



HUNTING FOR EXOTIC FEATURES IN COMPACT BINARY MERGERS WITH GRAVITATIONAL WAVES

ROBIN CHAN

STUDENT NUMBER: 02009110

Supervisor: Prof. Dr. Archisman Ghosh

Master's dissertation submitted in order to obtain the academic degree of Master of Science in Physics and Astronomy



Department of Physics and Astronomy
Chair: Prof. Dr. Jan Ryckebusch
Faculty of Sciences
Academic year 2024-2025

“Plurality is not to be posited without necessity.”
– William of Ockham

Acknowledgements

This thesis would not have been possible without the support of many people whom I would like to thank now and dedicate this work to. You have all contributed to this thesis, be it directly or indirectly, and made it possible for this work to be a labour of love.

First and foremost, I want to thank my promotor Prof. Dr. Archisman Ghosh for introducing me to the world of gravitational wave physics. Who knew that a small literature study would turn into an academic fascination? I want to thank him for suggesting this topic that plays into all of my interests in physics and for supporting me throughout the year by giving me the opportunity to attend PhD schools, for valuable feedback on my work and the friendly work environment he kept around Proeftuin. I also want to thank the rest of the Ghent Gravity Group for the group meetings and their help with some of the technical difficulties I encountered throughout the year. I want to extend this gratitude to the entire Elementary Particle Physics and Gravity group too as they provided a wonderful environment to work and socialise in. A special thank you goes out to the coffee machine in the break room/kitchen, I am convinced that you are the hardest working device on UGent premises and are deserving of a medal for your service.

I also want to thank my friends from the physics degree. Throughout our studies we have been through thick and thin and I could not wish for a better group of people to do that with. The other masters students working at Campus Proeftuin deserve a special mention too. We have seen each other at the highs and lows of the thesis work and despite the zillionth code error or miscalculation, we managed to keep our spirits up together.

The different student societies I have been a part of throughout the years are also much deserving of a thank you. It started with *WiNA* which introduced me to the student life in Ghent. I don't think I would have made it through my first year (which was a COVID year) without their activities to get to know fellow students when we were unable to go to campus. Next, the *VVN* has played such a large role throughout my studies and I want to thank the wonderful people there for all of the activities we organised and their friendship. Our silly in-jokes and terrible puns never fail to put a smile on my face. The singular pun in this thesis is dedicated to all of you. Last year I decided to join the university's student newspaper *Schamper*. Being active there has given me the opportunity to dive into non-physics-related subjects for fun. I want to thank everyone there for their fantastic friendships and hilarious evenings we spent at the *redactie* or *de Vooruit*¹.

I also want to thank the public observatories of Bruges and Ghent, *Cozmix* and *Armand Pien* respectively. I have had the opportunity to work/volunteer there for the past few years and they have enabled me to discover and maintain my passion for science communication. Also, who doesn't want to play around with big telescopes all evening?

I want to thank my family for their love throughout my entire life. I want to thank my parents in particular for their relentless support and for putting up with my (sometimes) incoherent physics ramblings. I know I don't say it enough, but I am incredibly lucky to have parents like you.

Lastly, I want to thank Zita-Luna for inspiring me to try and be the best version of myself every day.

¹Now Kunstencentrum VIERNULVIER, but this is an objectively worse name

Use of AI Tools

AI tools were used throughout the work for this thesis in several contexts. The main use was debugging code with ChatGPT. ChatGPT and chat100 were also used to convert images to $\text{\LaTeX}$, which saved a lot of time while writing equations for the final appendix. Finally, ChatGPT was also used sparsely to explore literature. Every reference provided by it was double-checked and read thoroughly. In the final stages of writing, ChatGPT was also used to analyse some paragraphs and make suggestions to improve the writing.

No real-life data was shared with any AI tool at any point during the work for this thesis.

Admission to loan

“The author(s) gives (give) permission to make this master dissertation available for consultation and to copy parts of this master’s dissertation for personal use. In the case of any other use, the copyright terms have to be respected, in particular with regard to the obligation to state expressly the source when quoting results from this master dissertation.”

Toegang tot bruikleen

“De auteur(s) geeft (geven) de toelating dit afstudeerwerk voor consultatie beschikbaar te stellen en delen van het afstudeerwerk te kopiëren voor persoonlijk gebruik. Elk ander gebruik valt onder de beperkingen van het auteursrecht, in het bijzonder met betrekking tot de verplichting de bron uitdrukkelijk te vermelden bij het aanhalen van resultaten uit dit afstudeerwerk.”

Robin Chan, June 2025

Hunting for Exotic Features in Compact Binary Mergers with Gravitational Waves

by
ROBIN CHAN

Master's dissertation submitted in order to obtain the academic degree of Master of Science in Physics and Astronomy

Supervisor: Prof. Dr. Archisman Ghosh

Faculty of Sciences
Ghent University

Department of Physics and Astronomy
Chair: Prof. Dr. Jan Rychkebusch

Abstract - The first detection of gravitational waves in September 2015 opened up a new window through which we can observe the universe. These waves carry information about the dynamics of mass-distributions in the universe. The signals we see with current LVK-observatories come from the inspiral of binary systems of compact objects. The analysis techniques for these signals have, so far, been separated in two categories: modelled and unmodelled.

This work serves as a proof of concept of combining modelled and unmodelled techniques for gravitational wave data-analysis. Such a combination could allow us to search for unmodelled *exotic* features, such as tidal resonances, in gravitational wave signals while simultaneously estimating the parameters of the binary system.

This thesis employs a recent transdimensional extension to the Bayesian inference package **Bilby**, named **tbilby**, to incorporate unmodelled features on top of modelled **Bilby** waveforms. Model-agnostic features are modelled by a superposition of N sine-Gaussian wavelets and added to a modelled waveform provided by **Bilby**. This joint inference method is then evaluated on a variety of scenarios to assess its feasibility. These are: random and hand-tuned modelled waveform + wavelet injections, a mock f -mode resonance and joint glitch + CBC parameter estimation.

Finally, suggestions for modifications and extensions are outlined in the final chapter.

Keywords - Gravitational waves, Bayesian inference, transdimensional, exotic compact objects, tidal resonances

Zoeken naar exotische fenomenen met zwaartekrachtgolven

door
ROBIN CHAN

Masterproef ingediend tot het behalen van de academische graad van Master of Science
in de Fysica en Sterrenkunde

Promotor: Prof. Dr. Archisman Ghosh

Faculteit Wetenschappen
Universiteit Gent

Vakgroep Fysica en Sterrenkunde
Voorzitter: Prof. Dr. Jan Ryckebusch

Samenvatting - De eerste detectie van zwaartekrachtgolven in september 2015 opende een nieuw venster waardoor we het heelal kunnen observeren. Deze golven dragen informatie over de dynamica van massaverdelingen in het heelal. De signalen die we zien met de huidige LVK-observatoria zijn afkomstig van de botsingen van compacte objecten in binaire systemen. De analysetechnieken van zulke signalen zijn tot nu toe verdeeld geweest in twee categorieën: gemodelleerd en niet-gemodelleerd.

Dit werk dient als een *proof of concept* voor het combineren van gemodelleerde en ongemodelleerde technieken voor de analyse van zwaartekrachtgolven. Zo'n combinatie zou ons in staat kunnen stellen om te zoeken naar ongemodelleerde exotische fenomenen, zoals getijdenresonanties, in zwaartekrachtgolfsignalen en tegelijkertijd de parameters van het binaire systeem te achterhalen.

Deze thesis maakt gebruik van een recente transdimensionale uitbreiding op het Bayesiaanse inferentiepakket *Bilby*, genaamd *tbilby*, om ongemodelleerde kenmerken toe te voegen aan gemodelleerde *Bilby*-golfvormen. Modelagnostische kenmerken worden gemodelleerd door een superpositie van N wavelets bovenop een gemodelleerde golfvorm geleverd door *Bilby*. Deze gezamenlijke inferentiemethode wordt vervolgens geëvalueerd op verschillende scenario's om de haalbaarheid ervan in te schatten. Deze scenario's zijn: willekeurige en handgemaakte gemodelleerde golfvormen + wavelet injecties, een namaak f -mode resonantie en gezamenlijke glitch + CBC parameterschatting. Tot slot worden in het laatste hoofdstuk suggesties gedaan voor wijzigingen en uitbreidingen.

Kernwoorden - Zwaartekrachtgolven, Bayesiaanse inferentie, exotische compacte objecten, getijdenresonanties

Table of Contents

Acknowledgements	v
Use of AI Tools	vii
Admission to loan	ix
Abstract	xi
Samenvatting	xiii
Table of Contents	xv
List of Figures	xvii
List of Tables	xix
List of Acronyms	xxi
1 Introduction	1
1.1 100 Years of Gravitational Wave Science	1
1.2 Gravitational Wave Sources	2
1.3 Exotic Compact Objects	3
1.4 Asteroseismology	5
1.4.1 Introduction	5
1.4.2 Dynamical tides	6
1.5 Tidal Resonances	7
1.5.1 Neutron Star Equation of State	8
2 Methods	11
2.1 Bayesian Inference	11
2.2 Sampling algorithms	11
2.2.1 Markov Chain Monte Carlo	12
2.2.2 Nested Sampling	12
2.3 Evidence and Occam’s Razor	15
2.4 Bilby	16
2.4.1 Corner Plots	16
2.4.2 Signal to Noise Ratio	16
2.4.3 Resonance Model	18
2.5 tBilby	19
2.5.1 Order Statistics	19
2.6 Inference pipeline	20
2.7 Quality of Reconstruction	20
2.7.1 Residuals	21
2.7.2 Bayes Factor	21
2.7.3 Significance σ	21
2.7.4 (Mis)match	21
2.8 Whitening Data	21

3	Testing (t)Bilby	23
3.1	Addition of Wavelets to Bilby	23
3.2	tBilby Tests	26
3.2.1	Bug in tBilby Source Code	26
3.2.2	Wavelet Superpositions	26
3.3	Discussion	28
4	Results	29
4.1	Double Wavelets with LIGO	29
4.2	ET Injections	31
4.2.1	Triple Wavelets	31
4.2.2	Wavelet in Merger	32
4.3	Random Injections	32
4.4	Mock f -mode Resonance	38
4.5	Glitch Identification	39
4.6	Remark on Sampling Priors	40
5	Conclusion and Discussion	43
5.1	Conclusion	43
5.2	Discussion	44
6	Future Work	45
6.1	Physics	45
6.2	Inference	45
A	Bilby Structure	51
A.1	core package	51
A.2	gw package	51
A.3	hyper package	52
B	Corner plots	53
C	Gravitational Backreaction	63
C.1	General Solution of Linearised Einstein Equations	63
C.2	GW Energy-Momentum Tensor	65
C.2.1	Energy Flux	67
C.3	Backreaction	67

List of Figures

1.1	Diagrams of a simple Michelson interferometer (left) (Brittanica, 2024) and the general setup used by state-of-the-art gravitational wave interferometers (right) (Davis et al., 2021). While the general structure of the Michelson interferometer is still recognisable, numerous improvements have been made to the design to achieve the sensitivity necessary to detect gravitational waves.	2
1.2	Measurements of the Hulse-Taylor binary periastron time (indicative of orbital period) shift compared to the GR prediction (solid line) (Weisberg et al., 1981). This provides indirect evidence for the existence of gravitational waves.	2
1.3	Currently observed sources of GWs: compact binary coalescences (CBCs) (left) and the stochastic GW background (right). The black arrows indicate what objects these detection methods observe	3
1.4	Fundamental and first two harmonic modes of string vibrations (top) compared to wind instrument vibrations (bottom). The string is fixed at both ends whereas one end of the wind instrument can move. For the latter the oscillations are in the density of the air inside the instrument.	5
1.5	Illustration of equilibrium (left) and dynamical (right) tides. Equilibrium tides deform the star due gravitational interactions. The left image in particular shows a quadrupolar deformation. The right image illustrates dynamical tides, which are due to the internal motions of the fluid. For more details on the equations shown, see (Suvorov et al., 2024)	6
1.6	Self-consistent mode amplitude evolutions of f - and $g_{1,2,3}$ -modes due to tidal forcing, together with the dephasing associated with the f - and g -modes. In green the frequency evolution of the binary is plotted, corresponding to the right-hand axis. The underlying simulation was carried out for a $1.35\text{-}1.35M_{\odot}$ binary NS merger. Figure taken from (Suvorov et al., 2024).	8
1.7	There are a wealth of EOSs compatible with current NS observations (Gonzalez et al., 2023).	9
2.1	Illustration of nested sampling. Darker orange regions signify a higher likelihood. If the sampler has found this peak, these darker regions should also correspond to a higher density of samples	13
2.2	Different sampling strategies in dynesty . Figure taken from Joshua Speagle’s presentation on nested sampling https://helper.ipam.ucla.edu/publications/gwaws3/gwaws3_17076.pdf	14
2.3	Illustration of univariate slice sampling. Figure taken from (Neal, 2003).	15
2.4	Samplers try to bring an initial distribution down to a smaller posterior distribution that is representative for the underlying likelihood	15
2.5	Bilbies are a type of marsupial native to Australia, where the Bilby codebase was developed (Rafferty, n.d.). They are currently endangered.	16
2.6	Results of a parameter inference for an eccentric binary black hole merger shown in a corner plot. It shows parameter distributions on the diagonal and parameter correlations off-diagonally. These do tend to “waste” a lot of space on a page, for which (Rowan and Roberts, 2025) have found an innovative solution	17
2.7	(a) The likelihood multimodality can be seen in the sampling results of the component masses in a CBC parameter inference. (b) An example of descending SNR priors. ρ_i is defined as $\text{SNR}_i/\text{SNR}_{\text{max}}$ (Tong et al., 2025)	19
2.8	Illustration of the composition of waveforms from a CBC of rotating BHs and a single wavelet. The signals are shown in frequency (left) and time (right) domain. The sharp peak at the start of the time domain data is due to wrapping of the time domain signal by LALSImulation , causing the ringdown of the IMR signal to occur at the start of the time series	22

3.1	Results of the single wavelet PEs. The left panel shows the relation between the log-Bayes factor and the average optimal detector SNR, which upper-bounds the SNR of the signal. The vertical line shows the SNR-cutoff. On the right a histogram of the log-Bayes factors is shown for $\text{SNR} > 8$. These indicate that Bilby is able to reliably find the implemented wavelets.	24
3.2	Sample corner plots of an $\text{SNR} \approx 6$ (left) and $\text{SNR} \approx 12$ (right) wavelet. The injection values are shown in orange. The phase shift ϕ is poorly constrained in both cases.	24
3.3	Top row: Histograms of the absolute relative deviations between the posterior and true parameter values. Middle row: Log of the absolute relative deviations plotted against average optimal SNR. The dashed line in red denotes $\text{SNR} = 8$. Bottom row: Same as top row, but with data filtered on $\text{SNR} > 8$	25
3.4		27
4.1	LIGO Hanford data and response to the reconstruction of an IMR + wavelet signal with 2 wavelets stacked on top of each other at $t = 0.5$ s to mimic a complex resonance. The blue and grey shaded regions represent the 90% CL of the total waveform and wavelet reconstruction, respectively.	30
4.2	Characteristic strain curves for current and future ground-based observatories. Figure generated from the web tool developed by (Moore et al., 2014).	31
4.3	Reconstruction of the injection of three wavelets at $t = 4, 5, 5.5$ s with simulated ET noise. Only the wavelet at $t = 5$ s was not found. This was a fairly quiet wavelet with an $\text{SNR} \approx 1.94$. Its exclusion is likely due to its minimal influence on the total signal.	32
4.4	Reconstruction of an $\text{SNR} \approx 45$ signal in ET with a single wavelet in the merger section of the waveform. The injected and reconstructed wavelets are shown as they may not be clear without indication in the merger section.	33
4.5	Overview of the injected waveforms. The indicated SNR is the average SNR of the injected wavelet(s) over the three detectors in ET. 4 waveforms were replaced by shorter duration signals due to the long inference time of the second signal. Note that wavelets with a large strain are not necessarily the loud wavelets.	35
4.6	Time domain reconstruction of waveform 5. The second wavelet was missed by the sampler. This was likely due to the use of the ordered SNR prior.	35
4.7		36
4.8	Reconstruction of the signals with wavelets near the apparent ET wavelet detection threshold of $\text{SNR} \approx 1.5$ as seen in the left panel of Fig. 4.7c.	37
4.9	Reconstruction of a mock f -mode resonance in a BNS merger for a 10s signal. The BNS was identified but the wavelet was missed despite its relatively large strain.	39
4.10	The observed GW170817 signal in the LIGO-Livingston detector containing a loud glitch.	39
4.11	Reconstruction of a BBH merger with a large glitch. Both the glitch and BBH were identified, despite the large difference in strain.	40
4.12	Comparison of the parameter posteriors of a BBH merger with and without the large glitch. The glitch and BBH parameters were inferred simultaneously. The glitch does not appear to influence the posteriors drastically.	41
A.1	Structure of the Bilby source code. Image taken from Bilby documentation: https://lscsoft.docs.ligo.org/bilby/index.html	51
B.1	Fig. 4.1a	53
B.2	Fig. 4.1b	54
B.3	Fig. 4.3	55
B.4	Fig. 4.4	56
B.5	Fig. 4.6	57
B.6	Fig. 4.8a	58
B.7	Fig. 4.8b	59
B.8	Fig. 4.9	60
B.9	Fig. 4.11	61

List of Tables

1.1	Typical values of the frequencies and the damping times of various families of modes for a polytropic star ($N = 1$) with $R = 8.86$ km and $M = 1.27M_{\odot}$ are given. A subscript 1 identifies a mode as the first of its family. w_{II} is the slowest damped <i>interface</i> mode. Values taken from (Kostas D. Kokkotas and Schmidt, 1999)	7
3.1	Injection windows for the wavelet parameters. The recovery priors are uniform with their bounds given by the injection windows.	23
3.2	Injection windows and priors for the wavelet superposition parameters. As there are possibly multiple wavelets injected, the injection amplitudes have been halved compared to Section 3.1. Priors have the same intervals as the injection intervals, except for the amplitude and N . Instead of amplitude priors, the SNR priors for the SNR ordering of <code>tBilby</code> are used.	26
4.1	Injection parameters and recovery priors of the BBH merger used in Sections 4.1 and 4.2. Parameter values marked with a “ $\sim$ ” indicate that these vary accross runs in these two sections.	29
4.2	Combined injection and prior values. The dashed line separates the IMR and wavelet parameters. $t_{\min}$ and $m_{\min}$ are calculated as follows: $t_{\min} = \max(0.2, t_{\max} - 6)$, $m_{\min} = \max(5, 15 - m_1)$	34
4.3	Injection parameters and recovery priors of the BNS merger with mock f -mode resonance.	38

List of Acronyms

GW	Gravitational wave
GR	General Relativity
LIGO	Laser Interferometer Gravitational wave Observatory
CBC	Compact binary coalescence
NS	Neutron star
KAGRA	KAmioka GRAvitational wave detector
LVK	LIGO-Virgo-KAGRA
CE	Cosmic Explorer
ET	Einstein Telescope
LISA	Laser Interferometer Space Antenna
QCD	Quantum chromodynamics
CP	Charge conjugation parity
EOS	Equation of state
TOV	Tolman-Oppenheimer-Volkov
MAP	Maximum a posteriori
SNR	Signal-to-noise ratio
IMR	Inspiral-merger-ringdown
PSD	Power spectral density
ASD	Amplitude spectral density
CL	Confidence level
BNS	Binary neutron star
BBH	Binary black hole
LHC	Large Hadron Collider
PE	Parameter estimation
MCMC	Markov chain monte carlo
PTMCMC	Parallel-tempered markov chain monte carlo
GMC	Galilean monte carlo
KL	Kullback-Leibler
BH	Black hole
DOF	Degree(s) of freedom
RJMC	Reversible-jump markov chain monte carlo

LVT	LIGO-Virgo Trigger
ISCO	Innermost stable circular orbit
LAL	LVK Algorithm Library
NR	Numerical relativity

Introduction

1.1 100 Years of Gravitational Wave Science

Every good gravitational wave (GW) study starts its introduction with Einstein's prediction of ripples in spacetime. In 1916, Einstein found wave solutions to a linearised form of the set of coupled differential equations that bear his name (Einstein, 1916). However, he remarked that they would be undetectable due to their small amplitude. This idea of waves in spacetime was met with much criticism as Einstein's Theory of General Relativity allows for much coordinate freedom. It is exactly this freedom that allows for a simple and elegant formulation of gravitational waves, but perhaps these waves themselves were an artifact of an improper choice of coordinates? Could they simply be "gauged away"? The detection of gravitational waves was not included in Einstein's classical tests of general relativity, perhaps because he deemed them undetectable. The three tests that he did posit for his theory were passed with flying colours.

