}$	$\log \mathcal{B}_N^{svt}$	$\log O_N^S$	$\log O_I^S$	$h_s^{95\%}$	$h_v^{95\%}$ $\times 10^{-26}$	$h_t^{95\%}$
J0023+0923	655.69	-0.27	-0.57	-0.56	-0.76	-0.79	-1.05	-1.24	-0.65	-0.89	3.30	2.62	2.76
J0024-7204AA	1083.79	0.07	-0.27	-0.17	-0.21	-0.14	-0.47	-0.46	-0.20	-0.55	10.44	6.56	7.33
J0024-7204AB	539.86	-0.25	-0.55	-0.49	-0.72	-0.70	-0.93	-1.08	-0.60	-0.84	3.38	2.53	2.78
J0024-7204C	347.42	-0.29	-0.68	-0.62	-0.86	-0.92	-1.24	-1.49	-0.73	-0.91	2.46	1.62	1.86
J0024-7204D	373.30	-0.05	-0.32	-0.69	-0.28	-0.76	-0.89	-0.92	-0.45	-0.69	4.53	2.55	1.49
J0024-7204E	565.56	-0.24	-0.57	-0.64	-0.74	-0.92	-1.12	-1.38	-0.67	-0.87	3.91	2.45	1.95
J0024-7204F	762.32	-0.25	-0.58	-0.53	-0.78	-0.78	-1.07	-1.23	-0.64	-0.86	3.60	2.72	2.86
J0024-7204G	495.00	-0.28	-0.38	-0.61	-0.57	-0.88	-0.84	-1.16	-0.59	-0.81	3.21	2.86	2.12
J0024-7204H	622.99	-0.24	-0.47	-0.37	-0.66	-0.60	-0.79	-1.01	-0.53	-0.78	3.59	2.94	3.16
J0024-7204I	573.89	-0.26	-0.51	-0.42	-0.73	-0.65	-0.89	-1.11	-0.58	-0.80	3.38	2.50	2.68
J0024-7204J	952.09	-0.22	-0.36	-0.22	-0.51	-0.49	-0.54	-0.69	-0.40	-0.70	5.44	4.40	4.60
J0024-7204L	460.18	-0.27	-0.64	-0.65	-0.85	-0.93	-1.22	-1.45	-0.72	-0.90	2.85	1.90	1.86
J0024-7204M	543.97	-0.28	-0.66	-0.39	-0.84	-0.63	-0.93	-1.19	-0.61	-0.81	2.84	1.90	2.86
J0024-7204N	654.89	0.02	-0.43	-0.56	-0.35	-0.53	-0.93	-0.92	-0.42	-0.68	6.51	3.20	2.49
J0024-7204O	756.62	-0.06	-0.63	-0.41	-0.58	-0.49	-0.95	-0.98	-0.48	-0.79	6.77	2.34	3.37
J0024-7204Q	495.89	0.02	-0.56	-0.56	-0.45	-0.56	-1.05	-0.97	-0.46	-0.81	3.92	1.70	1.77
J0024-7204R	574.64	-0.24	-0.56	-0.54	-0.72	-0.78	-0.99	-1.26	-0.63	-0.84	3.64	2.50	2.54
J0024-7204S	706.61	-0.19	-0.59	-0.46	-0.71	-0.64	-0.97	-1.13	-0.58	-0.79	4.95	2.55	3.10
J0024-7204T	263.56	-0.29	-0.71	-0.60	-0.90	-0.87	-1.17	-1.44	-0.72	-0.90	2.19	1.60	1.84
J0024-7204U	460.53	-0.11	-0.69	-0.53	-0.72	-0.65	-1.11	-1.17	-0.58	-0.79	4.60	1.67	2.30
J0024-7204W	850.22	-0.14	-0.59	-0.52	-0.65	-0.69	-1.04	-1.18	-0.58	-0.83	5.99	2.63	3.13
J0024-7204X	419.15	-0.20	-0.64	0.24	-0.79	-0.00	-0.35	-0.56	-0.19	-0.47	3.57	1.76	2.97
J0024-7204Y	910.47	-0.03	-0.56	-0.44	-0.56	-0.51	-0.96	-0.98	-0.47	-0.70	7.78	2.90	3.63
J0024-7204Z	439.13	-0.27	-0.53	-0.62	-0.74	-0.89	-1.09	-1.33	-0.67	-0.85	2.81	2.24	2.00

Continued on next page

