



Universitat
de les Illes Balears

DOCTORAL THESIS
2020

**SEARCHING FOR CONTINUOUS GRAVITATIONAL
WAVES WITH ADVANCED LIGO**

Josep Blai Covas Vidal

Doctor by the Universitat de les Illes Balears



Universitat
de les Illes Balears

DOCTORAL THESIS
2020

Doctoral Programme of *Physics*

**SEARCHING FOR CONTINUOUS GRAVITATIONAL
WAVES WITH ADVANCED LIGO**

Josep Blai Covas Vidal

Supervisor: Dr. Alicia Sintés Olives

Doctor by the Universitat de les Illes Balears



Universitat
de les Illes Balears

Dr Alicia Sintes, of the University of the Balearic Islands

I DECLARE:

That the thesis titled *Searching for continuous gravitational waves with Advanced LIGO*, presented by Josep Blai Covas Vidal to obtain a doctoral degree, has been completed under my supervision and meets the requirements to opt for an International Doctorate.

For all intents and purposes, I hereby sign this document.

Signature:

A handwritten signature in black ink, appearing to read 'Alicia Sintes', is written over a horizontal line. The signature is stylized and somewhat cursive.

Palma de Mallorca, 12 May 2020

La primera detecció d'ones gravitacionals (OGs), detectades el 14 de setembre de 2015, va marcar l'inici de l'astronomia d'ones gravitacionals. Totes les OGs que han estat detectades fins ara han vingut de fusions de sistemes binaris compactes, un tipus de senyal que és detectable durant segons o uns pocs minuts. Un tipus completament diferent d'ona gravitacional són les ones gravitacionals contínues, senyals de llarga duració majoritàriament produïdes per estrelles de neutrons (tant aïllades com a sistemes binaris) asimètriques que es troben rotant. Hi ha hagut moltes cerques que han intentat trobar aquest tipus de senyal, però cap d'elles ho ha aconseguit. Encara que les ones contínues poden ser monitoritzades durant molt de temps, es prediu que la seva amplitud és uns quants ordres de magnitud menor que l'amplitud de les ones provinents de fusions de sistemes binaris compactes: les ones ja detectades tenen una amplitud habitual al voltant de 10^{-21} , mentre que cerques d'ones contínues sense deteccions impliquen que aquesta amplitud hagi de ser menor que 10^{-26} per alguns púlsars, fet que remarca el gran repte que detectar aquestes ones representa.

Les cerques d'ones contínues poden dividir-se entre cerques d'ones provinents de púlsars coneguts o cerques d'ones provinents d'estrelles de neutrons desconegudes. Al contrari que per el primer tipus de cerca (per les quals la posició al cel, la freqüència de rotació, i el ritme amb que aquesta freqüència es redueix són coneguts), les cerques de fonts que no han estat detectades mitjançant radiació electromagnètica requereixen algorismes capaços d'analitzar espais de paràmetres gegants, ja que les dades han de ser correlacionades amb els models d'ona teòrics els quals depenen d'aquests paràmetres desconeguts, que han de ser inclosos per tenir en compte les diferents modulacions com la modulació Doppler produïda per la rotació de la Terra i la seva òrbita al voltant del sol. Desafortunadament, no hi ha prou potència computacional disponible per cercar un espai de paràmetres tan gran i gairebé continu en posició al cel, freqüència, i ritme de frenat. Per aquesta raó, el desenvolupament d'algorismes no òptims que puguin gestionar aquest espai de paràmetres és una tasca important dins el camp d'anàlisi de dades.

Aquesta tesi està separada en dues parts diferents. La primera part es compon de tres capítols que presenten introduccions a diferents temes que són necessaris per comprendre la recerca d'ones contínues: com es generen i propaguen les ones gravitacionals, què són les estrelles de neutrons i com poden generar ones gravitacionals, i quins mètodes estadístics s'han d'usar per detectar un senyal contínua i estimar els seus paràmetres.

La segona part consta de quatre capítols que resumeixen els resultats originals que han estat publicats en revistes d'alt impacte. El primer d'aquests capítols està dedicat a la millora d'un mètode de recerca ja existent per a senyals d'ones gravitacionals contínues de fonts desconegudes, anomenat SkyHough. Aquestes noves incorporacions milloren la sensibilitat d'aquest mètode sense incrementar el seu cost computacional, augmentant així la probabilitat de detectar un senyal d'ona contínua. Aquestes millores s'apliquen a una cerca utilitzant dades d'Advanced LIGO O2. Encara que no hi ha deteccions, els límits superiors de l'amplitud d'ona gravitacional milloren en comparació amb els resultats anteriors. El segon

capítol de resultats presenta una expansió del mètode SkyHough que permet la cerca d'ones contínues procedents d'estrelles de neutrons en sistemes binaris, ja que el mètode original només permetia la cerca de sistemes aïllats. Aquest tipus de cerca té el cost computacional més alt dins les cerques d'ones contínues, ja que cal tenir en compte més paràmetres que descriuen l'òrbita binària. Abans d'aquest nou desenvolupament només existia un algoritme capaç de fer aquest tipus de cerca. Els resultats presentats en aquest capítol mostren que el nou algoritme és més sensible que l'algoritme que existia prèviament. El següent capítol presenta una aplicació d'aquest nou algoritme, on també es realitza una cerca de fonts desconegudes utilitzant dades d'Advanced LIGO O2. No hi ha deteccions, però els nous resultats milloren els resultats anteriors en un factor 17. L'últim capítol presenta un estudi sobre com caracteritzar el soroll de llarga durada que pot estar present en els detectors interferomètrics. A més de presentar mètodes que permeten la detecció d'aquest soroll, mostrem una sèrie d'investigacions que van reduir o eliminar algunes fonts de soroll que eren presents al detector d'ones gravitacionals de Hanford.

La primera detección directa de ondas gravitacionales (OGs), detectadas el 14 de septiembre de 2015, marcó el comienzo de la astronomía de ondas gravitacionales. Todas las OGs que se han detectado hasta ahora provienen de fusiones de sistemas binarios compactos, un tipo de señal que es detectable durante segundos o unos pocos minutos. Un tipo completamente diferente de ondas gravitacionales son las ondas gravitacionales continuas (CWs), que son ondas de larga duración producidas principalmente por estrellas de neutrones (aisladas o en sistemas binarios) asimétricas en rotación. Aunque se han realizado muchas búsquedas de CWs, ninguna de ellas ha logrado una detección. A pesar de que las CWs se pueden rastrear durante tiempos mucho más largos que las fusiones de sistemas binarios compactos, se cree que la amplitud de estas ondas es muy menor: la amplitud de las OGs ya detectadas es de alrededor de 10^{-21} , mientras que las búsquedas de CWs están limitando esta amplitud a menos de 10^{-26} para algunos púlsares, hecho que remarca el enorme desafío de detectar CWs.

Las búsquedas de CWs se pueden dividir entre búsquedas de OGs de púlsares conocidos y búsquedas de OGs de estrellas de neutrones desconocidas. A diferencia de las búsquedas de OGs de los púlsares (cuyas ubicaciones, frecuencias de emisión de ondas gravitacionales, y ritmos a los que estas frecuencias se reducen son bien conocidas), las búsquedas de fuentes electromagnéticamente silenciosas requieren algoritmos que analicen espacios de parámetros mucho más grandes, porque los datos tienen que ser correlacionados con formas de onda teóricas que dependen de estos parámetros desconocidos, que deben incluirse para tener en cuenta las diferentes modulaciones que están presentes, como la modulación Doppler producida por la rotación de la Tierra y su órbita alrededor del sol. Desafortunadamente, no hay suficiente potencia de cómputo disponible para buscar un espacio de parámetros tan grande y casi continuo en posición del cielo, frecuencia y ritmo de frenado. Por esta razón, el desarrollo de algoritmos no óptimos que puedan manejar este enorme espacio de parámetros es una tarea importante dentro del campo de análisis de datos.

Esta tesis está separada en dos partes diferentes. La primera parte se compone de tres capítulos que presentan introducciones a diferentes temas que son necesarios para comprender las búsquedas de CWs: cómo se generan y propagan las ondas gravitacionales, qué son las estrellas de neutrones y cómo pueden generar ondas gravitacionales, y qué métodos estadísticos deben usarse para detectar una señal continua y estimar sus parámetros.

La segunda parte consta de cuatro capítulos que resumen los resultados originales que han sido publicados en revistas de alto impacto. El primero de estos capítulos está dedicado a la mejora de un método de búsqueda ya existente para señales de ondas gravitacionales continuas de fuentes desconocidas, llamado SkyHough. Estas nuevas incorporaciones mejoran la sensibilidad de este método sin incrementar su coste computacional, aumentando así las probabilidades de detectar una señal de onda continua. Estas mejoras se aplican a una búsqueda utilizando datos de Advanced LIGO O2. Aunque no hay detecciones, los límites superiores de la amplitud de onda gravitacional mejoran en comparación con los

resultados anteriores. El segundo capítulo de resultados presenta una expansión del método SkyHough que permite la búsqueda de ondas continuas procedentes de estrellas de neutrones en sistemas binarios, ya que el método original solo permitía la búsqueda de sistemas aislados. Este tipo de búsqueda tiene el coste computacional más alto dentro de las búsquedas de ondas continuas, ya que es necesario tener en cuenta más parámetros relacionados con la órbita binaria. Antes de este nuevo desarrollo solo existía un algoritmo capaz de hacer este tipo de búsqueda. Los resultados presentados en este capítulo muestran que el nuevo algoritmo es más sensible que el algoritmo que existía previamente. El siguiente capítulo presenta una aplicación de este nuevo algoritmo, donde también se realiza una búsqueda de fuentes desconocidas utilizando datos de Advanced LIGO O2. No hay detecciones, pero los nuevos resultados mejoran los resultados anteriores en un factor 17. El último capítulo presenta un estudio sobre cómo caracterizar el ruido de larga duración que puede estar presente en los detectores interferométricos. Además de presentar métodos que permiten la detección de este ruido, mostramos una serie de investigaciones que redujeron o eliminaron algunas fuentes de ruido que estaban presentes en el detector de ondas gravitacionales de Hanford.

The first direct detection of gravitational waves (GWs) on September 14 2015 marked the beginning of gravitational-wave astronomy. All of the GWs that have been detected until now came from compact binary coalescences, a type of signal that is detectable during seconds or a few minutes. A completely different type of gravitational waves are continuous gravitational waves (CWs), which are long-lasting waves mainly produced by asymmetric rotating neutron stars, either isolated or in binary systems. Although many searches for CWs have been done, none of them has reported a detection. Even though CWs can be tracked during much longer times than compact binary coalescences, the amplitude of these gravitational waves is expected to be many orders of magnitude smaller: the amplitude of the detected GWs is around 10^{-21} , whereas CW searches are constraining this amplitude to be less than 10^{-26} for some pulsars, which clearly underlines the challenge of detecting CWs.

CW searches can be divided between searches for GWs from known pulsars and searches for GWs from unknown neutron stars. Unlike searches for GWs from pulsars (whose locations, gravitational wave emission frequencies, and spin-down rates are well known), searches for electromagnetically quiet sources require algorithms that look at vastly larger parameter spaces, because the data has to be correlated with theoretical waveforms that depend on these unknown parameters, which have to be included to take into account the different modulations such as the Doppler modulation produced by Earth's rotation and orbit around the sun. Unfortunately, there is not enough computing power available to search such a large and nearly continuous parameter space in sky position, frequency, and spin-down rate. For this reason, developing non-optimal algorithms that can deal with this huge parameter space is an important task within the data analysis community.

This thesis is separated in two different parts. The first part is made up of three chapters that give introductions to different topics that are needed to understand CW searches: how gravitational waves are generated and propagated, what neutron stars are and how they can generate gravitational waves, and what statistical methods have to be used in order to detect a CW signal and estimate its parameters.

The second part is made up of four chapters that summarize original results that have been published in high-impact journals. The first of these chapters is devoted to the improvement of an already existing search method for continuous gravitational wave signals from unknown sources, called SkyHough. These new additions improve the sensitivity of this method without incrementing its computational cost, thus increasing the probabilities of detecting a continuous wave signal. These new improvements are applied to an all-sky search using Advanced LIGO O2 data. Although no detections are reported, the upper limits on the gravitational-wave amplitude are improved as compared to the previous results. The second results chapter presents an expansion of the SkyHough method that allows the search of CWs from neutron stars in binary systems, whereas the original method only allowed the search from isolated systems. This type of search has the highest computational cost within the CW searches, since more parameters related to the binary orbit need to be taken into account. Prior to this new development,

only one pipeline capable of doing this kind of search existed. The results presented in this chapter show that the new pipeline is more sensitive than the previous pipelines. The next chapter presents an application of this newly developed pipeline, where an all-sky search also using Advanced LIGO O2 data is done. No detections are reported, but the new results improve on the previous results by a factor of 17. The last chapter presents a study of how to characterize long-duration noise that can be present in interferometric detectors. Besides presenting methods that allow the detection of this noise, we show a number of investigations that mitigated some noise sources that were present in the Hanford gravitational-wave detector.

ACKNOWLEDGEMENTS

First of all I want to thank my thesis advisor, Alicia Sintés. A couple of years ago, she nicely agreed to be the director of my “Trell de fi de grau”. That was when I began to hear about gravitational waves and started to see how a research group works. She again agreed to tutor my master’s work, when I began to participate in real scientific research. Alicia always had time to solve the many questions that I had about continuous gravitational waves and scientific research, and having the opportunity of being a member of a collaboration such as the LSC was a very exciting start. Once more, she accepted being my thesis advisor and gave me one position in the GRG group. During this thesis, she has always provided good advice and has given me the confidence that is needed to pursue a career as a physics researcher.

I also want to mention Sascha Husa and Carlos Palenzuela, who were great teachers during the master’s courses and have always had time to answer questions and give advice whenever needed during this thesis. Furthermore, I also want to thank all the young people from the group that have been a part of this long journey: Rafel, Toni, David, Miquel, Cecilio, Héctor, Borja, Miguel, Daniele, Leila, Marta, Rodrigo. I have learned many things from all of you, and we have shared many discussions and travels that I always will remember with great joy.

During this thesis I have been fortunate enough to spend some time at other research institutions, an opportunity which has given me a broader perspective and the opportunity to meet new people from other countries. I’ve spent six months at the Hanford detector in Washington, USA, where I have met many great scientists that were always patient enough to answer any questions and to guide me when I was starting in new research adventures. I will always value having had the opportunity to actively participate in such a complex scientific facility, where I could witness in person many gravitational wave detections and how these detectors were actually operated and maintained. I want to particularly thank Rick Savage, Evan Goetz, Jeffrey Kissel, and Robert Schofield for their scientific guidance, all the commissioners and operators, and all the fellows whom I shared the fellows apartments and many hours of free time. I am also grateful to the Rome and AEI groups for hosting me during my visits and for their hospitality. I also want to thank all the members of the Continuous Waves group, a group which I have been part of for more than four years. Its members have always shown great disposition to answer questions and provide guidance. I’ve learned many things during the weekly teleconferences, and during the LSC meetings there was always an atmosphere of friendly cooperation and collaboration.

Finally, I want to thank my family and friends, which are the most important persons to me. My parents gave me the opportunity to enroll in the Physics degree and have always supported me, so this would not have been possible without them. I also want to thank my sister and my friends, because this would have been much harder without their support and the uncountable amount of good times that we have.

This thesis has been done as a member of the LIGO Scientific Collaboration and the Institute of Applied Computing and Community Code (IAC3) of the University of the Balearic Islands. This work

has been supported by the Vicepresidència i Conselleria d'Innovació, Recerca i Turisme del Govern de les Illes Balears (grant FPI-CAIB FPI/2134/2018) and the Fons Social Europeu 2014-2020 de les Illes Balears. The author also acknowledges the support by the Govern de les Illes Balears through the Direcció General de Política Universitària i Recerca with funds from the Tourist Stay Tax Law ITS 2017-006 (PRD2018/24), the European Union FEDER funds and EU COST Actions CA18108, CA17137, CA16214, and CA16104, the Ministry of Science, Innovation and Universities and the Spanish Agencia Estatal de Investigación grants FPA2016-76821-P, RED2018-102661-T, RED2018-102573-E, FPA2017-90566-REDC, FPA2017-90687-REDC, and the Generalitat Valenciana (PROMETEO/2019/071). The author is grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. The author thankfully acknowledges the computer resources at CTE-POWER and the technical support provided by *Barcelona Supercomputing Center - Centro Nacional de Supercomputación* (RES-AECT-2019-1-0021 and RES-AECT-2019-3-0011). This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org>) [1], a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

PUBLICATIONS DERIVED FROM THIS THESIS

- *Full band all-sky search for periodic gravitational waves in the O1 LIGO data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **97**, 102003 (2018)
- *Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO*, P. B. Covas et al. (LSC Instrument Authors), **Physical Review D** **97**, 082002 (2018)
- *All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **100**, 024004 (2019)
- *New method to search for continuous gravitational waves from unknown neutron stars in binary systems*, P. B. Covas and Alicia M. Sintes, **Physical Review D** **99**, 124019 (2019)
- *First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data*, P. B. Covas and Alicia M. Sintes, **Physical Review Letters** **124**, 191102 (2020)

LIST OF ABBREVIATIONS

BSH	BinarySkyHough
CBC	Compact Binary Coalescence
CW	Continuous wave
EM	Electromagnetic
EOS	Equation of state
GR	General Relativity
GW	Gravitational wave
H1	LIGO Hanford 4 km interferometer
LIGO	Laser Interferometer Gravitational-wave Observatory
LSC	LIGO Scientific Collaboration
LVC	LIGO-Virgo Collaboration
L1	LIGO Livingston 4 km interferometer
MCMC	Markov chain Monte Carlo
MDC	Mock Data Challenge
MLE	Maximum likelihood estimator
NS	Neutron star
O1	First Advanced LIGO science run
O2	Second Advanced LIGO-Virgo science run
O3	Third Advanced LIGO-Virgo science run
PSD	Power Spectral Density
RAM	Random Access Memory
SFT	Short Fourier Transform
SH	SkyHough
SNR	Signal to Noise Ratio
SSB	Solar System Barycenter
STD	Standard deviation
S6	Initial LIGO science run 6
V1	Virgo 3 km interferometer

Resum	I
Resumen	III
Abstract	V
Acknowledgements	VII
Publications derived from this thesis	IX
List of abbreviations	X
Contents	XI
I Theoretical introduction	1
1. Gravitational waves	2
1.1. Introduction to gravitational waves	3
1.1.1. Propagation of gravitational waves	4
1.1.2. Production of gravitational waves	7
1.2. Astrophysical sources of gravitational waves	10
1.3. Gravitational wave detectors	12
1.3.1. How does an interferometric gravitational wave detector work?	15
1.3.2. Noise in ground-based interferometric detectors	18
1.3.3. Antenna beam pattern functions of an interferometric detector	19
2. Neutron stars	23
2.1. Basics of neutron star physics	23
2.2. Demographics and parameters	24
2.3. Asymmetries of neutron stars	33
2.3.1. Magnetic deformations	35
2.3.2. Thermal deformations	36
2.3.3. Exotic deformations	36
2.4. Continuous gravitational waves from neutron stars	37
2.4.1. Gravitational waves from rotation around a principal axis	37
2.4.2. Gravitational waves from general rotation	38

2.4.3.	The signal at the detectors	40
3.	Data analysis and continuous wave searches	44
3.1.	Gravitational-wave data analysis	44
3.1.1.	Preliminary definitions	44
3.1.2.	Hypothesis testing	47
3.1.3.	Parameter estimation	49
3.1.4.	Mismatch and resolution	50
3.2.	Continuous waves searches: targets, methods, and past results	53
3.2.1.	Upper limits and estimation of sensitivity	58
3.2.2.	Parameters measurable with CWs	61
3.3.	SkyHough	63
3.3.1.	Input data	64
3.3.2.	Partial Hough map derivatives and look-up table approach	65
3.3.3.	Detection statistics and sensitivity	66
3.3.4.	Post-processing	70
II	Results	73
4.	All-sky search for continuous gravitational waves from isolated neutron stars using O2 Advanced LIGO data	74
4.1.	Introduction	74
4.2.	Preparing the search	75
4.2.1.	Number of valid frequency bins for the look-up table approach	76
4.2.2.	Size of the sky-patches and RAM formula	77
4.2.3.	Extra detection statistics	79
4.2.4.	Creating additional SFTs	80
4.3.	Advanced LIGO O2 observing run	81
4.4.	Description of the search methods	83
4.4.1.	SkyHough	83
4.4.2.	Comparison with the other pipelines	87
4.5.	Results	88
4.5.1.	SkyHough	88
4.5.2.	Global results	91
4.6.	Conclusions	93
5.	BinarySkyHough: a new method to search for continuous gravitational waves from unknown neutron stars in binary systems	96
5.1.	Introduction	96
5.2.	New method for CW all-sky binary searches	98
5.2.1.	BinarySkyHough	99
5.2.2.	Resolution, parameter space range and number of templates	100
5.2.3.	Maximum coherent time	102
5.2.4.	Maximum eccentricity	103
5.2.5.	Post-processing	104
5.2.6.	Computational model	106
5.3.	Sensitivity estimation	108
5.3.1.	Procedure	108
5.3.2.	Results	109

5.3.3. Comparison with other methods	111
5.4. Conclusions	113
6. First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data	115
6.1. Introduction	115
6.2. Search	117
6.2.1. Main search	117
6.2.2. Follow-up	118
6.3. Results	120
6.3.1. Sensitivity depth estimation	120
6.3.2. Upper limits	120
6.3.3. Excluded frequency bands from the upper limits calculation	122
6.4. Conclusions	124
7. Efforts to characterize and mitigate lines and combs during O2 and O3	126
7.1. Introduction	126
7.2. Effects of noise on CW and SGWB searches	127
7.3. LIGO data and noise sources for searches of persistent gravitational waves	128
7.4. Noise coupling mechanisms to the gravitational wave channel	131
7.4.1. Coupling through shared power and grounds	131
7.4.2. Coupling through magnetic or electrostatic fields	132
7.4.3. Mechanical coupling	133
7.4.4. Data acquisition artifacts and non-linear coupling	133
7.5. Data analysis tools	133
7.5.1. Finding lines in frequency domain	133
7.5.2. Finding combs	134
7.5.3. Lines and combs in time-domain: folding studies	134
7.5.4. Finding long-duration non-linearities	135
7.6. Results	137
7.6.1. 11.111 Hz comb (Vacuum sensors)	138
7.6.2. 86 Hz line (Pcal high frequency injections)	138
7.6.3. Lines from bounce and roll modes	139
7.6.4. Non-stationary line around 70 Hz	140
7.6.5. Non-linearities in O3a data	141
7.7. Producing a list of known lines and combs for CW searches	144
7.7.1. List of known lines and combs for CW searches	144
7.7.2. Known lines and combs for O1 and O2	145
7.8. Conclusions	145
8. Conclusions	148
8.1. Summary	148
8.2. Future work	149
References	150

Part I

Theoretical introduction

The first three chapters of this thesis present a theoretical introduction to the topics that are needed to understand the original results obtained and presented in the second part of the thesis. The first chapter presents a mathematical derivation of gravitational waves and their properties, a brief explanation of the different astrophysical sources of gravitational waves, and a summary of the different types of gravitational-wave detectors and a basic understanding of how they work. The second chapter gives an introduction to neutron stars, pulsars, and describes mathematically the continuous gravitational wave signals and how they can be generated. The third chapter introduces the basic statistical framework that is needed to understand how continuous gravitational waves can be detected and characterized, a summary of different continuous waves detection methods, and an introduction to the method used to obtain the results of this thesis, *SkyHough*.

In 1905, Einstein published his theory of special relativity [2]. In this theory, phenomena like time dilation or conversion of mass into energy were explained. The two main principles of special relativity are:

- **Principle of relativity:** the results of any experiment performed by any observer do not depend on the speed relative to other observers that are not involved in the experiment.
- **Universality of the speed of light:** the speed of light in vacuum has always the same value, regardless of the relative motion between the light's source and the observer.

Ten years later Einstein published the theory of general relativity (GR) [3], which explains the interaction between energy distributions and the curvature of spacetime, and it is a description of gravity as a geometric property of spacetime. The equation that explains this interaction is:

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad (1.1)$$

where $G_{\mu\nu}$ is the Einstein tensor that describes the geometry of spacetime, $T_{\mu\nu}$ is the stress-energy tensor that describes the matter and energy density content in a region of spacetime, G is the gravitational constant, and c is the speed of light in vacuum. GR presented a whole new theoretical interpretation of gravity, and was soon experimentally tested since two of its key predictions were confirmed: the correct value for the advance of the perihelion of Mercury's orbit, and the experimental verification that light could be bended by gravity.

One year after the publication of the theory of general relativity, Einstein published a paper describing gravitational waves (GWs) [4], presenting an equation that described the rate of energy emitted by a time-varying mass quadrupole. The actual physical existence of GWs was highly debated until the Chapel Hill conference in 1957, where a thought experiment known as the sticky bead argument was proposed, which showed that gravitational waves would change the proper distance between two objects and that energy would be subtracted from these waves. Furthermore, a calculation by Pirani showing the tidal effect that a GW has on a ring of particles also showed that GWs produced measurable effects [5].

After GWs were theoretically proven to exist, the hunt for their experimental detection began. The first detection of GW effects, although indirect, was due to the observed decay of the orbital period of the binary pulsar PSR 1913+16 (which was the first binary pulsar ever found), discovered by Hulse and Taylor in 1974 [6]. After several years of data-taking, they saw that the orbital period of the system (the time to complete a full orbit around the center of mass) was decreasing, and this decrease was perfectly fitted by the predictions of general relativity. This system has been used to test general relativity and constrain other theories of gravitation [7], as well as others detected after this first measurement.

Gravitational waves	Electromagnetic waves
Oscillations of spacetime	Oscillations propagated through spacetime
1 mass sign	2 charge signs
2 polarization states rotated by 45°	2 polarization states rotated by 90°
Frequencies: $\sim 10^{-9} - 10^{11}$ Hz	Frequencies: $\sim 10^4 - 10^{20}$ Hz
Hardly interact with matter	Strongly interact with matter
Propagate at speed of light	Propagate at speed of light
Amplitude decays with distance as $\propto 1/d$	Amplitude decays with distance as $\propto 1/d$

Table 1.1: Comparison between gravitational and electromagnetic waves.

The quest for a direct detection of a GW had begun in the 1960s with the design of the resonant-mass antennas pioneered by Joe Weber, who also attended the 1957 Chapel Hill conference. Although he claimed to have detected GWs, none of the other groups that had antennas could replicate his results. It was not until 2015 that GWs were directly detected by the LIGO collaboration. The first direct detection of a GW was made the 14th September 2015 [8], when a signal from a binary black hole merger located around 410 Mpc was found. Furthermore, in August 2017 the first detection of gravitational waves from a binary neutron star merger was made [9]. This was the first detection of gravitational waves and light from the same system: after the merger of the two neutron stars, light from the collision and the aftermath was collected [10]. Furthermore, this was the first time that the equation of state of neutron stars had been constrained by using gravitational waves. This event also confirmed the hypothesis that short gamma-ray bursts are produced by mergers of neutron stars. All the other detections made during the first two observing runs of the second generation ground-based detectors are reported in [11].

The rest of this chapter is organized as follows: in section 1.1 we introduce the mathematical foundation of gravitational waves; in section 1.2 we describe the different types of sources that can produce gravitational waves; in section 1.3 we present a brief review of gravitational wave detectors and how they work.

1.1 | Introduction to gravitational waves

Gravitational waves (GWs) are ripples or perturbations in the fabric of spacetime. When any energy distribution that has an asymmetry is accelerated, the curvature of spacetime is changed, and this change (which according to special relativity must propagate at a finite speed) is emitted from the source as gravitational radiation. General relativity predicts that gravitational waves travel at the speed of light, although other theories of gravitation speculate that this speed might be smaller due to a massive graviton. As shown in table 1.1, this form of radiation shares some properties with other kinds of waves, such as electromagnetic waves; however, one must notice that the latter are waves traveling within spacetime, while the former are perturbations of spacetime itself.

When GWs from astrophysical events reach the Earth they have an extremely small amplitude. This is related to the fact that gravitational interaction is the weakest of all known interactions and also due to the enormous distances to these systems. For this reason, we only expect to detect gravitational waves coming from the most violent and catastrophic events in the Universe, such as mergers of black holes or neutron stars, isolated spinning neutron stars, or supernova explosions.

In the next two subsections we present the mathematical formalism that describes the generation of gravitational waves and how they propagate through vacuum.

1.1.1 | Propagation of gravitational waves

Most sources of detectable gravitational waves are located very far away from us. The propagation of gravitational waves through a flat background is usually described by linearized general relativity.

Linearized theory is a weak-field approximation to GR: the Einstein equations are solved in a nearly flat spacetime, with a small perturbation acting on top of a flat background. With this formulation, we can separate the “static” and the oscillating parts of the metric. This approximation can be used for the description of gravitational waves passing through Earth since any likely source of detectable GWs will be very distant, and the amplitude of the GWs will be very small when they reach our detectors.

In linearized theory, the metric can be decomposed as a flat background plus a small perturbation [12]:

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

where $g_{\mu\nu}$ is the metric of spacetime, $\eta_{\mu\nu} = \text{diag}[-1, 1, 1, 1]$ is the Minkowski metric, and $h_{\mu\nu} \ll 1$ is a small perturbation of flat spacetime. All these quantities are symmetric tensors, and, for this reason, they only have ten independent components. Linearized theory is an approximation to general relativity that is correct to first order in the size of the metric perturbation $h_{\mu\nu}$. An example of a weak-field situation is the solar system, where $|h_{\mu\nu}| \approx |\Phi|/c^2 \propto GM_{\odot}/(R_{\odot}c^2) \sim 10^{-6}$, where Φ is the gravitational potential for a point mass and M_{\odot} and R_{\odot} are the mass and the radius of the Sun respectively.

The linearized Einstein equation is derived hereunder, following the derivation of [12]. The Einstein tensor is defined as:

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R, \quad (1.3)$$

where $R_{\mu\nu}$ is the Ricci tensor, defined as the contraction of the Riemann tensor (a comma indicates a partial derivative):

$$R_{\mu\nu} \equiv R^{\gamma}_{\mu\gamma\nu} = \Gamma^{\gamma}_{\mu\nu,\gamma} - \Gamma^{\gamma}_{\mu\gamma,\nu} + \Gamma^{\sigma}_{\mu\nu}\Gamma^{\gamma}_{\sigma\gamma} - \Gamma^{\sigma}_{\mu\gamma}\Gamma^{\gamma}_{\sigma\nu}, \quad (1.4)$$

where $\Gamma^{\alpha}_{\beta\gamma}$ are the Christoffel symbols, defined by:

$$\Gamma^{\alpha}_{\beta\gamma} = \frac{1}{2}g^{\alpha\mu}(g_{\mu\beta,\gamma} + g_{\mu\gamma,\beta} - g_{\beta\gamma,\mu}), \quad (1.5)$$

where $g_{\mu\nu}$ is the metric defined in equation (1.2), and R is the Ricci curvature scalar, defined by:

$$R = g^{\mu\nu}R_{\mu\nu}. \quad (1.6)$$

Equation (1.5) can be linearised using the metric defined in equation (1.2):

$$\begin{aligned} \Gamma^{\mu}_{\alpha\beta} &= \frac{1}{2}(\eta^{\mu\nu} + h^{\mu\nu})(\eta_{\alpha\nu,\beta} + \eta_{\beta\nu,\gamma} - \eta_{\alpha\beta,\nu} + h_{\alpha\nu,\beta} + h_{\beta\nu,\gamma} - h_{\alpha\beta,\nu}) \\ &= \frac{1}{2}\eta^{\mu\nu}(h_{\alpha\nu,\beta} + h_{\beta\nu,\gamma} - h_{\alpha\beta,\nu}) = \frac{1}{2}(h^{\mu}_{\alpha,\beta} + h^{\mu}_{\beta,\alpha} - h^{\mu}_{\alpha\beta}), \end{aligned} \quad (1.7)$$

where in the first line we have used the fact that $|h_{\mu\nu}| \ll 1$ and that the derivatives of the Minkowski metric $\eta_{\mu\nu}$ are zero. Also, we have used $\eta_{\mu\nu}$ instead of $g_{\mu\nu}$ to raise the indices of the components $h_{\mu\nu}$, because $h_{\mu\nu}$ behaves like a tensor defined on a flat spacetime, for which the metric is $\eta_{\mu\nu}$. From equations (1.4) and (1.7), we obtain:

$$R_{\mu\nu} = \Gamma^{\alpha}_{\mu\nu,\alpha} - \Gamma^{\alpha}_{\mu\alpha,\nu} = \frac{1}{2}(h^{\alpha}_{\mu,\nu\alpha} + h^{\alpha}_{\nu,\mu\alpha} - h^{\alpha}_{\mu\nu,\alpha} - h_{,\mu\nu}), \quad (1.8)$$

where the last term is $h_{,\mu\nu} \equiv h^{\alpha\beta}_{,\mu\nu} = \eta^{\alpha\beta}h_{\alpha\beta,\mu\nu}$. If $R_{\mu\nu}$ is contracted once more, the curvature scalar R is obtained, which can be combined with equation (1.8) and substituted in equation (1.3):

$$G_{\mu\nu} = \frac{1}{2}(h^{\alpha}_{\mu\alpha,\nu} + h^{\alpha}_{\nu\alpha,\mu} - h^{\alpha}_{\mu\nu,\alpha} - h_{,\mu\nu} - \eta_{\mu\nu}(h^{\alpha\beta}_{,\alpha\beta} - h^{\beta}_{,\beta})). \quad (1.9)$$

The last equation can be simplified by introducing a modified expression for the metric perturbation, called the trace-reversed metric perturbation:

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h. \quad (1.10)$$

The Einstein field equation is rewritten as:

$$-\bar{h}_{\mu\nu,\alpha} - \eta_{\mu\nu}\bar{h}_{\alpha\beta}^{\alpha\beta} + \bar{h}_{\mu\alpha,\nu} + \bar{h}_{\nu\alpha,\mu} = \frac{16\pi G}{c^4}T_{\mu\nu}. \quad (1.11)$$

The coordinate freedom of GR gives the possibility of choosing any coordinate system or gauge. This means that there might be other coordinate systems in which the metric can still be written as the Minkowski metric plus a perturbation as shown in equation (1.2), but this perturbation will be different. For this reason, the decomposition of the metric into a flat background plus a perturbation is not unique. Gauge transformations correspond to symmetries of the field equations, meaning that the field equations are invariant under such transformations. We will use the harmonic gauge (equivalent to the Lorentz gauge condition of electromagnetism), given by:

$$\square x^\mu = 0, \quad (1.12)$$

in which the following condition is true:

$$\bar{h}_{\mu\alpha}^{\mu} = 0. \quad (1.13)$$

There is still some gauge freedom remaining, since the coordinates can be shifted by (infinitesimal) harmonic functions:

$$x'^\mu \equiv x^\mu + \xi^\mu, \quad (1.14)$$

which induces a redefinition of the gravitational field tensor:

$$\bar{h}'_{\mu\nu} = \bar{h}_{\mu\nu} - \xi_{\nu,\mu} - \xi_{\mu,\nu} + \eta_{\mu\nu}\xi_{,\lambda}^{\lambda}. \quad (1.15)$$

It can be proved that ξ^μ must satisfy:

$$\square \xi^\mu = 0 \quad (1.16)$$

so that the new gravitational field is in agreement with the harmonic gauge condition.