Gravitational waves remained a topic of debate within the physics community until their direct detection in 2015. The first attempts to detect gravitational waves were made in the 1960s when Joseph Weber started experimenting with resonant bar detectors. These are large metal cylinders, carefully suspended to minimise dampening. The idea being that they would resonate when a gravitational wave passes by at the bar's resonant frequency, thereby amplifying a small, periodic displacement to a detectable oscillation of the bar. To illustrate this idea of resonance, you could imagine yourself sitting on a swing and swinging your legs back and forth. If you do this at the right frequency, the the swing's displacement will be much larger than that of your legs. Your legs here are acting as the gravitational wave, driving the swing (resonant bar) to move much more than the displacement you put in. I assume that, at this point in your life, you have already mastered the art of this playground apparatus. But if you have ever observed a small child on a swing, you may have noticed that they flail their legs back and forth in an attempt to copy the swinging motion their elders so effortlessly demonstrated. However, they only move slightly around the swing's low point. This is not too dissimilar to a gravitational wave passing by the resonant bar at the incorrect frequency. This makes resonant bars only work at the very specific frequencies that they are tuned to. Weber claimed several detections with them, but no one was able to reproduce his results and by now this technique has fallen out of favour.

Concurrently in the 1960s, researchers started experimenting with interferometers as detectors for gravitational waves. Such detectors are based on a 19th century piece of technology named the Michelson interferometer. Interestingly, it is this device that kickstarted the development of special, and subsequently general, relativity by disproving the existence of the aether, the medium that light was thought to propagate through during Michelson's time. The interferometric experiments from the sixties matured over decades and led to the eventual funding of the Laser Interferometer Gravitational wave Observatory (LIGO) in the 1990s. It consists of two observatories in Hanford, Washington and Livingston, Louisiana. Each is constructed out of two perpendicular, 4 km long arms. A laser beam is split, passed through both arms, reflected back at the end of each arm and recombined. Because light behaves like a wave, the two beams can interfere with each other, which allows such a setup to detect sub-wavelength variations in the interferometer arm length. A detailed description of these detectors is outside the scope of this thesis, but they are truly fascinating machines.

Before moving on to the direct detection of gravitational waves, this is perhaps the appropriate time to mention that gravitational waves had already been indirectly detected in 1981 (Weisberg et al., 1981), following the discovery of the first pulsar in a binary system by Hulse and Taylor in 1975 (Hulse and Taylor, 1975). It was observed that the orbital period of the binary decreased over time: a telltale sign of energy




Figure 1.1: Diagrams of a simple Michelson interferometer (left) (Brittanica, 2024) and the general setup used by state-of-the-art gravitational wave interferometers (right) (Davis et al., 2021). While the general structure of the Michelson interferometer is still recognisable, numerous improvements have been made to the design to achieve the sensitivity necessary to detect gravitational waves.

loss.

The comparison between this and predictions of energy emitted by the binary through gravitational waves showed a remarkable correspondence. It was for this discovery that Hulse and Taylor were awarded the Nobel prize in physics in 1993.

LIGO started listening for gravitational wave signals in 2002 and has been intermittently upgraded since then. When suddenly, during an engineering test on 14th September 2015, it detected gravitational waves from a colliding pair of black holes (BHs). This showed, unequivocally, that gravitational waves were real, almost 100 years after their prediction by Einstein. In 2017, Rainer Weiss, Barry C. Barish and Kip S. Thorne were awarded the Nobel prize in physics for their contributions to LIGO and the observation of gravitational waves. This first detection event was named GW150914, following the US date convention of year-month-day. All gravitational wave events discussed in this thesis will follow this convention.

Currently, there are several gravitational wave detectors active: the two LIGO sites in the USA (Aasi et al., 2015), Virgo near Pisa in Italy (Acernese et al., 2014) and the KAGRA GRAvitational wave detector (KAGRA) located in the Kamioka mine, Japan (Akutsu et al., 2020). Together they work under the LIGO-Virgo-KAGRA (LVK) collaboration. There are also several future detectors planned, such as Cosmic Explorer (CE) (Reitze et al., 2019), the Einstein Telescope (ET) (Punturo et al., 2010) and the Laser Interferometer Space Antenna (LISA) (Babak et al., 2021). This thesis will work with simulated data from ground-based detectors such as the current LVK detectors and ET. LISA is not considered as it is mainly designed for very low-frequency gravitational waves and its output data stream looks very different to that from ground-based detectors. In principle, the ideas from this work could be transferred to a data analysis pipeline for LISA. This could be interesting, for example, in the context of quasi-normal mode measurements of black holes (Deppe et al., 2024; Pitte et al., 2024).

1.2 Gravitational Wave Sources

Gravitational waves are very easy to produce. In fact, you are creating them right now while flicking/scrolling through this thesis. Any non-spherically symmetric acceleration will generate gravitational




Figure 1.2: Measurements of the Hulse-Taylor binary periastron time (indicative of orbital period) shift compared to the GR prediction (solid line) (Weisberg et al., 1981). This provides indirect evidence for the existence of gravitational waves.




Figure 1.3: Currently observed sources of GWs: compact binary coalescences (CBCs) (left) and the stochastic GW background (right). The black arrows indicate what objects these detection methods observe

waves. To put it in technical terms: the gravitational wave amplitude depends on the second time derivative of the mass-quadrupole moment of the system. A less technical explanation that I like a lot is one by Caplan and Horowitz:

*“All one needs is to put a mass on a stick and shake vigorously, although for the GWs to be detectable the mass must be very large and the stick very strong.”*¹ (Caplan and Horowitz, 2017)

So while you may be creating gravitational waves, they are undetectable as they are just too small. You need incredibly heavy and dense objects like black holes colliding to create the signals we see directly. Usually, electromagnetic observations probe the surface emission from objects, which makes it tricky to infer the interior properties of objects. Gravitational waves, on the other hand, depend on the dynamics of the mass-distribution of a system. Observing gravitational waves allows us to directly study these dynamics.

Currently, we detect two types of gravitational wave sources: on the one hand binary systems of compact objects and on the other the entire observable universe, of which the latter results in the stochastic gravitational wave background. It is indirectly detected as it uses the radio emission of pulsars to map the cumulative background from all gravitational wave emitting sources. This is done by using Pulsar Timing Arrays (PTAs) and forms a fascinating facet of gravitational wave science. Due to its indirect nature, it will not be discussed further in this work. In compact binary coalescences (CBCs) we see the final stages of the binary system before its two bodies coalesce. As these objects orbit each other, they lose energy through gravitational waves emitted at their orbital frequency which decreases their orbital separation and period. As the power radiated via gravitational waves by a binary system scales with frequency as $W \propto f^{10/3}$, a positive feedback loop is created in which the frequency of the emitted gravitational waves increases until coalescence, resulting a typical “chirp” signal. The binary systems that we observe like this contain black holes and neutron stars...

1.3 Exotic Compact Objects

...or at least we think that we only see black holes and neutron stars during compact binary coalescences. As we can only probe the masses of these objects, those in the $1\text{--}2\text{ }M_{\odot}$ range are categorised as neutron stars and heavier objects as black holes. However, one cannot fully exclude that these objects might be something more exotic. To date, a lot of the literature is unable to do so. Even with electromagnetic observations, it is nearly impossible to distinguish black holes or neutron stars from these hypothetical exotic objects. This is partially by design as any hypothetical object is constructed in such a way that its observational signatures will be consistent with current observations. So what are these objects and

¹This quote was published in the context of detecting “mountains” on neutron stars, where the mass-stick analogy makes more sense

why are they interesting?

Nuclear and condensed matter physicists might get very excited by prospect of new states of matter, so called *exotic matter*. These appear when matter is pushed into extreme regimes. One such example is the quark-gluon plasma. This may be found in the core of neutron stars, or in the bulk of *quark stars*: hypothetical objects made out of quark matter (Ivanenko and Kurdgelaidze, 1967). This state of matter consists of free quarks and gluons. Its best known example is the quark-gluon plasma.

Particle physicists have been looking for extensions to the Standard Model since its inception. One of its great successes was the prediction and detection of the Higgs boson in 2012 at the Large Hadron Collider (LHC) (Aad et al., 2012). But the Standard Model is most likely incomplete, with the nature of dark matter being one of its big question marks. Some families of exotic compact objects (ECOs), such as *dark stars*, can grow cores of dark matter in the early universe (Freese et al., 2016). This core could leave an imprint on the phase evolution of the coalescence of such objects.

In 1976, Gerard 't Hooft showed that quantum chromodynamics (QCD) permits the violation of charge conjugation and parity (CP) ('t Hooft, 1976). Experiments indicate that this violation must be extremely small (Abel et al., 2020), but why should this be the case? This is known as the strong CP problem in particle physics. Peccei and Quinn proposed a solution to this problem (Peccei and Quinn, 1977), which was later shown by Wilczek and Weinberg to produce a new bosonic particle, dubbed the axion. These are massive, weakly interacting particles which makes them a prime candidate for dark matter. They could condense into compact/star-like objects (Guerra et al., 2019) falling under the *Boson star* (Jetzer, 1992) category. Boson stars are made up of bosonic matter, contrary to the regular stars we know, which mainly consist of fermionic matter (namely, protons and electrons). When allowing for boson self-interaction and a boson-antiboson asymmetry, there are several possibilities for boson stars to form cosmologically (Madsen and Liddle, 1990).

General relativity itself, while a beautiful theory, still has some issues. For example: black holes are some of the simplest objects in the universe, needing only three intrinsic parameters (mass, spin and charge) to be described classically, but they are poorly understood at a fundamental level. The prediction of singularities has led some authors to propose alternatives to black holes. One such object is the *gravastar* (Mazur and Mottola, 2001): an expanding de Sitter spacetime contained within our own de Sitter spacetime by a thin shell of matter, a bit like a spacetime soap bubble. Its observational signatures are nearly identical to those of a black hole, possibly differing only in the quasinormal mode spectrum, which is as of yet unmeasured.

Black holes, while seemingly simple, are poorly understood at the quantum level. One of the first instances where classical black holes were combined with a quantum description of matter was in 1974 by Stephen Hawking. This led to the prediction that black holes emit thermal radiation characterised by a temperature T_H (Hawking, 1974), now called the Hawking temperature:

$$T_H = \frac{\hbar c^3}{8\pi G M k} \quad (1.1)$$

This poses an interesting paradox: Thermal radiation does not contain any information, however, everything that has fallen into a black hole carried a certain amount of information. According to the unitarity of quantum mechanics, information must be preserved. Hawking radiation violates this principle, leading to the information paradox. A possible resolution to this paradox has been proposed using a string theoretical description of black holes: so-called *fuzzballs* (Mathur, 2005). At the macroscopic scale fuzzballs look identical to classical black holes, but near the Planck scale they would be fuzzy, hence their name. This structure could allow information to escape in the radiation. Fuzzballs might be distinguishable from black holes through gravitational wave observations by looking at ringdown frequencies or tidal excitations (Mayerson, 2020; Ikeda et al., 2021).

There are many more exotic compact objects possible, but hopefully this brief overview has illustrated the wide range of physics that these objects cover and can shed light upon if discovered and studied. In what follows, we will look a bit more in depth into how we might be able to use gravitational waves to find such objects.

1.4 Asteroseismology

1.4.1 Introduction

As one of the envisioned applications of this thesis concerns itself with the internal dynamics of compact objects, a brief introduction to a relevant and older field is warranted: that of asteroseismology. Much of this section is based on (Aerts et al., 2010).




Figure 1.4: Fundamental and first two harmonic modes of string vibrations (top) compared to wind instrument vibrations (bottom). The string is fixed at both ends whereas one end of the wind instrument can move. For the latter the oscillations are in the density of the air inside the instrument.

When you pluck a guitar string, it vibrates at the frequency that it is tuned to. This is a one-dimensional object and therefore the waves it carries have nodes (zero-crossings) in one direction. Such oscillations can happen at the fundamental frequency, or integer multiples of it, also called overtones or harmonics. The fundamental frequency depends on the length of the string, its tension and its composition. A better analogy for stellar objects would be a wind instrument such as the humble recorder as one end of the wave is not constrained to a single location as can be seen in the bottom panel of Fig. 1.4. The resonant frequencies are once again dependent on the instrument length and the composition of the medium: the air in the pipe. Anyone who has inhaled helium, and experienced the subsequent hilarity, can attest to this.

Expanding to 2D-objects, such as the drum head, we find nodes in two orthogonal directions: *radial* and *nonradial*. Similarly, we can move to 3D objects, such as stars and compact objects. Here there are three independent modes of oscillation described by the distance to the centre r , co-latitude θ and longitude ϕ . For a spherically symmetric object, the displacements in these directions are given by:

$$\begin{aligned}\xi_r(r, \theta, \phi, t) &= a(r) Y_l^m(\theta, \phi) \exp(-i2\pi\nu t) \\ \xi_\theta(r, \theta, \phi, t) &= b(r) \frac{\partial Y_l^m(\theta, \phi)}{\partial \theta} \exp(-i2\pi\nu t) \\ \xi_\phi(r, \theta, \phi, t) &= \frac{b(r)}{\sin \theta} \frac{\partial Y_l^m(\theta, \phi)}{\partial \phi} \exp(-i2\pi\nu t)\end{aligned}$$

The $Y_l^m(\theta, \phi)$ are called spherical harmonics. They express nonradial dependencies in spherical symmetries. As these are quantised by three numbers (n, l, m) , these numbers are referred to as *quantum numbers*. The spherical harmonics are also usually first encountered by physicists as the eigenfunctions of the square angular momentum operator ($\hat{L}^2$) and azimuthal component of the angular momentum operator ($\hat{L}_z$) in quantum mechanics. This probably also explains, in part, why we like to refer to (n, l, m) as quantum numbers, although there is nothing inherently quantum mechanical about them. These numbers are labelled as the *overtone* n , the number of nodes in the radial direction, the *degree* l , the number of surface nodes present, and the *azimuthal order* m , where $|m|$ specifies how many of the surface nodes are lines of longitude.

With these three quantum numbers (n, l, m) , one can describe the vibrational modes of a nonrotating sphere. When applying this to stars, these modes can be linked to observations in the periodicity of light curves and used to infer properties about the interior of stars.

1.4.2 Dynamical tides

Dynamical tides are the internal fluid motions of stars, also referred to as quasi-normal modes (QNMs). There are also *equilibrium* tides that concern the geometric deformation of the star. Fig. 1.5 shows an illustration of both.




Figure 1.5: Illustration of equilibrium (left) and dynamical (right) tides. Equilibrium tides deform the star due to gravitational interactions. The left image in particular shows a quadrupolar deformation. The right image illustrates dynamical tides, which are due to the internal motions of the fluid. For more details on the equations shown, see (Suvorov et al., 2024)

There is a great wealth of physical origins that can be attributed to different modes. As a rule of thumb, each new piece of physics will add (at least) one new family of modes to one's model (Andersson, 2019). Generally, they are characterised by the restoring force that damps the oscillations. I will discuss some of the most common modes found in literature, especially in the context of GW asteroseismology.

Firstly, there are ***p*-modes** or pressure/acoustic modes. They are similar to the pressure waves seen in the recorder from the introduction and represent waves of varying pressure inside the star. *p*-modes with the lowest radial quantum number $n = 0$ are referred to as fundamental, or ***f*-modes** (Suvorov et al., 2024). They belong to the most studied modes in this field as the *f*-modes have the strongest tidal coupling (Pratten et al., 2020). For neutron stars, their frequencies typically lie in the 2-3 kHz range. For Newtonian stars an exact formula for the *f*-mode frequency can be derived:

$$\omega_f = \frac{2l(l-1)}{2l+1} \frac{M}{R^3} \quad (1.2)$$

In the case of relativistic stars, this relation is approximately correct (Andersson et al., 1996). Modes restored by buoyancy are called ***g*-modes**, or gravity modes (not to be confused with GWs). These modes are highly sensitive to the interior composition of the star and could allow us to infer information about temperature or composition gradients.

***r*-modes**, rotational/Rossby modes, have the Coriolis force as their restoring force. Therefore they can only occur in rotating stars. To first order in the rotational frequency of the star Ω , the frequency of the *r*-mode is given by (Kostas D. Kokkotas and Schmidt, 1999):

$$\omega_r = \frac{2m\Omega}{l(l+1)} \quad (1.3)$$

Where l, m are the degree and azimuthal order of the mode. For an inertial observer the observed frequency is:

$$\omega_i = \omega_r - m\Omega \quad (1.4)$$

This implies that a counter-rotating r -mode (w.r.t the star’s co-rotating frame) can be observed as co-rotating to an inertial observer. For all r -modes with $l \geq 2$ this will be the case, which makes them unstable to the emission of GWs due to the Chandrasekhar-Friedman-Schutz (CFS) mechanism (Chandrasekhar, 1970; Friedman and Schutz, 1978). This mechanism predicts that GW emission from the perturbed star triggers the transfer of stellar rotational energy into the oscillation energy, leading to the growth of the mode’s amplitude (Kraav et al., 2024). This has made r -modes the subject of increased interest as their potentially high amplitude could become a prominent feature in a GW waveform.

A common approximation made to calculate all of the above modes for Newtonian stars is the Cowling approximation² (Cowling, 1941). Within the Cowling approximation, variations in the gravitational potential due to varying densities are neglected as much of the star’s mass is expected to be concentrated near the centre. Of course, as this thesis deals with extremely compact objects, a relativistic extension is needed, which was provided several decades later and is known as the relativistic Cowling approximation (McDermott et al., 1983). This still neglects metric perturbations, so within the relativistic Cowling approximation, there is no gravitational backreaction, causing the modes to be undamped. The oscillation modes can also be derived in a relativistic framework, where they are close analogues to the Newtonian modes described above. However, there also are modes that have no Newtonian analogue. These have been christened **w -modes** (gravitational-wave or curvature modes) (K. D. Kokkotas and Schutz, 1992). They are strongly damped due to their strong coupling to GWs and are therefore expected to be short-lived.

A summary of the characteristic frequencies and damping times of a selection of the described modes can be found in Table 1.1

mode	frequency	damping time
f	2.87 kHz	0.11 s
p_1	6.57 kHz	0.61 s
g_1	19.85 Hz	years
w_1	12.84 kHz	0.024 ms
w_{II}	8.79 kHz	0.016 ms

Table 1.1: Typical values of the frequencies and the damping times of various families of modes for a polytropic star ($N = 1$) with $R = 8.86$ km and $M = 1.27M_\odot$ are given. A subscript 1 identifies a mode as the first of its family. w_{II} is the slowest damped *interface* mode. Values taken from (Kostas D. Kokkotas and Schmidt, 1999)

1.5 Tidal Resonances

What was the relevance of going into such painstaking detail describing resonant bar detectors in Section 1.1? Having an idea of how resonances work is essential to understand one of the key applications of this thesis.

Imagine placing one tuning fork next to another tuned to the same frequency. Then you hit one. This will result in the other resonating at the same frequency. In fact, any integer multiple frequency of the resonance will also be a resonant frequency. Not just tuning forks have resonant frequencies though, every object has them. Now imagine placing two neutron stars in orbit around each other. They will orbit and lose energy via GWs. In doing so, their orbital frequency, and consequently emitted GW frequency, will increase. If the orbital frequency equals the resonant frequency of one of these objects at any point during the inspiral, the resonance may get tidally excited and leave an imprint on the GW evolution. As such resonant frequencies are intimately related to the structure of their parent objects (see Section 1.4.2), they can be used to infer the nature of an object’s matter.

One expected effect is an overall dephasing of the GW signal as can be seen in the bottom panel of Fig. 1.6. A study has already been conducted by (Asali et al., 2020) to look for such dephasing effects in currently observed GWs. This thesis will try to model and detect such resonances explicitly by using a basis of sine-Gaussian wavelets (see Section 2.4.3).