	$f_{\rm GW}$ (Hz)	$\log \mathcal{B}^{\rm s}_{\rm N}$	$\log \mathcal{B}^{\rm v}_{\rm N}$	$\log \mathcal{B}^{\rm t}_{\rm N}$	$\log \mathcal{B}^{\rm sv}_{\rm N}$	$\log \mathcal{B}^{\rm st}_{\rm N}$	$\log \mathcal{B}^{\rm vt}_{\rm N}$	$\log \mathcal{B}^{\rm svt}_{\rm N}$	$\log O^{\rm S}_{\rm N}$	$\log O^{\rm S}_{\rm I}$	$h_{\rm s}^{95\%}$	$h_{\rm v}^{95\%}$ $\times 10^{-26}$	$h_{\rm t}^{95\%}$
J0030+0451	411.06	-0.29	-0.67	-0.53	-0.84	-0.85	-1.08	-1.34	-0.69	-0.86	2.27	1.81	2.24
J0034-0534	1065.43	-0.23	-0.32	-0.49	-0.48	-0.70	-0.77	-0.96	-0.50	-0.93	4.54	4.62	3.93
J0102+4839	674.74	-0.25	-0.18	-0.63	-0.43	-0.88	-0.82	-0.98	-0.50	-0.73	3.27	3.65	2.08
J0205+6449	30.41	0.06	-0.09	-0.42	-0.10	-0.30	-0.52	-0.53	-0.22	-0.57	18.98	10.65	6.87
J0218+4232	860.92	-0.15	-0.53	-0.48	-0.62	-0.60	-0.94	-1.10	-0.54	-0.80	5.73	3.24	3.43
J0340+4130	606.18	-0.20	-0.57	-0.54	-0.69	-0.72	-0.97	-1.17	-0.60	-0.88	4.16	2.74	2.66
J0348+0432	51.12	-0.26	-0.61	-0.35	-0.79	-0.64	-0.89	-1.09	-0.58	-0.79	3.38	2.42	3.79
J0407+1607	77.82	-0.27	-0.56	-0.68	-0.76	-0.94	-1.15	-1.39	-0.69	-0.89	2.79	2.35	1.86
J0437-4715	347.38	-0.30	-0.53	-0.63	-0.81	-0.93	-1.11	-1.39	-0.69	-0.88	1.98	2.30	1.85
J0453+1559	43.69	-0.23	-0.52	-0.37	-0.70	-0.61	-0.80	-0.99	-0.54	-0.80	5.02	3.64	5.25
J0534+2200	59.33	-0.25	-0.55	-0.48	-0.74	-0.68	-0.91	-1.16	-0.60	-0.82	3.36	3.50	3.77
J0605+37	733.15	-0.26	-0.56	-0.56	-0.79	-0.82	-1.05	-1.25	-0.65	-0.85	3.23	2.84	2.79
J0609+2130	35.91	0.25	-0.32	-0.38	-0.03	-0.10	-0.67	-0.38	-0.14	-0.67	24.86	11.57	9.43
J0610-2100	517.96	-0.08	-0.40	-0.59	-0.44	-0.67	-0.92	-0.99	-0.49	-0.72	4.56	3.26	2.54
J0613-0200	653.20	-0.25	-0.47	-0.52	-0.64	-0.73	-0.97	-1.15	-0.59	-0.81	3.25	2.80	3.11
J0614-3329	635.19	-0.21	-0.58	-0.39	-0.73	-0.62	-0.88	-1.08	-0.56	-0.83	3.98	2.32	3.44
J0621+1002	69.31	-0.22	-0.33	-0.60	-0.48	-0.84	-0.82	-1.02	-0.54	-0.78	4.15	3.60	3.04
J0636+5129	697.12	-0.21	-0.57	-0.54	-0.70	-0.74	-1.04	-1.26	-0.61	-0.83	4.30	2.78	2.71
J0645+5158	225.90	-0.30	-0.71	-0.41	-0.98	-0.72	-1.01	-1.30	-0.66	-0.85	2.24	1.39	1.98
J0711-6830	364.23	-0.28	-0.47	-0.58	-0.70	-0.92	-1.05	-1.24	-0.64	-0.84	2.35	2.48	2.14
J0721-2038	128.68	-0.28	-0.64	-0.27	-0.84	-0.53	-0.85	-1.11	-0.56	-0.78	2.53	1.81	2.70
J0737-3039A	88.11	-0.27	-0.54	-0.58	-0.80	-0.84	-1.08	-1.27	-0.66	-0.85	2.92	2.32	2.35
J0742+66	693.06	-0.26	-0.45	-0.39	-0.68	-0.67	-0.83	-1.11	-0.55	-0.77	3.43	3.15	3.09
J0751+1807	574.92	-0.23	-0.57	-0.35	-0.72	-0.54	-0.83	-1.10	-0.54	-0.86	3.71	2.53	3.34
J0835-4510	22.37	-0.07	-0.21	-0.18	-0.26	-0.28	-0.39	-0.44	-0.25	-0.69	34.36	26.64	27.79