Within the harmonic gauge, all the terms on the left side of equation (1.11) become zero, except for the first one (that represents the d'Alembert operator), and the linearized Einstein field equation becomes:

$$-\square \bar{h}_{\mu\nu} \equiv \left(-\frac{1}{c^2} \frac{\partial^2}{\partial t^2} + \nabla^2\right) \bar{h}_{\mu\nu} = \frac{16\pi G}{c^4} T_{\mu\nu}, \quad (1.17)$$

which in vacuum (absence of any sources) is:

$$\square \bar{h}_{\mu\nu} = 0, \quad (1.18)$$

which is the typical wave equation. For example, when there is one spatial dimension and one time dimension, this equation can be written as:

$$\left(-\frac{1}{c^2} \frac{\partial^2}{\partial t^2} + \frac{\partial^2}{\partial x^2}\right) \bar{h}_{\mu\nu} = 0. \quad (1.19)$$

A general solution to equation (1.19) is the plane wave solution (more generally, there could be a superposition of different plane waves that would still solve the wave equation):

$$\bar{h}_{\mu\nu}(t, \vec{x}) = \Re[A_{\mu\nu}(\vec{k})e^{ik_\alpha x^\alpha}] = A_{\mu\nu}(\vec{k}) \cos(k_\alpha x^\alpha), \quad (1.20)$$

where \Re denotes the real part, and the components $A_{\mu\nu}$ and k_μ are the amplitude tensor (whose components depend on the wave vector \vec{k} but are independent of \vec{x} and t) and the wave vector $k_\mu = [w, k_x, k_y, k_z]$

respectively, where w is the frequency of the wave and k_i are the wavenumbers, both constant. The amplitude gives the strength of the wave (the displacement of spacetime that it creates), and it is related to the source of the gravitational waves.

When the solution given by equation (1.20) is inserted in equation (1.18) the next condition is obtained:

$$\eta^{\rho\sigma}\partial_\rho\partial_\sigma\bar{h}_{\mu\nu} = \eta^{\rho\sigma}\partial_\rho(ik_\sigma\bar{h}_{\mu\nu}) = -\eta^{\rho\sigma}k_\rho k_\sigma\bar{h}_{\mu\nu} = -k_\alpha k^\alpha\bar{h}_{\mu\nu} = 0 \longrightarrow k_\alpha k^\alpha = 0, \quad (1.21)$$

which can be read as $-w^2 + k_x^2 + k_y^2 + k_z^2 = 0$, implying that the wave vector is light-like (meaning that gravitational waves travel at the speed of light). Applying the harmonic gauge condition given by equation (1.13) to the plane wave solution, we find:

$$\partial_\mu\bar{h}^{\mu\nu} = \partial_\mu\left(A^{\mu\nu}e^{ik_\sigma x^\sigma}\right) = iA^{\mu\nu}k_\mu e^{ik_\sigma x^\sigma} = 0 \longrightarrow A_{\mu\alpha}k^\alpha = 0, \quad (1.22)$$

implying that the amplitude tensor is orthogonal to the wave vector (the waves are transverse). This relation can be written as four equations that impose four conditions on $A_{\mu\alpha}$, which reduces the independent components of the symmetric amplitude tensor from ten to six.

Due to the coordinate freedom (the freedom in choosing the four components of the vector ξ_μ), the actual number of independent components of $A_{\mu\alpha}$ can be reduced to two. We will use the transverse-traceless (TT) gauge (a subgauge of the harmonic gauge, also called radiation gauge) to do this. The TT gauge consists of a choice of coordinates (t, x, y, z) corresponding to an inertial (Lorentz) frame in the unperturbed flat background, which makes explicit that the perturbation is transverse (orthogonal to the direction of propagation), and traceless ($A^\alpha_\alpha = 0$), meaning that the perturbation does not compress or expand elements of spacetime, but instead induces a (volume-preserving) strain only. Assuming that the waves propagate in the z direction ($k_x = k_y = 0$) and making use of the TT gauge, we can write $A_{\mu\nu}$ as:

$$A_{\mu\nu}^{TT}(t, z) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad h_{\mu\nu}^{TT}(t, z) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{xx} & h_{xy} & 0 \\ 0 & h_{xy} & -h_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (1.23)$$

where $h_{xx} = A_{xx} \cos(w(t-z))$ and $h_{xy} = A_{xy} \cos(w(t-z))$. In the TT -frame, $\bar{h}_{\mu\nu}^{TT} = h_{\mu\nu}^{TT}$.

In general, a gravitational wave propagating along the z -axis can be expressed in the TT -frame as a linear combination of two polarizations, usually named plus polarization $h_+ \equiv h_{xx}$ and cross polarization $h_\times \equiv h_{xy}$ (the two independent components of the amplitude tensor):

$$h = h_+e_+ + h_\times e_\times, \quad (1.24)$$

where $e_+ = \hat{x} \otimes \hat{x} - \hat{y} \otimes \hat{y}$ and $e_\times = \hat{x} \otimes \hat{y} + \hat{y} \otimes \hat{x}$ are the vector basis of the wave (\otimes indicates a tensor product):

$$e_+ = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad e_\times = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (1.25)$$

These two polarizations are not uniquely defined, as other combinations can be used. For example, the right-handed and left-handed circular polarizations basis can also be defined [14]:

$$e_R = \frac{1}{\sqrt{2}}(e_+ + ie_\times) \quad \text{and} \quad e_L = \frac{1}{\sqrt{2}}(e_+ - ie_\times), \quad (1.26)$$

$$e_+ = \frac{1}{\sqrt{2}}(e_R + e_L) \quad \text{and} \quad e_\times = \frac{1}{\sqrt{2}i}(e_R - e_L). \quad (1.27)$$

When only one of the two polarizations is active, the wave is said to be linearly polarized. If instead both polarizations are active and equal, the wave is said to be circularly polarized. Otherwise, if both polarizations are active but not equal, the wave is said to be elliptically polarized.

To obtain the expressions of this subsection we have assumed that no sources or additional perturbations exist between the source and the observer. The propagation of GWs would be modified if this were not the case, as for example happens in lensing of gravitational waves, which modifies the amplitude and phase of the emitted waveforms [15], opposite to what can be seen in equation (1.20), where the components of the amplitude tensor do not depend on time or the spatial coordinates. Even in an idealized Universe containing only GWs, the interaction between them would produce backscatter and tails [16], which would modify the previous expressions.

1.1.2 | Production of gravitational waves

In the previous subsection we described the propagation of gravitational waves, assuming an ideal situation where the stress-energy tensor is equal to 0 in all spacetime. In order to describe the generation of gravitational waves, this assumption cannot be made, and the equation that needs to be solved is (in linearized theory) [13]:

$$\square \bar{h}_{\mu\nu}(t, \vec{x}) = -\frac{16\pi G}{c^4} T_{\mu\nu}(t, \vec{x}). \quad (1.28)$$

This differential equation can be solved using a Green's function $G(t, \vec{x}, t', \vec{x}')$ (where \vec{x}' is a vector to position inside the source of GWs):

$$\bar{h}_{\mu\nu}(t, \vec{x}) = -\frac{16\pi G}{c^4} \int d^4x' G(t, \vec{x}, t', \vec{x}') T_{\mu\nu}(t', \vec{x}'). \quad (1.29)$$

Green's function can be found by solving the following equation:

$$\square G(t, \vec{x}, t', \vec{x}') = \delta(t - t') \delta(\vec{x} - \vec{x}'), \quad (1.30)$$

which depends on boundary conditions. This function describes the field generated at the point (t, \vec{x}) by the source point (t', \vec{x}') . For a radiation problem, the standard solution is the retarded Green's function:

$$G(t, \vec{x}, t', \vec{x}') = -\frac{1}{4\pi|\vec{x} - \vec{x}'|} \delta(t_{ret} - t') \Theta(t - t'), \quad (1.31)$$

where the Dirac delta term means that we are evaluating the function only at retarded time $t_{ret} = t - |\vec{x} - \vec{x}'|/c$ and the function $\Theta(t - t_{ret})$ represents the step function, meaning that only times at the future of t' are taken into account. For the Newtonian case, Green's function represents the potential produced by a point mass located at \vec{x}' . We can apply this expression to equation (1.29):

$$\bar{h}_{\mu\nu}(t, \vec{x}) = \frac{4G}{c^4} \int d^3x' dt' \frac{1}{|\vec{x} - \vec{x}'|} \delta(t_{ret} - t') T_{\mu\nu}(t', \vec{x}') = \frac{4G}{c^4} \int d^3x' \frac{1}{|\vec{x} - \vec{x}'|} T_{\mu\nu}(t - |\vec{x} - \vec{x}'|/c, \vec{x}'). \quad (1.32)$$

This equation shows that the disturbance of the metric at time t and position \vec{x} is produced by a sum (integral) of the influences from the energy and momentum sources given by the tensor $T_{\mu\nu}$ at time t_{ret} and positions \vec{x}' , belonging to the past light cone, as shown in figure 1.1. At an observing point far (at distance d) from a source of size R , we can expand the term $|\vec{x} - \vec{x}'| \approx d - \vec{x}' \cdot \hat{x} + \mathcal{O}(R^2/d)$.

The stress-energy tensor $T_{\mu\nu}$ can be expressed in terms that ease its integration over the source. We follow the multipolar expansion described in [13], which is based on a Taylor expansion in the parameter $\vec{x}' \cdot \vec{n}/c$ with the condition of $\omega_s R \ll c$, where ω_s is a typical frequency of the source (low-velocity expansion). After applying this Taylor expansion, the stress-energy tensor is:

$$T_{\mu\nu}(t - |\vec{x} - \vec{x}'|/c, \vec{x}') \approx T_{\mu\nu}(t - \frac{d}{c}, \vec{x}') + \frac{x'^i n^i}{c} \partial_0 T_{\mu\nu} + \frac{1}{2c^2} x'^i x'^j n^i n^j \partial_0^2 T_{\mu\nu} + \dots, \quad (1.33)$$

where derivatives are evaluated at $(t - d/c, \vec{x}')$. Substituting into (1.32):

$$\bar{h}_{\mu\nu}(t, \vec{x}) \approx \frac{4G}{dc^4} \int d^3x' T_{\mu\nu}(t - \frac{d}{c}, \vec{x}') + \frac{4G}{dc^4} \int d^3x' \frac{x'^i n^i}{c} \partial_0 T_{\mu\nu} + \dots \quad (1.34)$$

$$\equiv \frac{4G}{dc^4} \left(S_{\mu\nu}(t - \frac{d}{c}, \vec{x}') + \frac{n^i}{c} \dot{S}_{\mu\nu i}(t - \frac{d}{c}, \vec{x}') + \dots \right). \quad (1.35)$$




Figure 1.1: Schematic showing that the effect at time t and position \bar{x} comes from the integrated sources (red circle) at time t_{ret} and positions \bar{x}' , which are contained within the past light cone.

This expression is the basis for the multipole expansion: each new term has an extra factor of v/c compared to the previous one.

In order to clarify the physical meaning of the terms S of the previous equation, we will express them in terms of the mass M and momentum P^i moments, such as [13]:

$$M = \frac{1}{c^2} \int d^3 x' T^{00}(t, \bar{x}') \quad (1.36)$$

$$M^i = \frac{1}{c^2} \int d^3 x' T^{00}(t, \bar{x}') x^i \quad (1.37)$$

$$M^{ij} = \frac{1}{c^2} \int d^3 x' T^{00}(t, \bar{x}') x^i x^j \quad (1.38)$$

$$P^i = \frac{1}{c} \int d^3 x' T^{0i}(t, \bar{x}') \quad (1.39)$$

$$P^{ij} = \frac{1}{c} \int d^3 x' T^{0i}(t, \bar{x}') x^j \quad (1.40)$$

$$P^{ijk} = \frac{1}{c} \int d^3 x' T^{0i}(t, \bar{x}') x^j x^k. \quad (1.41)$$

From the conservation of the stress-energy tensor ($\nabla_\mu T_{\mu\nu} = 0$) it is possible to relate the time derivatives of these moments, as described in [13]:

$$\dot{M} = 0, \quad \dot{M}^i = P^i, \quad \dot{M}^{ij} = P^{ij} + P^{ji} \quad (1.42)$$

$$\dot{P}^i = 0, \quad \dot{P}^{ij} = S^{ij}, \quad \dot{P}^{ijk} = S^{ijk} + S^{ikj}. \quad (1.43)$$

The conservation of mass and momentum expressed in the previous equations is due to the linearized theory formalism, where back-reaction effects are neglected (these effects only appear at higher v/c orders). After some calculations, the following expressions are derived:

$$S^{ij} = \frac{1}{2} \ddot{M}^{ij} \quad (1.44)$$

$$\dot{S}^{ijk} = \frac{1}{6} \ddot{M}^{ijk} + \frac{1}{3} (\ddot{P}^{ijk} + \ddot{P}^{jik} - 2\ddot{P}^{kij}). \quad (1.45)$$

To lowest order, the gravitational wave given by equation (1.35) is:

$$\bar{h}_{ij}(t, \bar{x}) \approx \frac{2G}{dc^4} \ddot{M}^{ij}(t - \frac{d}{c}, \bar{x}'), \quad (1.46)$$

where (to lowest order) the second mass moment can be expressed as $M^{ij} = \int d^3 x' \rho(t, \bar{x}') x^i x^j$, where ρ is the Newtonian density of the source. Working in the TT-frame, the gravitational wave is:

$$\bar{h}_{ij}^{TT}(t, \bar{x}) \approx \frac{2G}{dc^4} \ddot{M}_{ij}^{TT}(t - \frac{d}{c}, \bar{x}'). \quad (1.47)$$

The previous expression shows that time-varying monopolar and dipolar moments do not emit gravitational waves. The dipolar moment of a collection of particles is $\vec{d}_m = \sum_i m_i \vec{x}_i$. As in electromagnetism,

the electric dipolar emission depends on the second time derivative of the dipolar moment. Taking the first derivative, we see that $\vec{d} = \sum_i m_i \vec{v}_i = \vec{p}$ is equal to the linear momentum, which is a conserved quantity. Furthermore, electromagnetism has the dipolar magnetic moment, which in the gravitational case can be expressed as $\vec{d}_c = \sum_i \vec{x}_i \times (m_i \vec{v}_i)$, equal to the angular momentum, which also is a conserved quantity. For this reason, no gravitational dipolar radiation emission can exist.

Since we work in the TT-frame, it does not matter whether the tensor that represents the mass moment has a trace or not, since it will be removed by the projection onto the TT frame:

$$Q_{ij} = I_{ij} - \frac{1}{3} \delta_{ij} I_{kk} \longrightarrow \Lambda_{ijkl} I^{kl} = \Lambda_{ijkl} Q^{kl}, \quad (1.48)$$

where the projection tensor Λ is explained below. This is why different definitions or quantities are usually used within the quadrupole formula, such as the second mass moment M , the quadrupole tensor Q , or the inertia tensor I . All of these tensors are rank-2 tensors representing the second-order moment of mass with respect to position, but with different projections of their diagonal ($i = j$) and off-diagonal ($i \neq j$) components. Canonically, the quadrupole tensor is used to quantify deviations from sphericity. For example, the moment of inertia of a sphere is $I = \text{diag}[0.4MR^2, 0.4MR^2, 0.4MR^2]$, with off-diagonal components equal to zero. It can be seen that the definition of the quadrupole, which is the traceless moment of inertia, is constructed to give a 0 quadrupole tensor for a sphere (i.e. non-zero quadrupole tensors indicate non-sphericity). The mass-quadrupole moment can also be defined as the coefficient of the $1/r^3$ term in an expansion in powers of $1/r$ of the Newtonian gravitational potential (far from the source), thus establishing a clear relationship with a measurable quantity.

Now we want to compute the amplitudes of a wave that propagates in a direction given by $\hat{n} = [\sin \theta \sin \phi, \sin \theta \cos \phi, \cos \theta]$ in the TT frame, where the angles denote the difference of orientation between the frame where the moments are calculated and the direction of the wave. To do this, the matrix \ddot{M}_{kl} has to be rotated to that specific frame, since the TT frame is defined differently for different directions (i.e. the transverse components in different directions are not equal). The TT projection tensor is given by [13]:

$$\Lambda_{ijkl} = P_{ik} P_{jl} - \frac{1}{2} P_{ij} P_{kl} \longrightarrow M_{ij}^{TT} = \Lambda_{ijkl} M_{kl} = (P_{ik} P_{jl} - \frac{1}{2} P_{ij} P_{kl}) M_{kl}, \quad (1.49)$$

where

$$P_{ij} = \delta_{ij} - n_i n_j = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} (\sin \theta \sin \phi)^2 & \sin^2 \theta \sin \phi \cos \phi & \cos \theta \sin \theta \sin \phi \\ \sin^2 \theta \sin \phi \cos \phi & (\sin \theta \cos \phi)^2 & \cos \theta \sin \theta \cos \phi \\ \cos \theta \sin \theta \sin \phi & \cos \theta \sin \theta \cos \phi & \cos^2 \theta \end{pmatrix}, \quad (1.50)$$

which gives:

$$\ddot{M}_{ij}^{TT} = \Lambda_{ijkl} \ddot{M}_{kl} = \begin{pmatrix} A & B & 0 \\ B & -A & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (1.51)$$

$$A = \ddot{M}_{11}(\cos^2 \phi - \sin^2 \phi \cos^2 \theta) + \ddot{M}_{22}(\sin^2 \phi - \cos^2 \phi \cos^2 \theta) - \ddot{M}_{33} \sin^2 \theta - \ddot{M}_{12} \sin 2\phi(1 + \cos^2 \theta) + \ddot{M}_{13} \sin \phi \sin 2\theta + \ddot{M}_{23} \cos \phi \sin 2\theta \quad (1.52)$$

$$B = (\ddot{M}_{11} - \ddot{M}_{22}) \sin 2\phi \cos \theta + 2\ddot{M}_{12} \cos 2\phi \cos \theta - 2\ddot{M}_{13} \cos \phi \sin \theta + 2\ddot{M}_{23} \sin \phi \sin \theta. \quad (1.53)$$

The two polarizations of the strain can be calculated from the previous equation and equations (1.47) and (1.25):

$$\begin{aligned} h_+(t, \theta, \phi) &= \frac{G}{c^4} \frac{1}{d} [\ddot{M}_{11}(\cos^2 \phi - \sin^2 \phi \cos^2 \theta) + \ddot{M}_{22}(\sin^2 \phi - \cos^2 \phi \cos^2 \theta) - \ddot{M}_{33} \sin^2 \theta \\ &\quad - \ddot{M}_{12} \sin 2\phi(1 + \cos^2 \theta) + \ddot{M}_{13} \sin \phi \sin 2\theta + \ddot{M}_{23} \cos \phi \sin 2\theta] \\ h_\times(t, \theta, \phi) &= \frac{G}{c^4} \frac{1}{d} [(\ddot{M}_{11} - \ddot{M}_{22}) \sin 2\phi \cos \theta + 2\ddot{M}_{12} \cos 2\phi \cos \theta \\ &\quad - 2\ddot{M}_{13} \cos \phi \sin \theta + 2\ddot{M}_{23} \sin \phi \sin \theta]. \end{aligned} \quad (1.54)$$

Thus, once the moments \ddot{M} are known it is easy to calculate the strain produced by a source. This quadrupolar formula shows that a gravitational wave is created only if the time derivative of a distribution of energy is not equal to zero. It can be seen that the gravitational wave emission is anisotropic, i.e. the strength depends on the direction. For example, for a black hole binary or an asymmetric rotating neutron star it is known that the radiation is stronger in the direction of the total angular momentum, i.e. when the system is observed face-on (the direction of angular momentum directed towards the observer).

We have obtained the canonical quadrupole formula starting from linearized general relativity. As explained in many references like [17], in principle this is not a correct procedure to obtain the GWs generated from objects such as neutron stars, since linearized GR assumes a weak-field gravity, a situation that is not true inside or very close to a neutron star. A more correct procedure is to use a Post-Newtonian (in powers of v/c) and Post-Minkowskian (in powers of G) expansion using the relaxed Einstein equations, where it is not assumed that $h_{\mu\nu} \ll 1$ (where now $h_{\mu\nu} = \eta_{\mu\nu} - (-g)^{1/2}g_{\mu\nu}$) and a pseudo-tensor that contains the stress-energy produced by the gravitational field itself (called Landau-Lifshitz pseudo-tensor), which depends on $h_{\mu\nu}$, is added to the usual stress-energy tensor. Even though this is true, several studies have shown that equation (1.47) is correct to lowest order since by using the correct formalism it is also obtained [18].

The energy emission rate (also called gravitational wave luminosity) is given by [13]:

$$L_{GW} = \frac{1}{5} \frac{G}{c^5} \langle \ddot{M}_{jk} \ddot{M}^{jk} \rangle, \quad (1.55)$$

where $\langle \dots \rangle$ denotes a temporal average over several periods. As an example, we can calculate the emission rate of a rotating neutron star. If it is perfectly symmetric, the derivative of the quadrupolar mass tensor is zero, but if it has an asymmetry, characterised by the ellipticity parameter $\epsilon = (I_{xx} - I_{yy})/I_{zz}$, the derivative is not equal zero, and the emission rate is approximately given by:

$$L_{GW} \propto \frac{G}{c^5} \epsilon^2 I_{zz}^2 \nu^6, \quad (1.56)$$

where $I_{zz} \sim MR^2$ is the moment of inertia with respect to the rotation axis, and ν is the frequency of rotation. Assigning specific values to these parameters, we can arrive at an estimate of the gravitational wave luminosity: $G/c^5 \sim 10^{-53}$ s/J, $I_{zz} \sim 10^{38}$ kg m² (the canonical moment of inertia for a neutron star), $\epsilon \sim 10^{-6}$, and $\nu \sim 100$ Hz, giving a value for the luminosity of $L_{GW} \sim 10^{23}$ W, which is lower than the luminosity of the Sun due to electromagnetic radiation ($L_{\odot} = 10^{26}$ W).

1.2 | Astrophysical sources of gravitational waves

As seen in the previous section, any accelerated mass emits gravitational waves. For example, humans can produce GWs, but the amplitude of such waves would be extremely weak [20]: $h \sim 10^{-43}$. From the wide variety of objects that we find in the Universe, we will only take into account those moving close to relativistic speeds and which have high compactness:

$$C = \frac{2GM}{c^2 R}, \quad (1.57)$$

where R is the radius of the source and M its mass. This compactness parameter attains its maximum of 1 for a Schwarzschild black hole, which has $R = 2GM/c^2$.

The astrophysical and cosmological sources of gravitational waves that we expect to detect with the current detectors are usually divided in four different groups, depending on the duration that the signal is detectable and on the knowledge that we have about the emitted waveform:

- **Compact Binary Coalescences (CBC)**: gravitational waves coming from binary systems with two black holes, two neutron stars, or a black hole and a neutron star. When these compact objects are orbiting each other, they emit gravitational waves due to the time-changing quadrupole

moment. This emission of gravitational waves gradually carries away some of the orbital energy, thus gradually decreasing the separation of the two objects. This happens slowly at first, but as the orbit gets tighter the gravitational wave amplitude and the frequency become higher and the process accelerates until eventually these objects collide and merge, creating the characteristic chirp in the frequency evolution. The gravitational waves emitted by binary mergers are not random bursts of energy, but highly structured waveforms that carry a lot of information about the emitting systems, such as the masses of the two inspiraling objects or their spins (how fast they are rotating around their own axes). These signals are only present in the ground-based detectors during seconds or minutes, because only a few orbits before the merger can be measured. This is the type of signal that the second generation ground-based interferometric detectors have detected until now: ten detections of binary black holes and one of a binary neutron star (prior to the beginning of the third observing run of the Advanced detectors) [11]. Because the waveform is analytically well understood, optimal detection methods such as matched filtering can be used.

- **Burst:** burst signals are expected to come mainly from supernovae and gravitational collapse, but other exotic objects like cosmic strings or unknown phenomena could produce them as well. Gravitational collapse is one of the most violent events known to astronomy, and probably, the least understood source of gravitational radiation. This is because we have little direct information about what is happening in the interior of the collapse, and accurate predictions about the shape of the radiation are not possible. These signals have a very short duration, and accurate theoretical models of the waveforms do not exist. Because there is no theoretical model for these types of signals, matched filtering cannot be used and other unmodeled methods such as coincidence tests between different detectors are routinely used.
- **Stochastic background:** the stochastic background of gravitational waves comes from the superposition of a large number of unresolved gravitational wave sources of astrophysical and cosmological [21] origin. Unresolvable sources are those which we cannot distinguish individually, either because they are too quiet or because there are simply too many occurring at once. Cosmological backgrounds are predicted to have been produced by sources that existed in the very early Universe, as for example during the inflation period (at $\sim 10^{-36}$ seconds after the Big Bang), while astrophysical backgrounds are predicted to have been produced by systems of massive stars such as the neutron stars and black holes that we see today. Thus, depending on the type of gravitational-wave background we detect, we may learn about the state of the Universe just a few moments after the Big Bang or how the Universe is evolving in more recent times. The strength of the gravitational-wave background at different frequencies strongly depends on the type of sources that produce it, but it is usually much lower than the already detected CBC signals. Similarly to the burst case, analytical models of the waveforms do not exist and the search pipelines mainly depend on correlating the data streams from different detectors in order to find a coincident signal. The stochastic background is assumed to be continuous and search methods can integrate data for longer durations.
- **Continuous gravitational waves (CWs):** the main topic of this thesis, CWs mainly come from rotating neutron stars with an asymmetry around their rotating axis. These waves are almost monochromatic (their frequency changes very little) and have a very long duration. The signal amplitude is very low and long observation times (up to months or even years) are required to build up enough signal-to-noise ratio to be detectable. More details about CWs are discussed in chapter 3.

Besides these four main groups, more sources can produce gravitational waves outside of the frequency band of the ground-based detectors. Figure 1.2 shows the different types of sources and the frequency of the GWs that they can emit, which is related to the size and the mass of the system emitting them.




Figure 1.2: Gravitational-wave spectrum, showing the different frequencies and their corresponding sources and detectors. Plot made with [22].

The frequency band of good sensitivity for the current ground-based detectors is located between ~ 20 Hz and several kHz, so with these detectors we will only be able to detect the sources that emit GWs in this range. For this reason, many sources such as black holes orbiting a supermassive black hole or galactic white dwarf binaries will only be detected with other types of gravitational-wave detectors such as LISA, which are discussed in the next section.

1.3 | Gravitational wave detectors

In order to detect gravitational waves, extremely sensitive detectors need to be designed and built to sense changes in length many orders of magnitude smaller than the diameter of the proton. In this section we review past, current and future gravitational wave detectors such as resonant-mass antennas and interferometers.

The first operative gravitational-wave detectors were resonant-mass antennas. The main idea behind these detectors is that a passing gravitational wave will excite the longitudinal modes of the antenna by absorbing energy from the GW, and by using a transducer this excitation can be measured. In the 1960s, Weber designed the first resonant-mass antenna, an aluminium cylindrical bar of 3 m long and 0.6 m of diameter, which weighed around a ton [23]. After this first antenna, several more antennas with very similar characteristics were built around the world. Although Weber claimed several times to have detected gravitational-wave events, none of the other experiments could confirm these detections. In the 80s and 90s new designs and improvements were proposed that formed the second generation of resonant-mass antennas, such as the implementation of cryogenic temperatures or different shapes such as spherical. Up until now, none of these resonant-mass detectors has detected any gravitational wave. Typical sensitivities of the first generation of detectors were around $10^{-17} \text{ Hz}^{-1/2}$, while the second generation antennas improved this sensitivity up to a thousand times, although such sensitivities are only attained within a narrow frequency bandwidth, of the order of tens of Hertz. A detailed review about resonant-mass detectors can be found at [23]. New developments in resonant-mass detectors are still being proposed, such as a new idea called TOBA (Torsion-Bar Antenna) that was proposed in 2010 [24], which already has a working prototype. This new instrument is made of two perpendicular bars, where instead of measuring internal excitations of each bar a relative rotation between both bars is sensed. TOBA is

supposed to detect gravitational waves in the 1 mHz to 1 Hz regime at a sensitivity only slightly worse than the proposed space-based detectors such as LISA, which are discussed below. Furthermore, another new idea consists of using superfluid helium as the resonant-mass in order to detect gravitational waves around a small bandwidth at kHz frequencies, where preliminary studies show that sensitivities similar or even better than interferometric ground-based detectors could be achieved [25].

Another type of instrument that can sense gravitational waves is the interferometric gravitational-wave detector. The seeds of the idea for interferometric detectors date from the 50's [26], and the first detailed descriptions of these detectors were produced in the 70's [27] by Rainer Weiss. The first interferometric detectors were constructed in the 90's, and began operations in the 00's. Some of the ground-based interferometric gravitational wave detectors that exist are (more details about how these interferometers work are given in the next subsection):

- **LIGO:** the Laser Interferometer Gravitational-Wave Observatory (LIGO) consists of two 4 km detectors, one in Hanford (WA) called H1 and another in Livingston (LA) called L1 (Initial LIGO had a third detector of 2 km called H2 inside H1). These two observatories are separated by ~ 3000 km, and they are operated by the LIGO Scientific Collaboration. Figure 1.3 shows an aerial picture of the Hanford interferometric detector. The Initial LIGO detectors were operating between 2002 and 2010, in six different science runs, not finding any gravitational wave signal. From ~ 2010 –2015 the detectors underwent a series of major upgrades to enhance their sensitivity, becoming the first “second generation” gravitational-wave detectors, called Advanced LIGO [28]. These enhancements have made Advanced LIGO almost an order of magnitude more sensitive than Initial LIGO. A third Advanced LIGO detector of similar characteristics is expected to be built in India and to begin operations around 2025.
- **Virgo:** the Virgo (V1) detector is a 3 km interferometer located in Cascina (Italy), operated by the Virgo Collaboration. Similarly to LIGO, this detector was active during its first generation phase, and it was upgraded to Advanced Virgo during the first years of the 2010s, joining the last month of the O2 observing run and participating in the first detection of a binary neutron star [29].
- **KAGRA:** a Japanese 3 km detector located in the same place as the Kamioka neutrino observatory (200 meters beneath Earth) [30]. KAGRA participated with GEO in a short observing run called O3KG that began just after the O3 observing run was finished. KAGRA test masses are maintained at very low temperatures with cryogenic technology, thus decreasing the Brownian/thermal noise and improving its sensitivity. Before KAGRA there was TAMA-300, which was a 300 m interferometer built in the city of Tokyo. The aim of the project was to develop advanced techniques needed for building a future interferometer. TAMA was followed by CLIO, a 100 m cryogenic detector.
- **GEO:** GEO-600 is a 600 m detector built as a collaboration between the United Kingdom and Germany, located near Hannover (Germany) [31]. Its short length makes it the least sensitive detector, and, because of this, it cannot compete with the other detectors. It has been mainly used as a test bank for the technology that can be implemented in the other interferometric detectors, and it operates in the so called astro-watch mode, which covers epochs when the other detectors are not operational in case that an unexpected but big event such as close-by supernova happens.

A science or observing run is a period of time during which the detectors are active, taking data that will be analysed to search for gravitational-wave signals. Due to periodic maintenance, seismic noise, equipment failures, and other problems, the duty cycle of a science run (the fraction of time where the detectors are recording data that will be analysed to search for gravitational waves) is always less than 100%, varying for each detector and each science run, usually having values between $\sim 50\%$ and $\sim 90\%$.

The Initial LIGO interferometers carried out six initial science runs called S1 to S6 of varying durations (Initial Virgo carried out four science runs called VSR1 to VSR4), while for the second generation




Figure 1.3: Aerial image of the LIGO Hanford gravitational wave observatory, located in Washington, EEUU, with 4 km arms. Source: Caltech/MIT/LIGO Lab, <https://www.ligo.caltech.edu/WA/image/ligo20150731f>.

of interferometers there have been three observing runs up until now, called O1 (which started on 12 September 2015 and ended on 19 January 2016, with only the two Advanced LIGO detectors participating), O2 (which began on 30 November 2016 and lasted until 25 August 2017, with Advanced Virgo joining for the last month), and O3 (which started on 1 April 2019 and ended on 27 March 2020). Plans for future observing runs are discussed in [32].

Currently there is an ongoing discussion about the planning and design of the third generation ground-based detectors, which aim to increment the sensitivity by at least an order of magnitude. There are two main proposals: Einstein Telescope, an European project currently at design phase, which will be underground and use cryogenic techniques, with arms of 10 km [33]; Cosmic Explorer, a North-American project which aims to build a 40 km detector at ground level also with cryogenic techniques [34]. The construction of both facilities is expected to start during the 2030s decade.

Besides resonant-mass detectors and classical interferometers, new ideas and techniques are being proposed in order to do ground-based gravitational-wave detection. MAGIS-100 is an atom interferometer of 100 meters that uses falling atoms instead of hanging mirrors as their test masses [35]. It is being constructed in order to test the performance and the possibilities of a future 1 km atom interferometer that would have a sensitivity around $10^{-19} \text{ Hz}^{-1/2}$ between 0.1 and 10 Hz, due to the absence of seismic noise, which would complement the frequency band of the current detectors. Another idea is the speedmeter, where the velocity of the test masses instead of their position is actively measured [36]. Different types of speedmeter designs exist which show promising improvements in sensitivity, although more detailed and thorough tests are pending [37].

Apart from the ground-based gravitational-wave detectors, there is a project called LISA that intends to send an interferometric detector to the outer space, with a planned launch around the 2030s [38]. LISA is a laser interferometer consisting of three spacecrafts in solar orbit, with near equilateral triangular configuration of around 10^6 km baselines. Due to unequal arm lengths, frequency noise from the laser is not exactly canceled when light is recombined at the beamsplitter. For this reason, a time-shifted combination of the output from the different spacecraft is combined to cancel this noise and achieve good sensitivity, in what is called Time Delay Interferometry [39]. Being outside Earth, LISA will not be affected by seismic noise as the ground-based detectors are, making it a perfect instrument to detect sources at frequencies approximately between 10^{-5} and 0.1 Hz, complementing the frequency range of

current detectors. A mission sent in 2015 to test LISA technologies called LISA Pathfinder has published the first set of results [40], attaining an even better than expected sensitivity. Besides LISA, other future space-based detectors such as DECIGO, TianQin/Taiji, Big Bang Observer, AMIGO, and AIGSO are also being planned, which may have slightly different configurations than LISA.