²While the approximation is named after Cowling, he mentions it was also used by Emden in his 1907 book *Gaskugeln* (Emden, 1907).




Figure 1.6: Self-consistent mode amplitude evolutions of f - and $g_{1,2,3}$ -modes due to tidal forcing, together with the dephasing associated with the f - and g -modes. In green the frequency evolution of the binary is plotted, corresponding to the right-hand axis. The underlying simulation was carried out for a $1.35\text{-}1.35M_{\odot}$ binary NS merger. Figure taken from (Suvorov et al., 2024).

1.5.1 Neutron Star Equation of State

Even if ECOs turn out to be a scientific fever dream, this research could still prove valuable in constraining the ever elusive neutron star (NS) equation of state (EOS). Such an equation links macroscopic variables such as pressure, volume and temperature for an ideal gas, or pressure and energy density for a NS (Prakash, 2013).

The EOS is interesting as it is usually derived from the microscopic interactions at the interior the NS. From there, a relation between mass and radius can be established using the Tolman-Oppenheimer-Volkov (TOV) equations (Oppenheimer and Volkoff, 1939; Tolman, 1939). These express constraints on a spherically symmetric, static body in gravitational equilibrium. As masses and radii³ can be measured, one can, in principle, find an EOS compatible with observations and learn something about the interior composition of the NS. Eq. (1.2), for example, gives a relation between M and R for a given f -mode frequency. This can then be compared to the M - R relations predicted by different EOSs to possibly eliminate those that are not in agreement with the observation of ω_f .

NSs have the potential to allow us to study matter in states that we cannot realise on earth, which makes them particularly interesting objects to study. For example, recently maximally massive NSs were shown to have an EOS consistent with quark matter (Annala et al., 2023). Currently, the NS EOS is poorly

³Defining the NS radius is rather tricky as there are at least three different definitions that I am aware of: (1) R_{∞} is based on the luminosity observed at infinity, which is influenced by gravitational lensing and the gravitational redshift of the observed temperature as $L_{\infty} = 4\pi R_{\infty}^2 \sigma T_{\infty}^4$. (2) One could use a ruler to define the distance from the centre of the NS to the outer layer as R , but due to the strong spacetime curvature this ruler will measure different lengths throughout the process. (3) One can also define the radius via the circumference $l = 2\pi R$, which circumvents the problem with the notion of length while moving through a strong gravitational field. In literature, it is usually not mentioned which definition is used, although R_{∞} would be most common as observations are photometric.




Figure 1.7: There are a wealth of EOSs compatible with current NS observations (Gonzalez et al., 2023).

constrained (see Fig. 1.7) and finding new ways to gain information about the NS interior could help in constraining it.

Methods

When a gravitational wave is detected, it can take on many different shapes, a *waveform*. From this waveform we would then like to infer the source parameters such as the component masses, spins, orbital parameters, etc. With increasing complexity of the model, we need to fit more parameters to describe the waveform. A parameter inference for fairly regular GW signals already encompasses 15 parameters. Classical methods such as gradient descent will rarely converge to a global optimum for such problems. To resolve this, *stochastic methods* have been developed. These sample over phase space with some degree of randomness in order to map the distribution of probable parameter combinations.

2.1 Bayesian Inference

Generally, a student's first encounter with statistics will be from a frequentist point of view. This approach views a parameter's value as fixed and its uncertainty a result of the limited number of measurements performed. In this thesis, however, we will approach the parameter estimation (PE) from a Bayesian perspective. The key difference lies in how uncertainty is treated: the Bayesian statistician will claim that a parameter's value is not fixed and that there is an underlying parameter *distribution* which your measurements sample. The error bars thus quantify the spread in the distribution of this parameter. Which school you prefer to believe is more of a philosophical question, but for practical purposes the Bayesian approach is quite natural. In real life we don't know the value of a parameter *a priori* and we will likely never know it exactly. However, we can quantify our uncertainty about this parameter by finding a distribution for it that is consistent with our observations.

Key to Bayesian statistics is *Bayes' Theorem*:

$$P(\Theta|\mathbf{D}, M) = \frac{P(\mathbf{D}|\Theta, M)P(\Theta|M)}{P(\mathbf{D}|M)} = \frac{\mathcal{L}_M(\mathbf{D}|\Theta)\pi_M(\Theta)}{\mathcal{Z}_M(\mathbf{D})} \quad (2.1)$$

Given a model M , this theorem expresses the conditional probability $P(\Theta|\mathbf{D})$ of obtaining the parameter set Θ for some data $\mathbf{D}$ as the likelihood $P(\mathbf{D}|\Theta)$ of obtaining that data with those parameters times the prior probability of those parameters occurring $P(\Theta)$ normalised to the probability of obtaining that data $P(\mathbf{D})$ for all parameter sets Θ , which is also called the evidence.

Bayes' theorem allows us to include prior knowledge about the system we study into our final probability distribution of Θ . We can then use this final posterior distribution as a new prior for future experiments. Thereby, we have essentially found a mathematical way of updating our prior beliefs about some system. For this initial discussion we will not look at the evidence ($\mathcal{Z}$) as it is simply a normalisation constant for a given model M , but it will play a very important role later. Because $P(\Theta|\mathbf{D}, M) \propto \mathcal{L}(\mathbf{D}|\Theta)\pi(\Theta)$ it is usually satisfactory to maximise this quantity, which is also called *a posteriori* maximisation (MAP). Maximising this means finding the parameters that have the highest probability of explaining the data that we observe while taking into account our prior knowledge about the system. Now, how should we go about maximising this?

2.2 Sampling algorithms

In general, there is no mathematically closed way of doing a posteriori maximisation. The objective function to maximise will usually have a lot of local maxima on a high-dimensional hypersurface. This is why we need to rely on stochastic techniques to reliably find a global maximum. The idea behind stochastic

sampling algorithms is to find a way to efficiently sample this high-dimensional hypersurface, with a certain degree of randomness, such that most samples will be concentrated around the highest maxima. The distribution of those samples then serves as an estimate of the unknown parameter distribution. Let's translate the following expression into something understandable¹:

$$P(\Theta|\mathbf{D}, M) \propto \mathcal{L}(\Theta)\pi(\Theta) \quad (2.2)$$

We can think of the priors as a “box” in phase space within which the high dimensional likelihood surface lives. Applying the appropriate sampling algorithm should then yield a proxy for the actual likelihood distribution that we do not know.

In the following subsections I will delve deeper into two popular sampling algorithms: Markov Chain Monte Carlo and Nested Sampling. Each of these has their extensions, but I will try to convey the main ideas of these base sampling methods.

2.2.1 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) samplers try to **solve a really hard problem once**: namely to find a distribution proportional to the posterior. There are many different flavours of MCMC sampling, so I will sketch one very popular algorithm: **Metropolis-Hastings**. It initially samples N random points, which we will call walkers, from the prior distribution and evaluates the likelihood for these walkers. It then moves all walkers to a different, random position. For each of them it checks the likelihood. If the likelihood is greater than at the previous location, the walker remains at its new location. If not, it remains at the new location with probability $\exp(-\Delta\mathcal{L}/kT)$. The temperature T in the denominator of the exponential is a measure for how likely a walker is to remain at an “unfavourable” location. Remaining at an unfavourable location is in this case a good thing as it allows the algorithm to explore more of the likelihood surface and prevents it from getting stuck in local maxima. If you crank up this virtual temperature, the algorithm will start to sample the phase space more randomly, which can become inefficient. Lowering it discourages the algorithm from exploring very much, which makes it more prone to getting stuck in local maxima. Tuning this temperature hyperparameter thus allows you to balance exploration against efficiency.

A nice way of illustrating this exploration-efficiency tradeoff is with a short bovine story: Imagine a herd of cows grazing in a mountainous landscape. They may be perfectly content in their current field, but you know what they say: *“The grass is always greener on the other side”*. Perhaps there is a particularly inquisitive cow who dreams of a field where the blades of grass are greener, bigger and juicier than those of his current field. So one day it decides to leave the herd and explore the mountains. It passes through dry, barren and rocky landscapes, barely surviving off the weeds it finds tucked between the rocks. But then it suddenly stumbles across the most beautiful field it has ever seen with greener, bigger and juicier grass than back home. There even are apple trees, something the cow had never dreamed of before! It calls for its herd which hurriedly moo-ves to its new utopia.

By exploring the mountainous landscape our hero cow had to endure unfavourable conditions, but eventually it found a better place for its herd. Perhaps if the herd had sent out more explorers, they would have found this place sooner, or maybe an even better one? Or maybe some cows would have died from starvation looking for a better field...

A nice extension to the Metropolis-Hastings MCMC method is Parallel-Tempered MCMC (PTMCMC). Here several MCMC samplers are run at different temperatures with walkers being exchanged between the samplers. The high-temperature samplers effectively smooth out the underlying likelihood distribution while the low-temperature ones probe local, small-scale features. Exchanging these walkers between the different-temperature samplers allows for a physics-inspired mapping of the likelihood surface that prevents the walkers from getting stuck in local maxima.

2.2.2 Nested Sampling

Nested sampling² was developed by John Skilling near the start of this century (Skilling, 2006) and takes a very different approach to MCMC samplers. Instead of solving one really hard problem once, it

¹Note that I am leaving out the $\mathbf{D}$ -dependence of $\mathcal{L}$. This notation is a matter of convenience and readability.

²This explanation is largely based on the `dynesty` documentation which can be found here: <https://dynesty.readthedocs.io/en/stable/index.html>

successively solves many easy problems by iteratively sampling within an iso-likelihood contour. It tries to come up with an estimate of the evidence $\mathcal{Z}$, which is usually very hard to compute. Let us rewrite the expression for $\mathcal{Z}$, which is an integral over all possible parameter configurations Ω_{Θ} :

$$\mathcal{Z} = \int_{\Omega_{\Theta}} \mathcal{L}(\Theta) \pi(\Theta) d\Theta \quad (2.3)$$


Instead of integrating over Θ , which can be high-dimensional, we can make use of the prior volume $X(\lambda)$ contained within an iso-likelihood contour $\mathcal{L}(\Theta) > \lambda$:

$$X(\lambda) = \int_{\Theta: \mathcal{L}(\Theta) > \lambda} \pi(\Theta) d\Theta \quad (2.4)$$

This allows to us to rewrite the evidence as:

$$\mathcal{Z} = \int_0^{+\infty} X(\lambda) d\lambda \quad (2.5)$$

$$= \int_0^1 \mathcal{L}(X) dX \quad (2.6)$$



We have transformed the really hard multidimensional integral from (2.3) into a simple 1D-integral! We can think of the integrand in (2.6) as a density $\mathcal{L}$ times a volume element dX , which yields a mass. Eq. (2.6) can then be interpreted as assigning more weight to the evidence from regions that have a higher “mass”.

During nested sampling we draw K *live points* from the prior $\pi(\Theta)$. We then remove the live point with the lowest likelihood and replace it with a new live point drawn from $\pi(\Theta)$ with the constraint that $\mathcal{L}_{i+1} \geq \mathcal{L}_i$. The removed live point is called a *dead point*.

We can approximate $\mathcal{Z}$ with our set of N dead points by writing (2.6) as a Riemann sum:

$$\mathcal{Z} \approx \sum_{i=1}^N f(\mathcal{L}_i) f(\Delta X_i) \equiv \sum_{i=1}^N w_i \quad (2.7)$$

Which expresses the evidence as a sum over each dead point’s weight w_i . f is a function that depends on the numerical integration scheme used. For a linear integration scheme using rectangles we can use: $f(\mathcal{L}_i) = \mathcal{L}_i$ and $f(\Delta X_i) = X_{i-1} - X_i$.

What’s really nice is that we can then use these weights to estimate the posteriors for free:

$$P(\Theta_i) \approx \frac{w_i}{\mathcal{Z}} \quad (2.8)$$

Nested sampling also allows for an easy stopping criterion. The remaining evidence at iteration i can be approximated as:

$$\Delta \mathcal{Z}_i \approx \mathcal{L}_{max} X_i \quad (2.9)$$

This assumes that the likelihood within the prior volume X_i is a flat slab with the value $\mathcal{L}_{max}$, where $\mathcal{L}_{max}$ is the maximum likelihood point contained within the set of K live points. With this, we can define a stopping criterion based on the current evidence estimation $\hat{\mathcal{Z}}_i$ and the remaining evidence $\Delta \mathcal{Z}_i$:

$$\Delta \ln \hat{\mathcal{Z}}_i \equiv \ln(\hat{\mathcal{Z}}_i + \Delta \mathcal{Z}_i) - \ln \hat{\mathcal{Z}}_i \quad (2.10)$$

This expresses the fraction of the evidence that is unaccounted for.

Figure 2.1: Illustration of nested sampling. Darker orange regions signify a higher likelihood. If the sampler has found this peak, these darker regions should also correspond to a higher density of samples

Dynesty

Dynesty (Speagle, 2020; Koposov et al., 2024) is a common nested sampler and the one that is mainly used in this work. Within the nested sampling framework there are several specific methods of sampling that are implemented in **dynesty**.

First of all, there are different ways of bounding the distribution within which to sample:

- **No bound:** This keeps on sampling within the priors provided. This decreases in efficiency as the sampling progresses.
- **Single ellipsoid:** (Mukherjee et al., 2006) Bounds the current samples by a single ellipsoid at each iteration, within which new samples are drawn.
- **Multiple ellipsoids:** (F. Feroz et al., 2009) Bounds the samples by multiple (overlapping) ellipsoids. This is useful, for example, when multimodalities start to appear in the distribution of samples. This is the standard in **dynesty** and what is used in this work.
- **Overlapping balls/cubes:** (Buchner, 2014; Buchner, 2019) Similar to multiple ellipsoids, but using balls/cubes.

“Live points” (i.e. “chains”)




Figure 2.2: Different sampling strategies in **dynesty**. Figure taken from Joshua Speagle’s presentation on nested sampling https://helper.ipam.ucla.edu/publications/gwaws3/gwaws3_17076.pdf.

Once a sampling strategy is chosen, there exist several methods of proposing a new live point conditioned on the bounding strategy:

- **Uniform:** Proposes a new live point based on a uniformly sampled point within the sampling bound.
- **Random walk:** Performs a random walk away from a current live point. The length of this walk can be set, for example, by specifying the number of autocorrelation times it should spend walking. This is what is used in this work, unless specified otherwise.
- **Random Slicing:** (Neal, 2003; Handley et al., 2015a; Handley et al., 2015b) is a non-rejection sampling method. The algorithm in 1D goes as follows: (1) choose a starting point $x_0 | f(x_0) > 0$ and sample a y -value uniformly between $[0, f(x_0)]$. (2) Draw the line segments within the curve at this y -value. (3) Sample a point (x, y) uniformly from these segments and repeat, starting from this x -value. This algorithm can be generalised to multivariate distributions.
- **Hamiltonian Slicing:** (Neal, 2003; Skilling, 2012; Farhan Feroz and Skilling, 2013; Speagle, 2020) combines slice sampling with Galilean Monte Carlo (GMC). In GMC a random point $\mathbf{x}_1$ is selected at each iteration i and given an initial velocity $\mathbf{v}$. A new point $\mathbf{x}_2 = \mathbf{x}_1 + \mathbf{v}$ is proposed and accepted if $\mathcal{L}(\mathbf{x}_2) > \mathcal{L}_i$. If rejected, a third point $\mathbf{x}_3$ is proposed by reflecting off $\mathbf{x}_3 = \mathbf{x}_2 + \mathbf{v}'$ where $\mathbf{v}' = \mathbf{v} - 2\mathbf{n}(\mathbf{n} \cdot \mathbf{v})$ and $\mathbf{n}$ is a unit vector perpendicular to $\nabla \mathcal{L}(\mathbf{x}_2)$. If $\mathcal{L}(\mathbf{x}_3) > \mathcal{L}_i$ it is accepted, otherwise the new point is reversed from $\mathbf{x}_1$ by giving it velocity $-\mathbf{v}$. This is not used in this thesis as we do not have access to the gradients of the likelihood.




Figure 2.3: Illustration of univariate slice sampling. Figure taken from (Neal, 2003).

2.3 Evidence and Occam's Razor

With all this talk about the evidence in nested sampling, yet its seeming irrelevance in MAP inference, you may be wondering why we should even care about it. It is simply a normalisation factor in (2.1) after all. While it is not too interesting for inference with a single model, it is very useful for model comparison because **Occam's razor** resides in $\mathcal{Z}$. Occam's razor is a "rule" that states that if two models have the same predictive power, but one is more complex than the other, then one should always pick the simpler model.

To describe how this mathematically slots into Bayes' theorem, we will first need to introduce the Kullback-Leibler (KL) divergence (Burnham and Anderson, 2004) of probability distributions f and g :

$$\mathcal{D}_{KL}(f, g) = \int f(x) \ln \left(\frac{f(x)}{g(x|\theta)} \right) dx \quad (2.11)$$

This denotes the information lost when g is used to approximate f . This can also be interpreted as the distance from g to f or as the *relative entropy* of f with respect to g . Writing out the log-evidence for a model M we get (Hergt et al., 2021):

$$\ln \mathcal{Z}_M = \ln \left(\int \mathcal{L}_M \pi_M d\Theta \right) = \int \mathcal{P}_M \ln \mathcal{L}_M d\Theta - \int \mathcal{P}_M \ln \left(\frac{\mathcal{P}_M}{\pi_M} \right) d\Theta \quad (2.12)$$

$$= \langle \ln \mathcal{L}_M \rangle_{\mathcal{P}} - \mathcal{D}_{KL}(\mathcal{P}_M, \pi_M) \quad (2.13)$$

For which the second term is exactly the KL divergence between the prior and posterior, which quantifies the overall compression in going from the prior to the posterior distribution.

The fantastic book by David MacKay (MacKay, 2002) provides a very nice intuitive understanding of the Occam factor. As the posterior distribution, which is proportional to the integrand of $\mathcal{Z}$, is peaked once the sampling is done, the integral (2.3) can be approximated by the integrand's height times its width $\sigma_{\Theta|\mathbf{D}}$. We can also think of the prior $\pi_M(\Theta)$ as the inverse of our initial uncertainty about the parameters σ_{Θ} , which allows us to write:

$$\begin{aligned} \mathcal{Z}_M &= \int \mathcal{L}_M(\mathbf{D}|\Theta) \pi_M(\Theta) d\Theta \\ &\approx \mathcal{L}_M(\mathbf{D}|\Theta) \pi_M(\Theta) \sigma_{\Theta|\mathbf{D}} = \mathcal{L}_M(\mathbf{D}|\Theta) \cdot \frac{\sigma_{\Theta|\mathbf{D}}}{\sigma_{\Theta}} \end{aligned}$$

The second factor in the final expression, also called the Occam factor, is the ratio between the uncertainty of the posterior $\sigma_{\Theta|\mathbf{D}}$ and the prior uncertainty σ_{Θ} . It can be thought of as the ratio between the posterior and prior volume and




Figure 2.4: Samplers try to bring an initial distribution down to a smaller posterior distribution that is representative for the underlying likelihood

expresses the following question: How much have we **reduced the uncertainty about our parameters** compared to our initial belief? A visual aid for this is given in Fig. 2.4.

With more parameters comes a larger prior volume, which is penalised by the denominator of the Occam factor. When comparing two models, the added complexity of one model needs to be compensated by a higher likelihood and/or a smaller posterior volume in order to be preferred to the simpler one.