J0900-3144	180.02	-0.17	-0.36	-0.40	-0.51	-0.61	-0.71	-0.88	-0.47	-0.84	10.42	14.95	11.27
J0908-4913	18.73	-0.06	-0.12	-0.11	-0.17	-0.17	-0.23	-0.29	-0.16	-0.68	65.88	67.51	62.92
J0931-1902	431.22	-0.30	-0.64	-0.68	-0.85	-0.91	-1.23	-1.51	-0.74	-0.92	2.40	1.84	1.86
J0940-5428	22.84	0.07	-0.17	0.26	-0.07	0.26	0.04	0.11	0.10	-0.46	46.73	28.61	35.09
J1012+5307	380.54	-0.31	-0.64	-0.22	-0.84	-0.49	-0.78	-1.07	-0.53	-0.81	2.15	1.86	2.52
J1016-5819	22.77	-0.11	-0.21	-0.16	-0.30	-0.28	-0.39	-0.49	-0.26	-0.83	26.39	27.41	25.31
J1016-5857	18.62	-0.05	-0.12	-0.16	-0.16	-0.21	-0.29	-0.32	-0.18	-0.73	75.57	64.39	48.88
J1017-7156	855.24	-0.21	-0.32	-0.53	-0.48	-0.73	-0.76	-0.94	-0.50	-0.78	5.49	4.58	3.13
J1022+1001	121.56	-0.27	-0.61	-0.68	-0.78	-0.95	-1.16	-1.43	-0.71	-0.89	2.52	1.90	1.52
J1024-0719	387.43	-0.23	-0.70	-0.45	-0.87	-0.67	-1.07	-1.29	-0.63	-0.92	3.12	1.43	2.59
J1028-5819	21.88	-0.10	-0.25	-0.15	-0.32	-0.27	-0.42	-0.49	-0.27	-0.68	31.41	23.66	28.95
J1038+0032	69.32	-0.15	-0.59	-0.60	-0.68	-0.73	-1.11	-1.22	-0.60	-0.81	4.72	2.32	2.36
J1045-4509	267.59	-0.30	-0.41	-0.64	-0.66	-0.95	-1.00	-1.25	-0.64	-0.82	1.94	2.13	1.67
J1055-6028	20.07	-0.07	-0.18	-0.03	-0.23	-0.12	-0.24	-0.27	-0.15	-0.68	50.87	39.90	46.08
J1105-6107	31.65	-0.17	-0.15	-0.42	-0.31	-0.61	-0.55	-0.71	-0.37	-0.71	11.59	10.69	6.60
J1112-6103	30.78	-0.14	-0.32	-0.16	-0.43	-0.29	-0.47	-0.53	-0.31	-0.86	13.70	8.79	10.14
J1122+78	476.01	-0.23	-0.12	-0.63	-0.29	-0.86	-0.67	-0.90	-0.43	-0.66	3.85	3.19	1.89
J1125-6014	760.35	-0.26	-0.50	-0.52	-0.67	-0.77	-0.93	-1.15	-0.61	-0.81	3.15	3.52	3.05
J1142+0119	394.07	-0.30	-0.46	-0.54	-0.73	-0.81	-0.93	-1.20	-0.62	-0.92	2.28	2.24	2.40
J1231-1411	542.91	-0.29	-0.52	-0.52	-0.76	-0.82	-0.98	-1.25	-0.64	-0.86	2.55	2.54	2.72
J1300+1240	321.62	-0.27	-0.53	-0.64	-0.73	-0.94	-1.06	-1.29	-0.67	-0.98	2.40	2.40	1.87
J1302-3258	530.38	-0.25	-0.46	-0.68	-0.71	-0.88	-1.00	-1.31	-0.64	-0.90	2.98	2.77	1.81
J1302-6350	41.87	-0.21	-0.45	-0.40	-0.58	-0.59	-0.76	-0.99	-0.51	-0.78	5.75	4.40	4.52
J1312+0051	473.03	-0.21	-0.65	-0.59	-0.82	-0.80	-1.14	-1.31	-0.66	-0.91	3.62	1.76	2.17
J1327-0755	746.85	-0.22	-0.58	-0.50	-0.75	-0.72	-1.01	-1.23	-0.61	-0.83	4.17	2.44	3.26

Continued on next page

	f_{GW} (Hz)	$\log \mathcal{B}^{s}_{N}$	$\log \mathcal{B}^{v}_{N}$	$\log \mathcal{B}^{t}_{N}$	$\log \mathcal{B}^{sv}_{N}$	$\log \mathcal{B}^{st}_{N}$	$\log \mathcal{B}^{vt}_{N}$	$\log \mathcal{B}^{svt}_{N}$	$\log O^{S}_{N}$	$\log O^{S}_{I}$	$h_{s}^{95\%}$	$h_{v}^{95\%}$ $\times 10^{-26}$	$h_{t}^{95\%}$