Two more techniques are used to detect gravitational waves at even lower frequency regimes, complementing the range of the ground-based and space-based detectors. Firstly, pulsar timing array (PTA) uses the precise arrival times of the electromagnetic signals from pulsars as the tool to measure the waves [41]. If a gravitational wave passes between the pulsar and the Earth, the time of arrival of the pulse will change, and this variation (correlated between all pulsars) will be seen in the arrival times of the pulses as a residual between the best fitted model and the data. The gravitational-wave frequencies that PTA can detect are around 10^{-9} and 10^{-7} Hz, and are produced mainly by mergers of supermassive black holes, although GWs from decaying cosmic strings and relic post-inflation are also a possibility. The latest data release consists of timing from 65 pulsars, with no detection of GWs up to date [42]. Secondly, at frequencies even lower than 10^{-9} Hz, another possibility is the indirect detection of gravitational waves in the very low frequency band based on measurements of the B+ polarization modes of the cosmic microwave background produced by primordial gravitational waves [43, 44]. These indirect detections could give information about the early Universe that is not accessible through electromagnetic radiation.

1.3.1 | How does an interferometric gravitational wave detector work?

As we explained before, in general relativity gravitational waves have two polarisations, with directions rotated by 45° and phases differing by 90° . A passing gravitational wave will stretch and squeeze space in the direction of every polarisation, and the proper distance between test masses will change. If this change of proper distance between the masses can be detected, we would be able to directly detect gravitational waves.

A typical example of the effect on matter of a passing gravitational wave is pictured in figure 1.4. It shows the effect of the passage of a plane gravitational wave, propagating along the z-axis, on a ring of test particles. The upper panel of the figure shows the case where the metric perturbation has $A_{xx}^{TT} \neq 0$ and $A_{xy}^{TT} = 0$. The lower panel shows the opposite case, where the metric perturbation has $A_{xx}^{TT} = 0$ and $A_{xy}^{TT} \neq 0$. Gravitational waves with $A_{xx}^{TT} \neq 0$ are identical to gravitational waves with $A_{xy}^{TT} \neq 0$, except for a rotation of $\pi/4$ radians. These two cases represent two gravitational wave polarization states, usually denoted by “+” and “×” respectively.

The change in the proper distance between the test particles during the passage of the gravitational wave is the physical quantity which gravitational wave detectors aim to measure, called ΔL . These detectors sense the difference between the travel time of a laser beam going back and forth along the arms of a Michelson interferometer, whose lengths are changing due to the passing of gravitational waves. Within these interferometric detectors, the ring of test particles is substituted by two test masses situated at the end of two perpendicular arms.

In figure 1.5 the composition of a dual-recycled Fabry-Perot Michelson interferometer is shown. Four masses (test masses) are hanging inside two perpendicular arms with an “L” shape, forming a 90° intersection. A laser shoots a beam of light that reaches the beamsplitter, and then half of the laser light is transmitted into one arm while the other half goes into the other arm. Afterwards, light travels through these two arms, reaching the end of the arms where very reflective mirrors are located (the test masses). The laser light travels back and forth many times: the arms are Fabry-Perot cavities which increase the power of the laser’s light. After bouncing, the light returns to the beamsplitter, where it interferes with the beam coming from the other arm. If no gravitational wave passes through the detector during the time that the laser’s light was traveling, the beams from the two arms will cancel out due to a perfect interference and no light will be detected by the photodetector. On the contrary,




Figure 1.4: Illustration of the effect of a gravitational wave traveling in the z direction (perpendicular to the page) on a ring of test particles. It shows the two polarisations, whose direction differs 45° . Time goes from left to right. Every panel shows five different phases (different values of ωt) of the oscillation of the wave, the first and the last one being the same (finishing a full cycle). Source: [45].

if a gravitational wave passed through the detector, the beams will not completely cancel out, and the photodetector will detect some light. The evolution of the pattern formed by the light received at the photodetector gives us information about the gravitational wave, like its amplitude, frequency, etc. Of course, current gravitational-wave detectors such as Advanced LIGO are very complicated apparatus with lots of interacting subsystems, and the previous explanation is a simplification that tries to give a general idea of how an interferometer can detect GWs.

Now, we derive an analytical model to calculate the changes of light travel time between the two arms when a GW has passed. As equation (1.2) shows, the perturbed metric is (for a wave traveling in the z direction) [12]:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} = \begin{pmatrix} -c^2 & 0 & 0 & 0 \\ 0 & 1 + h_+ & h_\times & 0 \\ 0 & h_\times & 1 - h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (1.58)$$

For the sake of simplicity, let us assume a linearly polarized wave $h_\times = 0$ with its wave vector perpendicular to the interferometer plane and its main oscillating directions aligned with the interferometer arms. The light from the laser follows a light-like geodesic, which can be expressed as:

$$ds^2 = 0 = g_{\mu\nu} dx^\mu dx^\nu = -c^2 dt^2 + (1 + h_+) dx^2 + (1 - h_+) dy^2 \quad (1.59)$$

From the last expression we can calculate the time it takes light to travel a certain distance in this perturbed spacetime. For example, when light follows the x -axis direction, the result is:

$$dt = \pm \frac{\sqrt{1 + h_+}}{c} dx, \quad (1.60)$$

where the \pm differentiates between the travel from the beamsplitter at x_0 to the mirror at $x_0 + L$ and the return trip, which has opposite sign. This equation can be integrated:

$$T_1 = \int_{x_0}^{x_0+L} \frac{\sqrt{1 + h_+}}{c} dx \approx \int_{x_0}^{x_0+L} \frac{1 + h_+/2 + \mathcal{O}(h_+^2)}{c} dx \approx \frac{L}{c} + \frac{Lh_+}{2c} = \frac{L}{c} \left(1 + \frac{h_+}{2}\right), \quad (1.61)$$

where in the first approximation we have used a Taylor expansion (when h_+ is of the order of 10^{-20} , it can be seen that higher order effects can be dismissed) and in the second one we have assumed the long wavelength approximation, which assumes that the perturbation h_+ does not depend neither on time or




Figure 1.5: Scheme of an interferometric detector: two masses located close together in the vertex of the “L” shaped structure and the other two at the end of each of the interferometer’s arms. The laser shoots a beam that is separated in the beamsplitter, and travels down through both arms, reaching the end and returning to the beamsplitter. When a gravitational wave passes by, light does not interfere completely and reaches the photodetector. The upper right inset shows the quadruple suspension system that is used to seismically isolate the test masses, and the complementary suspension chains that are used to actuate on the test masses. Source: [46].

space coordinates, since the wavelength of the wave is much longer than the arm and any finite length travel effects can be dismissed. It can be seen that when $h_+ = 0$, the unperturbed solution $T_u = L/c$ is recovered. The change from unperturbed to perturbed time is $\Delta T_1 = T_1 - L/c = Lh_+/(2c)$. This equation shows the importance of the length of the detector arms: for the same external perturbation h_{xx} , doubling the length of the arms means that the difference between perturbed and unperturbed times doubles.

The return trip is:

$$T_2 \approx \int_{x_0+L}^{x_0} -\frac{1+h_+/2}{c} dx = -\frac{L}{c} \left(1 + \frac{h_+}{2}\right) \quad (1.62)$$

$$\Delta T_2 = -\frac{L}{c} \left(1 + \frac{h_+}{2}\right) + L/c = -\frac{Lh_+}{2c}. \quad (1.63)$$

We define the variation on the travel time during a complete round-trip as $\Delta T_x = \Delta T_1 + \Delta T_2 = \frac{Lh_+}{c}$, and correspondingly for the light traveling through the y arm we have $\Delta T_y = -\frac{Lh_+}{c}$ (we have assumed that the unperturbed lengths of the two arms are equal).

Combining the strains produced in both x and y arms, we define the detected strain as:

$$h = \frac{\Delta L}{L} = \frac{\Delta L_x - \Delta L_y}{L} = c \frac{\Delta T_x - \Delta T_y}{L} = \frac{c}{L} \left(\frac{Lh_+}{c} + \frac{Lh_+}{c} \right) = 2h_+. \quad (1.64)$$

The variation in time/length produces a variation of the phase of the laser traveling within the arms of the detector, which is what the photodetector senses:

$$\Delta\phi = \frac{2\pi}{\lambda} \Delta L = \frac{2\pi}{\lambda} 2Lh_+. \quad (1.65)$$

In a realistic scenario, h_+ will be a sinusoidal function and $\Delta\phi$ will oscillate between positive and negative values around a central point.

1.3.2 | Noise in ground-based interferometric detectors

A gravitational wave with an amplitude of $h = 10^{-21}$ will only produce a change in the length of the arms of $\Delta L \propto hL = 4 \times 10^{-21}$ km (much shorter than the width of a proton), with $L = 4$ km being the length of the detector arms. The difficulty of performing such a measurement stems from the fact that the detector is not an ideal instrument: it has some sources of noise of instrumental and environmental origin that can mask the gravitational wave signals we are looking for. This makes the task more challenging: we have to identify these noise sources and understand how they behave in order to minimize them and to prevent confusions between noise and real gravitational wave signals.

The main sources of noise are:

- **Quantum noise:** due to the quantum properties of the photons of the laser. The quantum shot noise comes from the randomness of the times of arrival of the photons at the photodetector, which creates a fluctuation in the power received. Furthermore, the photons of the laser inflict a pressure on the mirror surfaces, called quantum radiation pressure noise, which generates a stochastic force that shakes the mirrors. Quantum noise is the main source of noise at frequencies higher than 10 Hz.
- **Thermal/Brownian noise:** due to movements of the atoms in the mirrors, the coatings of the mirrors, and the suspensions. Thermal noise can be analytically estimated using the fluctuation-dissipation theorem. Thermal noise has a high contribution at frequencies lower than 100 Hz.
- **Technical noise:** a wide variety of effects contribute to technical noise, such as laser frequency noise, noise coming from the feedback loops that are used to maintain the interferometer arms in resonance, or noise from the photodetectors. A discussion of technical noise sources and their contribution during O1 is given in [47].
- **Environmental noise:** due to external sources such as electromagnetic fields, wind and storms, or acoustic noises. The most important environmental noise sources are seismic and gravity-gradient noise, which come from the motion of Earth's ground and from the fluctuations of local gravity fields around the test masses. The use of pendulums (the test masses are hanged with a complex system of fiber suspensions as shown in figure 1.5) reduces the seismic noise by filtering out the frequencies above the natural frequency of the pendulum. For a 0.5 m pendulum, filtering above a few Hz is achieved. Seismic noise is the biggest source of noise at frequencies lower than ~ 10 Hz.

The sum of these noise sources make up a total contribution which is usually characterized as a random Gaussian process. This sum of noises is what limits the sensitivity of the detectors, and it can be characterized by the single-sided power spectral density (PSD) S_n , which is the Fourier transform of the auto-correlation function $\gamma(\tau)$:

$$S_n(f) = 2 \int_{-\infty}^{\infty} \gamma(\tau) e^{-i2\pi f\tau} d\tau. \quad (1.66)$$

The PSD can be estimated through different ways, one of them being the Welch method. It consists of dividing a series of data in shorter chunks, taking the Fourier transform of each of them, calculating the amplitudes for each frequency bin (thereby obtaining a periodogram) and finally combining these chunks by taking a mean between them. As can be seen in any spectrum of gravitational-wave data (such as figure 4.6), the noise can be described as nearly Gaussian (as explained below) but not white, since the value of the PSD is highly dependent on the frequency. This means that the different noise processes maintain a certain correlation.

Two of the main noise artifacts that contribute to the non-stationarity and non-Gaussianity of the detector are lines and glitches.

Lines are sinusoidal-like features that are coherent over long-duration stretches of time, and can be clearly seen in the PSD of a detector as peaks above the noise floor. Lines come from different sources (both known and unknown) such as: the power harmonics of the electricity components (60 or 50 Hz and its harmonics); the violin modes created by the suspensions of the mirrors; the calibration lines, deliberately inserted disturbances that are used to calibrate the detector; environmental disturbances, such as external magnetic fields. The frequency of these lines can be constant in time or it can be time-dependent. Chapter 7 will be devoted to further describe these lines and to list some efforts made during the O2 and O3 observing runs in order to characterize and eliminate them.

Besides lines, there are also many transient artifacts called glitches. Glitches appear with many different shapes, and the sources for most of them are unknown [48]. They can mimic the shape of some gravitational-wave events, thus decreasing the sensitivity of the pipelines. Short-duration glitches not only affect transient searches, since that each time there is a glitch the PSD of the detector is elevated from the noise floor over a wide range of frequencies, thereby also decreasing the sensitivity of longer-duration searches. Both lines and glitches contribute to deteriorate the efficiency of the different search algorithms that try to detect and characterize gravitational waves. For this reason, eliminating or reducing these noise sources is an important endeavour, done side by side to the other improvements which try to lower the other Gaussian noise sources.

In order to study all these different sources of noise and their coupling mechanisms, the output of the detectors not only consists of the main gravitational-wave channel $h(t)$. More than 200000 auxiliary channels are outputting data constantly, which can be used to localize sources of noise affecting different systems or locations of the interferometer [49]. Some examples of these auxiliary channels are magnetometers, microphones, cosmic ray detectors, or an array of photodetectors.

Periods of time in which there are significant problems with the quality of the data are vetoed. To record these vetoes, LIGO uses a system of flags and triggers to classify the data, which typically exclude periods of data on the order of seconds to hours. For example, there are times when some of the photodetectors used to sense the laser field in the detector are overflowing their analog-to-digital converters, and these times have to be vetoed. Data quality triggers are short duration vetoes generated by algorithms that identify significant statistical correlations between a transient in $h(t)$ and transient noise in auxiliary channels. There are three data quality categories, classifying the least and most usable data: CAT1, when a critical issue with a key detector component not operating in its nominal configuration has been identified; CAT2, when there is a known, understood physical coupling to the gravitational wave channel; CAT3, when there is statistical coupling to the gravitational wave channel which is not fully understood.

1.3.3 | Antenna beam pattern functions of an interferometric detector

Gravitational waves travel almost unaltered (because of the little interaction of gravitational waves with matter) from the source to the detector at the speed of light, as predicted by general relativity. As discussed in subsection 1.1, GWs are transverse waves. Because the plane of the detector will not coincide with the plane of the GW oscillation, the GW has to be projected onto the detector's plane in order to accurately describe the real physical effect of the gravitational wave on the time of travel of the laser.

As discussed in subsection 1.1.1, the traveling gravitational-wave H_w can be described in the wave frame (explained below) by:

$$H_w = \sum_p h_p e_p = h_+ e_+ + h_\times e_\times. \quad (1.67)$$

To transform the strain H_w described in the wave frame to the strain H_D in the detector frame, we have to concatenate three frame transformations: wave to celestial, celestial to cardinal, cardinal to detector. Each of these frames is an orthogonal Cartesian frame defined as:

- Wave frame: the z direction is the direction of propagation, while the x direction forms a counterclockwise angle ψ to the projection of the East axis in the sky.
- Celestial frame: also called equatorial frame, the x direction points from the centroid of the Earth to the vernal equinox, the z direction points to the celestial North pole and y axis forms a right-handed orthogonal basis.
- Cardinal frame: the x direction points to the East, the z direction points to the zenith and the y axis forms a right-handed orthogonal basis. In this way, the xy plane is tangent to the surface of the Earth.
- Detector frame: the z direction coincides with the cardinal frame, and the x direction is aligned with the bisector of the two arms of the detector.

For each of these transformations from frame 1 to 2, we perform three rotations given by three Euler angles, defined as: ϕ , the angle between \hat{x}_2 and \vec{N} ; θ , the angle between \hat{z}_2 and \hat{z}_1 ; ψ , the angle between \hat{x}_1 and \vec{N} , where \vec{N} is the line of nodes defined as the intersection between the planes $\hat{x}_1\text{-}\hat{y}_1$ and $\hat{x}_2\text{-}\hat{y}_2$, calculated as $\hat{z}_1 \times \hat{z}_2$. We apply the three Euler rotations following the zxz convention, i.e. a first rotation around the axis \hat{z}_2 by an angle ϕ , a second rotation around the axis \hat{x}_2 by an angle θ and a last rotation around the axis \hat{z}_2 by an angle ψ .

Firstly, we go from the wave frame to celestial frame. This first transformation does not involve any time dependence since the wave and celestial frames do not rotate. It is given by:

$$M_1 = R_z(\psi)R_x(\theta)R_z(\phi) = \begin{pmatrix} \sin(\phi)\sin(\psi) - \cos(\theta)\cos(\phi)\cos(\psi) & \cos(\psi)\sin(\phi) + \cos(\theta)\cos(\phi)\sin(\psi) & \cos(\phi)\sin(\theta) \\ \cos(\theta)\cos(\psi)\sin(\phi) + \cos(\phi)\sin(\psi) & \cos(\phi)\cos(\psi) - \cos(\theta)\sin(\phi)\sin(\psi) & -\sin(\theta)\sin(\phi) \\ -\cos(\psi)\sin(\theta) & \sin(\theta)\sin(\psi) & -\cos(\theta) \end{pmatrix}. \quad (1.68)$$

The Euler angles can be related to the sky positions (the right ascension α and declination δ) by:

$$\phi = \alpha + \pi/2, \quad \theta = \pi/2 - \delta. \quad (1.69)$$

This can be viewed graphically, or can be obtained by doing the dot product between the corresponding vectors. We define the polarisation angle ψ as [13]:

$$\cos(\psi) = \hat{n} \cdot (\hat{L} \times \hat{z}) \quad \text{and} \quad \sin(\psi) = |\hat{L} \cdot \hat{z} - (\hat{L} \cdot \hat{n})(\hat{z} \cdot \hat{n})|, \quad (1.70)$$

where \hat{n} is the unit vector pointing to the source from the detector and \hat{L} is the unit angular momentum of the source. ψ can be defined as the angle between the direction $\hat{n} \times Z$ and the x-axis of the TT wave-frame (corresponding to the “+” polarization), where $Z = (0, 0, 1)$ is the unit-vector pointing to the celestial north pole, or it can also be defined as the counterclockwise angle from the \hat{x}_w axis to the projection of the East axis in the sky.

The second transformation goes from the celestial frame to the cardinal frame. The three Euler angles can be related now to the latitude λ and a rotating longitude L by $\phi = L + \pi/2$ and $\theta = \pi/2 - \lambda$, and $\psi = 0$. The rotation matrix is:

$$M_2 = \begin{pmatrix} \sin \lambda \cos L & \sin \lambda \sin L & -\cos \lambda \\ -\sin L & \cos L & 0 \\ \cos \lambda \cos L & \cos \lambda \sin L & \sin \lambda \end{pmatrix}. \quad (1.71)$$

It can be seen that this transformation depends on time, since the angle $L = \phi_r + \Omega_r t$ changes because the Earth is rotating.

The last transformation goes from the cardinal frame to the detector frame, where a simple rotation of the x-y plane is needed since the z axes of both frames are the same. The rotation matrix, composed

by a rotation of γ (which gives the angle between the East axis or \hat{x}_C and the bisector of the arms) and a rotation of $-\zeta/2$ (where $\zeta/2$ is the angle between the bisector and one of the arms), is:

$$M_3 = \begin{pmatrix} -\sin(\gamma + \zeta/2) & \cos(\gamma + \zeta/2) & 0 \\ -\cos(\gamma + \zeta/2) & -\sin(\gamma + \zeta/2) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.72)$$

After these three transformations, the axis of the wave frame will coincide with the axis of the detector frame. The complete matrix of rotation is $M = M_3 M_2 M_1$. Now we can obtain H_D :

$$H_D = M \cdot H_w \cdot M^T. \quad (1.73)$$

For example, for the plus polarization $h_+ = A_+ e_+$ we have that $e_{+D} = M e_+ M^T$.

The equation describing the projection of a gravitational wave (already expressed in the detector frame) onto an interferometric detector whose response is given by $D = \hat{d}_1 \otimes \hat{d}_1 - \hat{d}_2 \otimes \hat{d}_2$, where \hat{d}_1 and \hat{d}_2 are the unit vectors parallel to the detector arms, is:

$$h(t) = H_D D = \frac{1}{2} \hat{d}_1^T H_D \hat{d}_1 - \frac{1}{2} \hat{d}_2^T H_D \hat{d}_2, \quad (1.74)$$

where H_D is the expression in equation (1.67) in the coordinate frame of the detector.

The unit vectors parallel to the arms can be written in the detector frame as $\hat{d}_1 = [1, 0, 0]$ and $\hat{d}_2 = [\cos \zeta, \sin \zeta, 0]$, which make the detector response:

$$D = \hat{d}_1 \otimes \hat{d}_1 - \hat{d}_2 \otimes \hat{d}_2 = \begin{pmatrix} 1 - \cos^2 \zeta & -\cos \zeta \sin \zeta & 0 \\ -\cos \zeta \sin \zeta & -\sin^2 \zeta & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (1.75)$$

With equations (1.73) and (1.75), the gravitational wave signal at the detector can be described as:

$$h(t) = h_+(t) F_+(t) + h_\times(t) F_\times(t), \quad (1.76)$$

where $F_+(t)$ and $F_\times(t)$ are the antenna pattern functions. They depend on the sky position of the source, the polarisation angle, the position of the detector and the angle between the detector arms. They can be expressed with these equations [13]:

$$\begin{aligned} F_+(t) &= \sin \zeta [a(t) \cos 2\psi + b(t) \sin 2\psi] \\ F_\times(t) &= \sin \zeta [b(t) \cos 2\psi - a(t) \sin 2\psi], \end{aligned} \quad (1.77)$$

where [50]:

$$\begin{aligned} a(t) &= \frac{1}{16} \sin 2\gamma (3 - \cos 2\lambda) (3 - \cos 2\delta) \cos [2(\alpha - \phi_r - \Omega_r t)] - \frac{1}{4} \cos 2\gamma \sin \lambda (3 - \cos 2\delta) \sin [2(\alpha - \phi_r - \Omega_r t)] \\ &+ \frac{1}{4} \sin 2\gamma \sin 2\lambda \sin 2\delta \cos [\alpha - \phi_r - \Omega_r t] - \frac{1}{2} \cos 2\gamma \cos \lambda \sin 2\delta \sin [\alpha - \phi_r - \Omega_r t] + \frac{3}{4} \sin 2\gamma \cos^2 \lambda \cos^2 \delta \\ &= A_1 \cos[\phi_2 - 2\Omega_r t] + A_2 \sin[\phi_2 - 2\Omega_r t] + A_3 \cos[\phi_1 - \Omega_r t] + A_4 \sin[\phi_1 - \Omega_r t] + A_5 \end{aligned} \quad (1.78)$$

$$\begin{aligned} b(t) &= \cos 2\gamma \sin \lambda \sin \delta \cos [2(\alpha - \phi_r - \Omega_r t)] + \frac{1}{4} \sin 2\gamma (3 - \cos 2\lambda) \sin \delta \sin [2(\alpha - \phi_r - \Omega_r t)] \\ &+ \cos 2\gamma \cos \lambda \cos \delta \cos [\alpha - \phi_r - \Omega_r t] + \frac{1}{2} \sin 2\gamma \sin 2\lambda \cos \delta \sin [\alpha - \phi_r - \Omega_r t] \\ &= B_1 \cos[\phi_2 - 2\Omega_r t] + B_2 \sin[\phi_2 - 2\Omega_r t] + B_3 \cos[\phi_1 - \Omega_r t] + B_4 \sin[\phi_1 - \Omega_r t]. \end{aligned} \quad (1.79)$$

These functions make the gravitational wave amplitude depend on time due to Earth's rotation. As can be seen, these functions have two different periodicities: one with a period equal to one sidereal day, and another with a half sidereal day period. Taking both into account, the antenna pattern functions have a periodicity of a sidereal day, as shown in figure 1.6. When a source of gravitational waves is in




Figure 1.6: Antenna patterns $a(t)$ and $b(t)$ for the H1 and L1 detectors. The starting time is 1126072080 GPS, while $\alpha = 1$ and $\delta = 0.5$.

a position of the sky that makes the antenna pattern functions equal to zero, the detector is unable to detect these waves.

It can be observed with equation (1.76) that the amplitude modulation produced by these antenna patterns will split the signal h_+ or h_\times at five different peaks in the frequency domain: one at the same frequency, and two for each of the two different modulations (the half sidereal day periodicity and the full sidereal day periodicity). This comes from the mixing of the frequencies, caused by the sinusoidal products:

$$\begin{aligned}
 \cos[\phi + 2\pi ft] \cos[\phi_r + 2\pi f_r t] &= \frac{1}{2} (\cos[\phi + \phi_r + 2\pi t(f + f_r)] + \cos[\phi - \phi_r + 2\pi t(f - f_r)]) \\
 \sin[\phi + 2\pi ft] \sin[\phi_r + 2\pi f_r t] &= \frac{1}{2} (\cos[\phi - \phi_r + 2\pi t(f - f_r)] - \cos[\phi + \phi_r + 2\pi t(f + f_r)]) \\
 \cos[\phi + 2\pi ft] \sin[\phi_r + 2\pi f_r t] &= \frac{1}{2} (\sin[\phi + \phi_r + 2\pi t(f + f_r)] - \sin[\phi - \phi_r + 2\pi t(f - f_r)]), \quad (1.80)
 \end{aligned}$$

where f is the gravitational-wave frequency and f_r is the inverse of Earth's rotational period.

The long wavelength approximation (LWA) assumes that the length of the detector arms is much smaller than the wavelength of the incoming gravitational wave. For this reason, the length of the arms can be neglected and changes of the gravitational wave itself during the round trip time of the photons are not taken into account. The full detector response without applying the LWA is presented and discussed in [52], where it can be seen that the functions F_+ and F_\times are modified to contain frequency dependent terms. It is shown that by using the LWA, errors are smaller than 1% for frequencies smaller than 2 kHz. These effects need to be taken into account for searches using LISA data, whose arms length are comparable to the wavelengths of the GWs it tries to detect.

The procedure to calculate the antenna patterns for other non-GR polarizations mirrors the calculations presented here, the only difference being the polarization tensor e_p given by equation (1.67), which differs between polarizations. A more detailed discussion can be found in [53] and references therein.

This chapter contains an introduction to neutron stars. In the first section we present a brief summary of the history and basic physics of neutron stars, while in the second section we describe the pulsar population. The third section discusses how neutron stars can build asymmetries that can trigger the emission of GWs, and the last section describes mathematically how different geometric configurations of the neutron star create different gravitational waveforms.

2.1 | Basics of neutron star physics

Neutron stars (NSs) were proposed in 1933 (one year after the discovery of the neutron) by Walter Baade and Fritz Zwicky [54]. This first idea consisted of an astrophysical ball made up of neutrons and protons.

When a star stops being able to burn its nuclear fuel, different scenarios can happen depending on several properties of the star, such as its mass and metallicity. If the mass is lower than a certain value, the star will not go supernova, but if the mass exceeds this threshold it will. After the supernova takes place, the final state depends also on the mass of the remnant: if it exceeds a certain value, the core that is left after the supernova explosion will become a black hole. Otherwise, the core will remain visible and it will become a neutron star. Stars with masses approximately between 8 and 30 solar masses (or with cores above the Chandrasekhar limit of $1.4 M_{\odot}$) will be converted to neutron stars. When the core of the progenitor star is collapsing enough free energy is available for the inverse beta decay reaction to occur:

$$p^+ + e^- + 1.36 \text{ MeV} \longrightarrow n + \bar{\nu}_e. \quad (2.1)$$

Neutrons generated from this reaction should be unstable and turn back into an electron and a proton within 10 minutes via beta decay. This does not happen because the degenerate electron gas in the star has filled all of the available electron states in the core. For this reason, no electrons of energies smaller than 1.36 MeV can be formed, which makes the neutrons stable.

The typical expected value for the mass of a neutron star is between 1 and 2 solar masses, and its radius is approximately between 10 and 20 km. The minimum and maximum mass value (and thus radius) are important for population synthesis models that try to understand the amount of neutron stars, black holes and other types of objects that are present in our and other galaxies at different moments of their evolution. Their mass and radius make them the second densest object in the Universe, with a compactness ($C = 2GM/Rc^2$) value of $C \sim 0.2$, placing them between black holes and white dwarfs. Due to this compactness, NSs attain density values at their center higher than the nuclear density $\rho_n = 1.4 \times 10^{14} \text{ g/cm}^3$. This supranuclear density makes neutron stars a very valuable cosmic

laboratory which allow us to study the behaviour of matter at such extreme environments, a possibility that does not exist at Earth.

One of the biggest uncertainties about neutron stars is their equation of state (a review can be found in [55]). The equation of state describes the relation between the energy density and the pressure inside of the star (alternatively, it relates the mass and the radius of the star), and at such high densities it is unknown. For this reason, the state and structure of matter near the core of neutron stars is one of the most important mysteries within nuclear physics that could be solved with astrophysics. The most external part of the star (without taking into account a possible atmosphere) is the crust, which most models predict as being solid and made up mostly of iron-like atoms ordered in a crystalline array (although non-solid crusts are also a theoretical possibility). Moving radially towards the center, the pressure rises rapidly and neutrons begin to drip outside of nucleus, in a process known as neutron drip. This fluid of neutrons and protons can be in a superfluid state, and the protons can be in a superconductor state. Between the external crust and the internal core, the nuclei start to lose their normal shape and form what is usually called nuclear pasta [56]. In the innermost core, the quarks that make up the neutrons could be located out of the neutron, a process known as quark deconfinement. The relative widths between the crust and the core are highly uncertain, with different models predicting different outcomes: the EOS determines the thickness of the crust, while the atomic composition of the crust determines its shear modulus [57]. These properties are schematically shown in figure 2.1. All of these extreme properties summarize why neutron stars represent an interesting object and why they are hard to model and simulate.

Neutron stars were just a theoretical object until 1967, when Shklovsky [59] correctly claimed that the first X-ray source detected outside of the solar system, discovered in 1962 by Giacconi (Scorpius X-1), was a neutron star [60]. Also in 1967 (after the identification of Scorpius X-1 as a neutron star) the first pulsar was discovered by Jocelyn Bell and Antony Hewish [61]. A pulsar is a neutron star with a fast rotational speed and an electromagnetic jet coming out from its poles (usually modeled with a dipole magnetic field) strong enough to make the star visible to us. The emission of electromagnetic radiation by the poles and the rotation of the star makes the pulsar look like a lighthouse: when it is pointing to us, we can see it, and a moment afterwards it is gone; this pattern is repeated with a very stable frequency. In fact, this frequency (and its evolution) is so stable that pulsars are one of the most accurate clocks that exist. This is used to search for gravitational waves with pulsar timing array techniques, as explained in the previous chapter [42]. This dipole model of the magnetic field may be too simple for some neutron stars, as new observations from the NICER experiment show: the observation of hot spots that emit X-rays seems to indicate that the magnetic field may be in a diquadrupolar configuration, where instead of having a simple north-south dipole the two hot spots are situated in the same hemisphere, and one of them follows a more elongated shape instead of a circular spot [62].

As discussed in the previous chapter, the first indirect evidence of gravitational radiation emission from neutron stars came from the Hulse and Taylor binary system. GR predicted that the orbital period would decay due to the emission of gravitational waves, which was afterwards observed and fitted almost exactly. This indirect detection won the Nobel Prize in 1993. In August 2017, the first direct detection of gravitational waves from a binary neutron star system was made [9]. This detection was groundbreaking, since it confirmed that short gamma ray bursts originate in binary neutron star mergers, and it was the first neutron star mass and radius measurement made with gravitational waves.

2.2 | Demographics and parameters

In this section we describe the different types of pulsars that have been discovered and the different parameters that characterize them. After pulsars were detected for the first time, pulsar astronomy has been improving its methods and thus incrementing the number of detected pulsars and their variety. Up




Figure 2.1: Theoretical structure of a neutron star. Source: [63].

Parameter	ν_0 [Hz]	ν_1 [Hz/s]	d [kpc]	τ_s [yr]	B_s [G]
Maximum	716.36	-3.8×10^{-10}	59.70	9.41×10^{10}	2.1×10^{15}
Minimum	0.04	-4.0×10^{-18}	0.09	218	4.5×10^7

Table 2.1: Observational maximum and minimum of some pulsar (isolated and in binary systems) parameters, taken from [64]. ν_0 is the rotational frequency, ν_1 is the rotational spin-down, d is the distance to Earth, $\tau_s = -\nu_0/(2\nu_1) = P/(2\dot{P})$ is the spin-down age (where $P = \nu_0^{-1}$ is the rotational period), B_s is the surface magnetic field.

to now, there are around 2800 pulsars that are catalogued by the Australia Telescope National Facility (ATNF), which maintains a list of all the discovered pulsars up to date and a database with all their properties [64]. Tables 2.1 and 2.2 show the maximum and minimum of some parameters of known pulsars. These pulsars can be separated in several subpopulations. Firstly, for isolated pulsars (about 2494), there are: radio/normal pulsars, millisecond pulsars, and magnetars (soft gamma-ray repeaters and anomalous X-ray pulsars). Secondly, for pulsars in binary systems (about 306), there are: low-mass and high-mass X-ray binaries (LMXB and HMXB), nuclear-powered pulsars (NXPs) and accreting millisecond X-ray pulsars (AMXPs), and double pulsars/neutron star systems.

There are probably around 10^8 neutron stars in the Milky Way [65]. This number is calculated by estimating the supernova rate in our galaxy and restricting it to stars in the adequate mass range, which gives an order of magnitude estimate of 1 neutron star per century, assuming an age of 10^{10} years [66]. From these 10^8 neutron stars, around 10^5 are estimated to be pulsars (with only around 2800 discovered up to now), and approximately 10^7 (unseen) dead magnetars.

One of the most important parameters of a neutron star for this thesis is the ellipticity, which is a quantity that measures the deviation from sphericity and was defined in equation (1.56). In section 2.3 we will study what processes can build and sustain such ellipticity and afterwards we will see how this ellipticity influences the gravitational wave emission. Hereunder we discuss the other parameters that characterize a neutron star.

Parameter	ν_0 [Hz]	ν_1 [Hz/s]	P_b [days]	a_p [s]	e	Mass [M_\odot]
Maximum	716.36	-9.9×10^{-13}	16835	13580	0.964	2.14
Minimum	0.36	-1.6×10^{-17}	0.065	3.0×10^{-6}	1.14×10^{-7}	1.174

Table 2.2: Observational maximum and minimum of some binary pulsar parameters, taken from [64]. ν_0 is the rotational frequency, ν_1 is the rotational spin-down, P_b is the orbital period, a_p is the projected semi-major axis amplitude, e is the eccentricity, M_\odot is the mass in solar masses.

Rotational frequency and spin-down

Figure 2.2 shows the derivative of the rotational period versus the rotational period for all known pulsars, usually called $P-\dot{P}$ diagram. Two clearly distinct populations can be seen: the normal pulsars and the millisecond pulsars. Normal pulsars have higher periods and higher spin-down rates, while millisecond pulsars spin faster and more steadily. The majority of the detected pulsars which are supposed to emit CWs in the frequency band of the ground-based detectors (from 50 to 1000 Hz, approximately) are millisecond pulsars, as this diagram shows. Furthermore, it can be seen that more than half of the millisecond population belongs to binary systems, while very few pulsars of the normal population are not isolated. Another distinct population can be seen in the top right corner: the magnetars. As the lines of constant magnetic field show, these are pulsars with very high magnetic fields. The orange shaded region marks the pulsar death line, below which it is believed that radio emission ceases to be viable.