To compare two different models (M_1, M_2) one can look at the ratio of evidences, or the Bayes factor:

$$\mathcal{B}_{M_2}^{M_1} \equiv \frac{\mathcal{Z}_{M_1}}{\mathcal{Z}_{M_2}} \quad (2.14)$$

A large Bayes factor is then indicative of M_1 being better-suited for the data than M_2 . When Bayes factors are mentioned without a second model in this work (simply as $\mathcal{B}$), the model is being evaluated against noise:

$$\mathcal{B} \equiv \frac{\mathcal{Z}_M}{\mathcal{Z}_{\text{noise}}} \quad (2.15)$$

2.4 Bilby

Bilby³ is a general purpose open-source Bayesian inference library, developed specifically with GW science in mind (Ashton et al., 2019). It allows its users to simulate GW waveforms from a variety of sources such as BH-BH and NS-NS mergers, supernovae and custom source models. These can then be injected into detector noise and recovered using a variety of samplers such as those discussed in Section 2.2. The standard likelihood function used in Bilby for GW parameter estimations (PEs) is a Gaussian noise likelihood. For data $\mathbf{D}$ and waveform $\mu(\Theta)$ in k frequency bins, the Gaussian noise log-likelihood can be written as:

$$\ln \mathcal{L}(\mathbf{D}|\Theta) = -\frac{1}{2} \sum_k \left\{ \frac{[\mathbf{D}_k - \mu_k(\Theta)]^2}{\sigma_k^2} - \ln(2\pi\sigma_k^2) \right\} \quad (2.16)$$

This checks for consistency between the residual signal and Gaussian noise. For more details on the exact structure of Bilby, see Appendix A.



Figure 2.5: Bilbies are a type of marsupial native to Australia, where the Bilby codebase was developed (Rafferty, n.d.). They are currently endangered.

2.4.1 Corner Plots

Most of the distributions inferred in this thesis live in a very high-dimensional space and are impossible to visualise in a single figure. To convey the main properties of such distributions, corner plots have been devised. These show the individual parameter distributions on the diagonal and correlations between parameters on the off-diagonal elements. An example is shown in Fig. 2.6.

2.4.2 Signal to Noise Ratio

Later on I will mention the signal to noise ratio (SNR) quite often, so let me briefly explain how this is calculated in Bilby. It follows a procedure called optimal matched filtering. For this explanation I will closely follow (Maggiore, 2007). Generally we can state that the data $d(t)$ we receive is composed of a template signal $h(t)$ and noise $n(t)$: $d(t) = h(t) + n(t)$. Now we would like to infer whether $h(t)$ is present in a given stretch of data. Let's say we have a linear filter $K(t)$ whose output is somewhat proportional to $h(t)$. As $h(t)$ is deterministic and the noise is zero-mean, we have that $\langle h(t) \rangle = h(t)$ and $\langle n(t) \rangle = 0$. We can now define the signal S and noise N :

³Bilby documentation: <https://lscsoft.docs.ligo.org/bilby/index.html>




Figure 2.6: Results of a parameter inference for an eccentric binary black hole merger shown in a corner plot. It shows parameter distributions on the diagonal and parameter correlations off-diagonally. These do tend to “waste” a lot of space on a page, for which (Rowan and Roberts, 2025) have found an innovative solution

$$S = \langle \hat{d} \rangle = \int_{-\infty}^{\infty} dt \langle d \rangle K(t) = \int_{-\infty}^{\infty} dt (\langle h(t) \rangle + \langle n(t) \rangle) K(t) = \int_{-\infty}^{\infty} df \tilde{h}(f) \tilde{K}^*(f) \quad (2.17)$$

$$\begin{aligned} N^2 &= [\langle d^2(t) \rangle - \langle d(t) \rangle^2]_{h=0} \\ &= \langle d^2(t) \rangle_{h=0} \\ &= \int_{-\infty}^{\infty} dt dt' K(t) K(t') \langle n(t) n(t') \rangle \\ &= \int_{-\infty}^{\infty} dt dt' K(t) K(t') \int_{-\infty}^{\infty} df df' e^{2\pi i f t - 2\pi i f' t'} \langle \tilde{n}^*(f) \tilde{n}(f') \rangle. \end{aligned} \quad (2.18)$$

We assume that the noise is stationary and the different Fourier components are uncorrelated, allowing us to write:

$$\langle \tilde{n}^*(f) \tilde{n}(f') \rangle = \frac{1}{2} \delta(f - f') S_n(f) \quad (2.19)$$

The function $S_n(f)$ is called the *noise spectral density*. This is what characterises the detector noise and,

usually, $\sqrt{S_n(f)}$ is what you see plotted when looking at LVK noise curves. Plugging this into Eq. (2.18), we finally obtain:

$$N^2 = \frac{1}{2} \int_{-\infty}^{\infty} df |\tilde{K}(f)|^2 S_n(f) \quad (2.20)$$

This allows us to define the SNR:

$$\frac{S}{N} = \frac{\int_{-\infty}^{\infty} df \tilde{h}(f) \tilde{K}^*(f)}{\left[\frac{1}{2} \int_{-\infty}^{\infty} df |\tilde{K}(f)|^2 S_n(f) \right]^{1/2}} \quad (2.21)$$

What exactly is the form of $\tilde{K}(f)$? If we introduce the noise-weighted inner product between two real functions $a(t)$ and $b(t)$ as (Jaranowski and Królak, 2012):

$$\langle a|b \rangle = 4\Re \int_{f_{min}}^{f_{max}} \frac{\tilde{a}^*(f) \tilde{b}(f)}{S_n(f)} \quad (2.22)$$

we can rewrite Eq. (2.21) as:

$$\frac{S}{N} = \frac{\langle u|h \rangle}{\langle u|u \rangle^{1/2}} \quad (2.23)$$

Where $u(t)$ is a function whose Fourier transform is:

$$\tilde{u}(f) = \frac{1}{2} \tilde{K}(f) S_n(f) \quad (2.24)$$

Maximising the SNR has now been reduced to finding the unit “vector” $\hat{n}$ for which the scalar product with h is maximum. This occurs when $\hat{n}$ and h are parallel, or $\tilde{u}(f)$ proportional to $\tilde{h}(f)$, which leads to:

$$\tilde{K}(f) \propto \frac{\tilde{h}(f)}{S_n(f)} \quad (2.25)$$

This matched filter, also called the Wiener filter, is proportional to the template signal weighted by the noise spectral density. This is quite natural as the filter will weigh the signal more when there is less noise present. The proportionality constant is completely arbitrary as rescaling $\hat{d}$ does not change its SNR.

2.4.3 Resonance Model

For this thesis a custom source function was implemented in **Bilby** allowing for a superposition of N wavelets on top of an inspiral-merger-ringdown (IMR) signal. The wavelet basis is used to reconstruct exotic/unmodelled features on top of an IMR signal. These could be astrophysical in origin, such as resonances, but also other features like detector glitches may be modelled in this manner.

A wavelet, more specifically a Morlet-Gabor wavelet, is a sinusoidal wave with a Gaussian envelope modulating its amplitude and was implemented with the following functional form:

$$\Psi(t) = \frac{\mathcal{A}}{\sqrt{2\sigma\sqrt{\pi}}} e^{-\frac{1}{2}\left(\frac{t-t_0}{\sigma}\right)^2} \cos(2\pi f(t-t_0) + \phi) \quad (2.26)$$

With the following 5 degrees of freedom (DOF):

- $\mathcal{A}$: amplitude of the wavelet
- σ : width of the Gaussian envelope
- t_0 : time at which the wavelet is centred
- f : frequency of the sinusoid
- ϕ : phase shift of the sinusoid

This basis is overcomplete, meaning that the reconstruction of a signal by wavelets is not unique. The basis functions have been normalised such that, for $\mathcal{A} = 1$:

$$\int_{-\infty}^{+\infty} |\Psi(t)|^2 dt = 1 \quad (2.27)$$

2.5 tBilby

We do not know the amount of wavelets necessary to reconstruct a signal a priori and every additional wavelet also adds 5 DOF to our phase space, meaning that the complexity of our inference ramps up quickly. For these reasons we would like to find the *minimal* number of wavelets necessary to *satisfactorily* reconstruct a given signal. In theory, you could fit an infinite number of wavelets to a signal and reconstruct it perfectly. Unfortunately, signals also contain noise and such a model would also fit to this noise. This is a well-known issue called **overfitting** and is particularly notorious in fields like machine-learning. We can combat this problem with a variety of **regularisation** techniques. The way tBilby⁴ tackles this issue is inspired by the **BayesWave** package (N. J. Cornish and Littenberg, 2015). **BayesWave** uses a technique called Reversible-Jump MCMC (RJMCMC) where the dimensionality of the phase space is actively varied during the MCMC sampling process. Such a sampler is called *transdimensional*. **BayesWave** has by now been well-established in the GW community but it is unable to mix modelled and unmodelled waveforms. For this we turn to tBilby, a transdimensional extension to Bilby (Tong et al., 2025). The technique it uses to incorporate transdimensionality is slightly different to RJMCMC as seen in **BayesWave**. It also varies the number of component functions N , but it samples their parameters in the full phase space instead of just the parameters of the used functions. The likelihood is then only evaluated over the used $\Theta_{k \leq N}$ component function parameters. The other $\Theta_{k > N}$ are referred to as “ghost parameters”. These samples are removed in post-processing as they are not a part of the final model. Drawing samples from the priors is a cheap operation, provided that the priors are not too complex. As such, the added cost of sampling ghost parameters is relatively small compared to the now cheaper likelihood evaluation, which is the costly step of the sampling algorithm. This construction allows tBilby to use the off the shelf samplers available in Bilby.

In principle, one could do a standard Bilby run for every value of N and compare the evidences of each run instead of using tBilby, but this approach can quickly become computationally prohibitive.

2.5.1 Order Statistics




Figure 2.7: (a) The likelihood multimodality can be seen in the sampling results of the component masses in a CBC parameter inference. (b) An example of descending SNR priors. ρ_i is defined as $\text{SNR}_i/\text{SNR}_{\max}$ (Tong et al., 2025)

The use of identical component functions introduces an unwanted symmetry: we can exchange parameter labels between the different component functions without changing the final likelihood. This introduces a

⁴tBilby github: <https://github.com/tBilby/tBilby/tree/main>

symmetric multimodality in the likelihood which makes it inefficient to sample over. A place where such a symmetry is very clear is in the GW waveform of CBCs. The waveform remains invariant under the exchange of the component masses which is reflected through a bimodal likelihood in both masses. This makes the sampling inefficient as an extra, uninformative, peak has been introduced. Fig. 2.7a shows this bimodality in the final distribution of component mass samples. This symmetry can be broken by sampling over the mass ratio $q = \frac{m_1}{m_2} \leq 1$ and the total mass $M = m_1 + m_2$ instead of the component masses (Buscicchio et al., 2019). In this work the total mass is not sampled over. Instead, the chirp mass $\mathcal{M}$ is used for sampling as it is more natural to use in a GW analysis context:

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad (2.28)$$

When sampling over N component functions, the degeneracy of a likelihood peak scales as $N!$. A nice way to break this symmetry is by assigning each component function an explicit label. Practically this means ordering the functions by some parameter such as time, frequency, amplitude, etc. All ordering schemes are mathematically equivalent, but for the problems encountered in this thesis, SNR ordering is most natural. As the sampler explores the likelihood landscape, it is likely to find the loudest contributions first, i.e. those with the highest SNR. This ordering scheme is incorporated into the priors as illustrated in Fig. 2.7b.

2.6 Inference pipeline

The first step of the inference is to create a simulated signal. In **Bilby** a base template is used from its source library on top of which N wavelets are added. The base signal depends on at least 15 parameters: 8 intrinsic and 7 extrinsic. Intrinsic are the component masses $m_{1,2}$, spins $a_{1,2}$, the tilt angle between their spins and the orbital angular momentum $\theta_{1,2}$ and the two spin vectors describing the azimuthal angle separating the spin vectors ϕ_{12} and the cone of precession about the system's angular momentum ϕ_{JL} . Extrinsic are luminosity distance d_L , orbital inclination angle ι , polarisation angle ψ , time of coalescence t_c , phase of coalescence ϕ_c , right ascension ra and declination dec . t_c is ignored in this work as we do not work with real data.

These base signals are usually generated by **Bilby** through **LALSimulation** (LVK Algorithm Library) (LIGO Scientific Collaboration, 2018). One peculiarity about this library is that it wraps the time series signal around the merger, which places the ringdown part of the IMR signal at the start of the time series. This leads to the sharp peak that can be observed at the start of the top right panel of Fig. 2.8.

The addition of a single wavelet then adds another 5 parameters as discussed in Section 2.4.3. The wavelets are sampled from a Gaussian centred at the same frequency as the IMR signal at the wavelet injection time. This is to model the resonance occurring at the frequency of the inspiral.

This composite signal is then injected into simulated detector noise.

This detector strain data, containing the injected signal and noise, is then passed to a sampler available in **Bilby** which returns the posterior parameter distributions. To ease computations, all extrinsic parameters were kept fixed. Additionally, in a real life setting one could imagine the following inference pipeline: A signal is detected using matched filtering, after which a “shallow” initial search is done using a waveform based on one of **Bilby**'s built in templates. If this still does not match the data well enough one could look to use more complex waveforms⁵, such as one containing wavelets. The initial search will have hopefully constrained the extrinsic parameters fairly well, which also motivates the choice to fix them in this work.

2.7 Quality of Reconstruction

With the methodology of how to perform PEs and an inference pipeline for joint modelled and unmodelled searches, we now need to find a way to assess the quality of a GW signal reconstruction. There are several metrics one can use.

⁵See also Matteo Vereetbrugghen's master thesis on an algorithm to see whether a waveform is complex enough to match the data

2.7.1 Residuals

The most obvious option is to subtract the injected and recovered signal and use the size of the residual signal as a figure of merit. This size could, for example, be the average or maximal strain of the residual. However, the residual is highly sensitive to slight offsets in phase, which will be hard to infer reliably using wavelets.

2.7.2 Bayes Factor

One number Bilby gives us automatically is the Bayes factor introduced in section 2.3. The information it encodes is twofold: both the qualities of fits as well as the model complexities are weighed. Unfortunately, there does not seem to be a consensus in literature on how large the Bayes factor needs to be in order to choose one model over another. When comparing against noise, the Bayes factor scales with the square of SNR (N. J. Cornish and Littenberg, 2015), which we will use as a sanity check to assess the implementation of new functionalities.

2.7.3 Significance σ

The trusty statistical significance is one many of us are used to. However, it can be a misleading metric. For example, the event GW151012 was at first classified as a LIGO-Virgo Trigger (LVT), due to its significance of 1.7σ . Its Bayes factor, however, was of the order of $\sim 10^{10}$ (D’Agostini, 2016)⁶. This is in part due to the significance solely looking at the SNR, while the Bayes factor also takes into account the detection of a similar signal accross multiple detectors. Moreover, this thesis works in a Bayesian setting, so the significance is not all that appropriate in this context.

2.7.4 (Mis)match

Using the the noise-weighted inner product from Eq. (2.22), we can define the match $\overline{M}$ between the recovered and injected waveforms, h_r and h_i respectively:

$$\overline{M} = \frac{\langle h_i | h_r \rangle}{\sqrt{\langle h_r | h_r \rangle \langle h_i | h_i \rangle}} \quad (2.29)$$

This allows us to define the mismatch (Janquart et al., 2023), which is simply $1 - \overline{M}$. The match can be used to define a metric on the waveform template space such that the mismatch is the metric distance between two nearby waveform templates (Owen, 1996):

$$1 - \overline{M} = g_{ij} \Delta x^i \Delta x^j \quad (2.30)$$

For a minimal mismatch, this distance (and therefore the mismatch) scales as SNR^{-2} (N. Cornish et al., 2011).

Of the figures of merit described above, the Bayes factor and mismatch are the ones that will be used to gauge the quality of reconstructions in this work.

2.8 Whitenning Data

In Chapter 4 the whitened strain will be shown often, so I will briefly describe what this is and how it is calculated. GW detectors’ noise can be described by a power spectral density (PSD) curve. This gives the square of the strain that the noise produces for each frequency. Its square root is called the amplitude spectral density (ASD). With the PSD, we can clean the data by trying to make the PSD flat. This process is called *whitenning* and allows us to see deviations from the noise more clearly.

In Bilby, whitening is done by Fourier transforming the time-domain detector data and dividing it by the square root of the PSD and a quarter of the duration:

$$d_{\text{whitened}}(f) = \frac{\tilde{d}(f)}{\sqrt{\text{PSD} \times \frac{t}{4}}} \quad (2.31)$$

⁶This is a rather sassily written paper, but it makes a few good points and the references therein are quite good. For example, it mentions that the American Statistical Association released a statement warning researchers about the improper use of p -values (Wasserstein and Lazar, 2016).




Figure 2.8: Illustration of the composition of waveforms from a CBC of rotating BHs and a single wavelet. The signals are shown in frequency (left) and time (right) domain. The sharp peak at the start of the time domain data is due to wrapping of the time domain signal by `LALSsimulation`, causing the ringdown of the IMR signal to occur at the start of the time series

Testing (t)Bilby

For this thesis I have made some modifications to the existing `Bilby` source code to accomodate wavelets into the analysis. In this chapter I will sketch these modifications and present the results of my tests. As `tBilby` is new and, as of writing this thesis, is only used in one paper (its release paper), some testing of this package is also warranted, especially with my modifications to the `Bilby` base in mind. I believe that I have also found a bug in the `tBilby` source code, which will be discussed in Section 3.2.1.

3.1 Addition of Wavelets to Bilby

`Bilby` was modified slightly for this work to incorporate wavelets into the analysis: source functions for wavelets in time and frequency domain as well as a superposition of wavelets on top of a base waveform were added to the `source.py` module of the `gw` package.

The wavelets are constructed such that their strain is independent of distance. As we do not know the amplitude of these features at the source, their amplitude is considered from the detector's point of view instead of the source's. This means that they do not scale with distance and should have their amplitudes manually rescaled if d_L is changed. Luckily, GW observatories do not measure power, but amplitude, which falls off as $1/r$, so rescaling is trivial.

To test whether `Bilby` is able to reconstruct wavelets well, 100 PEs are performed on random wavelets injected into the LIGO Hanford and Livingston detectors. The duration of the injection window was set to 10 s and the sampling frequency to 1024 Hz. Taking into account the Nyquist-Shannon sampling theorem, this means that the maximally inferable frequency is 512 Hz. The sampling prior distributions are uniform and bounded by the injection windows.

variable	unit	minimum	maximum
$\mathcal{A}$	-	1e-23	1e-22
σ	-	0.01	0.08
t_0	s	0	10
f	Hz	0	512
ϕ	rad.	0	2π

Table 3.1: Injection windows for the wavelet parameters. The recovery priors are uniform with their bounds given by the injection windows.

The PE was performed using the nested sampler `dynesty` on the *aalscholver* cluster of the Ghent Gravity Group, running in parallel on 50 CPUs with 500 live points. The sampling time took approximately 5.5 hours in total.

To evaluate whether a wavelet was detected, we look at the Bayes factor of the wavelet injection versus the noise. Only events with an average optimal SNR > 8 are selected, as is commonly done in LVK data analysis pipelines. In this way, 48 low-SNR events are filtered out of the analysis.

The Bayes factors scale with SNR in the expected quadratic fashion and, post-cutoff, have a median log-Bayes factor of 106.93. From this analysis, we can conclude that single wavelets can be reliably found with `Bilby`. Even without the SNR-cutoff, the median log-Bayes factor is still 41.11.

This only means that `Bilby` is able to detect wavelets, to infer the quality of the reconstructions, we look at the absolute relative deviations between the mean of the posterior parameter samples and their true values. The results of this are shown in Fig. 3.3. In case of $\mathcal{A}$, t_0 and f a fairly clear jump in the




Figure 3.1: Results of the single wavelet PEs. The left panel shows the relation between the log-Bayes factor and the average optimal detector SNR, which upper-bounds the SNR of the signal. The vertical line shows the SNR-cut-off. On the right a histogram of the log-Bayes factors is shown for SNR > 8. These indicate that **Bilby** is able to reliably find the implemented wavelets.

absolute relative deviation can be observed around the SNR = 8 threshold. σ posteriors do not seem to be as sensitive to SNR. The accuracy of the ϕ estimation is systematically quite poor. Fig. 3.2 shows two corner plots of reconstructed wavelets: one for a low and one for a high SNR injection. The low SNR samples appear random, indicating that the sampler struggled to converge to a likelihood maximum. The high SNR reconstruction is much better, but also struggles at reconstructing the phase.




Figure 3.2: Sample corner plots of an SNR ≈ 6 (left) and SNR ≈ 12 (right) wavelet. The injection values are shown in orange. The phase shift ϕ is poorly constrained in both cases.