J1410-6132	39.96	-0.08	-0.34	-0.29	-0.37	-0.38	-0.61	-0.67	-0.35	-1.05	10.13	6.10	5.12
J1418-6058	18.08	0.08	-0.10	0.00	-0.01	0.11	-0.10	-0.04	-0.00	-0.62	96.91	76.92	97.41
J1446-4701	911.29	-0.24	-0.41	-0.46	-0.60	-0.71	-0.82	-1.01	-0.54	-0.88	4.45	4.29	3.47
J1453+1902	345.29	-0.20	-0.45	-0.43	-0.61	-0.66	-0.76	-0.98	-0.52	-0.74	3.33	2.63	2.53
J1455-3330	250.40	-0.17	-0.69	-0.66	-0.82	-0.77	-1.18	-1.46	-0.66	-0.85	2.99	1.53	1.77
J1509-5850	22.49	-0.02	-0.23	-0.26	-0.24	-0.30	-0.48	-0.48	-0.26	-0.72	53.44	24.53	21.80
J1518+4904	48.86	-0.25	-0.46	-0.54	-0.66	-0.81	-0.91	-1.14	-0.60	-0.86	3.99	4.25	3.09
J1524-5625	25.57	-0.06	-0.22	-0.28	-0.27	-0.37	-0.53	-0.55	-0.30	-0.69	32.75	19.68	15.48
J1531-5610	23.75	-0.11	-0.26	-0.30	-0.33	-0.42	-0.54	-0.63	-0.34	-0.73	27.14	22.35	16.68
J1537+1155	52.76	-0.25	-0.55	-0.36	-0.74	-0.60	-0.86	-1.08	-0.56	-0.82	3.13	2.53	4.25
J1545-4550	559.40	-0.21	-0.55	-0.59	-0.72	-0.78	-1.05	-1.21	-0.63	-0.83	3.78	2.59	2.23
J1551-0658	281.94	-0.28	0.14	-0.49	-0.06	-0.80	-0.18	-0.52	-0.22	-0.51	2.35	2.91	2.02
J1600-3053	555.88	0.11	-0.55	-0.61	-0.43	-0.53	-1.11	-1.02	-0.42	-0.97	5.29	2.63	2.38
J1603-7202	134.75	-0.06	-0.73	-0.62	-0.71	-0.68	-1.30	-1.33	-0.59	-0.83	4.09	1.49	1.64
J1614-2230	634.76	-0.27	-0.57	-0.31	-0.78	-0.55	-0.83	-1.01	-0.55	-0.79	3.07	2.70	3.65
J1618-3921	166.84	-0.30	-0.69	-0.61	-0.91	-0.90	-1.23	-1.51	-0.73	-0.91	2.10	1.48	1.59
J1623-2631	180.57	-0.30	-0.68	-0.52	-0.94	-0.81	-1.06	-1.40	-0.69	-0.90	2.05	1.73	2.01
J1630+37	602.75	0.33	-0.61	-0.58	-0.23	-0.26	-1.12	-0.74	-0.24	-0.49	5.82	2.36	2.66
J1640+2224	632.25	-0.23	-0.54	0.37	-0.69	0.13	-0.23	-0.42	-0.09	-0.36	3.64	2.64	4.33
J1643-1224	432.75	-0.25	-0.48	-0.57	-0.63	-0.81	-0.92	-1.18	-0.61	-0.80	2.98	2.57	2.22
J1653-2054	484.36	-0.03	-0.57	-0.60	-0.53	-0.64	-1.14	-1.09	-0.51	-0.77	4.46	2.22	1.98
J1708-3506	443.94	-0.28	-0.48	-0.59	-0.70	-0.83	-0.94	-1.21	-0.63	-0.84	2.40	2.65	2.02
J1709+2313	431.85	-0.26	-0.62	-0.51	-0.80	-0.77	-1.06	-1.29	-0.65	-0.85	2.69	1.94	2.45
J1709-4429	19.51	0.01	-0.05	-0.15	-0.04	-0.12	-0.22	-0.20	-0.10	-0.61	90.99	66.52	49.59
J1710+49	621.07	-0.23	-0.61	-0.58	-0.78	-0.83	-1.15	-1.33	-0.66	-0.87	3.73	2.19	2.43
J1713+0747	437.62	-0.28	-0.64	-0.09	-0.83	-0.31	-0.66	-0.82	-0.43	-0.71	3.11	1.86	3.48