The observed values of the period derivatives are (to varying degrees) different than the true intrinsic values due to different processes: radial accelerations due to proper motion (the Shklovsky effect); differential Galactic rotation; and local forces (for pulsars in globular clusters). These effects should be taken into account if one wants to work with the intrinsic period derivatives instead of the observed ones. Pulsars with observed negative period derivatives (not shown in the diagram) can be due to these effects, or also due to accretion in a binary system.

It is believed that most millisecond pulsars are recycled pulsars: they are or were part of a binary system, and they were spun up with accretion of matter from a companion. This process may explain the lower absolute value of the first frequency derivative of binary pulsars, since the accretion balances the rotational energy that is lost through emission of electromagnetic or gravitational waves [67]. Furthermore, accretion of matter can bury the magnetic field of the NS [68]. Decreasing the magnetic field lowers the electromagnetic energy emission, which makes the rotation more stable (since the rate of change of frequency is related to the energy loss), thus reducing their spin-down. This lowered electromagnetic energy could also explain why there have been more electromagnetic detections of normal pulsars than millisecond pulsars.

An estimate of the initial rotational frequency of neutron stars can be obtained if the angular momentum $\vec{J} = I\vec{\omega}$ is conserved when the star is collapsing (we assume that the angular momentum and angular frequency are directed through the same axis, which does not change orientation during the collapse):

$$I_i \omega_i = I_f \omega_f \longrightarrow 2\pi f_f = 2\pi f_i \frac{I_i}{I_f} \longrightarrow P_f = P_i \frac{I_f}{I_i} = P_i \frac{M_f R_f^2}{M_i R_i^2}, \quad (2.2)$$

where we have used $I = 2/5MR^2$. As an example, if the mass of the neutron star is 0.1 the mass of the initial star and the radius is 0.001 the radius of the initial star, it can be seen that the frequency will be 10^7 times higher than the rotational frequency of the initial star.

The maximum rotational frequency found so far is from pulsar J1748-2446ad, which spins at 716 Hz. This means that at the equator, this pulsar has a linear velocity of $2\pi R/P = 62982596.3$ m/s for $R = 14$ km, which is around 0.21 times the speed of light, a fact that displays the extreme properties that neutron stars can possess. Different models try to explain why no other pulsars have been found at higher frequencies, since the Keplerian break-up frequency (also called mass-shedding limit, a theoretical maximum that highly depends on the equation of state), which is obtained by equating the centripetal

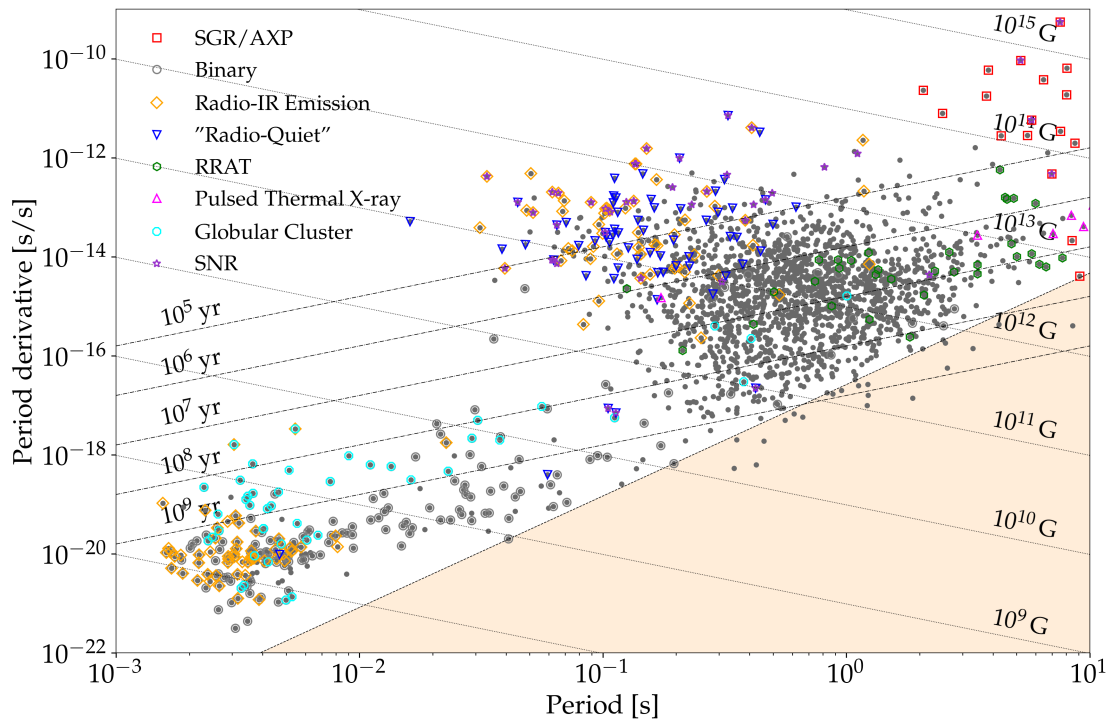


Figure 2.2: Rotational period and first period derivative for all discovered pulsars. Data taken from [64] and downloaded with [69].

acceleration to the gravitational acceleration at the equator and assuming perfect spherical shape [70]:

$$\frac{GM}{R^2} = R\omega_{max}^2 \longrightarrow \nu_{max} = \frac{1}{2\pi} \sqrt{\frac{GM}{R^3}}, \quad (2.3)$$

can be as high as 1500 Hz. A possible explanation is that the emission of gravitational waves acts as a braking torque which stops the spin-up of the neutron star. This would explain why no AMXP has been found at such high frequencies: the magnetic field from the neutron star builds a mountain (a process detailed in section 2.3) from the accreted matter and thus the gravitational wave emission becomes stronger. Another possible explanation is that the magnetosphere interacts with the accretion disk and a torque that removes rotational energy from the NS is created [71]. Furthermore, another explanation comes from the possible gravitational-wave emission through r-modes, which could become efficient at the observed maximum rotational frequency, even when no long-term mass quadrupole deformations are present in the neutron star [72]. Nevertheless, the maximum allowed rotational frequency is EOS dependent, as explained in [73], and the minimum value for minimal EOS assumptions seems to be around 1200 Hz, clearly above the maximum detected. For this reason, an observation of electromagnetic or gravitational radiation of a NS with unusually high rotational frequency could point us towards the right EOS (i.e. if a frequency is observed which is higher than the maximum allowed by some EOS, those EOS would be ruled out by that observation).

Another important parameter of pulsars is the relation between the frequency of rotation and its derivative. This parameter is called the braking index, which can provide information about the energy loss mechanisms of pulsars, including GW emission. The spin-down of pulsars is expected to follow this equation [74]:

$$\dot{\nu} = -K\nu^n, \quad (2.4)$$

where ν is the spin frequency of the pulsar, $\dot{\nu}$ is the frequency derivative, K is a constant of proportionality related to the pulsar's moment of inertia and magnetic field structure, and n is the braking index. More

specifically, this can be written as [75]:

$$\begin{aligned}\dot{P} &= \frac{32\pi^3 R^6 B_s^2}{3Ic^3\mu_0} P^{-1} + \frac{512\pi^4 GI\epsilon^2}{5c^5} P^{-3} \\ \dot{\nu} &= -\frac{32\pi^3 R^6 B_s^2}{3Ic^3\mu_0} \nu^3 - \frac{512\pi^4 GI\epsilon^2}{5c^5} \nu^5,\end{aligned}\tag{2.5}$$

where $\mu_0 = 4\pi \times 10^{-7}$ N A⁻² is the magnetic permeability of vacuum. Taking the time derivative of equation (2.4) results in the following equation:

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2},\tag{2.6}$$

where $\ddot{\nu}$ is the second-derivative of the spin frequency. If the torque is dominated by an outflowing wind, then $n = 1$; if magnetic dipole dominates, then $n = 3$; if magnetic quadrupole dominates, then $n = 5$; if gravitational-wave emission dominates, then $n = 5$; if the gravitational-wave emission from r-modes dominates, then $n = 7$. A measurement of n can be made only for the youngest pulsars for which $\ddot{\nu}$ is large enough to be detectable on human timescales, since for most other cases this value is dominated by other phenomena like timing noise or glitches. A recent survey has computed the braking index for 19 pulsars, with values ranging from 2.8 to 2890 (excluding two pulsars with negative braking index), showing that the simple dipolar spin-down model ($n = 3$) is clearly imperfect, although these measurements have a wide range of uncertainties [76].

The rotational phase of a neutron star can be described with very good accuracy by a Taylor expansion of a few orders around a reference time. Sometimes, pulsars show anomalous behaviour that produces sudden jumps of the rotation rate and the higher-order terms, which usually are called glitches. The first glitches were detected in 1969 for the Vela and Crab pulsars [77], and they showed a sudden increase of the rotational frequency (opposite sign as normal, since normally the period is decreasing due to energy emission). Since then, more than 300 glitches have been observed in hundreds of different pulsars. The timescale for the glitch is very rapid, around 30 s, although the exact timescale has not been resolved. The rotational frequency increases have ranged from 10^{-11} to 10^{-5} Hz, also accompanied by an increase of the spin-down rate.

Although the origin of glitches has not been confirmed, the strongest hypothesis proposes that although the neutron star is spinning down, the neutron superfluid in its interior is not since it is composed of several vortices that are pinned (whose positions are fixed). When the difference between rotation rates attains a certain threshold, some vortices may unpin and angular energy is liberated through glitches. Another hypothesis is related to starquakes, where due to the rigidity of the crust and because the star is spinning down, the crust maintains its shape until the strain on the crust reaches a critical level and energy is released. For a review of different glitching models see [78]. The effect that glitches produce on CW searches is discussed in [79]. A catalogue of glitches can be found in [80].

Besides glitches, regular timing of pulsars shows a different irregularity, called timing noise [81]. Timing noise is probably indicative of irregularities in pulsar rotation; its physical origin remains unclear. There appears to be a small correlation between timing noise and some pulsar spin parameters: pulsars with short periods (high spin frequencies) have less timing noise, and pulsars with low spin-down rates (small frequency derivatives, small \dot{P}) also have less timing noise.

Mass, radius, and moment of inertia

Masses of neutron stars have been measured for some pulsars, an updated catalogue and review of mass measuring methods can be consulted in [82]. Minimum and maximum masses can be seen in table 2.2, while the full population of galactic neutron stars can be fit with a mixture of two Gaussian probability distributions [83] (although only pulsars in binary systems are taken into account). There are theoretical bounds both for the maximum and minimum mass that NSs can have. The maximum mass

for a non-rotating star is given by the Tolman-Oppenheimer-Volkoff (TOV) limit, which is analogous to the Chandrasekhar limit for white dwarves, and is highly dependent on the equation of state. The maximum mass is also dependent on the rotational frequency ν of the star (faster rotations allow for higher maximum masses), a bound schematically given by $M_{\max} = M_{\text{TOV}} (1 + \alpha\nu^\beta)$, where α and β are parameters that depend on the equation of state and M_{TOV} is the maximum mass for a non-rotating star [58]. The increase in maximum mass due to rotation can be up to 20%, and this number can be even higher if differential rotation is taken into account [85]. The minimum mass of neutron stars depends on the supernova mechanism and the metallicity of the core of the progenitor star, with some models being able to accommodate the $1.17 M_\odot$ measurement [86]. Understanding the maximum and minimum masses has important consequences for stellar evolutionary and population models, and also to constrain the equation of state. New methods to measure NS masses such as gravitational waves or NICER (which can measure masses from isolated sources) will help to increment the amount of total measurements.

Radii of neutron stars are more difficult to measure (due to biases produced by unmodeled atmospheric re-processing and interstellar absorption), and for this reason only around 20 radii have been collected, mainly from NSs for which thermal X-ray and optical radiation has been observed. A minimum radius for a certain mass is imposed by general relativity, since the radius has to be larger than the Schwarzschild radius of $2GM/c^2$.

Although the moment of inertia of a simple body only depends on its mass, radius, and a geometrical factor, for a fully relativistic body with a realistic equation of state the relation is more complicated. In [88], an empirical fit is given by:

$$I \simeq (0.237 \pm 0.008)MR^2 \left[1 + 4.2 \frac{M}{M_\odot} \frac{\text{km}}{R} + 90 \left(\frac{M}{M_\odot} \frac{\text{km}}{R} \right)^4 \right], \quad (2.7)$$

which for normal neutron stars is almost universal (i.e. independent of the equation of state). The moment of inertia could be directly measured by studying the effects of spin-orbit coupling of neutron stars in binary systems. The precession induced by this coupling changes the angle ι_b between the line of sight and the total angular momentum $\vec{J} = \vec{L} + \vec{S}_A + \vec{S}_B$ [88]:

$$\frac{d\iota_b}{dt} = \frac{G}{ac^2} \frac{\pi}{(1-e^2)^{3/2}} \sum_i \frac{I_i (4M_i + 3M_{-i})}{M_i a^2 P_i} \sin \theta_i \cos \phi_i, \quad (2.8)$$

where $i = A, B$, $-i = B, A$, θ_i is the angle between \vec{S}_i and \vec{L} , ϕ_i is the angle between the line of sight to pulsar i and the projection of \vec{S}_i on the orbital plane, and a is the semi-major axis of the effective one-body orbital problem, which is equal to the sum of the semi-major axes of the two stellar orbits. A measurement of the moment of inertia can be used to constrain the equation of state of neutron stars as shown in [88]. This is very useful for systems where a radius measurement is not possible or less accurate than the moment of inertia measurement.

Magnetic field

Equation (2.5) shows the relation between the spin-down of a neutron star and the strength of its magnetic field. For this reason, the spin evolution will be highly coupled to the evolution of the magnetic field. When a star is collapsing, the magnetic flux through its surface is going to be constant. If the initial magnetic field and the initial and final radius are known, the final magnetic field can be calculated:

$$B_i 4\pi R_i^2 = B_f 4\pi R_f^2 \longrightarrow B_f = B_i \left(\frac{R_i}{R_f} \right)^2 \sim 4.84 \times 10^9 B_i, \quad (2.9)$$

where we have used the radius of the Sun and a radius of 10 km for the neutron star. If we assume that the magnetic field of the Sun is about 1 G, then the neutron star has a magnetic field of 4.84×10^9 G. This simple argument helps to understand why neutron stars reach such high values of magnetic field.

Magnetars are neutron stars that have magnetic fields higher than around 10^{14} G, and all of them possess rotational periods higher than a second. It is clear that the previous simple argument is not enough to explain these high values. The source of such high magnetic fields is uncertain, although two main processes are usually proposed: dynamo process in the core or magneto-rotational instability during the supernova collapse [89]. To generate their strong magnetic fields through a dynamo process, it is hypothesized that magnetars must be born rotating fast, with period less than 1 ms, but due to their high magnetic fields the spin-down is very high and they rapidly move to higher rotational periods. A catalogue of magnetars can be found in [90].

The magnetic field is usually decomposed in two different parts in polar coordinates: poloidal B_p (B_r and B_θ directions) and toroidal B_t (B_ϕ direction) components. The field is usually simulated with an internal toroidal component and an external poloidal component, where the internal field is much stronger than the external field. The internal magnetic field will evolve in time mainly due to Ohmic dissipation (i.e. Joule heating), and the external field may also change due to accretion processes. Some measurements suggest that some neutron stars have external magnetic field strengths several orders of magnitude larger than the fields predicted by the simple vacuum dipolar magnetic field model [91, 92].

Temperature

Neutron stars have typical surface temperature values around 10^6 K, whereas accreting neutron stars can have temperatures up to several 10^8 K. Although the initial temperature is much higher (around 10^{11} K), neutron stars rapidly cool down after their formation mainly through neutrino emission (direct Urca process), and they keep cooling down at a much slower rate throughout their evolution [93, 94]. The temperature also depends on the radial coordinate: closer to the core, neutron stars have higher temperatures than at their surface. The Fermi temperature (the temperature at which thermal effects are comparable to quantum effects associated with Fermi statistics, calculated as the energy difference between the highest and lowest occupied single-particle states divided by the Boltzmann constant) of a neutron star is around 10^{12} K, much higher than the actual temperature of the star: for this reason neutron stars are said to contain cold matter, regardless of their high temperatures.

Sky distribution

Figure 2.3 shows the distribution on the sky in equatorial coordinates of all the known pulsars. An over-density following the disk of the galaxy (a region where lots of new stars are born) can be clearly seen. It is believed that most neutron stars are born in the disk of the galaxy, and may be ejected out from it due to natal kicks from the supernova collapse. These natal kicks give neutron stars a fraction of their proper motion, which can be up to a few 1000 km/s [95]. While normal pulsars are usually in the disk of the galaxy, millisecond pulsars are more uniformly distributed around the halo of the galaxy, a zone where the oldest stars are expected to reside. Some studies mention that there might be a correlation between the direction of motion of pulsars and the orientation of their angular momentum [96]. The direction of the angular momentum has been measured only for a small number of pulsars. The main method consists of taking advantage of pulsar wind nebula around the pulsar, which has the shape of a torus [97].

Binary parameters

Almost half of the pulsars with gravitational-wave frequencies in the most sensitive band of the second generation ground-based detectors (approximately between 50 and 1000 Hz) are located in binary systems. Of the known pulsars in binary systems, four different types of companions can be distinguished: 177 are white dwarves, 24 are main-sequence stars, 20 are neutron stars (one of these systems is a double pulsar: PSR J0737-3039), and 54 are ultra-light objects (with masses less than 0.08 solar masses) [64].




Figure 2.3: Sky positions of all the known pulsars, where orange circles show isolated pulsars and blue circles show pulsars in binary systems. Data taken from [64] and downloaded with [69].

The movement of pulsars in binary systems can be described by a classical Keplerian orbit (when any kind of radiation is neglected). Five parameters are needed to describe it: the orbital period P_b , the projected semi-major axis amplitude a_p , the eccentricity e , the time of ascending node t_{asc} , and the argument of periastron ω . For a Keplerian orbit, the projected semi-major axis amplitude and the orbital period follow the relationship given by the third Keplerian law:

$$a_p \equiv \frac{a_{NS}}{c} \sin \iota_b = \left(\frac{G}{4\pi^2} \right)^{1/3} P_b^{2/3} \frac{M_T^{4/3}}{M_{NS}} \frac{\sin \iota_b}{c}, \quad (2.10)$$

where $a_{NS} = aM_T/M_{NS}$ is the semi-major axis amplitude of the pulsar's orbit, M_{NS} is the mass of the NS, and M_C is the mass of the companion star. The left plot of figure 2.4 shows values of these two quantities for the known pulsar population. The different companion masses and angles of inclination ι_b account for the spread in the vertical axis. Along with the observational data, we have plotted four lines which follow equation (2.10) for two different values of the companion mass and two values of the inclination angle (for a $1.4 M_\odot$ neutron star).

The eccentricity (defined as $e = \sqrt{1 - b^2/a^2}$, where a and b are respectively the semi-major and semi-minor axis of the binary orbit) of pulsars in binary systems is shown in figure 2.4. We observe that for most of the pulsars with measured eccentricity, it is smaller than 0.01 (for 167 out of 215). Pulsars in globular clusters are believed to have higher eccentricities, since the environments have higher stellar density and there are substantial gravitational interactions with nearby stars that can make an initially circular orbit acquire a substantial eccentricity.

Relativistic effects in binary motion can be expressed in terms of five post-Keplerian parameters that describe departures from Keplerian motion [99]. In any given theory of gravity, the post-Keplerian (pK) parameters can be written as functions of the pulsar and companion star masses and the Keplerian parameters (assuming point masses and negligible spin contribution). The first of these parameters to be measured was the precession of the periastron $\dot{\omega}$, given in GR (in fact, it is equivalent to the effect that causes the excess of advance of Mercury's perihelion) by [99]:

$$\dot{\omega} = 3 (P_b/2\pi)^{-5/3} (T_\odot M_T)^{2/3} (1 - e^2)^{-1}, \quad (2.11)$$

where ω is the longitude of periastron measured from the ascending node, $M_T = M_{NS} + M_C$, and $T_\odot \equiv GM_\odot/c^3 = 4.925490947 \mu\text{s}$. The next post-Keplerian parameter is γ , which describes the combination of

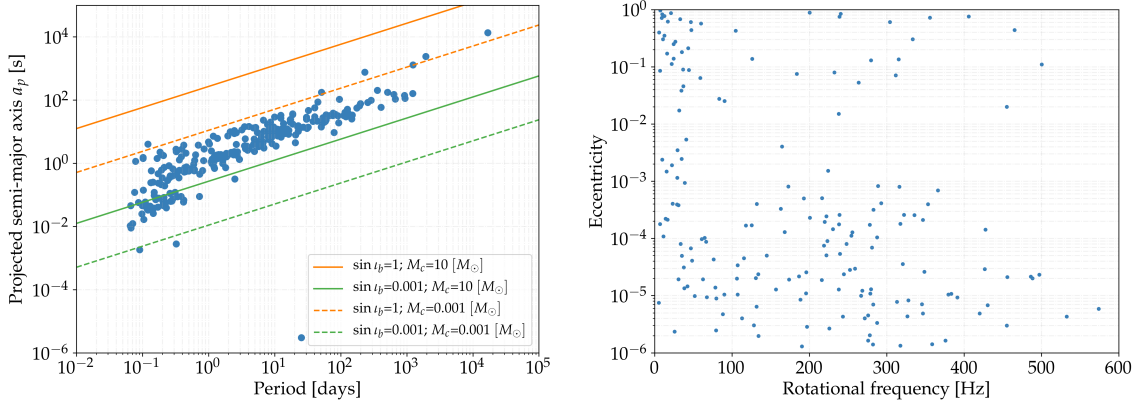


Figure 2.4: The left plot shows the projected semi-major axis amplitude and orbital period of known pulsars in binary systems. The four lines show different combinations of assumed companion mass and inclination angle for a $1.4 M_\odot$ neutron star, following equation (2.10). The right plot shows the eccentricity of pulsars in binary systems as a function of their rotational frequency. Data taken from [64] and downloaded with [69].

gravitational redshift and 2nd-order (or transverse) Doppler shift, given in GR by:

$$\gamma = e (P_b/2\pi)^{1/3} T_\odot^{2/3} M_T^{-4/3} M_C (M_{NS} + 2M_C). \quad (2.12)$$

Another post-Keplerian parameter is \dot{P}_b , the orbital decay due to the emission of gravitational waves from the system, given in GR by:

$$\dot{P}_b = -\frac{192\pi G^{5/3}}{5c^5} \left(\frac{P_b}{2\pi}\right)^{-5/3} (1-e^2)^{-7/2} \left(1 + \frac{73}{24}e^2 + \frac{37}{96}e^4\right) M_{NS} M_C M_T^{-1/3}. \quad (2.13)$$

As happened with the period derivative, the measurement of this parameter is also affected by radial accelerations due to proper motion, differential Galactic rotation, and local forces (for pulsars in globular clusters). The two final post-Keplerian parameters are related to the Shapiro delay (range and shape) suffered by the pulsar signal while passing through the curved spacetime surrounding the companion star. They are given by:

$$r = T_\odot M_C \quad (2.14)$$

$$s = a_p \left(\frac{P_b}{2\pi}\right)^{-2/3} T_\odot^{-1/3} M_T^{2/3} M_C^{-1}. \quad (2.15)$$

The Shapiro delay is generally only observable when the orbital inclination is relatively close to 90° (when the orbit is seen close to edge-on).

The measurement of any two post-Keplerian parameters determines the mass of the two objects in the system. If more than two pK parameters have been measured, they can be used to test and constrain theories of gravitation, i.e. by using the two masses of the system the other three pK parameters can be predicted and compared to the measurements, serving as a check for self-consistency. Other theories of gravity, such as those with one or more scalar parameters in addition to a tensor component, will have somewhat different mass dependencies for the pK parameters, thus producing different predictions for their values. Reference [99] shows a table with all the systems that have measured some pK parameters.

When the angular momentum of the NS is not aligned with the total angular momentum of the system, spin-orbit effects appear. This produces a precession (called geodetic precession) of the NS angular momentum around the total angular momentum given by (in GR):

$$\Omega_p = \frac{1}{2} \left(\frac{P_b}{2\pi}\right)^{-5/3} T_\odot^{2/3} \frac{M_C (4M_{NS} + 3M_C)}{(1-e^2) M_T^{4/3}}. \quad (2.16)$$

Neutron stars in binary systems may be accreting matter from its companion. If this is happening, the NS may experience the so-called spin wandering (SW), where the rotational frequency of the star changes accordingly to changes in the accretion rate [100]. The spin-down produced by the torque N generated by the accretion is given by:

$$\dot{f}_{SW} = \frac{N}{2\pi I} \simeq 1.6 \times 10^{-13} \left(\frac{\dot{M}}{10^{-10} M_{\odot} \text{yr}^{-1}} \right) \left(\frac{P}{s} \right)^{1/3} \left(\frac{r_m}{r_{co}} \right)^{1/2} \left(\frac{10^{38} \text{ kg m}^2}{I} \right), \quad (2.17)$$

where r_m is the magnetic radius and r_{co} is the co-rotation radius. The mass accretion rate can be decomposed in two terms:

$$\dot{M}(t) = \langle \dot{M}(t) \rangle + \Delta \dot{M}(t). \quad (2.18)$$

The instantaneous fluctuations given by $\Delta \dot{M}(t)$ are what create the SW. As explained in [100], the mass accretion rate presents different timescales.

2.3 | Asymmetries of neutron stars

In this section we discuss how NSs can generate and sustain the asymmetries that are needed for GW emission. We define the fiducial quadrupole ellipticity or asymmetry of the star as $\epsilon = (I_{xx} - I_{yy})/I_{zz}$, where I_{zz} is the moment of inertia of the star with respect to the principal axis aligned with the rotation axis. We can relate these quantities to the mass quadrupole moment Q_{22} of the star by [101]:

$$\epsilon = \sqrt{\frac{8\pi}{15}} \frac{Q_{22}}{I_{zz}}. \quad (2.19)$$

Q_{22} is the mass quadrupole moment associated to the distortion, i.e. excluding the contribution from rotation, which is necessarily axisymmetric and is not associated to GW emission.

This fiducial ellipticity can be related to a more physically measurable relativistic ellipticity ϵ_p , defined from quantities at the surface of the star: the star's equatorial circumference, s_e , and its longest polar circumference, s_p [102]. This new ellipticity is defined as:

$$\epsilon_p = \frac{s_p - s_e}{s_e}. \quad (2.20)$$

For a simple constant density ellipsoid, the fiducial ellipticity and the physical ellipticity are related by $\epsilon_p = \epsilon/8$.

Before describing actual mechanisms that generate asymmetries, we can try to quantify what is the maximum elastic asymmetry that a neutron star can sustain. The maximum ellipticity can be written schematically as [103]:

$$\epsilon_{max} = (\text{breaking strain}) \times (\text{shear modulus}) \times (\text{geometry}). \quad (2.21)$$

The geometry term depends on the equation of state and is uncertain by a factor of a few. The shear modulus increases with density and is given by the ratio between stress and strain (i.e. for a fixed stress, an object with higher shear modulus will suffer a smaller strain). Simulations of molecular dynamics show that the maximum strain that a solid crust can support before breaking (crustal breaking strain, u_b) is approximately 0.1 [104, 105]. These simulations are strictly applicable only to the outer crust (i.e., no neutron drip), but since the reason for this high breaking strain is that the extreme pressure simply crushes away defects that contribute to early fracture, it may apply to the inner crust (the major contributor to ellipticity) as well. A more specific estimation gives [106]:

$$\epsilon_{max} = 2 \times 10^{-5} \left(\frac{u_b}{0.1} \right). \quad (2.22)$$

If instead of normal matter, we assume a solid strange quark star the result is [107]:

$$\epsilon_{max;q1} = 6 \times 10^{-4} \left(\frac{u_b}{0.01} \right), \quad (2.23)$$

while for a crystalline color-superconducting quark phase [108]:

$$\epsilon_{max;q2} = 10^{-3} \left(\frac{u_b}{0.01} \right). \quad (2.24)$$

Note that for these exotic quark phases the molecular dynamics simulations previously cited do not apply. As can be seen, detections or upper limits can constrain the value of u_b .

The crust is suffering some strain due to the rotation quadrupole and more contributions from other multipole terms, so the ellipticity produced by a GW-emitting quadrupole can only be a percentage of the maximum strain that the crust can sustain. The asymmetry caused by the rotation of the NS is estimated as the ratio between rotational energy and gravitational energy: $\epsilon_\Omega \sim \frac{I\Omega^2}{GM^2/R} = 0.1 \left(\frac{f}{\text{kHz}} \right)^2$, and it makes the star oblate (although this quantity is not easy to evaluate, as the exact level of strain depends on the reference shape of the crust, which in turn depends on the history of the star). In this sense, neutron stars with higher rotational frequencies produce a larger $l, m = 2, 0$ quadrupolar strain (where l and m are spherical harmonics indices), and thus leave a smaller percentage of the maximum to be strained to the $2, 2$ quadrupole.

A fully relativistic estimation of the maximum quadrupole for different equations of state is presented in [101]. Results depend on the mass of the NS, its equation of state, and the thickness of the crust (thinner crusts giving smaller maximum deformations), but overall results show that the maximum ellipticity is smaller than 10^{-5} for non-exotic stars. For exotic stars, it is shown that the maximum ellipticities can be even higher than the previously presented values. A general property of these simulations is that the quadrupole decreases with the compactness of the star. A consequence of this fact is that CW searches are biased towards low-mass stars (for non-exotic stars), in the sense that for the same geometry, magnetic field and strain, they will have a larger quadrupole deformation, which makes them an easier target to observe [109].

Another way to estimate the maximum ellipticity that neutron stars can have is by using the spin-down upper limit on ellipticity. The argument is the following: if the spin-down of a neutron star is only caused by gravitational-wave emission (i.e. we assume that there is no electromagnetic radiation), we can obtain an upper limit to the existing ellipticity. Equating the gravitational-wave luminosity given by equation (1.56) to the derivative of the rotational energy ($E_{rot} = 2\pi^2\nu^2 I_{zz}$):

$$L_{GW} = -\dot{E}_{rot} = 2\pi^2(2I_{zz}\nu\dot{\nu} + \dot{I}_{zz}\nu^2)x, \quad (2.25)$$

where $0 \leq x \leq 1$ is a factor that represents the amount of gravitational-wave emission (1 for the spin-down limit, 0 for no gravitational-wave emission). Assuming a constant moment of inertia and $x = 1$, the result is:

$$\epsilon_{sd} = \sqrt{\frac{5(2c)^5}{2(4\pi)^4 G I_{zz}}} \sqrt{\frac{\dot{f}}{f^5}}. \quad (2.26)$$

By assuming a certain moment of inertia, we can plot the spin-down upper limit on ellipticity ϵ_{sd} for all the known pulsars. It should be noted that recent targeted searches have surpassed this value for some of the known pulsars, constraining the amount of gravitational-wave radiation that they emit [110].

Besides the maximum attainable ellipticity, there has been discussion of a minimum ellipticity which neutron stars might have due to the observed properties of the pulsar population [111]. We can plot in the $P-\dot{P}$ diagram lines of constant ellipticity, calculated by assuming an ellipticity and a period value and GW emission only. The line of $\epsilon = 10^{-9}$ at short period and low spin-down values seems to mark a region where no other pulsars have been detected beyond it, suggesting that this may be a minimum ellipticity. In [111], a Bayesian model selection analysis is done and this is highly favoured. The explanation for this

minimum ellipticity of millisecond pulsars can be due to a buried magnetic field of 10^{11} G (the observed external field is of the order of 10^8 G), or also due to asymmetric accretion.

In the next subsections we summarize three different scenarios that show how neutron stars can generate and sustain asymmetries. The asymmetry of a real neutron star could be due to any of these or to a combination of them. There are more mechanisms besides the three described below, such as asymmetries inherited from the supernova (the initial crust may be asymmetric if it forms on a timescale on which the neutron star is still perturbed by its violent formation and aftermath, including perturbations due to the fluid r-modes, and asymmetries may slowly relax due to mechanisms such as viscoelastic creep), but these three mechanisms are the ones that could sustain a long-lived asymmetry with more probability.

2.3.1 | Magnetic deformations

The idea that magnetic stresses can deform a star was originally proposed in [112]. As discussed previously, neutron stars possess an internal magnetic field, and the asymmetry that it creates can be approximated as the ratio between magnetic and gravitational energies:

$$\epsilon_B \sim \frac{E_B}{E_G} \sim \frac{R^3 \langle B^2 \rangle}{GM^2/R}, \quad (2.27)$$

although a more exact equation is:

$$\epsilon_B = -\frac{15}{4} \left(\frac{5R}{3GM^2} \right) \int \frac{1}{8\pi} B_t^2 dV, \quad (2.28)$$

where B_t stands for the internal toroidal field and the integral represents an average over the volume of the star.

The quadrupole ellipticity ϵ_B can be positive or negative: the shape of the star can be oblate or prolate, depending on the magnetic field structure. Neutron star magnetic fields can have poloidal or toroidal structure, and poloidal fields tend to deform the star to an oblate shape, while toroidal fields tend to deform it to a prolate shape. When poloidal and toroidal components are both present, the deformation has both positive and negative contributions. The deformation produced by an internal magnetic field is global, whereas for accreted mountains on the polar caps it is a local deformation.

If the magnetic field axis is not aligned with the rotation axis, this asymmetry will be the source of CWs. The actual value of this asymmetry is highly dependent on the magnetic field geometric configuration: most models assume a purely poloidal or toroidal magnetic fields, but analytical and numerical studies show that these types of fields are unstable. More realistic models include a mixture between poloidal and toroidal components, for which the asymmetry is [113]:

$$\epsilon_B \sim 4.5 \times 10^{-7} \frac{B_p^2}{10^{14} \text{ G}} \left(1 - 0.64 \frac{\langle B_t^2 \rangle}{B_p^2} \right). \quad (2.29)$$

Other models also include a sum of several multipoles (orders higher than the dipolar), for which the ellipticity generated can increment, as it is the sum of the ellipticity created by each magnetic multipole, as shown in [114]. The relative importance of each multipole depends on their relative energy.

The size of the asymmetry also depends on the state and structure of matter inside the neutron star. For example, by allowing a superconducting phase in the core, the estimate is (for a purely toroidal field):

$$\epsilon_B \sim 10^{-6} \langle B_{15} \rangle \frac{H_c}{10^{15} \text{ G}}, \quad (2.30)$$

where H_c is the critical field strength characterizing superconductivity and $B_{15} = 10^{15}$ G.

An interesting aspect of magnetic stresses is that they will evolve in time, since the magnetic field (both the strength and structure) will change. Furthermore, the angle between the angular momentum and the magnetic axis also evolves with time in a complicated manner that depends on the starting frequency and magnetic field strength, as shown in [115, 116]. These two aspects will make the asymmetry depend on time.