Figure 3.3: Top row: Histograms of the absolute relative deviations between the posterior and true parameter values. Middle row: Log of the absolute relative deviations plotted against optimal SNR. The dashed line in red denotes $\text{SNR} = 8$. Bottom row: Same as top row, but with data filtered on $\text{SNR} > 8$.

3.2 tBilby Tests

3.2.1 Bug in tBilby Source Code

One of the ways that **tBilby** adds transdimensionality to **Bilby**, is through its ordered priors. These can be understood by **Bilby**, but I believe that something got lost in translation, causing a bug. In the source code for the **tBilby** class `TransdimensionalConditionalDescendingOrderStatPrior` the array `self._prev_value` is set to length 50 000 by **Bilby**. Elements from this are then selected by a mask, whose length may vary, which causes an array indexing error. I started looking for a bug when I was unable to reproduce the GW150914 reconstruction with wavelets from the **tBilby** release paper (Tong et al., 2025) due to the program crashing. I fixed this bug by modifying the updating rule of `self._prev_value` to:

```
>>> self._prev_val = self._prev_val[:len(mask.index)]
```

Running this on the **tBilby** example code produces results consistent with those of the paper. I have contacted the authors in regards to this fix, but I have not received an answer from them.

3.2.2 Wavelet Superpositions


To assert **tBilby**’s performance at inferring wavelet parameters, we perform a similar test to Section 3.1. Now, 1 to 5 random wavelets are superimposed in every signal and injected into simulated LIGO noise. As there can be multiple wavelets in a single signal, the injection amplitudes were halved to ensure that the final SNRs are not too high. The priors are the same as those used for the wavelet injection, except for the amplitude and N . The SNR priors are ordered and defined on the interval $[0.02, 25]$. The prior for N is a discrete uniform prior over the interval $[0, 5]$. The direction of the source is now important as each detector has an antenna pattern, leading to “blind spots” which influence the conversion of SNR to detector strain. The injection directions were selected randomly and fixed for the inference as they are extrinsic parameters.

Injection Prior				Inference Prior			
variable	unit	minimum	maximum	variable	unit	minimum	maximum
$\mathcal{A}$	-	5e-24	5e-23	SNR	-	0.02	25
σ	-	0.01	0.08	σ	-	0.01	0.08
t_0	s	0	10	t_0	s	0	10
f	Hz	0	512	f	Hz	0	512
ϕ	rad.	0	2π	ϕ	rad.	0	2π
N	-	1	5	N	-	0	5
Extrinsic							
ra	rad.	0	2π				
dec	rad.	$-\pi/2$	$\pi/2$				


Table 3.2: Injection windows and priors for the wavelet superposition parameters. As there are possibly multiple wavelets injected, the injection amplitudes have been halved compared to Section 3.1. Priors have the same intervals as the injection intervals, except for the amplitude and N . Instead of amplitude priors, the SNR priors for the SNR ordering of **tBilby** are used.

The prior for N includes 0 to see whether **tBilby** will identify low SNR signals as being purely noise. 50 waveforms are generated and have a PE performed on them. This is less than in Section 3.1 as these PEs are computationally more demanding, taking a bit more than a day to run for all waveforms.


The left panel of Fig. 3.4a shows the resulting Bayes factors against the optimal SNR. **tBilby** appears to identify the low SNR signals as containing no wavelets. The Bayes factors are very high for the other signals. Interestingly enough, **tBilby** seems to systematically underestimate the number of wavelets needed to reconstruct a signal as is illustrated in Fig. 3.4b. The Bayes factors are still high in these cases, because the model is being evaluated against noise. Finding a single wavelet among 5 is still better than finding none. Fitting additional wavelets may be suppressed by the increased Occam penalty.



(a) Results of the wavelet superposition PEs with tBilby . The left panel shows the relation between the log-Bayes factor and the average optimal detector SNR per number of wavelets used for the reconstruction. The vertical line shows the SNR-cutoff. On the right a histogram of the log-Bayes factors is shown for $\text{SNR} > 8$.



(b) Analysis of the distribution of wavelets N used for the signal reconstruction. The left panel shows the distribution of N for the injection and recovery. The right panel shows the confusion matrix for the reconstruction. Exact matches are indicated by a red box.



(c) Example of the time and frequency strain of a superposition of 5 wavelets underestimated in its reconstruction with one wavelet.

Figure 3.4

3.3 Discussion

In this chapter we have tested the implementation of wavelets into `Bilby` and the ability of `tBilby` to infer the parameters of an unknown number of wavelets.

From the first test, we can state `Bilby` is able to reliably infer single wavelets for signals with an SNR that is not too low. It has given us a first idea of the systematics of the wavelet inferences. Namely that there are parameters whose inference quality is strongly linked to the SNR. However, σ and ϕ do not seem to benefit much from an increase in the SNR. Especially phases are poorly constrained by the sampling, but as the most interesting parameters for resonances are f and σ , this is not too much of an issue. The poor ϕ estimation does mean that studying residuals can be misleading as a slight phase shift may result in a relatively large residual signal.

The test of `tBilby` has shown us that the number of wavelets is usually underestimated and that zero-wavelet reconstructions are possible. Luckily, these zero-wavelet reconstructions seem to coincide with low-SNR signals where the reliability of the reconstruction could be questioned. The systematic underestimation of the number of wavelets is likely due to the increasing Occam penalty when a new wavelet is added, as this adds 5 DOF to the prior volume. The reconstruction of multiple wavelets seems to be difficult as the residual signal after reconstruction with 1 wavelet might be “consistent enough” with noise to not warrant the addition of more wavelets and the subsequent increase of the Occam penalty.

In both tests, the square scaling of the Bayes factor with SNR could be seen. This serves as a sanity check that the implemented wavelets follow the expected scaling relations from (N. J. Cornish and Littenberg, 2015).

Overall, the implemented wavelets work and we can move on to combining them with IMR signals. However, `tBilby` appears to systematically underestimate the number of wavelets in its reconstruction, which we will keep in mind during the further analyses.

Results


In this chapter, the results of several PEs will be shown and discussed. We will first look at injections into LIGO noise in Section 4.1 and then move on to ET noise in Section 4.2. The parameters of the base IMR signal remain the same between these runs, but the wavelet injections are different and hand-tuned. These mainly serve as an initial exploration of the joint inference method developed in this work. After dipping our toes into these waters, we move on to Section 4.3 where we recover the CBC and wavelet parameters for 20 injections of random binary black hole (BBH) + wavelet waveforms. In Section 4.4 a mock f -mode resonance of a neutron star is injected into a binary neutron star (BNS) merger. Finally, Section 4.5 shows the results of a joint glitch + CBC parameter estimation. All inferences are performed on simulated data. Corner plots for all results can be found in Appendix B.

4.1 Double Wavelets with LIGO


In this first section, we analyse a BBH merger signal with two wavelets added on top of each other at the same central time $t = 0.5$ s, which was injected into LIGO detector noise. This stacked wavelet serves as a model for a more complex resonance structure than a single wavelet. A PE was then performed for both a low- and high-SNR signal. The injection parameters and recovery priors for this PE are given in Table 4.1.

Injection Value		Recovery Prior				
variable	value	variable	unit	type	min	max
m_1	36.0 $M_\odot$	$\mathcal{M}$	$M_\odot$	UniformInComponentsChirpMass	20	80
m_2	29.0 $M_\odot$	q	-	UniformInComponentsMassRatio	0.1	1
a_1	0.4	$a_{1,2}$	-	Uniform	0	0.99
a_2	0.3					
θ_1	0.5	$\theta_{1,2}$	rad.	Uniform	0	π
θ_2	1.0					
ϕ_{12}	1.7	ϕ_{12}	rad.	Uniform	0	2π
ϕ_{jl}	0.3	ϕ_{jl}	rad.	Uniform	0	2π
SNR	$\sim$	SNR	-	Transdim.Cond.DescendingOrderStatPriorSNR	0	5
σ	$\sim$	σ	-	Uniform	0.01	0.08
t_0	$\sim$	t_0	s	Uniform	0	$t_{\max}$
f	$\sim$	f	Hz	Uniform	0	512
ϕ	$\sim$	ϕ	rad.	Uniform	0	2π
N	$\sim$	N	-	DiscreteUniform	$\sim$	5
Extrinsic						
d_L	$\sim$					
ra	1.375					
dec	-1.2108					
ι	0.4					
ψ	0.1					
ϕ_c	1.3					

Table 4.1: Injection parameters and recovery priors of the BBH merger used in Sections 4.1 and 4.2. Parameter values marked with a “ $\sim$ ” indicate that these vary accross runs in these two sections.



(a) Reconstruction of the low SNR (≈ 26) signal with $d_L = 1000$ Mpc. The log-Bayes factor is $\log \mathcal{B} = -22.7$, indicating that the wavelet + CBC model is disfavoured over noise.



(b) Reconstruction of the high SNR (≈ 260) signal with $d_L = 100$ Mpc. The log-Bayes factor is much higher now: $\log \mathcal{B} = 94779.6$

Figure 4.1: LIGO Hanford data and response to the reconstruction of an IMR + wavelet signal with 2 wavelets stacked on top of each other at $t = 0.5$ s to mimic a complex resonance. The blue and grey shaded regions represent the 90% CL of the total waveform and wavelet reconstruction, respectively.

The luminosity distance was changed from 1000 Mpc to 100 Mpc between the runs to increase the SNR by a factor of 10. To simulate a resonance, the injection frequencies of the wavelets are drawn from a Gaussian centred around the frequency of the waveform at the wavelet injection time. This is done for all runs throughout this work. In Fig. 4.1a the result for an injection with $\text{SNR} \approx 26$ is depicted. The sampler indicates that it has detected one wavelet, however the log-Bayes factor of -22.7 indicates that the reconstruction is not favoured over noise. There is also a very large uncertainty on the wavelet parameters. The inclusion of a wavelet here is most likely due to the prior used requiring a minimum of one wavelet. In future runs this will be remedied by allowing for zero-wavelet reconstructions. The second run with a 10 times higher SNR performed much better as can be seen in Fig. 4.1b. Unfortunately, SNRs as high as 260 are not observed in current detectors, with the loudest event to date being GW170817 with an SNR of 32.4 (Abbott et al., 2017). `tBilby` fit a single wavelet to the stacked signal that shares features with both injected wavelets.

4.2 ET Injections

We now move on to ET injections for two reasons: The SNRs are much higher due to the lower noise floor and the low frequency cutoff lies at much lower frequencies. This allows us to inject detectable wavelets into the low frequency inspiral part of BBH IMR signals. These signals are relatively short, which keeps their analysis cost tractable. Fig. 4.2 shows the characteristic strain curves for several current and future generation ground-based detectors. The mentioned improvements by ET on current detectors are clearly visible.




Figure 4.2: Characteristic strain curves for current and future ground-based observatories. Figure generated from the web tool developed by (Moore et al., 2014).

4.2.1 Triple Wavelets

For this run, the same base waveform as the previous section was used, but at a luminosity distance of $d_L = 500$ Mpc. Due ET's higher sensitivity, the signal remains in the detector for longer, giving us a ~ 6 s segment to work with. The optimal SNR was also increased to 294.3 by working with ET. Three wavelets were injected at $t = 4, 5, 5.5$ s: during the early inspiral, just before the late inspiral and during the late inspiral phases respectively.

`tBilby` was able to identify the wavelets in the early and late inspiral sections, but not the middle wavelet. This was a quieter wavelet with an SNR of ≈ 1.94 . Its influence on the signal would have been marginal, so its inclusion might not have warranted the increased Occam penalty.




Figure 4.3: Reconstruction of the injection of three wavelets at $t = 4, 5, 5.5$ s with simulated ET noise. Only the wavelet at $t = 5$ s was not found. This was a fairly quiet wavelet with an $\text{SNR} \approx 1.94$. Its exclusion is likely due to its minimal influence on the total signal.

4.2.2 Wavelet in Merger

The previous experiment was performed at a high SNR. In order to get an indication of the performance at low SNRs, we can use the same base CBC waveform as before and inject a single wavelet into the late inspiral phase. The SNR was lowered by setting d_L to 4000 Mpc, which reduces the SNR to ~ 45 . The wavelet was still recovered well at this SNR as can be seen in Fig. 4.4.

4.3 Random Injections

In this section, we will perform a similar analysis to (Asali et al., 2020), albeit on a much smaller scale and with a different method. In this paper the authors look at phase shifts in the GW signal as an indicator for resonances in LVK data, whereas this work attempts to detect resonances directly in simulated ET data.

One or two wavelets are superimposed on top of a BBH IMR signal. The masses of the binary are sampled uniformly from the $[5, 70] M_\odot$ range with the restriction that $m_1 + m_2 \in [15, 110] M_\odot$. This choice is consistent with the GWTC1 catalogue and was also used in (Asali et al., 2020). This leads to a maximum ISCO¹ frequency of $f_{\text{ISCO}, \text{max}} = 293$ Hz. To be able to capture this edge case, the sampling frequency was set to 1024 Hz. These waveforms are mainly low-frequency, especially during the early inspiral, which is why they were injected into ET noise instead of LIGO.

Wavelets are only injected at frequencies below f_{ISCO} . Due to the monotonically increasing frequency with time, this can be translated to a t_{ISCO} , after which no wavelets are injected. As we are mainly interested in resonances and not necessarily the post-merger ringdown, wavelets are injected in $t \in [0.2, t_{\text{ISCO}})$ s, or if the signal is longer than 6 s, $t \in [t_{\text{max}} - 6, t_{\text{ISCO}})$, where t_{max} is the duration of signal. The latter is done as signals in ET can be quite long, so to keep the computational cost tractable, wavelets are only injected in the final 6 s. This allows us to restrict the t_0 prior for the wavelets to a 6 s window at most. A full overview of the injection and recovery priors is shown in Table 4.2.

¹Innermost stable circular orbit




Figure 4.4: Reconstruction of an $\text{SNR} \approx 45$ signal in ET with a single wavelet in the merger section of the waveform. The injected and reconstructed wavelets are shown as they may not be clear without indication in the merger section.

It may seem strange that the samples of $\theta_{1,2}$ are drawn from a sine distribution whereas the sampling prior is uniform over the same interval. While it would be more efficient to draw samples from the same distribution, we do not know the true distribution of $\theta_{1,2}$ in a real setting. The uniform distribution is therefore chosen as it gives a minimally biased prior for the sampler.

During the inference, the second set of parameters took exceptionally long to run (~ 4 days) due to the 16 s duration of the signal. To finish this series of PEs in a timely manner, I decided to filter out the signals with a duration > 10 s and replace them with signals shorter than 10 s. This resulted in four injections being replaced. This means that this analysis mainly looks at mergers of more massive systems as these generally have a shorter duration. Due to the exploratory nature of this study, I feel that the restriction of signal durations is warranted. Future works should analyse signals of any length permissible by the parameter space of the injection priors. The final injected waveforms are shown in Fig. 4.5. In total, it took about 15 days to perform all PEs. All available resources on the *aalscholver* cluster were used. This cluster is shared with other users, so this runtime is not indicative of the minimal runtime.

Injection Prior					Recovery Prior				
variable	unit	type	min	max	variable	unit	type	min	max
SNR	-	Uniform	0.02	5	SNR	-	TransdimensionalConditionalDescendingOrderStatPriorSNR	0.02	5
σ	-	Uniform	0.01	0.08	σ	-	Uniform	0.01	0.08
t_0	s	Uniform	$t_{\min}$	t_{ISCO}	t_0	s	Uniform	$t_{\min}$	$t_{\max}$
f	Hz	Uniform	0	f_{ISCO}	f	Hz	Uniform	0	512
ϕ	rad.	Uniform	0	2π	ϕ	rad.	Uniform	0	2π
N	-	DiscreteUniform	1	2	N	-	DiscreteUniform	0	2
m_1	$M_{\odot}$	Uniform	5	70	$\mathcal{M}$	$M_{\odot}$	UniformInComponentsSchirrmass	4	50
m_2	-	Uniform	$m_{\min}$	$110 - m_1$	q	-	UniformInComponentsMassRatio	0.07	0.99
$a_{1,2}$	-	Uniform	0	0.99	$a_{1,2}$	-	Uniform	0	0.99
$\theta_{1,2}$	rad.	Sine	0	π	$\theta_{1,2}$	rad.	Uniform	0	π
ϕ_{12}	rad.	Uniform	0	2π	ϕ_{12}	rad.	Uniform	0	2π
ϕ_{ji}	rad.	Uniform	0	2π	ϕ_{ji}	rad.	Uniform	0	2π
Extrinsic									
d_L	Mpc	Uniform	500	2000					
ra	rad.	Uniform	0	2π					
dec	rad.	Cosine	$-\pi/2$	$\pi/2$					
ι	rad.	Sine	0	π					
ψ	rad.	Uniform	0	π					
ϕ_c	rad.	Uniform	0	2π					

Table 4.2: Combined injection and prior values. The dashed line separates the IMR and wavelet parameters. $t_{\min}$ and $m_{\min}$ are calculated as follows: $t_{\min} = \max(0.2, t_{\max} - 6)$, $m_{\min} = \max(5, 15 - m_1)$.




Figure 4.5: Overview of the injected waveforms. The indicated SNR is the average SNR of the injected wavelet(s) over the three detectors in ET. 4 waveforms were replaced by shorter duration signals due to the long inference time of the second signal. Note that wavelets with a large strain are not necessarily the loud wavelets.

From Fig. 4.7b, we can see that 13 of the 20 injections were reconstructed with the correct number of wavelets. The 7 others were reconstructed with one less wavelet than injected. Generally, the “lost” wavelets are those with a low SNR as is seen in Fig. 4.7c. Only waveform nr. 5 seems to be an exception here. Upon closer inspection, the undetected wavelet is actually the loudest, but it does lie inside the merger section of the waveform. The two injected wavelets are very similar in terms of SNR, having an SNR of ~ 2.32 and ~ 2.80 (in order of occurrence in Fig. 4.6).

From Fig. 4.7c it is not apparent that the detection of a wavelet depends on the ratio of the wavelet SNR to that of the total waveform. Instead, there appears to be a cutoff for wavelet detectability around $\text{SNR} \approx 1.5$ for this limited dataset. Two edge cases for this $\text{SNR} = 1.5$ rule of thumb are waveforms 11 and 13. Their reconstructions are shown in Figs. 4.8a and 4.8b, respectively. Waveforms 5 and 11 both have a wavelet inside the late inspiral section, which raises the question why the one in 11 was detected. Having found the quieter wavelet first in waveform 5, the sampler incorrectly assumes that this is the loudest wavelet and will preferentially search for a quieter wavelet next due to the SNR ordered prior. This can cause it to miss the second, louder wavelet.







Figure 4.6: Time domain reconstruction of waveform 5. The second wavelet was missed by the sampler. This was likely due to the use of the ordered SNR prior.



(a) *Left panel*: Distribution of the number of injected and recovered wavelets. The number of injected wavelets is, once again, underestimated in some cases. *Right panel*: Confusion matrix for the number of injected and recovered wavelets. Exact matches are highlighted by a red box.




(b) Mirror plot of the number of injected and recovered wavelets for every run. 13/20 waveforms were reconstructed with the correct number of wavelets.




(c) *Left panel*: Quietest wavelet SNR plotted against the total waveform SNR. The annotated numbers coincide with the indices from Fig. 4.5. There appears to be an SNR detectability cutoff for wavelets around ~ 1.5 . *Right panel*: Mismatch vs SNR. The data seems to follow the expected inverse quadratic scaling with SNR, which serves as a sanity check that the sampler is working correctly.

Figure 4.7



(a) Reconstruction of waveform 11. Contrary to waveform 5, the wavelet inside the late inspiral section was identified despite it being quieter than the undetected wavelet in 5.



(b) Reconstruction of waveform 13. The detected wavelet occurs at a similar point in the waveform as the missed wavelet in Fig. 4.3.

Figure 4.8: Reconstruction of the signals with wavelets near the apparent ET wavelet detection threshold of $\text{SNR} \approx 1.5$ as seen in the left panel of Fig. 4.7c.

The right panel of Fig. 4.7c shows the mismatches for the different runs. They are all quite low, indicating that both the IMR and wavelet components were reconstructed well. Waveforms with missing wavelets do not appear to differ in mismatch compared to correctly reconstructed waveforms. This is consistent with the discussion from earlier about the missing wavelets, generally, having a low SNR. The mismatch of the data also seems to generally obey the inverse quadratic scaling with waveform SNR.