J1718-3825	26.78	-0.12	-0.03	-0.32	-0.15	-0.42	-0.34	-0.44	-0.23	-0.75	18.84	22.36	14.69
J1719-1438	345.41	-0.25	-0.65	-0.57	-0.85	-0.78	-1.11	-1.33	-0.67	-0.86	2.96	1.74	2.17
J1721-2457	571.98	-0.27	-0.49	-0.54	-0.70	-0.81	-0.97	-1.20	-0.62	-0.88	2.67	2.83	2.39
J1727-2946	73.85	-0.14	-0.51	-0.60	-0.56	-0.72	-0.94	-1.11	-0.55	-0.83	4.64	3.24	2.52
J1729-2117	30.17	-0.17	0.03	-0.38	-0.12	-0.54	-0.32	-0.51	-0.24	-0.67	10.28	12.97	8.17
J1730-2304	246.22	-0.31	-0.66	-0.70	-0.88	-1.04	-1.23	-1.61	-0.76	-0.95	1.88	1.65	1.45
J1731-1847	853.04	-0.15	-0.48	-0.33	-0.61	-0.51	-0.81	-0.89	-0.47	-0.75	5.78	3.37	4.14
J1732-5049	376.47	-0.18	-0.66	-0.70	-0.81	-0.86	-1.26	-1.45	-0.68	-0.99	3.48	1.67	1.50
J1738+0333	341.87	-0.12	-0.52	-0.64	-0.60	-0.74	-1.00	-1.13	-0.57	-0.86	3.81	2.20	2.06
J1741+1351	533.74	-0.29	-0.69	-0.05	-0.88	-0.37	-0.63	-0.91	-0.44	-0.67	2.54	1.74	3.61
J1744-1134	490.85	-0.19	-0.36	-0.35	-0.52	-0.57	-0.70	-0.85	-0.46	-0.87	3.75	2.72	3.05
J1745+1017	754.11	-0.27	-0.59	-0.57	-0.84	-0.85	-1.14	-1.34	-0.68	-0.91	3.08	2.22	2.64
J1745-0952	103.22	-0.30	-0.48	-0.65	-0.74	-0.92	-1.01	-1.28	-0.66	-0.89	2.27	2.23	1.88
J1748-2446A	172.96	-0.29	-0.66	-0.63	-0.89	-0.85	-1.18	-1.46	-0.72	-0.91	2.21	1.56	1.67
J1748-3009	206.53	-0.28	-0.09	-0.61	-0.42	-0.90	-0.69	-0.91	-0.46	-0.79	2.28	2.47	1.62
J1750-2536	57.55	-0.21	-0.53	-0.47	-0.69	-0.69	-0.97	-1.16	-0.58	-0.82	4.53	2.95	3.17
J1751-2857	510.87	-0.22	-0.54	-0.37	-0.72	-0.60	-0.84	-1.08	-0.55	-0.82	4.08	4.02	4.58
J1753-1914	31.77	0.14	-0.33	-0.24	-0.22	-0.18	-0.55	-0.48	-0.21	-0.56	15.43	9.33	9.74
J1753-2240	21.02	0.08	-0.20	-0.14	-0.09	-0.07	-0.36	-0.25	-0.13	-0.64	63.40	31.89	36.62
J1756-2251	70.27	-0.19	-0.47	-0.49	-0.57	-0.71	-0.89	-1.09	-0.55	-0.81	4.24	3.52	3.05
J1757-27	113.08	0.02	-0.63	-0.40	-0.55	-0.41	-1.03	-0.92	-0.43	-0.69	3.80	2.15	2.65
J1801-1417	551.71	-0.18	-0.62	-0.41	-0.77	-0.53	-0.99	-1.04	-0.55	-0.84	4.42	2.17	3.26
J1801-3210	268.33	-0.24	-0.72	-0.51	-0.87	-0.73	-1.16	-1.35	-0.67	-0.86	2.66	1.45	2.01
J1802-2124	158.13	-0.31	-0.53	-0.72	-0.77	-0.99	-1.19	-1.39	-0.71	-0.90	1.96	1.85	1.41

Continued on next page

	$f_{\rm GW}$ (Hz)	$\log \mathcal{B}^{\rm s}_{\rm N}$	$\log \mathcal{B}^{\rm v}_{\rm N}$	$\log \mathcal{B}^{\rm t}_{\rm N}$	$\log \mathcal{B}^{\rm sv}_{\rm N}$	$\log \mathcal{B}^{\rm st}_{\rm N}$	$\log \mathcal{B}^{\rm vt}_{\rm N}$	$\log \mathcal{B}^{\rm svt}_{\rm N}$	$\log O^{\rm S}_{\rm N}$	$\log O^{\rm S}_{\rm I}$	$h_{\rm s}^{95\%}$	$h_{\rm v}^{95\%}$ $\times 10^{-26}$	$h_{\rm t}^{95\%}$