2.3.2 | Thermal deformations

When a neutron star is accreting matter from a companion, the infalling matter can get buried and compressed inside the NS. Due to nuclear reactions (electron capture processes, whose rate depends on temperature) that happen to this infalling matter, heat is released and a temperature gradient is formed inside the star. This temperature gradient gives rise to a quadrupolar deformation. The deformation due to a temperature quadrupole component δT_q is [117, 118]:

$$\epsilon_T \sim 10^{-10} R_6^4 \frac{\delta T_q}{10^5 \text{ K}} \left(\frac{Q}{30 \text{ MeV}} \right)^3, \quad (2.31)$$

where $R_6 = 10^6 \text{ cm}$ and Q is the threshold energy of electron capture by nuclei (a parameter controlling how deep the nuclear reactions start to take place). The quadrupolar thermal gradient is just a fraction of the total thermal gradient, expected to be at most around 10^6 K . The value of this fraction is highly uncertain and could be much smaller than the value 0.1 used previously, which would lower even more the size of the mountain. Since the mountains formed by these temperature gradients are elastic mountains, the elastic strain limits explained at the beginning of this section apply to them.

Recently, thermal deformations have been brought up in order to explain the state of PSR J1023+0038, which transitions between a radio millisecond pulsar state and a LMXB state, showing increased spin-down while in this second state. The increase in spin-down may be explained by a $\sim 10^{-10}$ deformation created in the accretion phase, which makes the NS emit CWs and increase its spin-down rate [119].

2.3.3 | Exotic deformations

If neutron stars have cores made of quarks in a color superconducting state, they could present larger quadrupolar deformations [120]. Examples of such cores made of quark matter are: the color-flavor-locked (CFL) state, which has top, down, and strange quarks; the two-flavor superconducting (2SC) state, which only has top and down quarks; the crystalline color-superconductor (CCS) state, which has a solid crystalline core. The properties of quark matter in neutron stars are largely uncertain and therefore the following theoretical predictions should be viewed as giving orders of magnitude estimates.

The cores from the two first states form an array of quantised vortices in a similar way that fluxtubes are formed in type II protonic superconductors, and these vortices are both magnetic and color conducting. The latter component is the dominant one and the energy per unit length of a color-magnetic vortex can be 2 or 3 orders of magnitude higher than that of a conventional protonic fluxtube. This produces an amplified vortex array tension that can create a large deformation in the CFL/2SC core of a neutron star. The ellipticity of this internal color-magnetic mountain can be estimated by means of the ratio of the vortex array tension energy to the gravitational energy, which for a CFL core is:

$$\epsilon_{cfl} \sim 10^{-7} \langle B_{12} \rangle \left(\frac{V_q}{V_{\text{star}}} \right) \left(\frac{\mu_q}{400 \text{ MeV}} \right)^2, \quad (2.32)$$

where μ_q is the strange quark chemical potential (normalised to a canonical value), V_q is the volume of the quark core, and $B_{12} = 10^{12} \text{ G}$.

For the CSS case, its shear modulus μ_{css} is estimated to be several orders of magnitude higher than the shear modulus of the crust. For this reason, the elastic deformability that this star could sustain is much higher, and an estimate of the ellipticity is:

$$\epsilon_{\text{css}} \sim \frac{\mu_{\text{css}} V_q u_b}{GM^2/R} \sim 6 \times 10^{-4} \left(\frac{V_q}{V_{\text{star}}} \right) \left(\frac{\mu_q}{400 \text{ MeV}} \right)^2 \left(\frac{\Delta_q}{10 \text{ MeV}} \right)^2 \left(\frac{u_b}{0.01} \right), \quad (2.33)$$

where Δ_q is a quark gap parameter.

2.4 | Continuous gravitational waves from neutron stars

After having described the general properties of neutron stars and the asymmetries that they can support, we explain how NSs can generate gravitational waves. We will focus on long-lived continuous gravitational waves, leaving aside other mechanisms (such as bursts of magnetars, mergers of two neutron stars, or long-duration transients) through which neutron stars can emit shorter-lived gravitational waves. Furthermore, we will not discuss CWs generated by fluid oscillations/perturbations like f-modes or r-modes, although most of the results from searches can be re-interpreted under the model explained in [121].

As we have seen in subsection 1.1.2 of the first chapter, in order to have emission of gravitational waves a time-varying quadrupole moment needs to be present. Deformations of rotating neutron stars were proposed as sources of continuous gravitational waves shortly after the discovery of pulsars and the realization that they are rotating neutron stars [122–125] (see [126] for an early review).

The following two subsections present two different models of continuous waves emission, the first one being more simple and the second one more general. As we will see, the CWs can be described in the Fourier domain by just a couple of frequencies, where the number of frequencies and their relative amplitude depends on the specific model of emission. These two models only take into account emission from the $l = 2$ mass quadrupole, without calculating the emission from the current quadrupole, mass octupole, or higher-order terms.

2.4.1 | Gravitational waves from rotation around a principal axis

The simplest type of CW emission generated by a rotating neutron star (modelled as a rigid body ellipsoid) is due to an asymmetry around its axis of rotation. We define two different coordinate systems, one which is fixed, given by (x, y, z) and another one that follows the rotating body given by (x', y', z') . If we pick z' as the rotation axis, the rotation matrix R (relating both coordinate systems) is given by:

$$R = \begin{pmatrix} \cos \phi(t) & \sin \phi(t) & 0 \\ -\sin \phi(t) & \cos \phi(t) & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (2.34)$$

where $\phi = \omega t$ (assuming no spin-down, which would add more terms to this expression) and ω is the angular rotational frequency. Since the rotation is around a principal axis, the inertia tensor in the rotating coordinate system is just made of diagonal entries, while the inertia tensor in the fixed coordinate system is, applying $I = R^T I' R$ where \mathcal{T} indicates a transpose:

$$\begin{aligned} I_{11} &= I'_{11} \cos^2 \omega t + I'_{22} \sin^2 \omega t = \frac{I'_{11} + I'_{22}}{2} + \frac{I'_{11} - I'_{22}}{2} \cos 2\omega t \\ I_{12} &= I_{21} = \frac{I'_{11} - I'_{22}}{2} \sin 2\omega t \\ I_{22} &= \frac{I'_{11} + I'_{22}}{2} - \frac{I'_{11} - I'_{22}}{2} \cos 2\omega t \\ I_{33} &= I'_{33} \\ I_{13} &= I_{31} = I_{23} = I_{32} = 0. \end{aligned} \quad (2.35)$$

From equation (1.54), we can calculate the gravitational-wave emission of this simple model:

$$\begin{aligned} h_+(t) &= \frac{4G}{c^4} \frac{(I'_{11} - I'_{22})\omega^2}{d} \frac{1 + \cos^2 \iota}{2} \cos 2\omega t \\ h_\times(t) &= \frac{4G}{c^4} \frac{(I'_{11} - I'_{22})\omega^2}{d} \cos \iota \sin 2\omega t. \end{aligned} \quad (2.36)$$

It can be seen that there is emission at twice the rotational frequency if $I'_{11} \neq I'_{22}$. The relative amplitude of the plus and cross polarisations depends on the inclination angle ι , and these two polarisations

are $\pi/2$ out of phase. The overall amplitude h_0 is:

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{\epsilon I'_{33} f^2}{d}, \quad (2.37)$$

where as explained in the previous section we have defined a fiducial ellipticity factor $\epsilon = (I'_{11} - I'_{22})/I'_{33}$. For the rigid-body ellipsoid case, the ellipticity is $\epsilon = (b^2 - a^2)/(b^2 + a^2)$, where a and b are the axes of the ellipse in the x-y plane.

This is the CW emission that most of the published searches refer to when discussing emission mechanisms, showing emission only at $f = 2\nu = \omega/\pi$. Including spin-down terms in the definition of the rotating angle $\phi = \phi_0 + \omega(t - t_0) + 0.5\omega_1(t - t_0)^2$ would only add extra terms proportional to the spin-down term, such as:

$$\ddot{I}_{11} = -4 \frac{I'_{11} - I'_{22}}{2} [\omega + \omega_1(t - t_0)]^2 \cos 2\phi(t) - 2\omega_1 \frac{I'_{11} - I'_{22}}{2} \sin 2\phi(t), \quad (2.38)$$

where the second term is present because of the addition of ω_1 , but for realistic values of this parameter this second term is always much smaller than the first term and can be dismissed. The physical intuition behind the emission at twice the rotational frequency is due to the fact that the ellipsoid will look equal to any observer after a 180° rotation.

From this simple model it is clear that a $l, m = 2, 0$ quadrupole moment produced by the centrifugal bulge will not emit CWs, since this quadrupole is symmetric around the rotation axis and thus $I_{xx} = I_{yy} \neq I_{zz}$. Looking at equation (2.36), it can be seen that the strain will be 0 in this case.

Instead of modeling the neutron star as a rigid-body ellipsoid with axes a, b, c , we can model the asymmetry as a small mountain on top of a spherically symmetric body, where the mountain is situated at angles (ϕ, θ) (that in the inertial frame will depend on time) with point mass density $\rho_m = m\delta(\vec{x} - \vec{\chi})$ with $\vec{\chi} = [R \sin \theta \cos \phi, R \sin \theta \sin \phi, R \cos \theta]$. The total moment of inertia is the sum of the spherical body and the mountain, and we have [130]:

$$I'_{11} - I'_{22} = -mR^2 \sin^2 \theta \cos 2\phi, \quad (2.39)$$

which results in an asymmetry given by:

$$\epsilon_m = -\frac{5}{2} \frac{m}{M} \sin^2 \theta \cos 2\phi \quad (2.40)$$

and the same emission as before. It can be seen that the ellipticity is proportional to the mass ratio and that it also depends on the position of the mountain relative to the principal axis: if the mountain is located at the poles (defined by the rotation axis), there will be no emission (the ellipticity grows as the mountain gets closer to the equator); also, the same will happen if the mountain is located halfway between the two principal axis perpendicular to the rotation axis, since the two moments of inertia will be equal (the ellipticity is maximum when the mountain is located at one of these two axes). Assuming optimal positioning, we see that the maximum ellipticity values presented in the previous section (around 10^{-6}) correspond to a mass ratio around 10^{-6} .

2.4.2 | Gravitational waves from general rotation

In a more general case, the neutron star will not be spinning around one of its principal axes. As in the previous case, we need to express the inertia tensors in an inertial frame, and to do so we have to perform several rotations. We follow the same conventions as in subsection 1.3.3, where the rotation matrix is given by $M = M_3 M_2 M_1$ following the zxz convention. The Euler angles are now a function of time, such as $\theta = \theta_0 + \dot{\theta}t$.

In the inertial frame, the angular momentum \vec{J} is conserved and we define it to be directed through the inertial z axis, i.e. $\vec{J}_I = [0, 0, J]$. In this way, the angle θ is the angle that is called ‘‘wobble angle’’, between the angular momentum and the third principal axis.

The angular velocities have the following directions: $\vec{\omega}_\theta$ is directed through the line of nodes; $\vec{\omega}_\psi$ is directed through the z_R axis; $\vec{\omega}_\phi$ is directed through the z_I axis. They can be expressed in the rotating frame as:

$$\begin{aligned}\vec{\omega}_\theta &= [\dot{\theta} \cos \psi, -\dot{\theta} \sin \psi, 0] \\ \vec{\omega}_\phi &= [\dot{\phi} \sin \theta \sin \psi, \dot{\phi} \sin \theta \cos \psi, \dot{\phi} \cos \theta] \\ \vec{\omega}_\psi &= [0, 0, \dot{\psi}].\end{aligned}\tag{2.41}$$

To find the time-dependence of the Euler angles, we need to solve the system $\vec{J} = I\vec{\omega}$ in the rotating frame, where the matrix I is diagonal and $\vec{\omega} = \vec{\omega}_\theta + \vec{\omega}_\psi + \vec{\omega}_\phi$:

$$\begin{aligned}I_{xx}(\dot{\theta} \cos \psi + \dot{\phi} \sin \theta \sin \psi) &= J \sin \theta \sin \psi \\ I_{yy}(-\dot{\theta} \sin \psi + \dot{\phi} \sin \theta \cos \psi) &= J \sin \theta \cos \psi \\ I_{zz}(\dot{\psi} + \dot{\phi} \cos \theta) &= J \cos \theta.\end{aligned}\tag{2.42}$$

These equations can be converted to these non-linear equations for the angles:

$$\begin{aligned}\dot{\phi} &= J \left(\frac{\cos^2 \psi}{I_{yy}} + \frac{\sin^2 \psi}{I_{xx}} \right) \\ \dot{\theta} &= J \left(\frac{1}{I_{xx}} - \frac{1}{I_{yy}} \right) \sin \theta \sin \psi \cos \psi \\ \dot{\psi} &= J \left(\frac{1}{I_{zz}} - \frac{\cos^2 \psi}{I_{yy}} - \frac{\sin^2 \psi}{I_{xx}} \right) \cos \theta,\end{aligned}\tag{2.43}$$

which do not have a simple analytical solution. Following the results from [129], the angles are described by complicated Jacobi elliptical functions. An approximated analytical result can be obtained if we expand in terms of these parameters:

$$\gamma = \frac{I_{xx} a}{I_{zz} b} \quad \text{and} \quad \kappa = \frac{1}{16} \frac{I_{zz} I_{yy} - I_{xx}}{I_{xx} I_{zz} - I_{yy}},\tag{2.44}$$

where $a = \omega_x(t=0)$ and $b = \omega_z(t=0)$. The parameter γ describes the precession angle while κ describes the non-axisymmetry $I_{yy} - I_{xx}$ relative to the axisymmetric non-sphericity $I_{zz} - I_{yy}$. The result, up to second order in these expansion parameters is:

$$h_+ = A_{+,0} \cos [2\omega_r t] + A_{+,1} \cos [(\omega_r + \omega_p)t] + A_{+,2} \cos [2(\omega_r + \omega_p)t]\tag{2.45}$$

$$h_\times = A_{\times,0} \sin [2\omega_r t] + A_{\times,1} \sin [(\omega_r + \omega_p)t] + A_{\times,2} \sin [2(\omega_r + \omega_p)t],\tag{2.46}$$

where

$$A_{+,0} = h_0(1/2) (1 + \cos^2 \iota)\tag{2.47}$$

$$A_{+,1} = 2h'_0 g(\theta_0) \sin \iota \cos \iota\tag{2.48}$$

$$A_{+,2} = 2h'_0 g^2(\theta_0) (1 + \cos^2 \iota)\tag{2.49}$$

$$A_{\times,0} = h_0 \cos \iota\tag{2.50}$$

$$A_{\times,1} = 2h'_0 g(\theta_0) \sin \iota\tag{2.51}$$

$$A_{\times,2} = 4h'_0 g^2(\theta_0) \cos \iota\tag{2.52}$$

and

$$h_0 = \frac{1}{d} \frac{4G\omega_r^2}{c^4} (I_{xx} - I_{yy})\tag{2.53}$$

$$h'_0 = \frac{1}{d} \frac{G(\omega_r + \omega_p)^2}{c^4} \left(I_{zz} - \frac{I_{xx} + I_{yy}}{2} \right)\tag{2.54}$$

$$g(\theta_0) \equiv \frac{I_{xx} a}{I_{zz} b} \approx \theta_0.\tag{2.55}$$

The precession frequency is defined as:

$$\omega_p = \frac{\pi b}{2K(m)} \sqrt{\frac{(I_{zz} - I_{yy})(I_{zz} - I_{xx})}{I_{xx}I_{yy}}}, \quad (2.56)$$

where $K(m)$ is the complete elliptic integral of the first kind and $m = \frac{(I_{yy} - I_{xx})I_{xx}a^2}{(I_{zz} - I_{yy})I_{zz}b^2}$.

The most striking difference between this model with precession and the first model is the appearance of new terms on the waveform. It can be seen that this model predicts emission of gravitational waves at three distinct frequencies: at twice the rotational frequency, which was already present in the simplified model of the previous subsection; at the rotational frequency plus the precession frequency, and at twice the rotational frequency plus twice the precession frequency.

If we set $\theta_0 = 0$ (the wobble angle is zero), we recover the simplified case of the previous subsection, where the rotation axis was aligned with a principal axis.

If $I_{xx} = I_{yy} = I_0$, then it can be seen that $\dot{\theta} = 0$, i.e. the angle between the angular momentum J and the third rotating axis is constant, and that $\dot{\phi} = J/I_0$ and $\dot{\psi} = J \cos \theta (1/I_{zz} - 1/I_0)$. We can compute the waveform for this simplified biaxial precessing model:

$$h_+ = A_{+,1} \cos(\omega t) + A_{+,2} \cos(2\omega t) \quad (2.57)$$

$$h_{\times} = A_{\times,1} \sin(\omega t) + A_{\times,2} \sin(2\omega t), \quad (2.58)$$

where

$$A_{+,1} = h_0'' \sin 2\theta \sin \iota \cos \iota \quad (2.59)$$

$$A_{+,2} = 2h_0'' \sin^2 \theta (1 + \cos^2 \iota) \quad (2.60)$$

$$A_{\times,1} = h_0'' \sin 2\theta \sin \iota \quad (2.61)$$

$$A_{\times,2} = 4h_0'' \sin^2 \theta \cos \iota \quad (2.62)$$

and $h_0'' = \frac{4\pi^2 G}{c^4} \frac{(I_{zz} - I_{xx})f^2}{d}$. This model is the same as the one described in [132], where instead of a precessing biaxial star, a non-precessing star with a pinned superfluid with a symmetry axis different from any principal axis is described.

2.4.3 | The signal at the detectors

The previous two subsections have presented two different models that describe how the signal looks like in the source frame, and section 1.3.3 has shown how to project a traveling GW onto the detector frame. In this section we join these two notions and also include other effects such as the Doppler modulation in order to describe the GW in the detector frame.

The CWs emitted by a neutron star can be generally described in the detector frame by:

$$h(t) = \sum_l A_{l,+}(t) F_+(t, \psi, \hat{n}) \cos l\phi(t) + A_{l,\times}(t) F_{\times}(t, \psi, \hat{n}) \sin l\phi(t), \quad (2.63)$$

where the sum over l implies a sum of different spectrum frequencies, where different models predict a different number as we have seen in the previous two subsections. In order to describe the GW in the detector frame, we need to relate the time at the source τ with the time at the detector t .

From electromagnetic observations of pulsars it is known that the rotational phase at the source frame can be described with a Taylor expansion in powers of frequency (usually one or two frequency derivatives are enough to describe the detected electromagnetic pulses, which are tied to the rotation of the star) with respect to a reference time τ_r (where τ refers to time in the source frame). Since the gravitational-wave phase is assumed to be locked to the rotational phase, we can also describe the GW phase with a Taylor expansion:

$$\phi_s(\tau) = \phi_0 + 2\pi[f_0(\tau - \tau_r) + \frac{f_1}{2!}(\tau - \tau_r)^2 + \dots], \quad (2.64)$$

where we define f_0 and f_1 as the frequency and first frequency derivative (spin-down/up), respectively, at τ_r and ϕ_0 as an initial phase. As discussed in section 6.3 of [76], it is known that higher-order frequency derivatives such as f_2 could be present due to a wide binary orbit.

To relate the phase in the source frame to the phase in the detector frame we need the timing relation between the source time and the detector time. The full relation is derived from the geometrical relation:

$$\vec{r}_{DS} = d\hat{n} - \vec{r}, \quad (2.65)$$

i.e. to go from the Solar System Barycenter (SSB) to the source is the same as to go from the SSB to the detector and then from the detector to the source (\vec{r}_{DS}). Usually, finite distance effects are not taken into account and this vectorial relation is translated to (for an isolated source, and neglecting relativistic wave propagation effects such as Einstein and Shapiro delays):

$$\tau = t + \frac{\vec{r}(t) \cdot \hat{n}}{c} - \frac{d}{c}, \quad (2.66)$$

where \vec{r} is the position of the detector with respect to the SSB and \hat{n} is the position of the source in sky (a full timing relation can be found in [133]). The second term in the right-hand side takes into account the Doppler modulation due to Earth's movement respective to the SSB, while the third term gives the time of travel between the source and detector. If the source is in a binary system, then:

$$\tau + \frac{R(\tau)}{c} = t + \frac{\vec{r}(t) \cdot \hat{n}}{c} - \frac{d}{c}, \quad (2.67)$$

where $R(\tau)$ is the radial distance of the binary orbit projected along the line of sight ($R > 0$ means that the NS is further away than the binary barycenter, $R < 0$ means otherwise). By neglecting the relativistic effects, we are assuming that the binary orbit can be described with a classical Keplerian orbit and that the relativistic effects such as the decay of orbital period do not produce a noticeable effect in our analysis. The decay of the orbital period is proportional to $P_b^{-5/3}$ (where P_b is the orbital period), shorter periods producing faster decays. For a period of 0.01 days and a binary of stars with solar masses, the orbital decay is of the order of 10^{-11} s/s.

The projected radial distance of an ellipse is given by [134]:

$$R = a(1 - e \cos E) \sin \iota_b \sin(\omega + \nu) = a(1 - e \cos E) \sin \iota_b (\sin \omega \cos \nu + \cos \omega \sin \nu), \quad (2.68)$$

where ν is the true anomaly, e is the eccentricity of the orbit, a is the semi-major axis amplitude, ι_b is the angle of inclination of the orbital angular momentum with respect to the observer's sky plane, the angle ω is the argument of periastron and the angle E is the eccentric anomaly, given by the transcendental equation:

$$\tau - \tau_p = \frac{E - e \sin E}{\Omega}, \quad (2.69)$$

where $\Omega = 2\pi/P_b$ is the angular frequency, τ_p is the time of periastron passage (where $E = 0$) and the true anomaly ν is related to the eccentric anomaly by:

$$\cos \nu = \frac{\cos E - e}{1 - e \cos E}. \quad (2.70)$$

Combining equations (2.68) and (2.70) we have:

$$\frac{R(\tau)}{c} = a_p [\sin \omega (\cos E - e) + \cos \omega \sin E \sqrt{1 - e^2}], \quad (2.71)$$

where we have defined $a_p = a \sin \iota_b / c$. This equation depends on five orbital binary parameters: a_p , ω , Ω , τ_p , and e , which fully describe the Keplerian elliptical orbit.

If we choose to consider only orbits with small eccentricity ($e \ll 1$), the ELL1 model gives an approximation to lowest order in e to equation (2.71), which is analytically tractable [134]. This model

substitutes the time of periastron passage τ_p with the time of passage through the ascending node τ_{asc} , a quantity that remains well-defined even for circular orbits. These two times are related by:

$$\tau_{\text{asc}} = \tau_p - \frac{\omega}{\Omega}. \quad (2.72)$$

For the ELL1 model the projected radial distance becomes:

$$\frac{R(\tau)}{c} = a_p \sin[\Omega(\tau - \tau_{\text{asc}})] + a_p \frac{e \cos \omega}{2} \sin[2\Omega(\tau - \tau_{\text{asc}})] - a_p \frac{e \sin \omega}{2} \cos[2\Omega(\tau - \tau_{\text{asc}})] + \mathcal{O}(e^2). \quad (2.73)$$

Now we want to express the phase evolution given by equation (2.64) in the detector frame. The first step that we take is to redefine the constant reference time τ_r to t_r . We also drop the constant d/c term, which is just changing again the reference time t_r by a constant offset (this virtually joins the SSB and the binary barycenter). Firstly, for an isolated source, we have:

$$\begin{aligned} \phi(t) &= \phi'_0 + 2\pi f_0 \left(t - t_r + \frac{\vec{r}(t) \cdot \hat{n}}{c} \right) + \pi f_1 \left(t - t_r + \frac{\vec{r}(t) \cdot \hat{n}}{c} \right)^2 + \mathcal{O}(f_2) \\ &\approx \phi'_0 + 2\pi f_0 \left(t - t_r + \frac{\vec{r}(t) \cdot \hat{n}}{c} \right) + \pi f_1 (t - t_r)^2 + 2\pi f_1 (t - t_r) \frac{\vec{r}(t) \cdot \hat{n}}{c} + \mathcal{O}(f_2). \end{aligned} \quad (2.74)$$

Since now we have the phase model in the detector frame, we can derive the frequency of the gravitational wave in the detector as:

$$f(t) = \frac{1}{2\pi} \frac{d\phi}{dt} \approx f_0 + f_0 \frac{\vec{v}(t) \cdot \hat{n}}{c} + f_1 t. \quad (2.75)$$

This frequency evolution only depends on four parameters: two sky positions, f_0 and f_1 .

For the binary case, with equations (2.67) and (2.73) we can derive a relation between the times in the two different frames. We redefine the constant τ_{asc} by t_{asc} , which again just changes the reference time by a constant factor. We express $R(\tau)$ as $R(t_{\text{SSB}})$. The error created by this approximation is shown in figure 2.5, where it can be seen that it becomes larger for higher frequencies and shorter periods. The timing equation now reads:

$$\tau = t + \frac{\vec{r}(t) \cdot \hat{n}}{c} - a_p \sin[\Omega(t - t_{\text{asc}})] - a_p \frac{e \cos \omega}{2} \sin[2\Omega(t - t_{\text{asc}})] + a_p \frac{e \sin \omega}{2} \cos[2\Omega(t - t_{\text{asc}})]. \quad (2.76)$$

With the previous equation, the phase model (without frequency derivatives in the source frame) at the detector frame now reads:

$$\begin{aligned} \phi(t) &= \phi'_0 + 2\pi f_0 \left(t - t_r + \frac{\vec{r}(t) \cdot \hat{n}}{c} \right. \\ &\quad \left. - a_p \sin[\Omega(t - t_{\text{asc}})] - a_p \frac{e \cos \omega}{2} \sin[2\Omega(t - t_{\text{asc}})] + a_p \frac{e \sin \omega}{2} \cos[2\Omega(t - t_{\text{asc}})] \right). \end{aligned} \quad (2.77)$$

Since now we have the phase model in the detector frame, we can derive the frequency of the gravitational wave in the detector as:

$$\begin{aligned} f(t) &= \frac{1}{2\pi} \frac{d\phi}{dt} = f_0 + f_0 \frac{\vec{v}(t) \cdot \hat{n}}{c} - f_0 a_p \Omega \cos[\Omega(t - t_{\text{asc}})] \\ &\quad - f_0 a_p \Omega e \cos \omega \cos[2\Omega(t - t_{\text{asc}})] - f_0 a_p \Omega e \sin \omega \sin[2\Omega(t - t_{\text{asc}})]. \end{aligned} \quad (2.78)$$

If we assume that $e = 0$, the frequency-time pattern is:

$$f(t) = f_0 + f_0 \frac{\vec{v}(t) \cdot \hat{n}}{c} - f_0 a_p \Omega \cos[\Omega(t - t_{\text{asc}})], \quad (2.79)$$

which depends on six parameters: two sky positions, f_0 , Ω , a_p and t_{asc} .

This model assumes that the neutron star does not suffer any glitches during the observing time, and that the effect of spin-wandering (stochastic variations on the rotational frequency due to the accretion process), if present, can be neglected.

The phase and frequency models given by equations (2.74, 2.75, 2.77, 2.78), jointly with equation (2.63), completely define the waveforms that we want to search in the gravitational-wave data. Methods to do this will be described in the next chapter. Figure 2.6 shows some examples of the frequency-time pattern given by equations (2.75, 2.79) with different zoom levels, where the different modulations can be observed.




Figure 2.5: Error in the frequency-time pattern during one month made when assuming that $R(\tau) = R(t_{SSB})$. The different colors show three different combinations of f_0 and orbital period P_b .




Figure 2.6: The left plot shows the frequency-time pattern of an isolated neutron star during 1.5 years. The yearly and daily (in the inset) Doppler modulations are shown, where the spin-down effect can be seen by noticing the different frequency in the maximum of the first year and the maximum of the second. The right plot shows a frequency-time pattern for a neutron star in a binary system without spin-down. The daily Doppler modulation is less clear due to the binary Doppler modulation.

CHAPTER 3

DATA ANALYSIS AND CONTINUOUS WAVE SEARCHES

This chapter presents an introduction to data analysis, and its application to the search of continuous waves. The first section gives a brief introduction to important statistical concepts such as decision theory, parameter estimation, and mismatch. The second section presents an overview of the different types of continuous waves searches, and the different types of search methods. The last section explains with more detail a particular search method, called *SkyHough*, which is the starting point for the results obtained in this thesis.

3.1 | Gravitational-wave data analysis

The problem of detecting a signal with unknown parameters buried in noise is classically studied in terms of hypothesis testing. After having detected the signal, the unknown parameters have to be estimated as accurately as possible.

In this section we give an introduction to both hypothesis testing and parameter estimation, and we also introduce the concept of mismatch.

3.1.1 | Preliminary definitions

Fourier transforms

The data $\vec{x} = [x_0, x_1, \dots, x_{N-1}]$ recorded from a gravitational-wave detector with a sampling frequency of $f_s = 16384$ Hz is usually described as being noise n (hypothesis H_0) or as a sum of noise and an astrophysical signal h (hypothesis H_1):

$$x_j = n_j + h_j, \quad (3.1)$$

where j indexes time ($t_j = t_0 + j\Delta t$) and Δt is the inverse of sampling frequency. The noise (known to be the result of many small independent effects that may change with time, as discussed in subsection 1.3.2) is modeled as a stochastic process with statistical properties given by the joint probability distribution $p(\vec{n})$ (usually modeled with a colored Gaussian distribution), while the signal is modeled as a deterministic process that depends on multiple parameters. No interactions (non-linear effects) between the noise and signal are considered.

The discrete Fourier Transform (DFT) is defined by:

$$\tilde{x}_k \equiv F(\vec{x}) = \Delta t \sum_{j=0}^{N-1} x_j e^{-2\pi i j k / N}, \quad (3.2)$$

where i is the imaginary unit and k is a frequency bin index. The resolved or bin-centered frequencies are given by:

$$f_k = k\Delta f, \quad (3.3)$$

where $\Delta f = 1/(N\Delta t)$ and k only needs to go from 0 to $N/2$, since the DFT of a real-valued signal is hermitian. Because x_j is dimensionless, it can be seen that \tilde{x}_k has dimensions of time. The Fourier transform is a linear operator, which implies that:

$$F(\vec{n} + \vec{h}) = \tilde{n}_k + \tilde{h}_k. \quad (3.4)$$

The DFT of a bin-centered sinusoidal signal $h_j = A \cos(2\pi f_s j \Delta t + \phi_0)$ is (where bin-centered means that $f_s = f_k$ for some k):

$$F(\vec{h}) = \frac{AN\Delta t}{2} e^{i\phi_0} [\delta(f_k - f_s) + \delta(f_k + f_s)], \quad (3.5)$$

where $\delta(f_k \pm f_s)$ is only nonzero when $f_k = \pm f_s$. It can be seen that only a single frequency bin (the same with positive and negative frequency) will be different than zero, with amplitude $\tilde{h}_k = \frac{AN\Delta t}{2} e^{i\phi_0}$ and power $\tilde{h}_k^2 = \frac{A^2 N^2 (\Delta t)^2}{4}$.

When the DFT is calculated, it is common to multiply the time-series by a window function w (a summary of different window functions can be found in [135]):

$$\tilde{x}_k = \frac{\Delta t}{C} \sum_{j=0}^{N-1} x_j w_j e^{-2\pi i j k / N}, \quad (3.6)$$

where $C = (1/N \sum_{j=0}^{N-1} w_j^2)^{1/2}$. This is done in order to alleviate spectral leakage of the signal to neighbouring frequency bins. The time-series has a finite duration, and this dataset without windowing can be understood as having an implicit rectangular window that sets the data to zero beyond the initial and ending times. The Fourier transform of the rectangular function R of duration T is proportional to the cardinal sine function:

$$F(R) = T \operatorname{sinc} \pi f T. \quad (3.7)$$

A multiplication in time-domain is equivalent to convolution in frequency-domain. For this reason, the convolution of $F(R)$ and the Fourier transform of the sinusoidal will be a cardinal sine function centered on the frequency f_s of the signal. Although windowing is widely used as a tool to alleviate leakage, it might not be the optimal way to achieve the best sensitivity, as discussed for example in [136].

It can be seen that when f_s is coincident with one of the frequencies of the DFT given by $f_k = k\Delta f$, the center of the cardinal sine function will fall at the center of one frequency bin, and the value of the cardinal sine will be exactly zero at all other bins (the zeroes are located at multiples of $1/T$). When the frequency f_s of the signal does not coincide with the resolved frequencies of the DFT, there will be leakage, since the values of the cardinal sine function at the other frequency bins will no longer be 0 (intuitively, this means that the number of cycles contained in the integration time are not an integer number, i.e. the point at t_0 is not equal to the point that would be obtained at t_N). The appearance of leakage is unrelated to the presence of noise from the detectors.

Defining the frequency f_l where $l = k - \delta$ and $\delta = [-1/2, 1/2]$, the Fourier transform of the non-centered sinusoidal windowed signal is given by:

$$\tilde{h}_k = AN\Delta t [e^{i\phi_0} D_N(k-l) + e^{-i\phi_0} D_N(k+l)] \quad (3.8)$$

$$D_N(z) = e^{i\pi z(1-1/N)} \frac{\sin(\pi z)}{N \sin(\pi z/N)}, \quad (3.9)$$

where D_N is the Dirichlet kernel, which is related to the sine cardinal function. The left plot of figure 3.1 shows three different examples: one with a resolved frequency (blue circles) and two with unresolved




Figure 3.1: DFT amplitudes of a 1 s signal with arbitrary amplitude. The left plot shows DFTs without spin-down/up at three different frequencies, while the right plot shows six DFTs with different frequencies and spin-down/up.

frequencies (orange and green circles). It can be seen that for $\delta = 1/2$ (orange circles), the reduction in amplitude can be as high as 0.36.

Even if by chance f_s is a resolved frequency, due to the different phase modulations affecting the signal (due to spin-down/up, Doppler modulation) the frequency will not stay resolved during the time of a single DFT. The amount of signal lost to neighbouring bins depends on the factor $\Delta f T$ (where Δf is the range of variation of the frequency): the higher this factor is, the more percentage of signal will be shared with other frequency bins. Furthermore, the amplitude modulation discussed in 1.3.3 will split the signal in five different frequencies with different amplitudes, so even if the central frequency is bin-centered the other four frequencies may not be.

As will be seen later, by design many CW algorithms restrict the product $\Delta f T$ to not be greater than 1, but as we can see even for smaller values the amplitudes are smeared. This is shown in the right plot of figure 3.1, where for example the brown crosses show that an amplitude factor of 0.37 can be lost. The brown and purple crosses are almost equal, showing that when the number of bins that are crossed by the signal are equal (in this case two bins), the quantity that is significant is the value of central frequency, not the absolute value of the spin-down/up. The green points marked with a + symbol show that if we took a long-duration Fourier transform with a CW signal in it (meaning that many frequency bins would be crossed), the signal would be spread in many frequency bins, thus complicating its detection above noise. Analytical estimations of the complete spectrum of CW signals are given in [137] and [138].