4.4 Mock f -mode Resonance

In this section an f -mode resonance will be simulated by injecting a loud, high-frequency wavelet at the peak of a NS-NS merger into LIGO detector noise. In reality, the BNS IMR signal stays in the detector band for ~ 3 minutes, but this would take far too long to perform a PE on for this thesis. The signal length is thus restricted to 10 s. As BNS mergers occur at high frequencies, the sampling rate was increased to 4096 Hz, which is the highest rate available for the LIGO detectors in Bilby. The injection parameters and recovery priors can be found in Table 4.3.

Injection Value		Recovery Prior				
variable	value	variable	unit	type	min	max
m_1	1.35 $M_\odot$	$\mathcal{M}$	$M_\odot$	UniformInComponentsChirpMass	1	5
m_2	1.35 $M_\odot$	q	-	UniformInComponentsMassRatio	0.1	1
a_1	0.58	$a_{1,2}$	-	Uniform	0	0.99
a_2	0.57					
Λ_1	300	$\Lambda_{1,2}$	-	Uniform	0	500
Λ_2	200					
θ_1	1.67	$\theta_{1,2}$	rad.	Uniform	0	π
θ_2	2.25					
ϕ_{12}	6.21	ϕ_{12}	rad.	Uniform	0	2π
ϕ_{jl}	2.49	ϕ_{jl}	rad.	Uniform	0	2π
SNR	2.0	SNR	-	Transdim.Cond.DescendingOrderStatPriorSNR	0	5
σ	0.01	σ	-	Uniform	0.01	0.08
t_0	10 s	t_0	s	Uniform	0	10
f	1545 Hz	f	Hz	Uniform	0	2048
ϕ	0	ϕ	rad.	Uniform	0	2π
Extrinsic						
d_L	50 Mpc					
ra	5.65					
dec	0.02					
ι	4.61					
ψ	1.02					
ϕ_c	3.25					

Table 4.3: Injection parameters and recovery priors of the BNS merger with mock f -mode resonance.

The wavelet was missed by the sampler, but the BNS parameters were constrained. This is an interesting result as the very first run I performed was able to find the resonance, but not the BNS parameters. The result shown here, however, was the only one I was able to reproduce accross several runs. The first run was also performed at an incorrect sampling rate of 8192 Hz and with 0.1, 0.2 as injections for Λ_1, Λ_2 , so I am not too sure about the reliability of the first result. These values are not too realistic when comparing them to the GW170817 prior I based the final injection values on. It is still remarkable that a signal with a relatively low SNR of 9.60 was reconstructed quite well as can be seen in Fig. 4.9.




Figure 4.9: Reconstruction of a mock f -mode resonance in a BNS merger for a 10s signal. The BNS was identified but the wavelet was missed despite its relatively large strain.

4.5 Glitch Identification

A common nuisance in GW data analysis is the presence of detector glitches. These are noise transients that cannot be attributed to a source. A particularly famous glitch occurred during the first detection of a BNS merger in the LIGO-Livingston detector (Abbott et al., 2017). This merger is known as GW170817. The GW170817 signal was cleaned by using **BayesWave** to model the glitch, after which it was subtracted from the data. Current glitch cleaning methods follow a similar procedure that separates the glitch and CBC inferences. Recently, machine-learning based approaches have been suggested to jointly infer glitch and CBC parameters (Malz and Veitch, 2025). **BayesWave** has also been used to simultaneously model glitches and CBC signals with wavelets (Chatziioannou et al., 2021). In this approach, the CBC signal wavelets are considered coherent across the detector network, whereas glitches do not correlate between detectors. The reconstructed coherent signal can then be used downstream for PE.

Inspired by these concepts, we could attempt to use the trans-dimensional model developed in this thesis to simultaneously model a simple glitch and perform a PE on the CBC signal. Currently, the code developed for this thesis is unable to consider different signal models across detectors, so having an incoherent component to a signal is not possible. The signal was thus only injected into LIGO-Livingston detector noise. An unrealistically large glitch, modelled by a single wavelet with an $\text{SNR} = 100$, was intentionally chosen to also test the joint method's ability to




Figure 4.10: The observed GW170817 signal in the LIGO-Livingston detector containing a loud glitch.

analyse signals of vastly different amplitudes at once. To ensure that the sampler would not content itself with finding the loud wavelet, the stopping criterion for **dynesty** was set to $\Delta \ln \hat{Z} = 0.001$ instead of the standard value of 0.1.

The sampler was able to correctly identify and reconstruct the glitch, as well as the BBH IMR signal. The amplitude of the glitch is underestimated by the reconstruction.




Figure 4.11: Reconstruction of a BBH merger with a large glitch. Both the glitch and BBH were identified, despite the large difference in strain.

To assess the influence of this glitch on the final posterior distribution, an additional PE with the same settings is performed on the BBH IMR signal without the glitch. The posteriors are compared in Fig. 4.12. The glitch has mainly influenced the chirp mass and orbital inclination posteriors with no appreciable difference in the other parameter posteriors.

4.6 Remark on Sampling Priors

As explained earlier, the extrinsic parameters were kept fixed for all of the above runs for two reasons. Firstly, this reduces the parameter space to search in, speeding up the computation. The main focus of this thesis is to combine modelled and unmodelled data-analysis techniques, not to perform highly realistic PEs. Secondly, in a real-world setting one could imagine an initial modelled search to have somewhat restricted the parameter space of the waveform. It could then be argued that the intrinsic parameters of the system should also be restricted instead of using a uniform prior distribution as I did in this work. I briefly experimented with this by using a truncated Gaussian distribution, centred on the injection values, within the bounds of the uniform priors used. However, the runtime of the PEs were increased instead of decreased. These truncated Gaussian priors are much more complex than simple uniform distributions, making them both more expensive to draw samples from and update throughout the sampling process. For this reason, I used uniform priors for the intrinsic parameters.




Figure 4.12: Comparison of the parameter posteriors of a BBH merger with and without the large glitch. The glitch and BBH parameters were inferred simultaneously. The glitch does not appear to influence the posteriors drastically.

Conclusion and Discussion

5.1 Conclusion

So far, the analysis of GW data has been split up into two main categories: **model-agnostic** searches, with **BayesWave** as the main tool, and **modelled** searches using **Bilby** or **LAL**. A modelled search library like **Bilby** can be used downstream of an initial **BayesWave** reconstruction for parameter estimation, but no simple method to jointly perform a modelled and unmodelled analysis of GW data exists, to my knowledge. A method to analyse GW signals in such a manner has been developed and tested in this thesis. Unmodelled features are folded into a modelled waveform as a supersposition of sine-Gaussian wavelets. The transdimensional toolkit for the analysis of these wavelets is provided by **tBilby**. This builds upon the existing **Bilby** framework, which I modified slightly to allow for the joint inference of these wavelets and the CBC parameters. This model was then applied to several cases where such a joint analysis may be useful. In all cases, the extrinsic parameters of the CBCs were fixed and only the intrinsic + wavelet parameters were sampled.

From the initial tests of the implemented wavelets, it appears that the phase and width of the wavelets are nearly insensitive to the wavelet SNR in LIGO detectors. The former's poor restriction is not too much of an issue in the context of resonances as the frequency is the most important parameter to infer. The width plays a secondary role for resonances, indicating the resonance's dampening time. Its apparent SNR independence might form an issue in LIGO, but the succesful ET PEs do restrict σ quite well, so we may stay optimistic for the future. Generally, **Bilby** is capable of identifying the injected wavelets. When testing the recovery of a random number of injected wavelets with **tBilby**, the number of injected wavelets seems to be systematically underestimated.

A limited number of IMR + wavelet injections into LIGO noise were briefly experimented with, but it appears that unrealistically high SNRs ($\gtrsim 260$) are necessary to reconstruct wavelets on top of an IMR signal. Most of the signals analysed in this thesis were massive BBH mergers as these signals are relatively short and easy to analyse. They do mostly occur at lower frequencies, near the bottom of the LIGO sensitivity band. For these reasons, the subsequent analyses were performed on ET noise. Here, the joint PE is able to reliably reconstruct wavelets with an SNR $\gtrsim 1.5$ on top of an IMR signal. The mismatches are small for all reconstructions, with the largest mismatch being ~ 0.025 . The missed wavelets do not have too much of an impact on the quality of the reconstruction, due to their low SNR. In general, it can be stated that the joint inference for simulated ET data works remarkably well. This analysis also showed a caveat of the ordered SNR priors: loud features may be missed if they are not detected before quieter ones are.

f-mode resonances have gained quite some interest in literature, which is why a mock version was tested in a BNS merger. As these signals can stay in the detector band for several minutes, the injection was restricted to the final 10 s. A single, loud wavelet was then injected at the peak of the merger. This wavelet was not recovered, but the BNS merger signal was.

Finally, a simulated glitch was injected into the LIGO-Livingston detector along with a BBH IMR signal. The glitch was modelled by a single wavelet with an SNR = 100. This is exceptionally loud, but it also serves as a test of the developed method's ability to simultaneously detect loud and quiet(er) signals. Both the glitch and BBH signals were detected, although the glitch amplitude was underestimated. The glitch did not appear to significantly influence the BBH parameter posteriors.

While this analysis mainly represents highly idealised scenarios and is far from complete, this work does serve as a proof of concept that joint modelled and unmodelled searches are possible with the current GW analysis tools. However, I believe that the robustness of this method still needs a thorough analysis.

5.2 Discussion

The joint method developed in this thesis is new and many different directions could have been taken to explore it. I have decided to test it in a controlled environment on purely simulated data. In order to minimise the amount of hand-tuning, the test on random injections was performed, but the injection priors for the wavelets were the result of previous tests that led to wavelets that looked reasonable. This is an unfortunate result of incorporating unmodelled features into waveforms: there is no model to tell us what they should look like! The set of random injections was also fairly limited in size, so the systematics of the detectability of wavelets folded into IMR signals are still quite rough, but the obtained results serve as an initial indication of the performance of this method.

Unfortunately, the mock f -mode resonance injected into a BNS merger signal was not detected. It would be interesting to see whether analysing the full in detector-band signal could make a difference. Such a computation would take up a lot of time, but there exist techniques to speed up this computation, such as relative binning (Zackay et al., 2018). These were not experimented with in this thesis as the focus was on developing the new inference method, not optimising it for specific use cases. I decided to show the result that missed the mock f -mode resonance as this was the only one I was able to reproduce. An initial experiment did identify it, but missed the BNS parameters. This was run with values of 0.1 and 0.2 for the dimensionless tidal deformabilities Λ_1 and Λ_2 . The final runs used 300 and 200, respectively. These are more realistic values and are consistent with the observations of GW170817 (Abbott et al., 2017; Chatziioannou, 2020), which I based the injection values on. Digging deeper into the detectability of wavelets in BNS IMR signals could be an interesting avenue to explore further as this thesis has mainly worked with BBH mergers for convenience.

The injected wavelet used to simulate a glitch did not overlap much with the merger section of the BBH IMR signal, meaning that the experiment does not represent the worst-case scenario. This experiment mainly served as an illustration of an alternate use case of the developed method.

In all of the experiments performed, the injected wavelets were fairly simple. In reality however, a resonance will likely have a much more complex structure. For example: it may be red- and blueshifted due to the orbital motion of binary. This was not taken into account in this work. This varying doppler shift could provide an interesting way to distinguish resonances from detector glitches, in addition to the coherence accross a detector network.

This work has also only used the nested sampler `dynesty` to perform PEs. There was a bit of experimentation with the sampler settings, such as changing the `rwalk` proposal to `rslice` (random slicing), but this led to abysmal results. Other samplers such as `ptemcee` could also have been used for this work, but as a nested sampler provides direct access to the evidence, MCMC samplers were not employed.

Future Work

While working on this thesis, I realised that there are many new avenues that can still be pursued with the developed method as well as additional tests that could be performed. In this chapter I will discuss these, and hopefully someone may find any one interesting and explore it further¹. These are rough ideas and will have to be worked out in more detail should someone decide to pursue any one of them.

6.1 Physics

Firstly, working towards the analysis of **real GW signals** would be the next obvious step with the developed method. Although I am sceptical about the ability to discern wavelets from noise in current ground-based detectors, it is definitely worthwhile trying and refining this technique for current detectors. The analysis of **numerical relativity** (NR) signals would be a good stepping stone to the analysis of real signals. These simulations are computationally very expensive, but they provide a realistic waveform that can include resonances (Gold et al., 2012). There is a database for NR waveforms named CoRe (Gonzalez et al., 2023), but I was unable to find suitable waveforms although waveforms with resonances should be available in it.

I have mainly injected single wavelets in this work, but **more complex wavelets** formed by a superposition around the same point in time could be interesting to study further. Real-world features like glitches are usually more complex than a single wavelet and a wavelet superposition could simulate these. Alternatively, an actual stretch of “GW-less” detector data containing a glitch could be used. A CBC signal can then be injected into this and a similar analysis to this work performed.

More **complex component functions** than wavelets could be used for injection and recovery. These could encode additional physics, such as a time-varying doppler shift of the resonance. I am not sure whether such an effect will be detectable, but it could be interesting to study. More complex component functions will likely require additional parameters, which will incur a larger Occam penalty. The model will likely also lose the generality that wavelets allow in favour of specificity.

6.2 Inference

Different likelihoods could also be experimented with. I have only used Bilby’s standard Gaussian noise likelihood in this work, but Bilby also contains others. There is, for example, a power spectral density (“burst”) likelihood that marginalises over phase and preserves frequency content (Ashton et al., 2019). As phase seemed hard to constrain with current detectors, this likelihood would allow for a simplification of the resonance model by fixing the phase shift. The “burst” likelihood could be employed for the unmodelled part of the signal together with a Gaussian noise likelihood that accounts for the modelled aspects.

A shortcoming of the way I coded the joint model, is that it is unable to employ **separate signal models per detector**. Allowing for this, would mean that coherent (e.g. resonances) and incoherent (glitches) features could be separated in a detector network.

Modified Bayes’ Theorem: A recurring issue in this work was the systematic underestimation of the number of wavelets in a signal. The regularisation imposed by the Occam factor is hardwired into the Bayesian framework, but we could modify Bayes Theorem to relax this regularisation. We could, for example, raise the likelihood to a power β . In (Grünwald, 2012) this is interpreted as a learning rate,

¹I will take this opportunity to finally say “*This is left as an exercise to the reader*”.

but as physicists we may be more used to thinking of it as an inverse temperature. Bayes' theorem now reads:

$$P_{\beta}(\Theta|\mathbf{D}) = \frac{\mathcal{L}^{\beta}(\mathbf{D}|\Theta)\pi(\Theta)}{Z_{\beta}} \quad (6.1)$$

with²

$$Z_{\beta} = \int \mathcal{L}^{\beta}(\mathbf{D}|\Theta)\pi(\Theta)d\Theta \quad (6.2)$$

This allows us to place more - or less - weight on the likelihood, which can relax the importance of the Occam factor's regularisation. In the limit for $\beta \rightarrow \infty$, the prior becomes irrelevant to the posterior and the MAP estimation reduces to maximum likelihood estimation (MLE).

There have been other modifications proposed to Bayes' theorem such as the Generalised Bayes Theorem by (Bissiri et al., 2016), where the likelihood is replaced by a general loss function $l(\mathbf{D}|\Theta)$:

$$P(\Theta|\mathbf{D}) = \frac{e^{-l(\mathbf{D}|\Theta)}\pi(\Theta)}{Z} \quad (6.3)$$

$$Z = \int e^{-l(\mathbf{D}|\Theta)}\pi(\Theta)d\Theta \quad (6.4)$$

For $l = -\beta \ln \mathcal{L}$ we recover the initial modified Bayes' theorem from Eq. (6.1).

A modified Bayes' theorem would allow for a more tunable solution, but the addition of a hyperparameter β brings its own problems such as finding its ideal value. Great care should also be taken when working with a modified Bayes' theorem in order to avoid overfitting.

²Note that the calligraphic $\mathcal{Z}$ is not used as Z no longer represents the evidence. Instead, it is more akin to a partition function.

Bibliography

- 't Hooft, G. (July 1976). “Symmetry Breaking through Bell-Jackiw Anomalies”. In: *Physical Review Letters* 37.1, pp. 8–11. ISSN: 0031-9007. DOI: [10.1103/physrevlett.37.8](https://doi.org/10.1103/physrevlett.37.8).
- Aad, G. et al. (Sept. 2012). “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716.1, pp. 1–29. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020).
- Aasi, J et al. (Mar. 2015). “Advanced LIGO”. In: *Classical and Quantum Gravity* 32.7, p. 074001. ISSN: 1361-6382. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001).
- Abbott, B.P. et al. (Oct. 2017). “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Physical Review Letters* 119.16, p. 161101. ISSN: 1079-7114. DOI: [10.1103/physrevlett.119.161101](https://doi.org/10.1103/physrevlett.119.161101).
- Abel, C. et al. (Feb. 2020). “Measurement of the Permanent Electric Dipole Moment of the Neutron”. In: *Physical Review Letters* 124.8, p. 081803. ISSN: 1079-7114. DOI: [10.1103/physrevlett.124.081803](https://doi.org/10.1103/physrevlett.124.081803).
- Acernese, F et al. (Dec. 2014). “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2, p. 024001. ISSN: 1361-6382. DOI: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001).
- Aerts, C. et al. (2010). *Asteroseismology*. Springer Netherlands. ISBN: 9781402058035. DOI: [10.1007/978-1-4020-5803-5](https://doi.org/10.1007/978-1-4020-5803-5).
- Akutsu, T et al. (Aug. 2020). “Overview of KAGRA: Detector design and construction history”. In: *Progress of Theoretical and Experimental Physics* 2021.5. ISSN: 2050-3911. DOI: [10.1093/ptep/ptaa125](https://doi.org/10.1093/ptep/ptaa125).
- Andersson, Nils (Nov. 2019). *Gravitational-Wave Astronomy: Exploring the Dark Side of the Universe- Exploring the Dark Side of the Universe*. ISBN: 9780198568032. DOI: [10.1093/oso/9780198568032.001.0001](https://doi.org/10.1093/oso/9780198568032.001.0001).
- Andersson, Nils et al. (May 1996). “On the Oscillation Spectra of Ultracompact Stars: an Extensive Survey of Gravitational-Wave Modes”. In: *The Astrophysical Journal* 462, p. 855. ISSN: 1538-4357. DOI: [10.1086/177199](https://doi.org/10.1086/177199).
- Annala, Eemeli et al. (Dec. 2023). “Strongly interacting matter exhibits deconfined behavior in massive neutron stars”. In: *Nature Communications* 14.1. ISSN: 2041-1723. DOI: [10.1038/s41467-023-44051-y](https://doi.org/10.1038/s41467-023-44051-y).
- Asali, Yasmeen et al. (July 2020). “Probing resonant excitations in exotic compact objects via gravitational waves”. In: *Physical Review D* 102.2, p. 024016. ISSN: 2470-0029. DOI: [10.1103/physrevd.102.024016](https://doi.org/10.1103/physrevd.102.024016).
- Ashton, Gregory et al. (Apr. 2019). “Bilby: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy”. In: *The Astrophysical Journal Supplement Series* 241.2, p. 27. ISSN: 1538-4365. DOI: [10.3847/1538-4365/ab06fc](https://doi.org/10.3847/1538-4365/ab06fc).
- Babak, Stanislav et al. (2021). *LISA Sensitivity and SNR Calculations*. DOI: [10.48550/ARXIV.2108.01167](https://doi.org/10.48550/ARXIV.2108.01167).
- Bissiri, P. G. et al. (Feb. 2016). “A General Framework for Updating Belief Distributions”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 78.5, pp. 1103–1130. ISSN: 1467-9868. DOI: [10.1111/rssb.12158](https://doi.org/10.1111/rssb.12158).
- Britannica, the Editors of Encyclopaedia (2024). “Michelson interferometer”. In: *Encyclopedia Britannica*.
- Buchner, Johannes (Sept. 2014). “A statistical test for Nested Sampling algorithms”. In: *Statistics and Computing* 26.1–2, pp. 383–392. ISSN: 1573-1375. DOI: [10.1007/s11222-014-9512-y](https://doi.org/10.1007/s11222-014-9512-y).
- (Aug. 2019). “Collaborative Nested Sampling: Big Data versus Complex Physical Models”. In: *Publications of the Astronomical Society of the Pacific* 131.1004, p. 108005. ISSN: 1538-3873. DOI: [10.1088/1538-3873/aae7fc](https://doi.org/10.1088/1538-3873/aae7fc).
- Burnham, K.P. and D.R. Anderson (2004). *Model Selection and Multimodel Inference*. Springer New York. ISBN: 9780387953649. DOI: [10.1007/b97636](https://doi.org/10.1007/b97636).
- Buscicchio, Riccardo et al. (Oct. 2019). “Label switching problem in Bayesian analysis for gravitational wave astronomy”. In: *Physical Review D* 100.8, p. 084041. ISSN: 2470-0029. DOI: [10.1103/physrevd.100.084041](https://doi.org/10.1103/physrevd.100.084041).