J1804-0735	86.58	-0.26	-0.63	-0.56	-0.85	-0.82	-1.10	-1.32	-0.67	-0.97	3.00	1.83	2.46
J1804-2717	214.06	-0.31	-0.44	-0.15	-0.65	-0.42	-0.56	-0.81	-0.43	-1.17	2.02	1.97	2.31
J1809-1917	24.17	-0.07	-0.24	-0.24	-0.29	-0.33	-0.47	-0.56	-0.29	-0.68	32.86	20.26	19.52
J1810+1744	1202.82	-0.20	-0.35	-0.52	-0.51	-0.72	-0.81	-1.02	-0.52	-0.79	6.01	4.86	4.04
J1811-2405	751.71	-0.21	-0.43	-0.55	-0.62	-0.81	-0.98	-1.19	-0.58	-0.78	4.30	3.37	2.61
J1813-1246	41.60	-0.19	-0.42	-0.55	-0.55	-0.71	-0.83	-1.02	-0.54	-0.80	5.52	4.90	4.02
J1813-2621	451.47	-0.25	-0.69	-0.59	-0.85	-0.88	-1.17	-1.42	-0.70	-0.91	2.83	1.85	2.24
J1823-3021A	367.65	-0.28	-0.57	-0.54	-0.81	-0.83	-1.09	-1.29	-0.66	-0.90	2.26	2.14	2.18
J1824-2452A	654.81	-0.25	-0.47	-0.34	-0.62	-0.56	-0.65	-0.87	-0.50	-0.72	3.76	3.33	3.38
J1825-0319	439.22	-0.24	-0.44	-0.58	-0.62	-0.80	-1.01	-1.17	-0.60	-0.82	3.12	2.53	2.04
J1826-1256	18.14	-0.03	-0.10	-0.01	-0.12	-0.02	-0.09	-0.14	-0.07	-0.64	87.34	83.25	100.73
J1826-1334	19.71	-0.06	-0.10	-0.02	-0.16	-0.10	-0.13	-0.19	-0.11	-0.64	65.37	59.45	85.41
J1828-1101	27.76	-0.11	-0.06	-0.12	-0.17	-0.21	-0.19	-0.27	-0.16	-0.58	19.63	18.77	18.38
J1832-0836	735.53	-0.20	-0.55	-0.55	-0.72	-0.78	-1.02	-1.20	-0.61	-0.83	4.28	2.76	2.80
J1833-0827	23.45	-0.08	-0.28	-0.20	-0.36	-0.29	-0.49	-0.54	-0.30	-0.70	29.20	17.97	27.18
J1837-0604	20.77	-0.08	-0.25	-0.14	-0.28	-0.22	-0.38	-0.45	-0.24	-0.68	49.15	25.10	39.34
J1840-0643	56.21	-0.26	-0.45	-0.58	-0.66	-0.82	-0.94	-1.18	-0.61	-0.83	3.55	3.11	3.02
J1843-1113	1083.62	-0.22	-0.44	-0.52	-0.61	-0.71	-0.90	-1.11	-0.57	-0.86	5.49	4.94	3.84
J1845-0743	19.10	-0.06	-0.14	-0.14	-0.19	-0.21	-0.27	-0.33	-0.18	-0.74	68.50	49.13	57.96
J1853+1303	488.78	-0.29	-0.58	-0.49	-0.79	-0.79	-0.99	-1.26	-0.65	-0.86	2.70	2.19	2.84
J1853-0004	19.72	-0.07	0.00	-0.12	-0.08	-0.19	-0.14	-0.21	-0.11	-0.65	58.67	57.30	64.98
J1856+0245	24.72	-0.09	-0.27	-0.23	-0.33	-0.30	-0.51	-0.57	-0.30	-0.69	33.13	16.98	24.27
J1857+0943	372.99	-0.26	-0.47	-0.09	-0.66	-0.38	-0.51	-0.75	-0.39	-0.76	2.70	2.29	3.27
J1903+0327	930.27	-0.17	-0.39	-0.52	-0.51	-0.69	-0.88	-0.97	-0.52	-0.75	5.41	3.77	3.33
J1903-7051	555.88	-0.23	-0.60	-0.38	-0.78	-0.63	-0.94	-1.15	-0.58	-0.78	3.89	2.17	2.89
J1909-3744	678.63	-0.25	-0.56	-0.61	-0.74	-0.83	-1.10	-1.35	-0.66	-0.85	3.68	2.77	2.24