Likelihoods

The likelihood of measuring a single value x_j when hypothesis H_0 is true is given by [208]:

$$P(x_j|H_0) = \frac{1}{\sqrt{2\pi}\sigma_j} e^{-\frac{x_j^2}{2\sigma_j^2}}, \quad (3.10)$$

where we have assumed that the noise distribution is given by a Gaussian with 0 mean and standard deviation σ_j . The probability of measuring an entire dataset \vec{x} of N points is then given by:

$$P(\vec{x}|H_0) = \frac{1}{[2\pi\det(\gamma)]^{N/2}} \exp\left[-\frac{1}{2} \sum_{jl} x_j \gamma_{jl}^{-1} x_l\right], \quad (3.11)$$

where $\det(\cdot)$ denotes the determinant, and the non-diagonal elements of the correlation matrix $\gamma_{jl} = E[n_j n_l]$ quantify the correlation between samples j and l and the diagonal elements represent the variance for each sample (when no correlation exists and the variance is the same for all samples, γ^{-1} can be

substituted for a constant σ^2). When γ only depends on the time difference and not on the absolute value of time, the noise is said to be wide-sense stationary.

When data from more than one detector is included, the expression is:

$$P(\vec{x}|H_0) = \frac{1}{[2\pi\det(\gamma)]^{N/2}} \exp \left[-\frac{1}{2} \sum_{IKjl} x_{Ij} \gamma_{IjKl}^{-1} x_{Kl} \right], \quad (3.12)$$

where I and K are indices that go over the different detectors (i.e. for two detectors, $I = 0, 1$). If the noise is uncorrelated between detectors $\gamma_{IjKl} = \delta_{IK} S_{n;I,jl}$. The single-sided power spectral density (PSD) S_n is defined as the Fourier transform of the noise auto-correlation function:

$$S_{n;IK}(f) = 2 \int_{-\infty}^{\infty} \gamma_{IK}(\tau) e^{-i2\pi f\tau} d\tau. \quad (3.13)$$

The PSD can also be approximated by [158]:

$$S_n(f) \approx \frac{2}{T_{obs}} E [|\tilde{n}(f)|^2], \quad (3.14)$$

which holds exactly in the limit of infinite duration T_{obs} of the time-series. The PSD is estimated from the data itself, either by calculating the mean of different groups of samples (more than one group is needed) or by fitting the frequency-domain values to a physical model of the spectrum.

The probability (likelihood) of data \vec{x} when hypothesis H_1 is true is given by:

$$P(\vec{x}|H_1) = \frac{1}{[2\pi\det(\gamma)]^{N/2}} \exp \left[-\frac{1}{2} \sum_{jl} (x_j - h_j) \gamma_{jl}^{-1} (x_l - h_l) \right], \quad (3.15)$$

since when we subtract the exact h from x only the noise remains. We define an inner product given by [208]:

$$(x|y) \equiv \int_{-\infty}^{\infty} \frac{\tilde{x}(f)\tilde{y}^*(f) + \tilde{x}^*(f)\tilde{y}(f)}{S_n(f)} df = 4 \int_0^{\infty} \frac{\Re[\tilde{x}(f)\tilde{y}^*(f)]}{S_n(f)} df, \quad (3.16)$$

where \Re means taking the real part of a complex quantity and the last equality is only true when x and y are real-valued signals. For more than one detector, the inner product is:

$$(x|y) \equiv 4 \int_0^{\infty} \Re[\tilde{x}^I(f) S_{n;IK}^{-1}(f) \tilde{y}^{*K}(f)] df = 4 \sum_I \int_0^{\infty} \Re[\tilde{x}^I(f) S_{n;I}^{-1}(f) \tilde{y}^{*I}(f)] df, \quad (3.17)$$

where in the last equality we have assumed uncorrelated noise between the different detectors. With this inner product, the likelihoods can be expressed as:

$$P(\vec{x}|H_0) = \frac{1}{[2\pi\det(\gamma)]^{N/2}} e^{-\frac{1}{2}(x|x)} \quad (3.18)$$

$$P(\vec{x}|H_1) = \frac{1}{[2\pi\det(\gamma)]^{N/2}} e^{-\frac{1}{2}(x-h|x-h)}. \quad (3.19)$$

The signal-to-noise ratio (SNR) of a signal is defined as:

$$\text{SNR} = \sqrt{(h|h)}. \quad (3.20)$$

3.1.2 | Hypothesis testing

Before explaining what hypothesis testing is, we summarize the main differences between the two major schools of statistical inference: frequentism and Bayesianism. The most basic difference is related to the concept of probability. While frequentism understands probability as a measure of the long-term (ideally infinite) frequency of events, Bayesianism understands probability as a measure of the degree of confidence

in a statement. Another major conceptual difference is related to the data and the parameters of a model: while frequentism assumes that the parameters of a model are fixed and the data is one realization of an ensemble of possible data realizations, Bayesianism thinks of the data as fixed and parameters as unknown variables described by prior and posterior distributions. The main mathematical tool for Bayesianism is Bayes' formula, which relates the posterior distribution to the prior distribution and the likelihood function. Since frequentism does not rely on prior distributions, it only uses the likelihood function.

Hypothesis testing consists of quantifying the relative probabilities between the different proposed hypotheses. The most basic form of hypothesis testing is a simple test between two different hypotheses, when all the parameters that define the signals in the hypotheses are known. In our case, H_0 represents the hypothesis that there is no signal in the data (i.e. $h_j = 0$ for all j), while H_1 stands for the hypothesis that there is a nonzero signal $h(\vec{\Theta}, t)$, where $\vec{\Theta}$ is a vector that contains all the parameters of the signal.

The frequentist procedure to decide between the two hypotheses given a dataset \vec{x} is defined by computing a scalar detection statistic $\Lambda(\vec{x})$, and comparing it to a pre-defined threshold Λ_{th} such that H_0 is accepted if $\Lambda(\vec{x}) < \Lambda_{th}$, while H_1 is accepted otherwise. From the probability distributions of \vec{x} under both hypotheses, the two conditional probabilities $P(\Lambda|H_0)$ and $P(\Lambda|H_1)$ can be calculated. When deciding which hypothesis is true, two different types of error can be made: type I error α , also called false alarm probability, which happens when H_0 is true but H_1 is selected; type II error β , also called false dismissal probability, which happens when H_1 is true but H_0 is selected. They are defined [208]:

$$\alpha(\Lambda_{th}) \equiv \int_{\Lambda_{th}}^{\infty} P(\Lambda|H_0) d\Lambda \quad (3.21)$$

$$\beta(\Lambda_{th}) \equiv \langle \beta(\Lambda_{th}, \Theta) \rangle_{\Theta} = \int d\Theta p(\Theta) \beta(\Lambda_{th}, \Theta) = \int d\Theta p(\Theta) \int_{-\infty}^{\Lambda_{th}} P(\Lambda|H_1, \Theta) d\Lambda. \quad (3.22)$$

The last line represents the global false dismissal probability, which is an average over the parameters of the population. The detection probability is defined as $\eta = 1 - \beta$. Another useful quantity that can be defined is $\eta' = 1 - \beta - \alpha = \eta - \alpha$, which can be interpreted as the detection probability only due to the presence of a signal, with the contribution from noise removed (i.e. the probability of the threshold being exceeded because the signal was present, rather than because of the detector noise) [136].

The only thing that is left to choose is the scalar function Λ . The Neyman-Pearson criterion for the optimality of a hypothesis test $\Lambda(x)$ is that the test should maximize the detection probability at a given fixed false alarm probability. This is given by the likelihood ratio, which is defined as:

$$\Lambda(\vec{x}) \equiv \frac{P(\vec{x}|H_1)}{P(\vec{x}|H_0)} = \frac{e^{-(\vec{x}-\vec{h}|\vec{x}-\vec{h})}}{e^{-(\vec{x}|\vec{x})}} \quad (3.23)$$

$$\ln \Lambda(\vec{x}) \propto (x|h) - \frac{1}{2}(h|h), \quad (3.24)$$

where it can be seen that the only term that depends on the data is the inner product $(x|h)$, which is known as the matched filter.

In case that we have a composite hypothesis (when any of the hypotheses present unknown parameters), the likelihood depends on these unknown parameters. This happens for example if the signal in H_1 depends on an unknown amplitude, which is the case for gravitational-wave searches. The common procedure is to use the generalized likelihood ratio or maximum likelihood estimation, where the likelihood $P(\vec{x}|H_1)$ is maximized with respect to the unknown parameters. Then, these maximum likelihood estimators are substituted back into the likelihood, and the maximum likelihood is found, which is independent from the maximized parameters. Furthermore, when the analytical maximization is not possible, a search over a range of values for these unknown parameters has to be done by using some algorithm, like a deterministic search within a defined grid or through stochastic sampling. An example of such generalized likelihood ratio test is the \mathcal{F} -statistic [50], where the four amplitude parameters that define the triaxial aligned model of subsection 2.4.1 are analytically maximized, and the four phase parameters have to be numerically searched.

The generalized likelihood ratio does not provide an optimal detection statistic, in the sense of it having the best detection probability at a certain false alarm probability across the entire parameter space spanned by the unknown parameters. For the CW detection case, an example of this can be seen in [139], where it is shown that the frequentist maximum likelihood ratio test carries an implicit prior distribution of the signal parameters, which does not produce the most powerful test across the full range of parameter space. The performance of a composite hypothesis detection statistic can be compared with the ideal performance of a simple hypothesis test where the parameters are known, thus providing an upper bound on the best possible performance. An example of this comparison for the continuous wave case can be seen in [140].

A high value of the likelihood ratio test does not immediately imply that there is an astrophysical signal in the data: the hypothesis ratio test only says that the signal model fits the data better than the Gaussian noise model. A more complete hypothesis testing would include more hypotheses that would cover other non-Gaussian noise hypotheses or non-astrophysical signals. Instead of proposing an a priori model for the H_0 hypothesis, the background could be better modeled by estimating it from the data. A way to do this without being biased by possible signals is to use time-slides, where the data samples are reordered at random and the background distribution of the likelihood ratio test is estimated.

The overall false alarm probability α_T of crossing the threshold in \mathcal{N} trials is related to the individual false alarm probability for a single trial (assuming independence between the different trials):

$$\alpha_T = 1 - (1 - \alpha)^{\mathcal{N}} \approx \alpha\mathcal{N}. \quad (3.25)$$

where for the approximation we have assumed that $\alpha\mathcal{N} \ll 1$ and only kept the first-order term. Repeating a search over different independent trials can be seen as repeating the single trial search over different noise realizations.

The main difference with the Bayesian framework (besides the usage of priors) is that for the composite hypothesis test, the unknown parameters are integrated out (marginalized) instead of being maximized. This procedure is optimal if the signal parameters of the underlying population of signals are correctly described by the used prior distributions. A disadvantage of this method is that the marginalization usually cannot be performed analytically, and the integrals have to be solved or estimated by an algorithm that highly increases the computational cost. As will be seen later, CW searches use both the Bayesian and frequentist methodologies, depending on the type of search.

3.1.3 | Parameter estimation

After having detected a signal, or having strong indications that there might be a signal in the data, we want to estimate its parameters as accurately as possible. The accuracy of any estimation is limited due to the presence of noise in the detectors. Different point estimators $\hat{\Theta}$ (functions of the data) can be proposed, and their accuracy can be measured using different loss functions, such as the mean-squared-error (MSE) [141]:

$$\text{MSE} \equiv E[(\hat{\Theta} - \Theta)^2] = E\left[\left((\hat{\Theta} - E[\hat{\Theta}]) + (E[\hat{\Theta}] - \Theta)\right)^2\right] = \text{var}(\hat{\Theta}) + (E[\hat{\Theta}] - \Theta)^2 = \text{var}(\hat{\Theta}) + \text{bias}^2(\hat{\Theta}), \quad (3.26)$$

where it can be seen that there are two contributions, one coming from the variance of the estimator and another from the bias of the estimator.

One of the most important frequentist point estimators are the maximum likelihood estimators (MLE), already mentioned previously when discussing the generalized likelihood ratio test. An example of these maximum likelihood estimators can be seen in the continuous wave \mathcal{F} -statistic [50], where the MLE for the four amplitude parameters are found by analytically maximizing the likelihood ratio. The MLE are known to reach the Cramer-Rao lower bound (explained below) when the number of data samples

tends to infinity. Moreover, it can be demonstrated that when a uniform prior distribution is used, the maximum likelihood estimation is equivalent to the maximum a posteriori Bayesian estimation [141].

In the frequentist case, the best precision achievable by an unbiased estimator is given by the Cramer-Rao bound, which gives a lower bound on the variance of the estimator. These bounds can be obtained from the Fisher information matrix, which describes the extent to which detector noise introduces random errors in the estimations:

$$F_{ij} \equiv - \langle \partial_i \partial_j \log p(\vec{x}|H_1, \Theta_0) \rangle_n = \left\langle \left(\frac{\partial h}{\partial \Theta^i} \middle| \frac{\partial h}{\partial \Theta^j} \right) \right\rangle_n. \quad (3.27)$$

It can be seen that it depends on the second derivative of the likelihood around the true parameters Θ_0 , meaning that if the likelihood is Gaussian the Fisher matrix calculates the curvature around the maximum. The inverse of the Fisher information matrix gives the lower bounds on the variance of the estimators, and they can be interpreted in three different ways, as explained in [142, 143]:

1. Asymptotically (when the SNR of the signal tends to infinity or equivalently in the limit in which the waveforms can be considered as linear functions of source parameters) it is equal to the variance-covariance of the posterior distribution (i.e. it gives the error for a singular data realization) when priors are not important.
2. Asymptotically (when the SNR of the signal tends to infinity or equivalently in the limit in which the waveforms can be considered as linear functions of source parameters) it is the frequentist error covariance for the maximum-likelihood parameter estimators (which in this limit are unbiased), assuming Gaussian noise.
3. It gives the Cramer-Rao bound on the expected variance of any unbiased estimator over repeated measurements while keeping the parameter fixed.

For low SNR signals, the non-Gaussian shapes (such as multi-modalities, present in the likelihoods of continuous waves signals as discussed in [144]) of the posterior distribution contribute to lowering the validity of these bounds, making the real variances larger than predicted. Even for high SNR signals, the bounds from the Fisher matrix may not be accurate due to the usage of prior distributions, as seen in [143]. This is because prior distributions may produce biased estimators, but their correspondent variance can be lower than the variance of unbiased estimators, thus lowering the MSE defined in equation (3.26), as seen in example 10.1 of [141]. As happened in the previous hypothesis testing subsection, although the Bayesian framework can produce better results (besides producing full posterior distributions instead of point estimators), the multi-dimensional integrals that need to be solved in order to produce the posterior distributions increase the computational cost and may not be practical. Similarly to the hypothesis testing case, CW searches use both Bayesian and frequentist parameter estimators, depending on the search.

Besides the noise present in the detectors, another limitation to the parameter estimation precision comes from systematic calibration errors present in the current gravitational-wave data [46]. Even if the SNR is very high, an error in the estimations of the calibration parameters will produce a shift on the estimation of the astrophysical parameters, which can be greater than the shift produced by the noise, as shown in [146].

3.1.4 | Mismatch and resolution

The previous subsection described the best precision that can be expected when estimating signal parameters, which is limited by the noise present in the detectors. A different but related notion of precision or resolution arises from the difference between the shape of two waveform templates, produced by a difference of their signal parameters. This difference produces a change on the value of the likelihood ratio or the inner product present in the signal-to-noise ratio, given by equation (3.20), creating a notion of resolution between different parameters.




Figure 3.2: Matched filter between a sinusoidal signal of $f = 100$ Hz and another sinusoidal signal of $f = 100 + \sqrt{m/(T^2\pi^2)}$, where m is the mismatch parameter shown in the x -axis. I represents the matched filter defined in equation (3.28). The integration time T is 3600 s.

This notion is independent of the noise present in the detectors, and can be exemplified by studying how the value of a matched filter operation between two distinct waveforms changes as the parameters describing them are changed, and how this depends on the coherent integration time. A simple example is shown in figure 3.2, where the value of the matched filter given by (where we have assumed the noise to be constant and factored out):

$$I \equiv (x|y) = 4 \int_0^\infty \frac{\Re[\tilde{x}(f)\tilde{y}^*(f)]}{S_n(f)} df \longrightarrow \frac{1}{S_n(f)} \int_0^T x(t)y(t)dt, \quad (3.28)$$

is calculated as a function of the frequency difference between two normalized sinusoidal signals.

The mismatch μ (complementary to the fitting factor FF or overlap used in other fields) can be defined by the loss of signal-to-noise ratio produced when filtering with an incorrect waveform [149]:

$$\mu = \frac{\text{SNR}^2 - \text{SNR}_r^2}{\text{SNR}^2} = 1 - \frac{\text{SNR}_r^2}{\text{SNR}^2} = 1 - \frac{(h|h')}{(h|h)}. \quad (3.29)$$

The mismatch is a number between 0 (fully recovered SNR) and 1 (no recovered SNR), where SNR_r is the recovered signal-to-noise ratio and SNR is the optimal signal-to-noise that would be obtained if the matching was performed with a template calculated with the true signal parameters (when noise is not present in the data). This expression could be used by methods that do not rely on the matched filter, such as the Hough method later explained.

As previously discussed, an exploration of the parameter space of a signal might be necessary in order to detect a signal in the data. Two main methods exist to do this: a gridded search, where the parameter space is divided in different cells in a deterministic way, or a stochastic method such as Markov chain Monte Carlo (MCMC) sampling where the parameter space is explored in a non-deterministic way. For the deterministic methods, we need to define the resolution of parameter space in the different dimensions, which is going to decide the spacing between the different waveform models in order to construct the template bank that contains the templates that are going to be searched. This spacing can be estimated from the mismatch.

To estimate this mismatch, a local (in the sense that global correlations such as those explained in [147] and [148] between parameters are not taken into account) approximation is used, which can be

obtained by doing a Taylor expansion around the true signal parameters and only keeping second-order terms (the zero-order term is zero since the mismatch is 0 at the true parameters and the first-order term is zero since at that location there is a minimum/maximum):

$$\mu \approx g_{ab}(\Theta)d\Theta^a d\Theta^b + \mathcal{O}(d\Theta^3), \quad (3.30)$$

where g_{ab} is the parameter space metric (a and b go from 0 to the number of parameters), Θ represents all the different parameters such as amplitude, frequency, or sky position, and $d\Theta = \Theta_s - \Theta$, where Θ_s are the parameters of the signal. In [150] it is discussed that this may not be the best local approximation, since a sinusoidal function gives a better local approximation.

This metric depends both on the amplitude and phase parameters. When the amplitude parameters have been analytically maximized or are not explicitly searched over, we need to project the general metric into the subspace of phase parameters. In [149] it is shown that for high-frequency signals (where a high number of cycles are present in the integration time, i.e. $fT \gg 1$) this is equivalent to using the phase metric, which consists of assuming a constant amplitude and only taking into account the phase parameters λ :

$$g_{ab} = \langle \partial_a \phi(\lambda) \partial_b \phi(\lambda) \rangle - \langle \partial_a \phi(\lambda) \rangle \langle \partial_b \phi(\lambda) \rangle, \quad (3.31)$$

where $\langle g \rangle = 1/T \int_0^T g(t) dt$, ∂ denotes a partial derivative, and ϕ is given by equation (2.77). This approximated mismatch is unbounded and can be higher than 1, and from previous studies it is known that it highly overestimates the actual mismatch for mismatches higher than 0.5 [149].

As a quick and simple example, if $\phi = \phi_0$, it can be seen that $g_{\phi_0\phi_0} = 0$, meaning that no SNR is lost when using a wrong initial phase, or that no search over this parameter is needed. For a simple sinusoidal, $\phi = 2\pi f_0 t$, so $g_{f_0 f_0} = 4\pi^2(T^2/3 - T^2/4) = T^2\pi^2/3$. If the required mismatch is 0.1, then the needed resolution is $df_0 = \sqrt{m/g_{f_0 f_0}} = \sqrt{0.3/(T^2\pi^2)}$. This can be reversed, and we can get the mismatch as $m = df_0^2 T^2 \pi^2 / 3$. This can be compared to the resolution of the DFT, which is $1/T$:

$$\frac{df_0(m)}{df_0^{\text{DFT}}} = \frac{\sqrt{3m}T}{\pi T} = 0.17\sqrt{\frac{m}{0.1}}. \quad (3.32)$$

From this comparison it can be seen that the resolution obtained from mismatch arguments is much finer than the resolution given by the DFT algorithm. Many CW search algorithms use the resolution obtained from the mismatch metric to construct their template banks, such as [202] or [186]. The mismatch metric approach is also used for other types of gravitational-wave searches, such as Compact Binary Coalescence searches [151], or even for searches of gamma rays from unknown neutron stars [152].

The semi-coherent (the difference between a coherent and semi-coherent search is explained in section 3.2) metric is obtained by averaging the different coherent metrics. Search methods such as *SkyHough* do not use the metric to obtain the resolution needed to limit the maximum mismatch in their template bank, as shown in subsection 3.3.3. Other arguments can be used to estimate the resolution equations, such as calculating the minimum measurable shift in the waveform's frequency-time path, as done in [158].

A related problem to the per-dimension template (parameter space points) placement is how to optimally place the templates in more than one dimension in order to cover as optimally as possible a given range of parameter space [153]. This is a complementary problem to the description of the distance between templates at each dimension, governed by the mismatch as just explained. The optimal placement of parameter space points is similar to the sphere covering problem, where a given volume of space has to be covered with the smallest possible amount of spheres. CW search methods are able to use the optimal lattice when the parameter space metric given by equation (5.8) is independent of the signal parameters, as in the case of directed searches, whereas for all-sky searches non-optimal lattices have to be used [154]. The parameter space points can also be drawn in a stochastic way, as shown in [155].

3.2 | Continuous waves searches: targets, methods, and past results

Continuous waves (CWs) are nearly monochromatic long-lasting gravitational waves that have not yet been detected. They can be generated by multiple sources, such as asymmetric rotating neutron stars (as discussed in the previous chapter), which are the focus of this thesis, white dwarves in individual [159] and binary systems [160], supermassive black hole binaries [161], boson clouds around black holes [162], or dark matter inside stars or planets [163,164]. The main instrument to detect CWs in the frequency range around 20 to 2000 Hz are ground-based interferometer detectors such as LIGO or Virgo. Furthermore, as discussed before, CWs could also be detected with the resonant-mass detectors described in section 1.3 or other new ideas such as [25].

There are several types of CW searches for signals from neutron stars, differentiated by the available information about their targets. From highest to lowest sensitivity:

- Targeted searches: searching for CWs from known pulsars (in isolated and binary systems). Since the rotational phase parameters of the targets are known and are assumed to be equal to the gravitational-wave phase parameters, only the amplitude parameters need to be estimated, and the most sensitive search methods can be used with small computational cost. More recently, narrow-band searches have been developed to complement the traditional targeted methods. These searches allow for a small mismatch between the electromagnetic and gravitational-wave phase parameters such as the initial frequency, and do a search over a small range around the known value. A brief list of past targeted searches is: [165–175].
- Directed or spotlight searches: searching for CWs from sources at known sky positions, such as globular clusters, the center of the galaxy, supernova remnants, or low-mass X-ray binaries like Scorpius X-1. Although the sky positions are known, the initial frequency and frequency derivatives are not, and a search over these parameters has to be done. For this reason, the most sensitive methods used for targeted searches cannot be used (or they have to be used with a reduced set of data). A brief list of past directed searches is: [176–189].
- All-sky searches: searching for CWs from unknown neutron stars (in isolated and binary systems), these searches have the highest computational cost since both the sky positions and frequency parameters are unknown. A brief list of past all-sky searches is: [190–205].

All-sky searches, which are the focus of this thesis, possess interesting particular properties such as:

1. A very small percentage of the estimated neutron star population in our galaxy has been detected as pulsars [65]. We expect that these unseen neutron stars (perhaps with more extreme properties than the detected pulsar population) may emit detectable CWs. Furthermore, all-sky searches inspect regions of parameter space that targeted and directed searches leave unexplored, making all-sky searches a valuable endeavour.
2. Even if the rotational evolution of the electromagnetic poles and the gravitational quadrupole rotational evolution are different (possibly due to an internal asymmetry combined with differential rotation), the semi-coherent methods used for all-sky searches (explained below) are robust to these kinds of mismatches, since they do not use electromagnetic information. Furthermore, these semi-coherent methods are also more robust against other phase fluctuations such as the ones produced by glitches or timing noise [79].
3. Due to the possible high mismatch in the initial search stage and/or the usage of a signal model that is not completely accurate (for example, due to not including higher-order spin-down/up terms, proper motion, or spin-wandering), when a signal is found in the first stage, the SNR that can be

recovered after lowering the mismatch and correcting the signal model is much higher, possibly even more than an order of magnitude higher than the initial SNR. For this reason, the detection confidence could be increased even more by searching the signal at other frequencies (which might be present as explained in subsection 2.4.2) that at the initial stage might have been undetectable (depending on the geometry of the neutron star), or with other less sensitive detectors. This argument is even stronger for all-sky searches from neutron stars in binary systems, since their initial stage mismatch is expected to be higher.

4. All-sky searches are able to constrain physical parameters of the full galactic population within a single search, instead of just constraining parameters from individual targets.

There are two main ways to analyse the data from the gravitational-wave detectors: coherent and semi-coherent methods. Coherent methods demand amplitude and phase coherence between the templates and the data during all the observation time, while semi-coherent methods break up the data in shorter segments and only demand coherence within each of these segments, but not between them. The semi-coherent \mathcal{F} -statistic method is an example of this: a coherent \mathcal{F} -statistic is calculated for each segment by finding the analytical maximum likelihood amplitude estimators, but it is not demanded that these amplitude parameters (h_0 , $\cos\iota$, ψ and ϕ_0) are the same for all segments. Between segments, semi-coherent methods only demand coherence of the frequency-time pattern.

Since coherence is only required within each segment, semi-coherent methods do not need search grids to be as fine as coherent methods do. This is related to what was explained in subsection 3.1.4: the longer the comparison between template and data is, the finer the spacing between templates needs to be in order to maintain a constant mismatch level. For these reasons, semi-coherent methods have a smaller computational cost when compared to coherent methods, and they can be used to do the more expensive directed and all-sky searches. To perform all-sky searches with a limited computational budget, semi-coherent methods have been proven to be more sensitive than coherent methods [157].

The difference in sensitivity between coherent and semi-coherent procedures stems from the fact that the background distributions of semi-coherent searches have a much higher number of degrees of freedom (N times higher, where N is the number of segments), and for this reason at the same threshold level many more candidates from the background distribution will cross it, so to attain the same false alarm probability this threshold has to be increased for semi-coherent searches. Although the number of templates for the coherent search would be much higher, even when this higher number of trials is taken into account the higher number of degrees of freedom of semi-coherent methods has more weight.

The enormous difference in computational cost between coherent and semi-coherent methods can be understood by looking at the scaling of the required number of coherent \mathcal{N}_c and semi-coherent \mathcal{N}_s templates per parameter-space volume $d^4\lambda$ of a 4-dimensional all-sky search for isolated neutron stars with one spin-down/up [208]:

$$d\mathcal{N}_c \propto \sqrt{|\det g_{ab}|} d^4\lambda \propto T_{obs}^5 f^2 d^4\lambda \quad (3.33)$$

$$d\mathcal{N}_s \propto T_{obs} T_c^4 f^2 d^4\lambda. \quad (3.34)$$

One example of a coherent method used for targeted searches is the time-domain Bayesian targeted pipeline [209], which is the most sensitive CW pipeline. This method is able to heterodyne (multiply the data by a complex exponential to remove the phase modulations) the data for each of the searched pulsars, since the phase parameters are known. After heterodyning only the amplitude modulation is left, and downsampling the data to one sample per minute reduces the computational cost. If the phase parameters were not known, the heterodyning procedure could not be applied, and a Bayesian nested sampling search over at least an 8-dimensional parameter space and using the full time-domain data would need to be used, which would increase by many orders of magnitude the computational cost of this search.

Another coherent method relies on the \mathcal{F} -statistic, which is the frequentist likelihood ratio maximized over the amplitude parameters. The way this detection statistic is calculated is detailed in [210]. It can be calculated by heterodyning the data or by resampling as explained in [211]. Several modifications of this coherent statistic exist. For example, when the two orientation parameters of the neutron stars are known, it is substituted by the more sensitive \mathcal{G} -statistic, defined in [140]. Moreover, in [139] it was shown that the \mathcal{F} -statistic can be derived from a Bayesian point of view by proposing an unphysical prior (which was implicit in the derivation of the \mathcal{F} -statistic), and a more sensitive statistic name the \mathcal{B} -statistic with a physical prior was developed. Although CW signals from neutron stars in binary systems can be searched with the \mathcal{F} -statistic, further modifications (which are not a maximum-likelihood estimator) can be done in order to lower the computational cost of these high-dimensional searches. These are the \mathcal{C} -statistic, the Bessel weighted \mathcal{F} -statistic, and the \mathcal{J} -statistic, explained in detail in [212].

Regarding semi-coherent methods, many different pipelines exist and they differ between both the coherent detection statistic that they calculate and the way that these results are combined between segments. Usually, the coherent step consists of either calculating the \mathcal{F} -statistic (or one of its variants) for a given T_c , or on calculating the power in frequency bins of DFTs from the time-domain data. The advantage of using the \mathcal{F} -statistic is that it takes into account both the amplitude and phase modulations that the signal has, and for this reason T_c can be arbitrarily long, which for the DFT powers case is not true, since the phase modulations would distribute the power of the signal to neighbouring bins. This increases the sensitivity when compared to using power from Fourier transforms of the data (also increasing the computational cost, so the overall sensitivity might not be better if the mismatch parameters are not low enough, as recent results have shown [213]). Some methods allow T_c to be longer even when using DFT powers, such as the Cross-Correlation [214] method or the Loosely-Coherent approach [215]. Although these semi-coherent methods can achieve higher sensitivity due to the usage of longer segments, their computational cost is much higher and are usually only used for directed searches or for follow-ups of candidates from all-sky searches, although a new implementation of the Loosely-Coherent approach called Falcon has been recently applied to an all-sky search [205]. A different semi-coherent method is used by the TwoSpect pipeline [98], which searches for all-sky signals from neutron stars in binary systems. This pipeline starts from a time-series of frequency bin powers, but then applies a further Fourier transformation (after having corrected for the amplitude and Earth's modulations) for all time-series of each different frequency bin, in order to reveal excess power at frequencies that are harmonics of the binary orbital period.

For the combination of the results from different segments, there are four different mechanisms:

- Sum the coherent statistics (possibly after weighting them by the noise and antenna patterns of the detectors) along the frequency-time pattern of the signal, as for example done by the PowerFlux pipeline (a technique usually called “stack and slide”). The most sensitive pipeline that uses this technique relies on the usage of the Einstein@Home infrastructure, which uses idle CPU time from the computers of volunteers. The parameter space to be searched is divided in smaller portions, and each of these portions can be independently done by different volunteers. For this reason, a huge computational power is available, and longer coherent times (while maintaining small mismatch parameters) can be used.
- Apply a threshold to the (weighted) coherent detection statistics to convert them to 1s and 0s, and sum them along the frequency-time pattern, done by the Hough pipelines. Although the sensitivity that can be achieved is smaller, this method is more robust to huge power contributions from noise artifacts and its computational cost is smaller, so for a limited computational power its sensitivity can be the same or better than the previous method.
- Search for coincidences between the results from the different coherent segments, as done by the all-sky Time-Domain \mathcal{F} -statistic pipeline.

- Sum the coherent statistics but instead of following a deterministic frequency-time pattern (different for each of the different searched templates), use a stochastic technique to predict the most probable path of the signal. This technique can be applied with the Viterbi algorithm, as explained in [216] or [217].

As mentioned previously, all-sky searches have to search over (at least) a 4-dimensional parameter space, and their computational costs are the highest ones. For this reason, all-sky searches use the less computationally demanding semi-coherent methods.

The first proposed method to do all-sky searches dealt with resampling (stroboscopic or nearest-neighbour) the data in order for it to contain a monochromatic sinusoidal signal, and then doing a FFT. This had to be repeated for all different points of parameter space, so the computational cost was very high and only a small amount of data could be used [218]. Later [156] it was studied that instead of performing a coherent search over all different templates, if the data was divided in smaller segments and combined in an incoherent way, and afterwards candidates from this first stage were followed up by incrementing the coherent time (a strategy called hierarchical search) the total amount of data could be incremented and the sensitivity would be higher. This was the first time that a semi-coherent approach was proposed. In that study it was already mentioned the need for a finer resolution in the incoherent combination step, since residual errors may cause the combination of the incorrect frequency powers (even if the coherent demodulation correctly combined all the power in a single frequency bin). A further study showed that the optimal hierarchical procedure consisted on doing 3 semi-coherent stages and a subsequent final fully coherent stage [219].

The first search for CWs was reported in [220], where seismic waves on Earth produced by CWs were used to put an upper limit around 10^{-17} for 4 different pulsars. The first search for CWs with interferometers was done in 1972 [221], where data from a 30 m laser interferometer was used to search for CWs from the Crab pulsar, setting an upper limit of 3×10^{-17} .

The first CW search using large-scale interferometers was reported in [222], where 50 days of TAMA300 data were used to put an upper limit of 5×10^{-23} at 934.9 Hz for CWs from the supernova remnant SN1987A. The first CW search using data from the LIGO detectors was [165], where GWs from the pulsar J1939+2134 were searched, obtaining an upper limit around 10^{-22} using around 17 days of S1 (the first science run of Initial LIGO) data. These results show a great improvement from the previous upper limits due to the usage of large scale interferometers. As can be seen, the first searches were all targeted towards known pulsars due to their lower computational cost and more straightforward data analysis techniques. The first published all-sky search using more than one day of LIGO data was reported in [190], where the best upper limits were around 4×10^{-23} at 200 Hz using S2 data. A more detailed summary of CW searches up to S4 can be found in [208], and a public webpage listing all the published CW searches from the LVC can be found in [223]. A summary of results of O1 searches is given in [224].

The difference in sensitivity between using a coherent or semi-coherent method can be clearly seen when results from a targeted search and an all-sky search are compared, as shown in figure 3.3, where the upper limits on the gravitational-wave amplitude h_0 for the O2 observing run are presented. The difference in sensitivity can be more than an order of magnitude. This figure also conveys the huge improvement in sensitivity between the first searches with LIGO data and the current level, of more than two orders of magnitude, both due to improvements in the detectors and also to the data analysis pipelines.

Due to the high computational costs of directed and all-sky searches, and since their sensitivity is limited by the available computational power, an important area of research and improvement is related to how to optimally distribute the available computational resources between the different regions of parameter space, and how to select the best parameters defining a search, such as the maximum allowed mismatches for each dimension or the amount and length of the segments. If a sensitivity function depending on all these factors can be defined, and if the computational cost of the search (which is also




Figure 3.3: Comparison of the O2 upper limits between the targeted search [110] and the best results from the all-sky search [213].

a function of these parameters) can be analytically predicted, an optimization procedure like the one provided by Lagrange multipliers can be used to find the optimal solution. Attempts have been made for how to best select the input data for searches [227, 228, 232], and a more complex optimization procedure dealing with how to select the optimal parameter space regions and mismatch parameters has also been proposed, both for all-sky searches [156, 157] and for directed searches [225, 226].