- Caplan, M. E. and C. J. Horowitz (Oct. 2017). “Colloquium: Astromaterial science and nuclear pasta”. In: *Reviews of Modern Physics* 89.4, p. 041002. ISSN: 1539-0756. DOI: [10.1103/revmodphys.89.041002](https://doi.org/10.1103/revmodphys.89.041002).
- Chandrasekhar, S. (Mar. 1970). “Solutions of Two Problems in the Theory of Gravitational Radiation”. In: *Physical Review Letters* 24.11, pp. 611–615. ISSN: 0031-9007. DOI: [10.1103/physrevlett.24.611](https://doi.org/10.1103/physrevlett.24.611).
- Chatziioannou, Katerina (Nov. 2020). “Neutron-star tidal deformability and equation-of-state constraints”. In: *General Relativity and Gravitation* 52.11. ISSN: 1572-9532. DOI: [10.1007/s10714-020-02754-3](https://doi.org/10.1007/s10714-020-02754-3).
- Chatziioannou, Katerina et al. (Feb. 2021). “Modeling compact binary signals and instrumental glitches in gravitational wave data”. In: *Physical Review D* 103.4, p. 044013. ISSN: 2470-0029. DOI: [10.1103/physrevd.103.044013](https://doi.org/10.1103/physrevd.103.044013).
- Cornish, Neil et al. (Sept. 2011). “Gravitational wave tests of general relativity with the parameterized post-Einsteinian framework”. In: *Physical Review D* 84.6, p. 062003. ISSN: 1550-2368. DOI: [10.1103/physrevd.84.062003](https://doi.org/10.1103/physrevd.84.062003).
- Cornish, Neil J and Tyson B Littenberg (June 2015). “Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches”. In: *Classical and Quantum Gravity* 32.13, p. 135012. ISSN: 1361-6382. DOI: [10.1088/0264-9381/32/13/135012](https://doi.org/10.1088/0264-9381/32/13/135012).
- Cowling, T. G. (Dec. 1941). “The Non-radial Oscillations of Polytropic Stars”. In: *Monthly Notices of the Royal Astronomical Society* 101.8, pp. 367–375. ISSN: 1365-2966. DOI: [10.1093/mnras/101.8.367](https://doi.org/10.1093/mnras/101.8.367).
- D’Agostini, Giulio (2016). *The Waves and the Sigmas (To Say Nothing of the 750 GeV Mirage)*. DOI: [10.48550/ARXIV.1609.01668](https://doi.org/10.48550/ARXIV.1609.01668).
- Davis, D et al. (June 2021). “LIGO detector characterization in the second and third observing runs”. In: *Classical and Quantum Gravity* 38.13, p. 135014. ISSN: 1361-6382. DOI: [10.1088/1361-6382/abfd85](https://doi.org/10.1088/1361-6382/abfd85).
- Deppe, Nils et al. (2024). *Echoes from Beyond: Detecting Gravitational Wave Quantum Imprints with LISA*. DOI: [10.48550/ARXIV.2411.05645](https://doi.org/10.48550/ARXIV.2411.05645).
- Einstein, Albert (Jan. 1916). “Näherungsweise Integration der Feldgleichungen der Gravitation”. In: *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pp. 688–696. URL: <https://ui.adsabs.harvard.edu/abs/1916SPAW.....688E>.
- Emden, R. (1907). *Gaskugeln: anwendungen der mechanischen wärmetheorie auf kosmologische und meteorologische probleme*. Gaskugeln: Anwendungen der mechanischen wärmetheorie auf kosmologische und meteorologische probleme. B.G. Teubner. URL: <https://books.google.be/books?id=nMwxAAAAAAAJ>.
- Feroz, F. et al. (Oct. 2009). “MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics”. In: *Monthly Notices of the Royal Astronomical Society* 398.4, pp. 1601–1614. ISSN: 1365-2966. DOI: [10.1111/j.1365-2966.2009.14548.x](https://doi.org/10.1111/j.1365-2966.2009.14548.x).
- Feroz, Farhan and John Skilling (2013). “Exploring multi-modal distributions with nested sampling”. In: *AIP Conference Proceedings*. AIP. DOI: [10.1063/1.4819989](https://doi.org/10.1063/1.4819989).
- Freese, Katherine et al. (May 2016). “Dark stars: a review”. In: *Reports on Progress in Physics* 79.6, p. 066902. ISSN: 1361-6633. DOI: [10.1088/0034-4885/79/6/066902](https://doi.org/10.1088/0034-4885/79/6/066902).
- Friedman, J. L. and B. F. Schutz (May 1978). “Secular instability of rotating Newtonian stars”. In: *The Astrophysical Journal* 222, p. 281. ISSN: 1538-4357. DOI: [10.1086/156143](https://doi.org/10.1086/156143).
- Gold, Roman et al. (Dec. 2012). “Eccentric binary neutron star mergers”. In: *Physical Review D* 86.12, p. 121501. ISSN: 1550-2368. DOI: [10.1103/physrevd.86.121501](https://doi.org/10.1103/physrevd.86.121501).
- Gonzalez, Alejandra et al. (Mar. 2023). “Second release of the CoRe database of binary neutron star merger waveforms”. In: *Classical and Quantum Gravity* 40.8, p. 085011. ISSN: 1361-6382. DOI: [10.1088/1361-6382/acc231](https://doi.org/10.1088/1361-6382/acc231).
- Grünwald, Peter (2012). “The Safe Bayesian: Learning the Learning Rate via the Mixability Gap”. In: *Algorithmic Learning Theory*. Springer Berlin Heidelberg, pp. 169–183. ISBN: 9783642341069. DOI: [10.1007/978-3-642-34106-9_16](https://doi.org/10.1007/978-3-642-34106-9_16).
- Guerra, Davide et al. (Sept. 2019). “Axion boson stars”. In: *Journal of Cosmology and Astroparticle Physics* 2019.09, pp. 061–061. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2019/09/061](https://doi.org/10.1088/1475-7516/2019/09/061).
- Handley, W. J. et al. (Apr. 2015a). “polychord: nested sampling for cosmology”. In: *Monthly Notices of the Royal Astronomical Society: Letters* 450.1, pp. L61–L65. ISSN: 1745-3925. DOI: [10.1093/mnrasl/slv047](https://doi.org/10.1093/mnrasl/slv047).
- (Sept. 2015b). “polychord: next-generation nested sampling”. In: *Monthly Notices of the Royal Astronomical Society* 453.4, pp. 4385–4399. ISSN: 1365-2966. DOI: [10.1093/mnras/stv1911](https://doi.org/10.1093/mnras/stv1911).
- Hawking, S. W. (Mar. 1974). “Black hole explosions?” In: *Nature* 248.5443, pp. 30–31. ISSN: 1476-4687. DOI: [10.1038/248030a0](https://doi.org/10.1038/248030a0).
- Hergt, L. T. et al. (June 2021). “Bayesian evidence for the tensor-to-scalar ratio and neutrino masses m_ν : Effects of uniform versus logarithmic priors”. In: *Physical Review D* 103.12, p. 123511. ISSN: 2470-0029. DOI: [10.1103/physrevd.103.123511](https://doi.org/10.1103/physrevd.103.123511).

- Hulse, R. A. and J. H. Taylor (Jan. 1975). “Discovery of a pulsar in a binary system”. In: *The Astrophysical Journal* 195, p. L51. ISSN: 1538-4357. DOI: [10.1086/181708](https://doi.org/10.1086/181708).
- Ikeda, Taishi et al. (Sept. 2021). “Black-hole microstate spectroscopy: Ringdown, quasinormal modes, and echoes”. In: *Physical Review D* 104.6, p. 066021. ISSN: 2470-0029. DOI: [10.1103/physrevd.104.066021](https://doi.org/10.1103/physrevd.104.066021).
- Ivanenko, D. D. and D. F. Kurdgelaidze (1967). “Hypothesis concerning quark stars”. In: *Astrophysics* 1.4, pp. 251–252. ISSN: 1573-8191. DOI: [10.1007/bf01042830](https://doi.org/10.1007/bf01042830).
- Janquart, Justin et al. (May 2023). “Analyses of overlapping gravitational wave signals using hierarchical subtraction and joint parameter estimation”. In: *Monthly Notices of the Royal Astronomical Society* 523.2, pp. 1699–1710. ISSN: 1365-2966. DOI: [10.1093/mnras/stad1542](https://doi.org/10.1093/mnras/stad1542).
- Jaranowski, Piotr and Andrzej Królak (Mar. 2012). “Gravitational-Wave Data Analysis. Formalism and Sample Applications: The Gaussian Case”. In: *Living Reviews in Relativity* 15.1. ISSN: 1433-8351. DOI: [10.12942/lrr-2012-4](https://doi.org/10.12942/lrr-2012-4).
- Jetzer, P (Nov. 1992). “Boson stars”. In: *Physics Reports* 220.4, pp. 163–227. ISSN: 0370-1573. DOI: [10.1016/0370-1573\(92\)90123-h](https://doi.org/10.1016/0370-1573(92)90123-h).
- Kokkotas, K. D. and B. F. Schutz (Mar. 1992). “W-modes: a new family of normal modes of pulsating relativistic stars”. In: *Monthly Notices of the Royal Astronomical Society* 255.1, pp. 119–128. ISSN: 1365-2966. DOI: [10.1093/mnras/255.1.119](https://doi.org/10.1093/mnras/255.1.119).
- Kokkotas, Kostas D. and Bernd G. Schmidt (Sept. 1999). “Quasi-Normal Modes of Stars and Black Holes”. In: *Living Reviews in Relativity* 2.1. ISSN: 1433-8351. DOI: [10.12942/lrr-1999-2](https://doi.org/10.12942/lrr-1999-2).
- Koposov, Sergey et al. (2024). *joshspeagle/dynesty: v2.1.4*. DOI: [10.5281/ZENODO.12537467](https://doi.org/10.5281/ZENODO.12537467).
- Kraav, Kirill Y. et al. (Feb. 2024). “Instability windows of relativistic r -modes”. In: *Physical Review D* 109.4, p. 043012. ISSN: 2470-0029. DOI: [10.1103/physrevd.109.043012](https://doi.org/10.1103/physrevd.109.043012).
- LIGO Scientific Collaboration (2018). *LIGO Algorithm Library*. DOI: [10.7935/GT1W-FZ16](https://doi.org/10.7935/GT1W-FZ16).
- MacKay, David J. C. (2002). *Information Theory, Inference & Learning Algorithms*. USA: Cambridge University Press. ISBN: 0521642981.
- Madsen, Mark S. and Andrew R. Liddle (Nov. 1990). “The cosmological formation of boson stars”. In: *Physics Letters B* 251.4, pp. 507–510. ISSN: 0370-2693. DOI: [10.1016/0370-2693\(90\)90788-8](https://doi.org/10.1016/0370-2693(90)90788-8).
- Maggiore, Michele (Oct. 2007). *Gravitational Waves: Volume 1: Theory and Experiments*. Oxford University Press/Oxford. ISBN: 9780191717666. DOI: [10.1093/acprof:oso/9780198570745.001.0001](https://doi.org/10.1093/acprof:oso/9780198570745.001.0001).
- Malz, Ann-Kristin and John Veitch (2025). *Joint inference for gravitational wave signals and glitches using a data-informed glitch model*. DOI: [10.48550/ARXIV.2505.00657](https://doi.org/10.48550/ARXIV.2505.00657).
- Mathur, S.D. (June 2005). “The fuzzball proposal for black holes: an elementary review”. In: *Fortschritte der Physik* 53.7–8, pp. 793–827. ISSN: 1521-3978. DOI: [10.1002/prop.200410203](https://doi.org/10.1002/prop.200410203).
- Mayerson, Daniel R. (Nov. 2020). “Fuzzballs and observations”. In: *General Relativity and Gravitation* 52.12. ISSN: 1572-9532. DOI: [10.1007/s10714-020-02769-w](https://doi.org/10.1007/s10714-020-02769-w).
- Mazur, Pawel O. and Emil Mottola (2001). “Gravitational Condensate Stars: An Alternative to Black Holes”. In: DOI: [10.48550/ARXIV.GR-QC/0109035](https://doi.org/10.48550/ARXIV.GR-QC/0109035).
- McDermott, P. N. et al. (May 1983). “Nonradial g-mode oscillations of warm neutron stars”. In: *The Astrophysical Journal* 268, p. 837. ISSN: 1538-4357. DOI: [10.1086/161006](https://doi.org/10.1086/161006).
- Moore, Christopher J. et al. (2014). “Gravitational-wave sensitivity curves”. In: DOI: [10.48550/ARXIV.1408.0740](https://doi.org/10.48550/ARXIV.1408.0740).
- Mukherjee, Pia et al. (Jan. 2006). “A Nested Sampling Algorithm for Cosmological Model Selection”. In: *The Astrophysical Journal* 638.2, pp. L51–L54. ISSN: 1538-4357. DOI: [10.1086/501068](https://doi.org/10.1086/501068).
- Neal, Radford M. (June 2003). “Slice sampling”. In: *The Annals of Statistics* 31.3. ISSN: 0090-5364. DOI: [10.1214/aos/1056562461](https://doi.org/10.1214/aos/1056562461).
- Oppenheimer, J. R. and G. M. Volkoff (Feb. 1939). “On Massive Neutron Cores”. In: *Physical Review* 55.4, pp. 374–381. ISSN: 0031-899X. DOI: [10.1103/physrev.55.374](https://doi.org/10.1103/physrev.55.374).
- Owen, Benjamin J. (June 1996). “Search templates for gravitational waves from inspiraling binaries: Choice of template spacing”. In: *Physical Review D* 53.12, pp. 6749–6761. ISSN: 1089-4918. DOI: [10.1103/physrevd.53.6749](https://doi.org/10.1103/physrevd.53.6749).
- Peccei, R. D. and Helen R. Quinn (June 1977). “CPConservation in the Presence of Pseudoparticles”. In: *Physical Review Letters* 38.25, pp. 1440–1443. ISSN: 0031-9007. DOI: [10.1103/physrevlett.38.1440](https://doi.org/10.1103/physrevlett.38.1440).
- Pitte, Chantal et al. (2024). *Exploring the no-hair theorem with LISA*. DOI: [10.48550/ARXIV.2406.14552](https://doi.org/10.48550/ARXIV.2406.14552).
- Prakash, Madappa (2013). *Neutron Stars and the EOS*. DOI: [10.48550/ARXIV.1307.0397](https://doi.org/10.48550/ARXIV.1307.0397).
- Pratten, Geraint et al. (May 2020). “Gravitational-wave asteroseismology with fundamental modes from compact binary inspirals”. In: *Nature Communications* 11.1. ISSN: 2041-1723. DOI: [10.1038/s41467-020-15984-5](https://doi.org/10.1038/s41467-020-15984-5).

- Punturo, M et al. (Apr. 2010). “The third generation of gravitational wave observatories and their science reach”. In: *Classical and Quantum Gravity* 27.8, p. 084007. ISSN: 1361-6382. DOI: [10.1088/0264-9381/27/8/084007](https://doi.org/10.1088/0264-9381/27/8/084007).
- Rafferty, John P. (n.d.). “bilby”. In: *Encyclopedia Britannica* ().
- Reitze, David et al. (2019). “Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO”. In: DOI: [10.48550/ARXIV.1907.04833](https://doi.org/10.48550/ARXIV.1907.04833).
- Rowan, Dominick M. and John D. Roberts (2025). *PromoPlot: Covering open-access fees by filling wasted space in corner plots*. DOI: [10.48550/ARXIV.2503.24254](https://doi.org/10.48550/ARXIV.2503.24254).
- Skilling, John (Dec. 2006). “Nested sampling for general Bayesian computation”. In: *Bayesian Analysis* 1.4. ISSN: 1936-0975. DOI: [10.1214/06-ba127](https://doi.org/10.1214/06-ba127).
- (2012). “Bayesian computation in big spaces-nested sampling and Galilean Monte Carlo”. In: *AIP Conference Proceedings*. AIP. DOI: [10.1063/1.3703630](https://doi.org/10.1063/1.3703630).
- Speagle, Joshua S (Feb. 2020). “dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences”. In: *Monthly Notices of the Royal Astronomical Society* 493.3, pp. 3132–3158. ISSN: 1365-2966. DOI: [10.1093/mnras/staa278](https://doi.org/10.1093/mnras/staa278).
- Suvorov, Arthur G. et al. (Nov. 2024). “Premerger Phenomena in Neutron Star Binary Coalescences”. In: *Universe* 10.12, p. 441. ISSN: 2218-1997. DOI: [10.3390/universe10120441](https://doi.org/10.3390/universe10120441).
- Tolman, Richard C. (Feb. 1939). “Static Solutions of Einstein’s Field Equations for Spheres of Fluid”. In: *Physical Review* 55.4, pp. 364–373. ISSN: 0031-899X. DOI: [10.1103/physrev.55.364](https://doi.org/10.1103/physrev.55.364).
- Tong, Hui et al. (Jan. 2025). “Transdimensional Inference for Gravitational-wave Astronomy with Bilby”. In: *The Astrophysical Journal Supplement Series* 276.2, p. 50. ISSN: 1538-4365. DOI: [10.3847/1538-4365/ad9deb](https://doi.org/10.3847/1538-4365/ad9deb).
- Wasserstein, Ronald L. and Nicole A. Lazar (Apr. 2016). “The ASA Statement on p-Values: Context, Process, and Purpose”. In: *The American Statistician* 70.2, pp. 129–133. ISSN: 1537-2731. DOI: [10.1080/00031305.2016.1154108](https://doi.org/10.1080/00031305.2016.1154108).
- Weisberg, Joel M. et al. (Oct. 1981). “Gravitational Waves from an Orbiting Pulsar”. In: *Scientific American* 245.4, pp. 74–82. ISSN: 0036-8733. DOI: [10.1038/scientificamerican1081-74](https://doi.org/10.1038/scientificamerican1081-74).
- Zackay, Barak et al. (2018). *Relative Binning and Fast Likelihood Evaluation for Gravitational Wave Parameter Estimation*. DOI: [10.48550/ARXIV.1806.08792](https://doi.org/10.48550/ARXIV.1806.08792).

Bilby Structure




Figure A.1: Structure of the Bilby source code. Image taken from Bilby documentation: <https://lscsoft.docs.ligo.org/bilby/index.html>

In this section, I want to delve a bit deeper into Bilby’s code structure. Most of this information is taken from (Ashton et al., 2019). Bilby is composed of three main modules: **core**, **gw** and **hyper**.

A.1 core package

core contains the code for general Bayesian inference. One can pass a data and a model into the **Likelihood** class. The model’s parameters that need to be inferred are put into priors of the **Prior** class. These define the prior volume within which the sampler will perform the inference. They need not be uniform as the user can define their own priors next to the already implemented ones such as **Normal**, **LogUniform**, **Sine**, ... Constraints can also be added between different priors. Similarly, there are many different **Likelihood** subclasses. The prior and likelihood are then passed to the **run_sampler** function where the user can select what sampler to use as well as set the different hyperparameters for the inference problem. This function outputs an **hdf5** file which contains all output (posterior samples, injected parameters, likelihood evaluations, ...) and can be loaded in using the **Result** class. This class can also perform common tasks such as generating corner plots.