J1910+1256	401.32	-0.28	-0.49	-0.52	-0.67	-0.80	-0.92	-1.16	-0.61	-0.87	2.60	2.34	2.45
J1910-5959A	612.33	-0.17	-0.57	-0.61	-0.73	-0.81	-1.10	-1.25	-0.62	-0.87	4.80	2.44	2.22
J1910-5959C	378.98	-0.04	-0.56	-0.62	-0.54	-0.63	-1.05	-1.15	-0.52	-0.76	4.30	2.12	1.90
J1910-5959D	221.35	-0.26	-0.54	-0.68	-0.76	-0.96	-1.12	-1.38	-0.68	-0.89	2.61	1.93	1.48
J1911+1347	432.34	-0.26	-0.60	-0.49	-0.77	-0.74	-0.99	-1.21	-0.63	-0.84	2.73	2.20	2.33
J1911-1114	551.61	-0.23	-0.25	-0.59	-0.39	-0.78	-0.75	-0.97	-0.49	-0.75	3.83	3.61	2.60
J1915+1606	33.88	-0.17	-0.34	-0.38	-0.46	-0.49	-0.65	-0.82	-0.43	-0.76	10.05	6.78	7.07
J1918-0642	261.58	-0.30	-0.54	-0.73	-0.77	-0.99	-1.21	-1.51	-0.72	-0.89	1.93	1.97	1.35
J1923+2515	527.96	-0.29	-0.66	-0.63	-0.87	-0.92	-1.21	-1.47	-0.72	-0.92	2.52	1.86	2.03
J1925+1721	26.43	-0.13	-0.26	-0.18	-0.36	-0.35	-0.42	-0.54	-0.30	-0.67	19.36	18.54	22.34
J1928+1746	29.10	-0.15	-0.37	-0.30	-0.49	-0.46	-0.61	-0.78	-0.41	-0.73	10.80	9.07	11.19
J1932+17	47.81	-0.20	-0.52	0.78	-0.71	0.68	0.30	0.25	0.35	0.01	6.01	3.06	6.80
J1935+2025	24.96	-0.09	-0.29	-0.27	-0.33	-0.34	-0.54	-0.61	-0.32	-0.74	30.68	16.69	21.37
J1939+2134	1283.86	-0.23	-0.44	-0.49	-0.60	-0.73	-0.89	-1.10	-0.56	-0.83	5.17	4.63	4.37
J1943+2210	393.38	-0.29	-0.66	-0.57	-0.90	-0.87	-1.20	-1.41	-0.71	-0.90	2.38	1.64	1.95
J1944+0907	385.71	-0.25	-0.60	-0.53	-0.80	-0.77	-0.97	-1.28	-0.64	-0.82	2.74	1.94	2.26
J1946+3417	630.89	-0.05	-0.58	-0.55	-0.57	-0.61	-1.04	-1.12	-0.52	-0.76	5.13	2.53	2.68
J1949+3106	152.23	-0.32	-0.68	-0.66	-0.90	-0.97	-1.19	-1.53	-0.75	-0.92	1.76	1.52	1.67
J1950+2414	464.60	-0.21	-0.34	-0.49	-0.48	-0.68	-0.73	-0.97	-0.50	-0.72	3.62	2.87	2.55
J1952+3252	50.59	-0.23	-0.58	-0.52	-0.77	-0.77	-1.07	-1.24	-0.63	-0.89	4.07	2.75	3.03
J1955+2527	410.44	-0.30	-0.62	-0.38	-0.83	-0.69	-0.94	-1.17	-0.62	-0.87	2.23	2.24	2.88
J1955+2908	326.10	-0.28	-0.26	-0.68	-0.51	-0.92	-0.91	-1.16	-0.56	-0.77	2.47	2.37	1.62
J1959+2048	1244.24	-0.22	-0.48	-0.47	-0.63	-0.70	-0.91	-1.03	-0.56	-0.82	5.53	4.51	4.34
J2007+2722	81.64	-0.30	-0.26	-0.46	-0.48	-0.75	-0.55	-0.85	-0.48	-0.78	2.56	3.77	3.61

Continued on next page

	$f_{\rm GW}$ (Hz)	$\log \mathcal{B}^{\rm s}_{\rm N}$	$\log \mathcal{B}^{\rm v}_{\rm N}$	$\log \mathcal{B}^{\rm t}_{\rm N}$	$\log \mathcal{B}^{\rm sv}_{\rm N}$	$\log \mathcal{B}^{\rm st}_{\rm N}$	$\log \mathcal{B}^{\rm vt}_{\rm N}$	$\log \mathcal{B}^{\rm svt}_{\rm N}$	$\log O^{\rm S}_{\rm N}$	$\log O^{\rm S}_{\rm I}$	$h_{\rm s}^{95\%}$	$h_{\rm v}^{95\%}$ $\times 10^{-26}$	$h_{\rm t}^{95\%}$