None of the CW searches done until now has detected a gravitational wave signal. Even though most of the directed and all-sky searches have many candidates crossing the thresholds at several standard deviations, this is not enough to confidently claim a detection, since non-Gaussian features in the noise are known to be able to generate these candidates. In order to increase the confidence of candidates truly being generated from an astrophysical signal, several tests or vetoes are routinely used (many of them summarized in [228]):

- Lines veto: if the frequency-time pattern of the candidate intersects any of the frequencies of known artifacts (discussed and listed in chapter 7), the candidate can be dismissed as there is a high probability that it is generated by the detector.
- Consistency veto: when a search over multiple datasets is done, the significance values have to be consistent between these datasets. For example, for two datasets from two detectors with similar amount of data and noise floors, the significance should be at the same level. Enormous differences between significances are normally caused by artifacts present in only one of the detectors or during only a specific epoch.
- Permanence veto: the coherent detection statistics for the multiple segments can be inspected, and it can be checked if the final value is accumulated in a smooth way, or if a great portion of it is accumulated during a short time. The former is true for a real CW signal, while the latter can be generated by a transient noise artifact present in one of the detectors. Highly related to these veto is the χ^2 veto, which quantifies how the detection statistic of a signal should be distributed within each of these segments, and compares the obtained value with the theoretical prediction.
- Coincidences [213]: if the search has been separated in multiple datasets (for example multiple detectors or datasets from different epochs), a candidate must be present in the output from all datasets, since the opposite case hints at the candidate not being truly astrophysical. Usually, a

range around the parameters of the candidate is selected, and this range is searched for candidates in the complementary datasets. Searching for coincidences can be related to the consistency and permanence vetoes.

- Population veto: it is known that for a true astrophysical signal not only a single candidate will attain a high value of the detection statistic: a neighbourhood of high significance values around the true signal parameters will be present. If this is not the case, the candidate can be dismissed as not being astrophysical.
- “Doppler Modulation off” veto [229]: the obtained detection statistic value can be compared with the result of a new search done with the Doppler modulations produced by the Earth assumed to be zero (i.e. as if the Earth was located at the Solar System Barycenter and not rotating). If the significance of the detection statistic value increases, it is more probable that the candidate is generated from an artifact of the detectors than from a real astrophysical signal.
- Follow-up [144, 253]: if all previous tests are passed, the coherent time can be increased and/or the mismatch parameters can be decreased, in order to increase the statistical confidence on the signal. In order to keep a low computational cost, the amount of refinement is limited by the number of outliers to be followed-up. This fact governs the number of follow-up stages than have to be employed to reach a final fully coherent stage.

If a candidate passes all the previous tests, we can still try to increase the confidence (both statistical and subjective) in it. One of the first steps would be to check all the auxiliary channels. Although the lines veto has been passed, some noise artifacts may not be present in the lines/combs list used to veto outliers, because that list is generated from short Fourier transforms averaged over all an observing run, and a signal with these phase modulations would not appear above the noise and would not be noted. If a signal with the same parameters and modulations is found in any auxiliary channel (which are not sensitive to gravitational waves) such as a magnetometer, the confidence on the signal being astrophysical would be much lower. Due to the continuous nature of CW signals, another strategy would be to wait for more data to be obtained and analysed, to further increase the SNR and to narrow the uncertainty on the signal parameters. A further check that the candidate is not being generated by a software bug could be done by repeating the last coherent stage with a different pipeline. For example, the analysis could be done both by a search that calculates the coherent \mathcal{F} -statistic and by a search that calculates the Bayesian Bayes factor with a stochastic algorithm such as nested sampling. Furthermore, for an outlier coming from an all-sky search, if not all available detectors have been used in the last coherent follow-up stage (possibly due to different sensitivities), these additional detectors can be used since at this stage they might increment the SNR. A similar procedure would be to use data from previous observing runs. Moreover, if the signal has been found only at a single frequency, the same signal at other frequencies (given by the models explained in subsection 2.4.2) could be searched. A final check would be to try to find an electromagnetic counterpart in the radio, X-ray, or gamma-ray bands, situated at the same sky position of the detected signal.

3.2.1 | Upper limits and estimation of sensitivity

When a search pipeline is designed, modified, or improved, its sensitivity needs to be estimated in order to compare it to other pipelines or to other possible configurations. The sensitivity is defined as the smallest gravitational-wave amplitude that an analysis can detect with a certain statistical confidence (i.e. it will detect at least a certain percentage of signals over a population in one observation, or it will detect a signal over an ensemble of observations with the same statistical properties) at a given fixed false alarm probability.

Although the sensitivity is a useful quantity, it only takes into account how deep into the noise-floor of a dataset a pipeline can detect a signal. A more complete description would take into account other important aspects such as the parameter estimation accuracy, the size of the parameter space to be analysed, the computational cost of the pipeline per unit of volume, or its robustness against deviations from the model. Mock Data Challenges (MDCs) are designed to provide a controlled environment where different properties (such as sensitivity, parameter estimation accuracy, robustness, or computational cost) of many different pipelines can be compared when analysing the same parameter space. In these MDCs, many signals (whose parameters can be shared with the participants or can be kept unknown) are simulated and added to real or simulated noise with diverse statistical properties. Examples of Mock Data Challenges applied to CW searches are [230] or [231].

In order to obtain an analytical estimation of the sensitivity of semi-coherent pipelines, many approximations need to be taken, such as assuming that the number of segments is large, that the noise is Gaussian, or that the astrophysical signal is small. An example of such a sensitivity estimation is developed in subsection 3.3.3. These analytical estimations can serve as a first order comparison between similar pipelines, but they usually overestimate the absolute sensitivity of the pipeline. Besides being used to compare different pipelines, an analytical equation that relates the sensitivity of a pipeline to its different tuneable parameters (such as mismatch or duration of segments) can be used to optimize the sensitivity, as was previously discussed.

To derive the analytical sensitivity, the statistical properties of the detection statistic have to be calculated. Firstly, the probability distribution function of the detection statistic needs to be known. Many detection statistics follow a non-central χ^2 distribution of a given number of degrees of freedom (such as the \mathcal{F} -statistic), while others may follow a Gaussian or binomial distribution. Secondly, the false alarm probability is calculated: for a given desired false alarm probability, what is the threshold that is needed? The false alarm probability equation is inverted, and the threshold can be calculated. The last step consists of calculating the detection probability for the given threshold, which includes an integration of a function that depends on the non-centrality parameter. Since this parameter depends on many parameters of the astrophysical signal, an integration over these parameters needs to be done too, in order to obtain an average detection probability over the complete parameter space.

The optimal SNR can be defined by using the inner product, given by equation (3.16), and assuming a triaxial aligned model:

$$\begin{aligned} \text{SNR}^2 &= (h(t)|h(t)) = \frac{2}{S_n(f)} \int_{-T_0/2}^{T_0/2} h(t)^2 dt \\ &= \frac{h_0^2}{S_n(f)} \left[\frac{1}{4} (1 + \cos^2 \iota)^2 \int_{-T_0/2}^{T_0/2} F_+^2 dt + \cos^2 \iota \int_{-T_0/2}^{T_0/2} F_\times^2 dt \right], \end{aligned} \quad (3.35)$$

where the term proportional to the product $F_+ F_\times$ has been dropped out since its integral over time equals to zero. For coherent integration times longer than a day, as shown in [50], the squared SNR depends linearly on the integration time, whereas for shorter times the expression has an extra term with a different oscillatory dependence on T_c .

For a semi-coherent search, the total SNR^2 is given by the sum of the SNR^2 of each segment. If the coherent times of each segment are much less than the sidereal periodicity of the antenna patterns (such as 1800 s), the expressions F_+ and F_\times can be taken as constant during each segment:

$$\begin{aligned} \text{SNR}_S^2 &= h_0^2 T_c \sum_{J=1}^N \left[\frac{1}{4} (1 + \cos^2 \iota)^2 \frac{F_{+;J}^2}{S_{n;J}(f)} + \cos^2 \iota \frac{F_{\times;J}^2}{S_{n;J}} \right] \\ &= h_0^2 T_c \left[\frac{1}{4} (1 + \cos^2 \iota)^2 \sum_{J=1}^N \frac{F_{+;J}^2}{S_{n;J}} + \cos^2 \iota \sum_{J=1}^N \frac{F_{\times;J}^2}{S_{n;J}} \right], \end{aligned} \quad (3.36)$$

where the index J can belong to multiple detectors (we have assumed uncorrelated noise between detectors). A well-known result regarding semi-coherent searches is that their sensitivity is proportional to

$N^{-0.25}$ (in the limit of large N and false alarm probability), but this scaling may be incorrect for certain setups, as discussed in [233].

A quantity that is currently widely used to measure the sensitivity of a pipeline is the sensitivity depth \mathcal{D} , firstly proposed in [228]:

$$\mathcal{D} = \frac{\sqrt{S_n}}{h_0}, \quad (3.37)$$

where in this case S_n is the harmonic mean over all segments used in the search. The advantage of this quantity is that it is independent of frequency, since the frequency dependence from S_n is removed, i.e. it is a measure of sensitivity independent of the noise floor. In [207], a comparison of the sensitivity between different searches in terms of their sensitivity depth is presented.

When no detections are reported in a search, the common procedure is to present upper limits on the gravitational wave amplitude with a certain confidence probability. An upper limit is a probabilistic bound over a certain parameter. The strict definition of what an upper limit represents is different within the frequentist and Bayesian frameworks, due to the different meaning of frequentist's confidence intervals and Bayesian's credible regions, as discussed in subsection 3.1.3, and explained in [234]. The Bayesian upper limit h_0^η corresponds to the interval $[0, h_0^\eta]$ that contains the true value of h_0 with probability η , i.e. it is the point of the distribution where the credible region reaches a total of η of its area. On the other side, the frequentist upper limit is defined as the weakest signal amplitude h_0^η that can be detected (meaning that it is above a threshold d_{th} on a statistic $d(x)$ selected for a given false alarm probability) with a given detection probability η (typically chosen as 90% or 95%).

In order to calculate these upper limits, the analytical methods that are used to estimate the sensitivity are not precise enough, since they assume Gaussian noise (besides other approximations), which is not the true case for a real search. This is empirically known from previous searches, where the analytical sensitivity and the obtained upper limits are very different. Usually, frequentist upper limits are estimated through the usage of a massive set of Monte-Carlo injections (a procedure called population-averaged upper limits), where signals taken from an underlying distribution (usually isotropic in position and orientation, while uniform in the other parameters) are added to the data from the observing run, and for each of these signals the search is repeated with the same configuration as the real search.

To find the upper limits at a particular detection confidence η , signals are added with varying h_0 values, and the h_0 point where a fraction η (such as 0.9 or 0.95) of the injections is recovered is selected as the upper limit value. Usually, this point is not within the h_0 values that have been used, and an interpolation procedure has to be used. A common function to use for this fitting is a sigmoidal. This is related to the fact that many analytical estimations find that the sensitivity depends on the false dismissal through an inverse complementary error function (which has a sigmoidal shape), as the example in subsection 3.3.3 shows.

Since running again the full search over a great number of injections has a non-negligible computational cost, a common procedure is that injections are only added at a reduced number of frequencies, and then a mean value obtained from these frequencies is generalized to the full frequency range. This can easily be done with the sensitivity depth, which are obtained (with a certain confidence value) at these frequency bands, and the h_0 upper limits for the full frequency range are easily calculated by using the amplitude spectral density in equation (3.37). This is the procedure that we use to calculate upper limits in chapter 4.

Bayesian upper limits do not need to perform large scale MCMC simulations, since they can just marginalize the full (usually over the 4 amplitude parameters) posterior distribution and obtain a 1-dimensional posterior distribution over h_0 . From there, obtaining the value with η confidence is trivial. Besides the Bayesian upper limits, there are other methods that do not rely on adding simulated signals to the data, and instead compute frequentist upper limits, such as the Feldman-Cousins method [235] or the Universal Statistics procedure [236]. The advantage of these methods is that they do not require

adding simulated signals, thus saving up the required computational cost and researcher's time. Another advantage of the Universal Statistics procedure is that it can produce upper limits even in the most contaminated frequency bands.

3.2.2 | Parameters measurable with CWs

In this subsection we detail the astrophysical information that can be obtained after the detection of a CW signal. This information depends on the type of CW signal that is detected, and on if an electromagnetic counterpart is also detected. Detecting a CW from a NS in a binary system gives the possibility not only of measuring the five Keplerian parameters, but also of trying to measure post-Keplerian parameters and using that signal to test general relativity, as has been done with electromagnetic signals from pulsars in binary systems (such as the Hulse-Taylor pulsar).

Furthermore, the information to be learned also depends on the geometric configuration of the NS, like its wobble angle θ or if the star is precessing. For the simplest signal model, eight different parameters can be measured: four amplitude parameters and four phase parameters. The amplitude parameter h_0 depends on the combination $I_{zz}\epsilon f^2/d \propto Q_{22}f^2/d$, where f will have been measured. If no EM counterpart is detected, then a disentangled measurement of the distance and the quadrupole moment is not possible, but if an EM measurement of the neutron star is achieved, then its distance could be estimated, and the quadrupole Q_{22} would be known. This would be possible even without an EM detection, if the sky position of the found CW signal lies within a known globular cluster, since then the distance to the NS could be assumed to be the distance to the globular cluster.

If from the EM detection an independent measurement of the moment of inertia can be done, then the ellipticity ϵ could be directly measured, and from its value different mountain building mechanisms could be compared. For example, by knowing the frequency, the first frequency derivative, and the ellipticity, equation (2.5) can be used to infer values of the product $R^6 B_s^2$. If the radius has also been measured with EM information, the value of the external magnetic field could be estimated. This approximate relation can also be seen from equation (2.27), which estimates the ellipticity by comparing the magnetic and gravitational energy (by substituting the mass with $M \propto I/R^2$). It is clear that if an EM detection is not done, the amount of parameters that can be estimated is reduced.

When the signal is detected with multiple frequencies, geometric quantities like the wobble angle can be obtained. If the signal is detected at a single frequency, these geometric quantities can be constrained by comparing the SNR at the different frequencies. Assuming that the SNR is given by $\text{SNR} = h_0\sqrt{T}(a_+F_+ + a_\times F_\times)/\sqrt{S_n}$, we can compare the optimal SNR at two different frequencies:

$$\frac{\text{SNR}_{1f}}{\text{SNR}_{2f}} = \sqrt{\frac{S_n(f_2)}{S_n(f_1)} \frac{a_{+;1}F_{+;1} + a_{\times;1}F_{\times;1}}{a_{+;2}F_{+;2} + a_{\times;2}F_{\times;2}}}, \quad (3.38)$$

which averaging over sky positions and polarization angle ψ and using equation (2.62) gives:

$$\left\langle \frac{\text{SNR}_{1f}^2}{\text{SNR}_{2f}^2} \right\rangle_{\alpha, \delta, \psi} = \frac{S_n(f_2)}{S_n(f_1)} \frac{\sin^2 2\theta}{\sin^4 \theta} \frac{\frac{1}{20} \sin^2 \iota \cos^2 \iota + \frac{1}{5} \sin^2 \iota}{\frac{4}{5}(1 + \cos^2 \iota)^2 + \frac{16}{5} \cos^2 \iota} \equiv ABC, \quad (3.39)$$

where the last two terms (named B and C) depend on geometric angles θ and ι and the first term A depends on the ratio of noise background at the two different frequencies. With this equation an estimate can be made about the relative signal-to-noise ratio, or inversely, from SNR measurements at two different frequencies an estimate of the wobble angle can be made (since ι can be estimated from an individual frequency). Figure 3.4 shows the product BC , showing a very similar shape to figure 5 of [237], while figure 3.5 shows the square root of the noise ratio term A . These figures also show that for an all-sky search, if the initial SNR can be increased by an order of magnitude, other frequencies which had 0.1 of that SNR could be detected (signals with parameters $\sqrt{\left\langle \frac{\text{SNR}_{1f}^2}{\text{SNR}_{2f}^2} \right\rangle} > 0.1$).

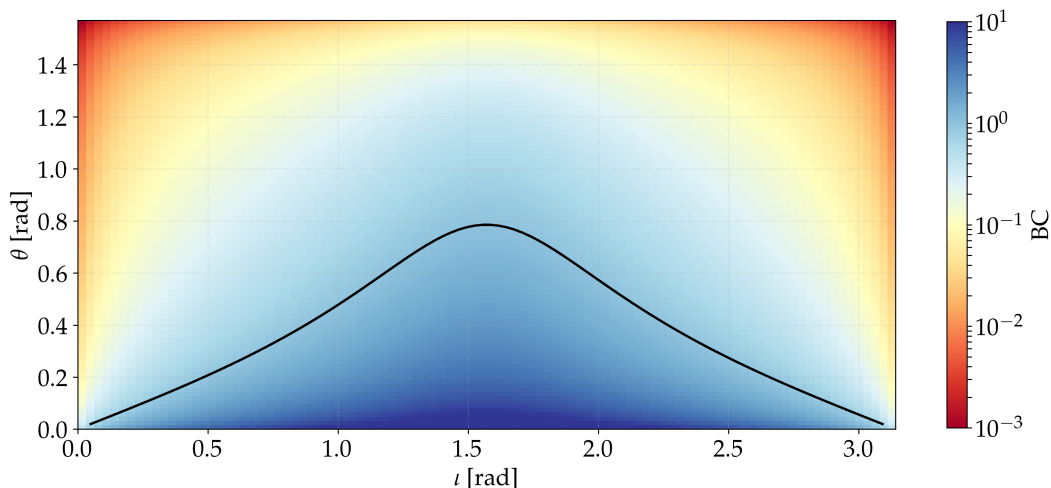


Figure 3.4: The colormap shows the product BC given by equation (3.39). The black line marks the limit between regions where SNR_{1f}^2 is higher (below the line) or lower than SNR_{2f}^2 (assuming equal power spectral densities).

When multiple frequencies are detected, a more complex (and realistic) distribution of unequal moments of inertia could be estimated, following equation (2.52). Furthermore, the method described in [84] explains that when detecting two different frequencies, the Shapiro delay due to the mass of the NS is different for each of them, and an estimation of the mass of the NS can be done, with accuracy depending on the SNR of the signal. Even if this measurement is not possible when the signal is first detected, due to the continuous nature of CW signals, more observation time can be accumulated and the accuracy of these types of measurements can be improved, although the accuracy in the long term is limited by calibration systematic errors discussed in [238].

Besides parameters of individual neutron stars, when multiple detections exist global parameters of the galactic NS population could be estimated as proposed in [239], as now is routinely done for the binary black hole population [240]. Many properties could be investigated with more detail after a detection of multiple sources, such as the distribution of ellipticities, a possible dependency of the orientation vector with proper motion direction, the sky distribution of galactic neutron stars, or the unknown distribution of rotational frequencies.

Continuous waves can also be used to test theories of gravitation. For example, [241] estimates how accurately a possible difference between the speed of gravitational waves and the speed of light could be measured. Furthermore, [242] proposes a method to measure polarizations that are not predicted by general relativity, such as vector and scalar polarizations.

Finally, it can also be mentioned that even not detecting CWs can help to constrain astrophysical parameters of the galactic population such as the ellipticity. For all-sky searches, the upper limits on the gravitational-wave amplitude can be converted to upper limits on the ellipticity at a certain distance, as shown in the results sections of chapters 4 and 6. Thus, we can claim that up to a certain distance (assuming the canonical moment of inertia and a certain detection confidence, usually 95%) there are no neutron stars with an ellipticity higher than a certain value. This constraint can also be used to constrain the internal magnetic fields of these unseen neutron stars, by assuming a certain magnetic field model.

It is important to notice (as discussed in [103]) that the ellipticity upper limits obtained by an all-sky search do not constrain the maximum ellipticity ϵ_{max} of the galactic NS population, which is a common quantity to all neutron stars and is a function of the properties of dense matter. What is really constrained are the current ϵ values of the population, which may not be at their maximum possible values.




Figure 3.5: Square root of the ratio between power spectral densities at frequencies $2f$ and f for the H1 detector during O2.

3.3 | SkyHough

In this section we review the *SkyHough* (SH) method used up to and including O1, and leave the explanation of new added improvements to the respective results chapters.

The Hough transform is a pattern recognition algorithm that was originally invented to analyse bubble chamber pictures from CERN [243]. It was later patented by IBM [244], and it has found many applications in the analysis of digital images. Although initially it was only used to recognise lines, it was later generalised (named “Generalized Hough Transform”) to more complex patterns such as circles and ellipses [245].

The main idea behind the Hough transform is to try different combinations of the parameters describing the shape that is being searched and calculate which one fits better the input data/image. Imagine a discrete input dataset, where only values of 1 or 0 are allowed (such as an image with only black and white pixels). If we want to find a single line in this image, different combinations of slope and initial value of the vertical axis are tried (also assumed to be suitably digitized), and the combination that collects or touches more of the black pixels is selected as the most probable combination of parameters describing the line. For our CW case, we substitute the parameters describing a line by parameters describing the frequency-time pattern.

The *SkyHough* pipeline is a semi-coherent method designed to do all-sky searches of continuous gravitational waves from isolated neutron stars, and was initially described in [158]. It has been used in past searches using data from the Initial LIGO S2, S4 and S5 and Advanced LIGO O1 observing runs [190, 191, 195, 201, 203], and it is the pipeline that we have used (after making some changes and improvements) to obtain the results of this thesis. The Hough pipeline substitutes the frequency bin powers to 1s and 0s after applying a threshold, and performs a weighted sum of these 1s and 0s along the frequency-time pattern of the signal, which is different for every different template. The main advantage of this pipeline is its low (relative to other all-sky CW pipelines) computational cost. This is achieved by reusing the same Doppler modulation for several frequency bins, which makes the *SkyHough* pipeline the cheapest all-sky semi-coherent method (more details in subsection 3.3.2). This pipeline is also more robust against transient spectral disturbances because no matter how large a spectral disturbance is in a single SFT, it will contribute at the most with a 1 (instead to the full power) to the final semi-coherent

detection statistic. The code for the main part of the search is called *lalapps_DriveHoughMulti* and is part of the publicly available LALSuite package [51]. The Hough method of substituting coherent results by 1s and 0s has also been used for an Einstein@Home search with S5 data, where in this case the threshold is applied to \mathcal{F} -statistic values instead of SFT powers [194].

3.3.1 | Input data

SkyHough starts by splitting the data from an observation run into smaller chunks of duration T_c (a quantity that might depend on frequency) and then produces Short Fourier Transforms (SFTs) from the calibrated and windowed $x(t)$ data produced by the different ground-based detectors such as H1, L1 or V1.

The sensitivity of a semi-coherent method increases with the coherent time T_c , so in principle one should aim to use coherent times as long as possible. On the other side, the computational cost increases with longer coherent times, setting a limit to this value. Furthermore, the spread of power to neighbouring bins (as discussed in the beginning of this chapter) limits the maximum SFT time baseline that can be used (which for our method is equal to the coherent time). To calculate the maximum possible coherent time, we demand that the signal is contained in half a single frequency bin:

$$\frac{\Delta f}{2} = \frac{1}{2T_c} \geq \dot{f}T_c \rightarrow T_c \leq \frac{1}{\sqrt{2|\dot{f}|_{max}}}. \quad (3.40)$$

We can estimate the maximum frequency derivative through the frequency evolution model from equation (2.75) (the binary case will be treated in chapter 5):

$$\dot{f} = f_0 \frac{\vec{a} \cdot \hat{n}}{c} + f_1, \quad (3.41)$$

where \vec{a} is the acceleration vector of the detector in the SSB frame. The highest contribution to the acceleration due to detector motion comes from the daily rotation of the Earth [158], which simplifies the previous equation to:

$$|\dot{f}|_{max} = \frac{f_0 v^2}{c R_e} + f_1 = \frac{f_0 4\pi^2 R_e}{c T_e^2} + f_1, \quad (3.42)$$

where R_e is the radius of the Earth and T_e is a sidereal day. Introducing this last equation in (5.17), we have:

$$T_c < \frac{1}{\sqrt{2f_1 + 2.246 \times 10^{-10} f_0}}. \quad (3.43)$$

It can be seen that the maximum coherent time depends on the frequency f_0 and on the spin-down/up f_1 (the spin-down/up term will only be significant when compared to frequencies less than ~ 100 Hz). The optimal search strategy should use SFTs with different coherent times (the maximum allowed, following equation (3.43)) in different regions of the parameter space. In realistic set-ups, only a small amount of different coherent times are selected, mostly to ease the later calculation of upper limits (which will need different characterization for each different coherent time).

The data x_i obtained by the detectors and its DFT \tilde{x}_k is described at the beginning of this chapter. The dataset to be analysed is split in a number N_{SFTs} of different chunks, which are used to produce a spectrogram (a matrix of frequency-time bin powers). These powers are afterwards normalized by the estimated background noise $\langle n_{Jk} \rangle$:

$$\rho_{Jk} = \frac{|\tilde{x}_{Jk}^2|}{\langle n_{Jk} \rangle^2}, \quad (3.44)$$

where J is the SFT index, k is the frequency bin index, and $\langle n_{Jk} \rangle^2$ is calculated with a running median over a number of frequency bins (usually 101 bins are used). The estimated noise is related to the PSD $S_{n;J}$ by:

$$\langle n_{Jk} \rangle^2 \approx \frac{T_c}{2} S_{n;J}(f_k). \quad (3.45)$$

Other pipelines estimate the background noise $\langle n_{Jk} \rangle^2$ in a different way. For example, the Frequency-Hough pipeline estimates it using an auto-regressive algorithm.

The spectrogram is replaced by 1s and 0s by defining a power threshold ρ_{th} (if the power in a bin is above this threshold, it is substituted by a 1, otherwise it is substituted by a 0). The FrequencyHough pipeline also demands that the peak is a local peak, meaning that 1s are only placed in frequency bins with more power than their two neighbours, although in [136] it is shown that this is not an optimal procedure in terms of sensitivity, although the computational cost savings achieved by this requirement should also be taken into account. In [158] it was found that for a signal embedded in Gaussian noise (without taking into account the weights) the optimal choice for this threshold is $\rho_{th} = 1.6$. These 1s and 0s are multiplied by per-SFT weights (calculated at the mid-time of each SFT), given by:

$$w_J \propto \frac{a_J^2 + b_J^2}{S_{n;J}}, \quad (3.46)$$

where $a(t)$ and $b(t)$ are explained in subsection 1.3.3 and given by equation (1.79). This weighting scheme aims to give more importance to times when the data has lower noise and when the detectors are optimally oriented to the specific sky position being searched. These weights were derived in [246].

3.3.2 | Partial Hough map derivatives and look-up table approach

The *SkyHough* pipeline calculates the so-called ‘‘Partial Hough Map Derivatives’’ (PHMDs) at each timestamp and frequency bin [158]. These structures contain the weighted 1s and 0s and are calculated by using the fact that at a given time, a circle of sky positions produces the same Doppler modulation, given by:

$$\cos \theta = \frac{\vec{v}(t) \cdot \hat{n}}{v(t)} = \frac{c}{v} \frac{f(t) - \hat{f}(t)}{\hat{f}(t)}, \quad (3.47)$$

where θ is the angle between \vec{v} (the velocity vector of a detector) and \hat{n} (the position of the source on the sky), and \hat{f} is given by equation 2.5 of [158].

Due to the limited resolution in frequency, the group of sky positions producing the same modulation is given by an annulus (centered on the velocity vector of the given detector) of a certain width $\Delta\theta$ instead of a circle:

$$\cos \Delta\theta = \frac{c}{v} \frac{n \delta f}{\hat{f}(t)} = \frac{n}{n_0}, \quad (3.48)$$

where δf is the width of a frequency bin and n is a number which indexes the frequency modulation, from 0 to a maximum of $n_0 = fv/(c\delta f)$. At a given time and for an observed frequency, all the sky is covered by a finite amount n_0 of annuli produced by the Doppler modulation.

Each PHMD contains all the possible annuli for a certain sky-patch (set of sky positions) compatible with a given time and observed frequency. In other words, one sky position will produce different modulations at different timestamps, and for this reason it will be mapped to different frequency bins at different timestamps.

After calculating all the PHMDs, as pictured in figure 3.6, the pipeline adds one PHMD for each timestamp by following the frequency path created by the source frequency variation (given by $f_0 + f_1 t$), which is the change produced by the spin-down/up of the source. This produces a final Hough map for each combination of f_0 , f_1 and sky-patch.




Figure 3.6: An example showing the PHMD scheme for the SkyHough pipelines. The dotted line shows the type of tracks which the BinarySkyHough pipeline (developed in chapter 5) uses to combine PHMDs at different timestamps, while the dashed lines show some tracks given by different spin-down/up parameters, which are used by the SkyHough pipeline.

The usage of the PHMDs greatly reduces the computational cost of the pipeline since many sky positions are analysed at once (the ones in the sky-patch), because only the borders of the selected annulus need to be tracked.

Furthermore, *SkyHough* uses the look-up table (LUT) approach. This method reuses the same Doppler modulation for contiguous searched frequency bins. This means that it calculates the PHMDs once instead of calculating them for each searched frequency. The same frequency which appears in the denominator of equation (3.47) is used for a number of searched frequencies (up to a maximum, given by equation 4.24 of [158]), changing only the starting frequency f_0 to calculate the source frequency variation. The LUT approach produces computational savings in exchange of searching for modulations which are not exact, which lowers the sensitivity of the search (a quantitative estimation of the effect of the LUT in sensitivity has not been published).

3.3.3 | Detection statistics and sensitivity

The analysis of the full parameter space is split in smaller frequency bands and in smaller regions of the sky called sky-patches (whose size may depend on frequency). The resolution for the different dimensions is derived in [158]:

$$\text{Frequency resolution: } df_0 = \frac{1}{T_c} \quad (3.49)$$

$$\text{Frequency derivative resolution: } df_1 = \frac{1}{T_c T_{obs}} \quad (3.50)$$

$$\text{Sky resolution: } d\theta = \frac{1}{T_c f_0 P_F \frac{v}{c}} \quad (3.51)$$

where T_c is the coherent time of the SFTs, T_{obs} is the observation time (i.e. from the highest analysed epoch minus the lowest analysed epoch, not the true time duration of the dataset, which is $N_{SFTs} T_c$), P_F is the pixel-factor, v is the speed of the detector and c is the speed of light in vacuum. For simplicity of the template-bank construction a hyper-cubic lattice is used to place the templates, although it is known that these lattices are not the optimal ones.

The frequency resolution comes from the resolution of the SFT, but a higher resolution could be

used in order to decrease the mismatch, since when the frequency-time pattern produced by $f_0 + f_1 t$ is used to locate the pertaining PHMDs, a higher amount of f_0 values would produce different tracks. The same happens with f_1 : although its resolution is obtained by demanding that at least there is a shift of one frequency bin when the track reaches the last SFT (in order to distinguish the track from the zero spin-down/up track), a higher amount of f_1 values would contribute to lower the mismatch, so the best resolution achievable is not given by (3.50).

For each of these regions a toplist is produced, which contains information about the most significant candidates in that region. The detection statistic that ranks the searched templates is the number count significance, given by:

$$s_H = \frac{n - \langle n \rangle}{\sigma_H}, \quad (3.52)$$

where $\langle n \rangle$ and σ_H are the expected mean and standard deviation of the Hough number count n (the weighted sum of 1s and 0s summed along the frequency-time pattern of the signal) when only noise is present:

$$\langle n \rangle = e^{-\rho_{th}} \sum_{J=1}^{N_{SFTs}} w_J \quad (3.53)$$

$$\sigma_H = e^{-\rho_{th}} (1 - e^{-\rho_{th}}) \sum_{J=1}^{N_{SFTs}} w_J^2. \quad (3.54)$$

This detection statistic is used to rank all the searched templates and put them ordered by significance in a toplist. Afterwards, only the top candidates (a very small percentage in order to deal with computational costs) are passed to the post-processing and follow-up stages. Since we have assumed Gaussian stationary noise, these significance values cannot be understood as a faithful account of significance, only as a ranking statistic. In order to really believe these significance values, we should obtain the noise distribution from the data itself and not from an a priori analytical Gaussian distribution.

With these ingredients, we can try to estimate the sensitivity (defined to be the smallest h_0 that would cross the threshold n_{th} for a given false alarm and false dismissal probabilities) of this basic setup (i.e. without taking into account posterior vetoes, post-processing, or follow-up, and without taking mismatches into account). In order to estimate the sensitivity of the search, we follow the procedure outlined in subsection 3.2.1: relate the SNR to the false dismissal probability and invert this equation to obtain the dependence of h_0 on the false alarm and dismissal probabilities.

First of all, we will calculate the false alarm and false dismissal. The power defined in equation (3.44) can be seen as the sum of two random variables (assuming stationary and Gaussian noise) [158]:

$$2\rho_{Jk} = z_1^2 + z_2^2, \quad (3.55)$$

where the two z 's are given by:

$$z_1 = \frac{\sqrt{2}\Re[\tilde{x}_{Jk}]}{\sqrt{\langle |\tilde{n}_{Jk}|^2 \rangle}} \quad \text{and} \quad z_2 = \frac{\sqrt{2}\Im[\tilde{x}_{Jk}]}{\sqrt{\langle |\tilde{n}_{Jk}|^2 \rangle}}, \quad (3.56)$$

where \Re denotes the real part and \Im the imaginary part. Therefore the random variables z_1 and z_2 are normally distributed and have unit variance (but nonzero mean). Thus $2\rho_{Jk}$ is distributed according to a non-central χ^2 distribution with 2 degrees of freedom with non-centrality parameter λ_{Jk} (equal to the optimal squared SNR, which changes at different frequency bins k and different SFTs J due to the non-stationarity of the noise and to the amplitude modulations caused by the antenna patterns):

$$\lambda_{Jk} = (E[z_1])^2 + (E[z_2])^2 = \frac{4|\tilde{h}_J(f_k)|^2}{T_c S_{n;J}(f_k)}. \quad (3.57)$$

The distribution of ρ_{Jk} is given by:

$$p(\rho_{Jk}|\lambda_{Jk}) = 2\chi^2(2\rho_{Jk}|2, \lambda_{Jk}) = \exp\left(-\rho_{Jk} - \frac{\lambda_{Jk}}{2}\right) I_0(\sqrt{2\lambda_{Jk}\rho_{Jk}}), \quad (3.58)$$

where I_0 is the modified Bessel function of zeroth order. The mean and variance of this distribution are:

$$E[\rho_{Jk}] = 1 + \frac{\lambda_{Jk}}{2} \quad \text{and} \quad \sigma^2[\rho_{Jk}] = 1 + \lambda_{Jk}. \quad (3.59)$$

When there is no signal, this is just an exponential distribution. *SkyHough* substitutes these powers for 1s and 0s. The false alarm and false dismissal probabilities for the bin selection are:

$$\alpha(\rho_{\text{th}}) = \int_{\rho_{\text{th}}}^{\infty} p(\rho|0)d\rho = e^{-\rho_{\text{th}}} \quad (3.60)$$

$$\beta(\rho_{\text{th}}|\lambda) = \int_0^{\rho_{\text{th}}} p(\rho|\lambda)d\rho \quad (3.61)$$

$$\eta = 1 - \beta \approx \alpha \left\{ 1 + \frac{\rho_{\text{th}}}{2} \lambda + \mathcal{O}(\lambda^2) \right\}, \quad (3.62)$$

where in the last approximation we have assumed a small signal (i.e. $\lambda \ll 1$):

$$\begin{aligned} p(\rho_{Jk}|\lambda_{Jk}) &\approx e^{-\rho_{Jk}} \left(1 - \frac{\lambda_{Jk}}{2} + \frac{\lambda_{Jk}^2}{8} + \mathcal{O}(\lambda^3) \right) \left(1 + \frac{\rho_{Jk}\lambda_{Jk}}{2} + \frac{\rho_{Jk}^2\lambda_{Jk}^2}{16} + \mathcal{O}(\lambda^3) \right) \\ &= e^{-\rho_{Jk}} \left(1 + \frac{\lambda_{Jk}}{2}(\rho_{Jk} - 1) + \frac{\lambda_{Jk}^2}{4} \left(\frac{\rho_{Jk}^2}{4} - \frac{\rho_{Jk}}{1} + \frac{1}{2} \right) + \mathcal{O}(\lambda^3) \right). \end{aligned} \quad (3.63)$$

The false alarm probability of the bin selection is used to obtain the optimal threshold $\rho_{\text{th}} = 1.6$ [158].