A.2 gw package

This package adds additional functionality to the **core** package through priors and likelihoods that are adapted especially to GW-related inference problems (for example the Gaussian noise likelihood from Eq. (2.16)). It also incorporates current and future detectors as objects of the **Interferometer** class,

into which signals constructed with the `WaveformGenerator` class can be injected. This class takes a source model and parameters as inputs and allows one to calculate strains in both time and frequency domain.

A.3 hyper package

This package was not used in this work, but it contains all tools necessary to perform Bayesian inference of populations.

Corner plots

As corner plots take up a lot of space, I have placed them in this appendix in case you want to reference them. The caption of a plot refers back to the figure that they relate to. The red lines denote the true injected values.




Figure B.1: Fig. 4.1a




Figure B.2: Fig. 4.1b




Figure B.3: Fig. 4.3




Figure B.4: Fig. 4.4




Figure B.5: Fig. 4.6




Figure B.6: Fig. 4.8a




Figure B.7: Fig. 4.8b




Figure B.8: Fig. 4.9




Figure B.9: Fig. 4.11

Gravitational Backreaction

Throughout this thesis I have mentioned energy loss by GWs and posited some formulas for the chirp mass and radiated power. In this appendix I would like to show the derivation of these, starting from the linearised Einstein equations and perturbations on a curved background. This will, much like the SNR derivation, closely follow (Maggiore, 2007).

C.1 General Solution of Linearised Einstein Equations

We will start from the linearised Einstein equations:

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu} \quad (\text{C.1})$$

where $\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h$.

Introducing the Green's function for which $\square_x G(x - x') = \delta^{(4)}(x - x')$, we can write the solution for the linearised equations as:

$$\bar{h}_{\mu\nu}(x) = -\frac{16\pi G}{c^4} \int d^4x' G(x - x') T_{\mu\nu}(x') \quad (\text{C.2})$$

Just as in electromagnetism, the *retarded Green's function* is an appropriate solution:

$$G(x - x') = -\frac{1}{4\pi|\mathbf{x} - \mathbf{x}'|} \delta(x_{\text{ret}}^0 - x'^0) \quad (\text{C.3})$$

where $x'^0 = ct'$, $x_{\text{ret}}^0 = ct_{\text{ret}}$ and

$$t_{\text{ret}} = 1 - \frac{|\mathbf{x} - \mathbf{x}'|}{c} \quad (\text{C.4})$$

Plugging this into Eq. (C.2) and integrating over the history of the source, we find that:

$$\bar{h}_{\mu\nu}(t, \mathbf{x}) = \frac{4G}{c^4} \int d^3x' \frac{1}{|\mathbf{x} - \mathbf{x}'|} T_{\mu\nu} \left(t - \frac{|\mathbf{x} - \mathbf{x}'|}{c}, \mathbf{x}' \right) \quad (\text{C.5})$$

Using the lambda-tensor, we can transform this expression into the Transverse-Traceless (TT) gauge. Far away from the , at a distance r , we can make the following approximation:

$$|\mathbf{x} - \mathbf{x}'| \approx r \left(1 - \frac{\mathbf{x}' \cdot \hat{\mathbf{n}}}{r} \right) \quad (\text{C.6})$$

where $\mathbf{n}$ expresses the line-of-sight vector to the source. Together with the lambda-tensor this approximation can now be used to obtain:

$$h_{ij}^{\text{TT}}(t, \mathbf{x}) = \frac{1}{r} \frac{4G}{c^4} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \int d^3x' T_{kl} \left(t - \frac{r}{c} + \frac{\mathbf{x}' \cdot \hat{\mathbf{n}}}{c}, \mathbf{x}' \right) \quad (\text{C.7})$$

Note that the GW amplitude falls off as $1/r$. In a non-relativistic expansion, we can expand T_{kl} as follows:

$$T_{kl} \left(t - \frac{r}{c} + \frac{\mathbf{x}' \cdot \hat{\mathbf{n}}}{c}, \mathbf{x}' \right) \simeq T_{kl} \left(t - \frac{r}{c}, \mathbf{x}' \right) + \frac{x'^i n^i}{c} \partial_t T_{kl} + \frac{1}{2c^2} x'^i x'^j n^i n^j \partial_t^2 T_{kl} + \dots \quad (\text{C.8})$$

Introducing the moments of the energy-momentum tensor:

$$S^{ij}(t) = \int d^3x T^{ij}(t, \mathbf{x}) \quad (\text{C.9})$$

$$S^{ij,k}(t) = \int d^3x T^{ij}(t, \mathbf{x}) x^k \quad (\text{C.10})$$

$$S^{ij,kl}(t) = \int d^3x T^{ij}(t, \mathbf{x}) x^k x^l \quad (\text{C.11})$$

we can rewrite h_{ij}^{TT} in terms of these moments:

$$h_{ij}^{TT}(t, \mathbf{x}) = \frac{1}{r} \frac{4G}{c^4} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \left[S^{kl} + \frac{1}{c} n_m \dot{S}^{kl,m} + \frac{1}{2c^2} n_m n_p \ddot{S}^{kl,mp} + \dots \right]_{\text{ret}} \quad (\text{C.12})$$

Each additional term in this series contributes an extra

$$\frac{x'^i}{c} \partial_t \sim \frac{1}{c} O\left(\frac{d}{t}\right) \sim O\left(\frac{v}{c}\right)$$

which, for our non-relativistic source, converges rapidly. We can therefore cut the expansion short after the first order correction:

$$h_{ij}^{TT}(t, \mathbf{x}) \approx \frac{1}{r} \frac{4G}{c^4} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \left[S^{kl} + \frac{1}{c} n_m \dot{S}^{kl,m} \right]_{\text{ret}} \quad (\text{C.13})$$

From the energy-momentum tensor moments, we can define energy- and linear momentum-density moments:

$$\begin{aligned} M &= \frac{1}{c^2} \int d^3x T^{00}(t, \mathbf{x}) & P^i &= \frac{1}{c^2} \int d^3x T^{0i}(t, \mathbf{x}) \\ M^i &= \frac{1}{c^2} \int d^3x T^{00}(t, \mathbf{x}) x^i & P^{i,j} &= \frac{1}{c^2} \int d^3x T^{0i}(t, \mathbf{x}) x^j \\ M^{ij} &= \frac{1}{c^2} \int d^3x T^{00}(t, \mathbf{x}) x^i x^j & P^{i,jk} &= \frac{1}{c^2} \int d^3x T^{0i}(t, \mathbf{x}) x^j x^k \end{aligned}$$

Integrating the coservation of energy-momentum $\partial^0 T_{00} = -\partial^i T_{i0}$ over a volume V :

$$c\dot{M} = c^2 \partial_0 M = \int_V d^3x \partial_0 T^{00} = - \int_V d^3x \partial_i T^{0i} = - \int_S dS^i T^{0i} = 0 \quad (\text{C.14})$$

$$c\dot{M}^i = \int_V d^3x x^i \partial_0 T^{00} = - \int_V d^3x \partial_j x^i T^{0j} = \int_V d^3x \delta_j^i T^{0j} = cP^i \quad (\text{C.15})$$

In a similar manner we can derive the following identities:

$$\dot{M}^{ij} = P^{i,j} + P^{j,i} \quad (\text{C.16})$$

$$\dot{P}^{i,j} = S^{ij} \quad (\text{C.17})$$

We can combine these to find:

$$\ddot{M}^{ij} = \dot{P}^{i,j} + \dot{P}^{j,i} = S^{ij} + S^{ji} = 2S^{ij} \quad (\text{C.18})$$

This can be substituted into our truncated expansion, giving us:

$$h_{ij}^{TT}(t, \mathbf{x}) \approx \frac{1}{r} \frac{4G}{c^4} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \ddot{M}^{kl} \quad (\text{C.19})$$

If we introduce the (traceless) quadrupole moment $Q^{kl} \equiv (M^{kl} - \frac{1}{3}\delta^{kl}M_{ii})$, the GW's amplitude dependence on the second time derivative of the mass-quadrupole moment, as mentioned in the introduction, becomes apparent:

$$h_{ij}^{\text{TT}}(t, \mathbf{x}) \approx \frac{1}{r} \frac{2G}{c^4} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \ddot{Q}^{kl} \left(t - \frac{r}{c} \right) \equiv \frac{1}{r} \frac{2G}{c^4} \ddot{Q}_{ij}^{\text{TT}} \left(t - \frac{r}{c} \right) \quad (\text{C.20})$$

For a wave propagating along the $\hat{n} = (0, 0, 1)$ direction, we find that:

$$h_+ = \frac{1}{r} \frac{G}{c^4} (\ddot{M}_{11} - \ddot{M}_{22}) \quad (\text{C.21})$$

$$h_\times = \frac{2}{r} \frac{G}{c^4} \ddot{M}_{12} \quad (\text{C.22})$$

For a binary system in the xy -plane, we find:

$$M_{11} = \frac{\mu R^2 (1 - \cos 2\omega_s t)}{2} \quad (\text{C.23})$$

$$M_{22} = \frac{\mu R^2 (1 + \cos 2\omega_s t)}{2} \quad (\text{C.24})$$

$$M_{12} = \frac{\mu R^2 \sin 2\omega_s t}{2} \quad (\text{C.25})$$

$$(\text{C.26})$$

where R denotes the orbital separation of the binary, ω_s the orbital frequency and $\mu = m_1 m_2 / (m_1 + m_2)$ the reduced mass of the system.

For an arbitrary direction of propagation of GWs from a binary, we can express h_+ and $h_\times$ as:

$$h_+(t; \theta, \phi) = \frac{1}{r} \frac{4G\mu\omega_s^2 R^2}{c^4} \left(\frac{1 + \cos^2 \theta}{2} \right) \cos(2\omega_s t_{\text{ret}} + 2\phi) \quad (\text{C.27})$$

$$h_\times(t; \theta, \phi) = \frac{1}{r} \frac{4G\mu\omega_s^2 R^2}{c^4} \cos \theta \sin(2\omega_s t_{\text{ret}} + 2\phi) \quad (\text{C.28})$$

C.2 GW Energy-Momentum Tensor

For this analysis we must first admit a curved background:

$$g_{\mu\nu}(x) = \bar{g}_{\mu\nu}(x) + h_{\mu\nu}(x) \quad (\text{C.29})$$

But how can we distinguish between background and perturbation? Here we will use the short-wavelength approximation. It assumes that $\bar{g}_{\mu\nu}(x)$ varies spatially or temporally slowly compared to $h_{\mu\nu}(x)$. Often times one of these assumptions holds in real scenarios. This means that we can separate the perturbation (GW) from the curved background. We now recast the Einstein equations into their less popular form and expand the Ricci tensor to second order in $h_{\mu\nu}$:

$$\frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right) = R_{\mu\nu} \approx \bar{R}_{\mu\nu} + R_{\mu\nu}^{(1)} + R_{\mu\nu}^{(2)} \quad (\text{C.30})$$

with $T_{\mu\nu}$ the energy-momentum tensor and T its trace. From Eq. (C.30) we can deduce that $\bar{R}_{\mu\nu}$ only depends on $\bar{g}_{\mu\nu}$ and contains only low-frequency modes. $R_{\mu\nu}^{(1)}$ is linear in $h_{\mu\nu}$ and is independent of $h_{\mu\nu}$, containing only high frequency modes. $R_{\mu\nu}^{(2)}$ is quadratic in $h_{\mu\nu}$ and depends on both $h_{\mu\nu}$ and $\bar{g}_{\mu\nu}$. We can thus split Eq. (C.30) into low- and high-frequency parts:

$$\bar{R}_{\mu\nu} = -[R_{\mu\nu}^{(2)}]^{\text{Low}} + \frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right)^{\text{Low}} \quad (\text{C.31})$$

$$R_{\mu\nu}^{(1)} = -[R_{\mu\nu}^{(2)}]^{\text{High}} + \frac{8\pi G}{c^4} \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right)^{\text{High}} \quad (\text{C.32})$$

When there is a clear-cut separation between the length scale λ of the GWs and that of the background L_B , we can introduce a new scale $\bar{l}$, such that $\lambda \ll \bar{l} \ll L_B$. A similar separation can be made for time scales. We now coarse-grain the Einstein equations through a low-frequency projection by averaging over a spatial volume with side $\bar{l}$. This allows us to rewrite the low-frequency field equation as:

$$\bar{R}_{\mu\nu} = -\langle R_{\mu\nu}^{(2)} \rangle + \frac{8\pi G}{c^4} \left\langle T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right\rangle \quad (\text{C.33})$$

$$= -\langle R_{\mu\nu}^{(2)} \rangle + \frac{8\pi G}{c^4} \left(\bar{T}_{\mu\nu} - \frac{1}{2} g_{\mu\nu} \bar{T} \right) \quad (\text{C.34})$$

$\bar{T}_{\mu\nu}$ is a purely low-frequency quantity and represents the smoothed/coarse-grained energy-momentum tensor. Let us now introduce $t_{\mu\nu}$:

$$t_{\mu\nu} = -\frac{c^4}{8\pi G} \left\langle R_{\mu\nu}^{(2)} - \frac{1}{2} \bar{g}_{\mu\nu} R^{(2)} \right\rangle \quad (\text{C.35})$$

and its trace

$$t = \bar{g}^{\mu\nu} t_{\mu\nu} \quad (\text{C.36})$$

$$= \frac{c^4}{8\pi G} \langle R^{(2)} \rangle \quad (\text{C.37})$$

which allows us to rewrite Eq. (C.35) by inserting Eq. (C.37) into it:

$$-\langle R_{\mu\nu}^{(2)} \rangle = \frac{8\pi G}{c^4} \left(t_{\mu\nu} - \frac{1}{2} \bar{g}_{\mu\nu} t \right) \quad (\text{C.38})$$

After some reshuffling, we can now rewrite Eq. (C.33):

$$\bar{R}_{\mu\nu} - \frac{1}{2} \bar{g}_{\mu\nu} \bar{R} = \frac{8\pi G}{c^4} (\bar{T}_{\mu\nu} + t_{\mu\nu}) \quad (\text{C.39})$$

This can be called the coarse-grained version of the Einstein equations. $\bar{T}_{\mu\nu}$ represents the smoothed energy-momentum tensor and $t_{\mu\nu}$ the energy in the gravitational field, independent of the external matter, and is quadratic in $h_{\mu\nu}$.

It can be shown that $R_{\mu\nu}^{(2)}$ is given by:

$$\begin{aligned} R_{\mu\nu}^{(2)} = \frac{1}{2} \bigg[& \frac{1}{2} \partial_\mu h_{\alpha\beta} \partial_\nu h^{\alpha\beta} + h^{\alpha\beta} \partial_\mu \partial_\nu h_{\alpha\beta} - h^{\alpha\beta} \partial_\nu \partial_\beta h_{\alpha\mu} - h^{\alpha\beta} \partial_\mu \partial_\beta h_{\alpha\nu} \\ & + h^{\alpha\beta} \partial_\alpha \partial_\beta h_{\mu\nu} + \partial^\beta h^\alpha_\nu \partial_\beta h_{\alpha\mu} - \partial^\beta h^\alpha_\nu \partial_\alpha h_{\beta\mu} - \partial_\beta h^{\alpha\beta} \partial_\nu h_{\alpha\mu} \\ & + \partial_\beta h^{\alpha\beta} \partial_\alpha h_{\mu\nu} - \partial_\beta h^{\alpha\beta} \partial_\mu h_{\alpha\nu} - \frac{1}{2} \partial^\alpha h \partial_\alpha h_{\mu\nu} + \frac{1}{2} \partial^\alpha h \partial_\nu h_{\alpha\mu} \\ & + \frac{1}{2} \partial^\alpha h \partial_\mu h_{\alpha\nu} \bigg] \end{aligned}$$

This equation looks daunting, but we can use some of GR's famous gauge freedom to drastically simplify it. Using the gauge conditions $\partial^\mu h_{\mu\nu} = 0$ and $h = 0$, and the equation of motion $h_{\alpha\beta} = 0$, all terms except the first two of the above equation become zero. The first two first are related by an integration by parts, leading to:

$$\langle R_{\mu\nu}^{(2)} \rangle = -\frac{1}{4} \langle \partial_\mu h_{\alpha\beta} \partial_\nu h^{\alpha\beta} \rangle \quad (\text{C.40})$$

This allows us to rewrite $t_{\mu\nu}$ as:

$$t_{\mu\nu} = \frac{c^4}{32\pi G} \langle \partial_\mu h_{\alpha\beta} \partial_\nu h^{\alpha\beta} \rangle \quad (\text{C.41})$$

This quantity is gauge-invariant, so can only depend upon the two physical DOF of the TT-gauge, which are the GW polarisations. Introducing the dot notation for $\partial_t = c\partial_0$, we can express the gauge-invariant energy density as:

$$t^{00} = \frac{c^4}{32\pi G} \langle \dot{h}_{ij}^{TT} \dot{h}_{ij}^{TT} \rangle = \frac{c^4}{32\pi G} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle \quad (\text{C.42})$$

C.2.1 Energy Flux

We can state that the energy carried away from a volume V through a surface dA in a time dt is given by the outward-propagating GW flux:

$$\frac{dE}{dAdt} = ct^{00} = \frac{c^3}{32\pi G} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle \quad (\text{C.43})$$

or when expressed with the surface element $dA = r^2 d\Omega$:

$$\frac{dE}{dt} = \frac{c^3 r^2}{32\pi G} \int d\Omega \langle \dot{h}_{ij}^{TT} \dot{h}_{ij}^{TT} \rangle \quad (\text{C.44})$$

Filling in the expression for $\dot{h}_{ij}^{TT}$ from Eq. (C.20):

$$\frac{dE}{dt} = \frac{G}{8\pi c^5} \int d\Omega \Lambda_{ij,kl} \langle \ddot{Q}^{ij} \ddot{Q}^{kl} \rangle = \frac{G}{8\pi c^5} \langle \ddot{Q}^{ij} \ddot{Q}^{kl} \rangle \int d\Omega \Lambda_{ij,kl} \quad (\text{C.45})$$

$\langle \ddot{Q}^{ij} \ddot{Q}^{kl} \rangle$ could be pulled out of the integral as it is independent of direction. It can be shown that:

$$\frac{1}{4\pi} \int d\Omega \Lambda_{ij,kl} = \frac{1}{30} (11\delta_{ik}\delta_{jl} - 4\delta_{ij}\delta_{kl} + \delta_{il}\delta_{jk}) \quad (\text{C.46})$$

from which we find:

$$\frac{dE}{dt} = \frac{G}{5c^5} \langle \ddot{Q}_{ij} \ddot{Q}_{ij} \rangle = \frac{G}{5c^5} \left\langle \ddot{M}_{ij} \ddot{M}_{ij} - \frac{1}{3} (\ddot{M}_{kk})^2 \right\rangle \quad (\text{C.47})$$

C.3 Backreaction

From Kepler's third law, we know that:

$$\omega_s^2 = \frac{Gm}{R^3}$$

We can use this and the chirps mass $\mathcal{M}$:

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{(1/5)}} \quad (\text{C.48})$$

to rewrite Eqs. (C.27) and (C.28) to:

$$h_+(t) = \frac{4}{r} \left(\frac{G\mathcal{M}}{c^2} \right)^{5/3} \left(\frac{\pi f_{\text{gw}}}{c} \right)^{2/3} \frac{1 + \cos^2 \theta}{2} \cos(2\pi f_{\text{gw}} t_{\text{ret}} + 2\phi) \quad (\text{C.49})$$

$$h_\times(t) = \frac{4}{r} \left(\frac{G\mathcal{M}}{c^2} \right)^{5/3} \left(\frac{\pi f_{\text{gw}}}{c} \right)^{2/3} \cos \theta \sin(2\pi f_{\text{gw}} t_{\text{ret}} + 2\phi) \quad (\text{C.50})$$

Note that the frequency of GWs is twice the orbital frequency of the binary. Combining Eq. (C.47) with the mass moments, we can express the power radiated through GWs on a circular orbit as:

$$\frac{dE}{dt} = \frac{32}{4} \frac{c^5}{G} \left(\frac{G\mathcal{M} f_{\text{gw}}}{2c^3} \right)^{10/3} \quad (\text{C.51})$$

Which gives us the scaling of radiated power with frequency as mentioned in the introduction.