J2010-1323	382.90	-0.22	-0.58	-0.54	-0.76	-0.76	-1.05	-1.31	-0.63	-0.82	2.91	1.94	2.24
J2017+0603	690.56	-0.03	-0.44	-0.43	-0.39	-0.40	-0.80	-0.76	-0.39	-1.08	5.57	3.34	3.39
J2019+2425	508.32	-0.22	-0.45	-0.42	-0.60	-0.63	-0.82	-1.02	-0.53	-0.82	5.92	6.74	6.83
J2033+1734	336.19	-0.29	-0.57	-0.52	-0.79	-0.81	-1.08	-1.27	-0.66	-0.87	2.68	1.99	2.20
J2043+1711	840.38	-0.26	-0.08	-0.58	-0.32	-0.83	-0.60	-0.86	-0.42	-0.68	3.59	4.53	2.64
J2043+2740	20.81	-0.05	-0.07	-0.13	-0.13	-0.18	-0.19	-0.25	-0.14	-0.65	46.55	47.49	42.33
J2047+1053	466.64	-0.27	-0.54	-0.65	-0.75	-0.97	-1.15	-1.31	-0.68	-0.89	2.57	2.34	1.70
J2051-0827	443.59	-0.29	-0.60	-0.69	-0.79	-0.93	-1.19	-1.42	-0.71	-0.90	2.46	2.07	1.76
J2124-3358	405.59	-0.24	-0.53	-0.60	-0.69	-0.84	-1.04	-1.27	-0.64	-0.86	3.05	2.27	1.90
J2129+1210A	18.07	-0.00	-0.12	-0.03	-0.12	-0.05	-0.13	-0.15	-0.08	-0.63	91.59	73.71	98.25
J2129+1210B	35.63	-0.18	-0.42	-0.37	-0.56	-0.56	-0.74	-0.94	-0.48	-0.76	9.66	6.67	7.57
J2129+1210C	65.51	-0.25	-0.57	-0.58	-0.75	-0.76	-1.08	-1.25	-0.65	-0.86	3.91	2.52	3.07
J2129+1210D	416.42	-0.22	-0.67	-0.64	-0.80	-0.84	-1.21	-1.34	-0.68	-0.87	3.39	1.72	2.09
J2129+1210E	429.97	-0.28	-0.43	-0.49	-0.62	-0.68	-0.83	-1.09	-0.56	-0.81	2.41	2.43	2.61
J2129-5721	536.72	-0.23	-0.53	-0.55	-0.72	-0.81	-1.02	-1.19	-0.62	-0.81	3.62	2.48	2.42
J2145-0750	124.59	-0.30	-0.53	-0.68	-0.76	-0.98	-1.15	-1.42	-0.70	-0.88	2.14	1.99	1.85
J2214+3000	641.18	-0.27	-0.48	-0.14	-0.70	-0.42	-0.64	-0.77	-0.43	-0.69	3.07	2.85	3.91
J2222-0137	60.94	-0.19	-0.51	-0.61	-0.62	-0.77	-1.06	-1.22	-0.60	-0.81	5.16	2.72	2.72
J2229+2643	671.63	-0.27	-0.41	-0.51	-0.66	-0.78	-0.88	-1.10	-0.58	-1.01	2.87	3.12	2.96
J2229+6114	38.72	-0.19	-0.45	-0.48	-0.61	-0.67	-0.90	-1.06	-0.54	-0.86	8.63	6.46	4.41
J2234+06	559.19	-0.28	-0.60	-0.55	-0.81	-0.82	-1.08	-1.24	-0.67	-0.86	2.66	2.22	2.54
J2235+1506	33.46	-0.21	-0.47	-0.41	-0.60	-0.58	-0.84	-0.98	-0.52	-0.79	7.22	5.17	7.54
J2241-5236	914.62	-0.23	-0.54	-0.51	-0.73	-0.71	-0.95	-1.17	-0.60	-0.83	4.49	3.08	3.69
J2302+4442	385.18	0.21	-0.48	-0.72	-0.27	-0.52	-1.03	-0.90	-0.34	-1.20	4.40	2.40	1.57
J2317+1439	580.51	-0.23	-0.63	-0.43	-0.76	-0.71	-0.99	-1.19	-0.61	-0.85	3.59	2.03	2.91
J2322+2057	415.94	-0.28	-0.52	-0.65	-0.78	-0.86	-1.08	-1.36	-0.67	-0.86	2.45	2.24	2.03