The probability of the full dataset of weighted 1s and 0s is given by a binomial distribution with mean and variance [246]:

$$\bar{n} = \sum_{J=1}^{N_{\text{SFTs}}} \eta_J w_J = A\alpha + \frac{\alpha\rho_{\text{th}}}{2} \sum_{J=1}^{N_{\text{SFTs}}} w_J \lambda_J \quad (3.64)$$

$$\sigma_H^2 = \sum_{J=1}^{N_{\text{SFTs}}} w_J^2 \eta_J (1 - \eta_J) = \alpha(1 - \alpha) \|\mathbf{w}\|^2 \left(1 + \frac{\rho_{\text{th}}}{2\|\mathbf{w}\|^2} \frac{1 - 2\alpha}{1 - \alpha} \sum_{J=1}^{N_{\text{SFTs}}} w_J^2 \lambda_J \right) \quad (3.65)$$

where in the second equalities we have applied the small signal approximation, the weights are given by equation (3.46), and

$$A = \sum_{J=1}^{N_{\text{SFTs}}} w_J \quad \text{and} \quad \|\mathbf{w}\|^2 = \sum_{J=1}^{N_{\text{SFTs}}} w_J^2. \quad (3.66)$$

When the number of SFTs is large (i.e. $N_{\text{SFTs}} > 1000$), the binomial distribution function is well approximated by a Gaussian with the same mean and variance:

$$p(n|\rho_{\text{th}}, \lambda) = \frac{1}{\sqrt{2\pi\sigma_H^2}} e^{-(n-\bar{n})^2/2\sigma_H^2}. \quad (3.67)$$

Then, we can calculate the false alarm and false dismissal probabilities of the full Hough pipeline:

$$\alpha_H = \int_{n_{\text{th}}}^{\infty} p(n|\rho_{\text{th}}, 0) dn = \frac{1}{2} \text{erfc} \left(\frac{n_{\text{th}} - A\alpha}{\sqrt{2\alpha(1-\alpha)}\|\mathbf{w}\|^2} \right) \quad (3.68)$$

$$\beta_H = \frac{1}{2} \text{erfc} \left(\frac{\bar{n} - n_{\text{th}}}{\sqrt{2}\sigma_H} \right). \quad (3.69)$$

To derive the previous false dismissal equation, a population of constant SNR has been assumed, which results in an overestimation of the sensitivity as discussed in [233]. From the false alarm we can obtain the number count threshold needed to have a certain false alarm:

$$n_{\text{th}} = \alpha A + \sqrt{2\|\mathbf{w}\|^2\alpha(1-\alpha)} \text{erfc}^{-1}(2\alpha_H), \quad (3.70)$$

which corresponds to a threshold in significance of $s_{\text{th}} = \sqrt{2} \operatorname{erfc}^{-1}(2\alpha_H)$, a similar expression to equation 3.5 of [233]. Notice that α_H is the false alarm probability for a single trial, whereas the total false alarm probability is $\alpha_H \mathcal{N}$ (only true when the different trials are statistically independent), where \mathcal{N} is the number of templates that are analysed.

Expanding β_H in powers of λ_J (which may not be appropriate for small false dismissal values, since the signals needed to achieve such low values and cross the threshold, for realistic false alarm probabilities, will be strong) we obtain this approximation [246]:

$$S = \operatorname{erfc}^{-1}(\alpha_H) + \operatorname{erfc}^{-1}(2\beta_H) = \sqrt{\frac{\alpha \rho_{\text{th}}^2 \sum_{J=1}^{N_{SFTs}} w_J \lambda_J}{8(1-\alpha) \|\mathbf{w}\|}} + \frac{\rho_{\text{th}}}{4} \frac{1-2\alpha}{1-\alpha} \frac{\sum_{J=1}^{N_{SFTs}} w_J \lambda_J}{\|\mathbf{w}\|^2} \operatorname{erfc}^{-1}(2\alpha). \quad (3.71)$$

Usually only the first term is kept since it is proportional to $\sqrt{N_{SFTs}}$, while the second term does not grow with N_{SFTs} when N_{SFTs} is large.

The only thing left to do is to relate the non-centrality parameter to the amplitude h_0 :

$$\lambda_J = \frac{4|\tilde{h}(f_k)|^2}{T_c S_{n;J}} \approx \frac{h_0 T_c}{S_{n;J}} (A_+^2 F_{+;J}^2 + A_x^2 F_{x;J}^2) \left(\frac{\sin[\pi(f-f_k)T_c]}{\pi(f-f_k)T_c} \right)^2. \quad (3.72)$$

We can obtain a population average by using these expressions:

$$\langle (F_{+;J})^2 \rangle_\psi = \langle (F_{x;J})^2 \rangle_\psi = \frac{a_J^2 + b_J^2}{2} \quad (3.73)$$

$$\langle \lambda_J \rangle_{\nu, \psi, \Delta f} = 0.7737 \times \frac{2h_0^2 T_c (a_J^2 + b_J^2)}{5S_{n;J}} \quad (3.74)$$

$$\langle \lambda_J \rangle_{\nu, \psi, \alpha, \delta, \Delta f} = 0.7737 \times \frac{4}{25} \frac{h_0^2 T_c}{S_{n;J}}, \quad (3.75)$$

where the 0.7737 comes from averaging the leakage (i.e. taking a uniform distribution of the signal frequency between $-1/2$ and $1/2$ of the center of the bin). Joining the last equation and equation (3.71) we obtain the final sensitivity to the average population (assuming zero mismatch between template and signal):

$$h_0 = 3.38 \sqrt{\frac{S}{T_c}} \sqrt{\frac{\|\mathbf{w}\|}{\bar{w} \cdot \bar{X}}} \quad (3.76)$$

$$X_J = \frac{a_J^2 + b_J^2}{S_{n;J}}. \quad (3.77)$$

When all weights are unity, we recover the classic result derived in [158]:

$$h_0 = 5.34 \frac{\sqrt{S}}{N_{SFTs}^{1/4}} \sqrt{\frac{S_n}{T_c}} = \frac{8.54}{N_{SFTs}^{1/4}} \sqrt{\frac{S_n}{T_c}} = 8.54 N_{SFTs}^{1/4} \sqrt{\frac{S_n}{T_{\text{obs}}}} = \frac{\sqrt{S_n}}{\mathcal{D}}, \quad (3.78)$$

where \mathcal{D} is the sensitivity depth, and in the second equality we have used a false alarm probability of 0.01 and a false dismissal probability of 0.1. Since this calculation does not take into account the loss in SNR due to mismatch and leakage, and it is assuming stationary Gaussian noise, we can treat it as an upper bound of the real search sensitivity.

Using the previous equation (unweighted case) we can obtain the number of SFTs needed to detect a signal of a certain amplitude with a given false alarm and false dismissal probabilities:

$$N_{SFTs} = \frac{5.34^4 S^2 S_n^2}{h_0^4 T_c^2} = 5319 \frac{\mathcal{D}^4}{T_c^2}, \quad (3.79)$$

where in the last equality we have used a false alarm probability of 0.01 and a false dismissal probability of 0.1.

3.3.4 | Post-processing

As we mentioned before, high significance values do not mean a 100% confidence on a true CW signal. More tests have to be passed in order to eliminate candidates coming from pure noise and increase the confidence, as described in section 3.2.

The final product of the main *SkyHough* search is one toplist per analysed dataset for each region in parameter space (e.g. one toplist for H1 for each 0.1 Hz band), which contains the top templates ordered by significance. These lists can be huge, each of them containing around 10^4 or more candidates (the larger the allowed size of the list, the more sensitive the analysis will be, but the computational cost will increase accordingly).

If more than one dataset has been analysed, such as from two detectors or from different epochs, the first step followed by the *SkyHough* pipeline is to search for coincident candidates between these different datasets. In order to search for coincidental pairs between the toplist, we calculate the distance in parameter space and select the pairs that are closer than a certain threshold. The optimal value for this threshold is not analytically known, since its value depends on a balance between being more sensitive and having too many outliers. A reasonable value can be found by doing simulations. For each coincidental pair, its centroid (average locations in parameter space weighted by significance) is calculated. The distance d in parameter space is given by:

$$d^2 = \left(\frac{\Delta f_0}{\delta f_0}\right)^2 + \left(\frac{\Delta f_1}{\delta f_1}\right)^2 + \left(\frac{\Delta x}{\delta \theta}\right)^2 + \left(\frac{\Delta y}{\delta \theta}\right)^2, \quad (3.80)$$

where the numbers in the numerators represent the difference between two templates and the numbers in the denominators represent the parameter resolution. This distance is dimensionless and is given as a number of bins. The quantities x and y represent the cartesian ecliptical sky coordinates projected into the ecliptical plane.

After searching for coincidences, clusters are found within the list of coincident candidates in order to reduce the amount of candidates. Real astrophysical signals are thought to produce multiple nearby candidates with elevated significance values, and clustering these candidates can help in improving the parameter estimation and also reducing the amount of computational cost that will be needed to use in further follow-up stages. In order to find clusters, we set a threshold in parameter space distance and find candidates which are closer than this distance. Clusters are found by analysing the distance of each template from all other templates, and keeping a list of indices of members with distances below the threshold. Each template can only be part of a cluster, so if a template was already in a cluster its newly generated cluster and the old one will be joined to form a unique cluster. Afterwards, the centroid of each cluster is found. This is calculated as a weighted (by significance) average among all the members of the cluster. *SkyHough* usually keeps the most significant cluster per 0.1 Hz band (if any), selected by the summed significance of all its members. This produces the final list of outliers of the search.

SkyHough tries to find coincidences first and then searches for clusters, although other pipelines may only use one of the two approaches or do them in the inverse order. If only clusters are searched, another distance threshold between the top template in a cluster and all the other members of the clusters is calculated, which eliminates all members that are further away than a certain threshold. This step is not needed after finding coincidences since many candidates are already eliminated, but if this is not the case the clusters can grow too wide and the final parameter estimation might be deteriorated unless a cut is made, due to a high number of cluster members being too distant from the true signal values.

Even after these two post-processing steps, we cannot claim that the found clusters represent a real astrophysical signal, because from previous searches it is known that instrumental noise or spurious coincidences can also end up in the final list. For this reason additional vetoes and follow-up procedures that remove spurious candidates, increase the significance of CW signals, and enhance their parameter estimation are used, as briefly discussed in section 3.2. Some veto procedures routinely used by the *SkyHough* pipeline include the χ^2 veto, the population veto, and the lines veto. *SkyHough* searches prior

to the ones presented in this thesis did not need a follow-up stage, since the previous vetoes already removed all candidates minus the hardware injections. Since we wanted to increase the sensitivity of the search, we allowed more candidates by increasing the false alarm probability. For this reason, not all the candidates were eliminated by the vetoes, and new veto procedures (described in their corresponding results chapter) had to be developed, together with two different follow-up procedures, also described in the results chapters.

Part II

Results

The following chapters describe the different original results that have been obtained during the development of this thesis. Chapter 4 presents results from an all-sky search for CWs from neutron stars in isolated systems using Advanced LIGO O2 data, while chapter 6 uses the same the data for an all-sky search of CWs but from neutron stars in binary systems. Chapter 5 introduces a new method developed to do all-sky searches from neutron stars in binary systems. Lastly, chapter 7 gives an introduction to the characterization and mitigation of long-duration noise artifacts present in the interferometric gravitational-wave detectors.

CHAPTER 4

ALL-SKY SEARCH FOR CONTINUOUS GRAVITATIONAL WAVES FROM ISOLATED NEUTRON STARS USING O2 ADVANCED LIGO DATA

In this chapter we present results of an all-sky search for continuous gravitational waves (CWs) from isolated neutron stars using data from the second observing run of the Advanced LIGO detectors. Three different semi-coherent methods (*FrequencyHough*, *SkyHough*, and *Time-Domain \mathcal{F} -statistic*) are used, although we focus on the description of the search design and results from the *SkyHough* pipeline. *SkyHough* searches a gravitational-wave frequency band from 50 to 1500 Hz and a first frequency derivative from -1×10^{-8} to 1×10^{-9} Hz/s. None of these searches has found clear evidence for a CW signal, so we present upper limits on the gravitational-wave strain amplitude h_0 (the lowest upper limit on h_0 is 1.7×10^{-25} in the 123-124 Hz region) and discuss the astrophysical implications of this result. This is the most sensitive all-sky search ever performed over the broad range of parameters explored in this study. The material in this chapter is adapted from [213].

4.1 | Introduction

Fast-spinning neutron stars in the Milky Way can generate continuous gravitational waves via various processes which produce an asymmetry. Crustal distortions from cooling or from binary accretion, or magnetic field energy buried below the crust could lead to the non-axisymmetry necessary for detectable emission. Recently, some evidence for a limiting minimum ellipticity was discussed in [111]. A comprehensive review of continuous gravitational wave emission mechanisms from neutron stars can be found in [247]. The detection of a CW, possibly combined with electromagnetic observations of the same source, could yield insight into the structure of neutron stars and into the equation of state of matter under extreme conditions, as discussed in the previous chapters.

In this chapter we present the results of an all-sky search of CWs by three different semi-coherent pipelines (*FrequencyHough* [248], *SkyHough* [158], *Time-Domain \mathcal{F} -statistic* [50]) using O2 data from the Advanced LIGO detectors. Each pipeline uses different data analysis methods and covers different regions of the parameter space, although there exists some overlap between them. Overall, we search the whole sky for gravitational wave frequencies from 20 to 1922 Hz (this number was chosen in order to avoid the violin modes of the test masses found at higher frequencies) and a first frequency derivative from -1×10^{-8} to 2×10^{-9} Hz/s (the different regions of parameter space searched by each pipeline are shown in figure 4.9). No detection has been made, and upper limits on the gravitational wave amplitude are presented.

Furthermore, we present a detailed account of how some parameters of the *SkyHough* pipeline were selected, such as the coherent time, and we present several improvements to the pipeline that helped to improve its sensitivity with a negligible increase of the computational cost, such as:

- The coherent time is not the same for all frequencies: we use four different coherent times.
- The size of the toplist (number of candidates per 0.1 Hz band that go to the post-processing stage) is incremented by at least an order of magnitude (this can be done because a new follow-up strategy is used).
- The size of the sky-patch changes with frequency, which produces savings in computational cost.
- Addition of two extra detection statistics, calculated in the second stage of the main search (where just the top 4000 templates per patch are used).
- Usage of an extra second set of complementary SFTs in the second stage.
- New follow-up based on the \mathcal{F} -statistic: the characterization and usage of a sensitive follow-up method is very important to improve the sensitivity of the pipeline. Previously, after the main search only veto methods were used, and for this reason only a small number of outliers were allowed to reach the final stages. With a follow-up method, more outliers are allowed to arrive to the final stage, implying that the false alarm threshold can be lowered, and thus the possibility of detecting smaller signals increases.

The signal model has been described in detail in subsection 2.4.3. We remark again that the frequency-time pattern that we search is:

$$f(t) = \frac{1}{2\pi} \frac{d\phi}{dt} \simeq f_0 + f_0 \frac{\vec{v}(t) \cdot \hat{n}}{c} + f_1 t. \quad (4.1)$$

This assumes that the neutron star is isolated. In case it is part of a binary system, the frequency evolution is complicated by the binary system orbital motion, which introduces an additional frequency modulation. Such modulation, on a signal of frequency f and neglecting the binary system ellipticity, is given by:

$$\Delta f_{\text{orb}} \simeq \frac{2\pi}{P} a_p f, \quad (4.2)$$

where P is the binary orbital period and a_p is the projected orbital semi-major axis (in light-seconds). By imposing that the orbital frequency modulation is contained into a frequency bin $\delta f = 1/T_c$, where T_c is the duration of the data chunks which are incoherently combined in the analysis (see section 4.4), we find that two of the search pipelines (*FrequencyHough* and *SkyHough*) used in this chapter would be fully sensitive to a CW signal from a NS in a binary system if:

$$a_p \ll 0.076 \left(\frac{P_b}{1\text{day}} \right) \left(\frac{f}{100\text{Hz}} \right)^{-1} \left(\frac{T_c}{1800\text{s}} \right)^{-1} \text{ s}. \quad (4.3)$$

For larger orbital frequency modulations the pipelines would start to lose signal-to-noise ratio but a detailed study of this issue is outside the scope of this chapter.

4.2 | Preparing the search

In this section we present some preliminary studies that helped to design the search and the different improvements that were added between the previous O1 search and the current O2 search.

4.2.1 | Number of valid frequency bins for the look-up table approach

The *SkyHough* pipeline uses the so called look-up tables, an approach that saves a lot of computational time but introduces a small error in the annuli that are calculated this way. The look-up tables are a way to re-use the Partial Hough Map Derivatives (PHMD) for several frequency bins at once in the Hough map computation.

The maximum number of valid frequency bins that f_0 can be changed by so that the annuli changes by only a fraction r of the pixel size $d\theta$ is (taken from line 309 of the DriveHoughMulti code available within LALSuite [51]):

$$B = \frac{r}{P_F \frac{v}{c}} \sqrt{\frac{(f_0 T_c \frac{v}{c})^2}{(f_0 T_c \frac{v}{c} - 1)^2} - 1}, \quad (4.4)$$

where P_F is the pixel-factor, and r is the fraction of the pixel size that we allow the annuli to be moved. In the current version of the code, r is hard-coded to be equal to 0.5.

The PHMD structure has a size of `nOfSize`, which is equal to `nSpinDown + nSpinUp + 1`. The frequency bin for which the look-up table is constructed is f_0 , and this is the bin around which the PHMD is constructed. Since we are using more templates with negative frequency derivative than positive, the maximum difference in frequency bins will be for the negative derivatives. The minimum bin of the PHMD is:

$$f_{binmin} = f_0 - nOfSize + nSpinUp + 1 = f_0 - nSpinDown. \quad (4.5)$$

Thus, the maximum difference in frequency bins from the look-up table bin to the PHMD is:

$$\Delta f_{max} = \text{abs}(f_{binmin} - f_0) = nSpinDown. \quad (4.6)$$

This maximum difference has to be smaller than the allowed number of valid frequency bins for the look-up table:

$$nSpinDown \leq \frac{r}{P_F \frac{v}{c}} \sqrt{\frac{(f_0 T_c \frac{v}{c})^2}{(f_0 T_c \frac{v}{c} - 1)^2} - 1}. \quad (4.7)$$

If we know the number of spin-down values, the frequency, the coherent time, and the pixel-factor, we can calculate the fraction r that is needed for this equation to be valid:

$$r > \frac{nSpinDown P_F \frac{v}{c}}{\sqrt{\frac{(f_0 T_c \frac{v}{c})^2}{(f_0 T_c \frac{v}{c} - 1)^2} - 1}}. \quad (4.8)$$

For this O2 search, we use four different coherent times (as explained in subsection 4.4.1). The number of spin-down values is 835, 625, 418 and 209 (for 3600 s, 2700 s, 1800 s and 900 s), the pixel-factor is 2 and v/c is 10^{-4} . We can calculate the maximum r that we need in each of the four different T_c stretches:

$$\begin{aligned} T_c = 3600 \text{ s} \quad \text{and} \quad f = 300 \text{ Hz} &\longrightarrow r > 1.22 \\ T_c = 2700 \text{ s} \quad \text{and} \quad f = 550 \text{ Hz} &\longrightarrow r > 1.07 \\ T_c = 1800 \text{ s} \quad \text{and} \quad f = 1300 \text{ Hz} &\longrightarrow r > 0.90 \\ T_c = 900 \text{ s} \quad \text{and} \quad f = 1500 \text{ Hz} &\longrightarrow r > 0.34. \end{aligned} \quad (4.9)$$

These are the maximum values needed in each stretch. A more optimal design would be to calculate r for each 0.1 Hz frequency band, and use that value, which is a more optimal choice than setting a constant value for the variable `nFreqValid` in the code. The left plot of figure 4.1 shows the minimum value of r that is needed at each frequency. The dashed trace of this figure shows the result for the parameters from the previous O1 search: `nSpinDown` \sim 200 and a constant coherent time of 1800 s. We can see that r never crosses 0.5 (as opposed to the O2 search case).




Figure 4.1: The left plot shows the minimum value of r that is required for the look-up table approach to be valid, computed from equation (4.8). The black, green, yellow, and red lines correspond to the different (incrementing) coherent times, and the blue line marks the path of our choice of coherent time across frequencies. The dashed blue trace shows the result for the O1 parameters, where the number of spin-down values was around 200 and the constant coherent time was 1800 s. The right plot shows the corrected traces when the observing run has been split in two datasets of approximately 90 days each.

A study of the decrease in sensitivity when r is increased is needed if we use this approach. Another possibility to avoid having to increase r would be to decrease the range of spin-down values to be searched, but this is not ideal since we do not want to decrease the range of spin-down values if possible.

Another alternative is to construct two different datasets by joining data from H1 and L1, but using half of the observing time for each detector in each dataset, thus decreasing the total observation time of each dataset and decreasing the needed resolution in the first frequency derivative. This can be seen in the right plot of figure 4.1, where the value of r is always less than 0.5 if we take this approach.

4.2.2 | Size of the sky-patches and RAM formula

As mentioned in the introduction, we want to make the size $d\Theta$ of the sky-patches frequency dependent, since the *SkyHough* code is more efficient when the sky-patches are bigger. When the size of the sky-patches is increased, the Random Access Memory (RAM) used by each job is increased, since the number of templates in each patch is increased. This imposes a maximum sky-patch size. Furthermore, because we are using weights calculated at the center of each sky-patch, the smaller that these sky-patches are the more accurate the weights will be, slightly improving the sensitivity of the search.

If we have an analytical formula for the RAM that depends on the sky-patch size, we can invert that equation and choose the maximum sky-patch size that is allowed at each frequency by a certain available RAM. To estimate the amount of RAM used by a job, we have to know the size in bytes of the structures that occupy the higher amount of memory (we do not have to take into account each individual variable, since compared to the big structures that depend on the number of SFTs or frequency derivative values they will not be important). We want an equation for the memory M that has the following structure:

$$M = (a d\Theta^2 + b d\Theta + c)N_{SFTs} \rightarrow a d\Theta^2 + b d\Theta + K = 0, \quad (4.10)$$

where $K = c - M/N_{SFTs}$, and a , b and c are some constants. This has the classical solution for the quadratic equation:

$$d\Theta = \frac{-b \pm \sqrt{b^2 - 4aK}}{2a}, \quad (4.11)$$

of which we only use the $+$ term, since the other one is negative.

By reading the code, the different structures are called: LoadSFTs, VelV, timeV, timediff, weightsV,

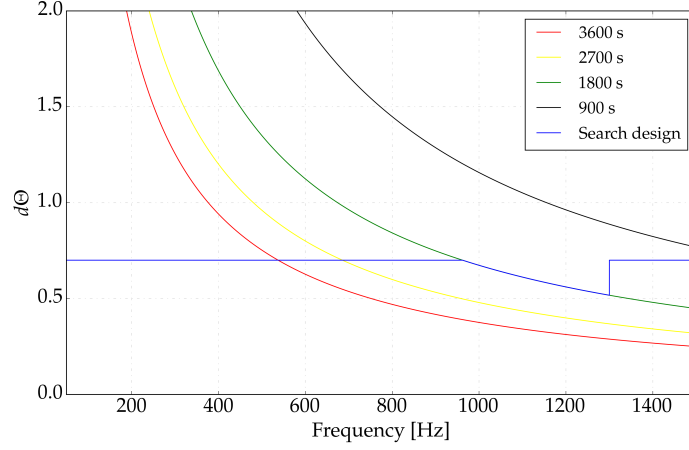


Figure 4.2: Maximum size of sky-patches for 8 GB of memory. The number of SFTs used for 3600 s, 2700 s, 1800 s and 900 s is, respectively: 2544, 3460, 5283, 10801.

pgV, upgV, best, LUT, PHMDVs, patches. Of all these structures, the biggest ones are:

$$\text{best} = 20 + 8S + N_{SFTs}(70 + S + S \text{nPeaks}) \quad (4.12)$$

$$\text{LUT} = 4 + S + N_{SFTs}(42 + 2S + \text{maxBorders}(14 + S + 2 \text{ySide}) + 16 \text{maxBins}) \quad (4.13)$$

$$\text{PHMDVs} = 28 + S + N_{SFTs} \text{nfSize}(17 + 3S + \text{ySide} + 2(S \text{maxBorders} + 14 + S + 2 \text{ySide})), \quad (4.14)$$

where S is the size of the pointers (8 bytes on a 64-bit machine), and:

$$\text{ySide} = \frac{4 \tan(0.25d\Theta)}{1/(\frac{v}{c} P_F T_c f_0)} \approx \frac{v}{c} P_F T_c f_0 d\Theta \quad (4.15)$$

$$\text{xSide} = \text{ySide} \quad (4.16)$$

$$\text{maxBorders} = 1 + 2 \min(1 + 2f_0 T_c \frac{v}{c}, 1.5 \max(\text{xSide}, \text{ySide})/P_F) \approx 3 \frac{v}{c} T_c f_0 d\Theta \quad (4.17)$$

$$\text{maxBins} = 1 + 2f_0 T_c \frac{v}{c} \approx 2f_0 T_c \frac{v}{c} \quad (4.18)$$

$$\text{nPeaks} = 0.2(\Delta f T_c + 2[(f_0 + \Delta f) \frac{v}{c} T_c + 101 + \text{nfSize}]) \approx 0.4(f_0 \frac{v}{c} T_c + \text{nfSize}). \quad (4.19)$$

After some algebraic operations, the constant factors are:

$$\begin{aligned} a &= 6P_F \left(f_0 T_c \frac{v}{c} \right)^2 \\ b &= (5P_F + 6S) \text{nfSize} f_0 T_c \frac{v}{c} \\ c &= (45 + 5.4S) \text{nfSize} + 0.4(80 + S) f_0 T_c \frac{v}{c} \\ K &= c - \frac{M}{N_{SFTs}}. \end{aligned} \quad (4.20)$$

Now, we can solve equation (4.11) with the parameters from the O2 search and a maximum memory of 7.5 GB. We have to be aware that the real memory used by the jobs will usually be higher than this, since the previous equations were a simplification which only took into account the biggest terms. For this reason, we input a memory of 7.5 GB instead of 8 GB (which will be the RAM of the machines that we will use in the search) in the equation. The results are shown in figure 4.2. It can be seen that for low frequencies ($f < 200$ Hz), we could use a sky-patch size bigger than 0.7 rad, but we do not obtain much computational cost savings and the sky weighting becomes less accurate. These results have been obtained by using two datasets separated by different epochs, as explained in the last subsection.

4.2.3 | Extra detection statistics

The previous searches that used the *SkyHough* algorithm used the weighted Hough number count as their detection statistic (a detection statistic is a function which depends on the data, and it is used to sort the templates being searched depending on the probability of containing an astrophysical signal). In this subsection we present two new detection statistics, and test them with injections in O1 data.

The *SkyHough* pipeline can be divided in two stages. At the first step, the pipeline uses the so-called “look-up table approach”. It allows to use the same Doppler modulation pattern for several contiguous frequency bins, speeding up the calculation. At the end of this first stage, an approximated number count significance is calculated for each template. Only N_{cand} are then passed to the second stage of the pipeline, which recalculates the significance (and calculates values for the χ^2 veto) using the correct frequency-time path, and calculates the weights at each exact sky position. More than 90% of the computational cost comes from the first stage, since $N_{Cand}/N_{Total} \ll 1$.

The new idea is that at the second stage two detection statistics (besides the number count significance) are calculated: the weighted power and the combined (from n frequency bins around the central bin) weighted power, which should be more sensitive than the Hough number count. The increase in efficiency is produced because more templates than those that will be used in the post-processing are recalculated, thus reordering the final toplist and producing three different ones (one for each detection statistic) with different templates in each of them. Furthermore, the ability to calculate these three detection statistics incorporates the benefits of all of them at once, such as the robustness of the Hough number count and the higher sensitivity of the power. In order to calculate these new detection statistics, few modifications to the code have to be done, and the new additions do not increment the computational cost in a significative way, since the second stage of the search is much faster.

The weighted power is defined as:

$$P = \sum_{J=1}^{N_{SFTs}} w_J \rho_J, \quad (4.21)$$

where $\rho_{Jk} = \frac{2|\tilde{x}_{Jk}|^2}{T_c S_{n,J}}$ is the whitened power and $w_J \propto \frac{a_J^2 + b_J^2}{S_{n,J}}$ are the weights at each SFT. This detection statistic is very similar to the one used by the PowerFlux pipeline. The final statistic that we report is the significance of the power, defined by:

$$s_P = \frac{P - \bar{P}}{\sigma_P}, \quad (4.22)$$

where $\bar{P} = N_{SFTs}$ is the mean of the Erlang distribution and $\sigma_P = \sqrt{\sum_{J=1}^{N_{SFTs}} w_J^2}$ is its standard deviation.

The weighted power is near optimal when the frequency of the signal is resolved, i.e. if the frequency of the signal coincides with the center of a Fourier frequency bin. If this is not the case, the signal power will be spread into more than one frequency bin. The weighted mixed power is a detection statistic that combines the complex values of $2P$ frequency bins around the central bin, and it was derived in [136]. The weighted mixed power is given by:

$$\tau = x^T M x, \quad (4.23)$$

where x is a vector of complex values of frequency bins, x^T is the complex transpose and M is given by:

$$M_{nm} = (-1)^{n-m} \int_{-\Delta/2}^{\Delta/2} j_0(\pi(n+\delta)) j_0(\pi(m+\delta)) d\delta \quad (4.24)$$

where n and m are indices that run from $-P$ to P , δ is the distance to the frequency of the central bin and j_0 is the spherical Bessel function.

As an example, for $\Delta = 1$ and $P = 1$, M and x are:

$$x = \begin{bmatrix} x_{n-1} \\ x_n \\ x_{n+1} \end{bmatrix}$$

$$M = \begin{bmatrix} 0.0786983 & -0.0588504 & -0.0588504 \\ -0.0588504 & 0.773695 & -0.0588504 \\ -0.0588504 & -0.0588504 & 0.0786983 \end{bmatrix}.$$

As it is shown in [136], this detection statistic can be understood as the average power in a continuous interpolated spectrum within a frequency band of width $\pm\Delta/2$ around the central bin. For this detection statistic we also report its significance, defined in a similar way to equation (4.22).

The sensitivity of the full search is limited by the sensitivity of the first stage. Only N_{cand} are recalculated in the second stage, where the sensitivity of the search can be increased by using the new detection statistics. Since the first stage is an approximation, the best strategy to follow is to select enough N_{cand} to be recalculated, more than the final candidates that we will keep for the post-processing stages. In this way, the loss in sensitivity due to the approximation done in the first stage can be “surpassed”, and we obtain a cheap procedure with nearly the same sensitivity as the pipelines which use the power since the first stage, but with a reduced computational cost. With this strategy, the first stage can be understood as a first look at the parameter space to investigate which portions are interesting enough to be passed to the second stage.

We have tested the new detection statistics with injections in O1 data. We have injected 500 signals at 14 different sensitivity depths, making a total of 7000 injections. The signals have been injected at three different frequency bands, and are isotropically distributed in the sky, uniformly random distributed in spin-down between -10^{-8} and 10^{-9} Hz/s and have uniform random distributions in $\cos \iota$, ϕ_0 and ψ .

The results can be seen in figure 4.3 (we show the results at 872.1 Hz, a frequency band which was used for the O1 upper limits calculation), where the green trace represents the efficiency using the weighted mixed power statistic, the red trace represents the weighted power and the blue trace represents the usual Hough number count significance. The results for the mixed statistic that are shown were obtained using $P = 3$ and $\Delta = 1$, i.e. using the information from three bins at each side of the central bin and allowing the maximum distance in frequency to be half a bin. An increase in efficiency can be seen for the two new statistics, with a greater increase for the weighted mixed power. An improvement over the weighted power statistic was also seen for $P = 1$ and for $P = 2$, but the improvement was greater for $P = 3$.

4.2.4 | Creating additional SFTs

For previous *SkyHough* searches, the detector data was divided in non-overlapping SFTs. Following an idea from [206], we predict that an improvement in sensitivity may come from creating additional SFTs that combine data in a coherent way which was not coherently combined in the original set of SFTs.

The idea consists of combining data using the same coherent time as before, but with SFT starting times separated by an amount T_c/q s, where q is an integer that specifies the number of times in which we will recombine the segments. For example, if $T_c = 1800$ s and $q = 2$, each new extra SFT will begin 900 s after a previous SFT. This is bound to the condition that the new times must be contained in the “science” segments, i.e. times in which the data is not vetoed. For this reason, the new amount of SFTs will usually be less than q times the amount of previous SFTs. We can use any integer q , but with increasing q the computational cost of the search increases, and we note that the sensitivity improvement that can be achieved is bounded. A visualization of this idea can be seen in figure 4.4, where the procedure to obtain the additional SFTs from the initial set is exemplified. We note that this idea is different than overlapping windowed time-series, as for example done with the 50% overlapping Hann




Figure 4.3: The plot shows the efficiency curve at 14 different sensitivity depth points evaluated with different detection statistics. The blue points are from the Hough number count significance, the red points are from the weighted power significance, and the green points are from the weighted mixed power significance.

windowing. As happened with the extra detection statistics, this new addition will not highly increment the total computational cost since the extra SFTs are only used in the second stage of the search.

We have tested this idea by adding injections in the O1 data (the same injections that were used in the previous subsection). The results for $q = 2$ can be seen in figure 4.5. We observe an improvement in efficiency over just using the original SFTs. For example, at sensitivity depth 51 $\text{Hz}^{-1/2}$ the difference in efficiency between the original number count significance and the weighted mixed power significance with additional SFTs is $0.25558 - 0.13360 = 0.12198$.

4.3 | Advanced LIGO O2 observing run

The LIGO second observing run (called O2) started on November 30 2016 and finished on August 25 2017. During this time, three different gravitational-wave detectors of second generation were active and

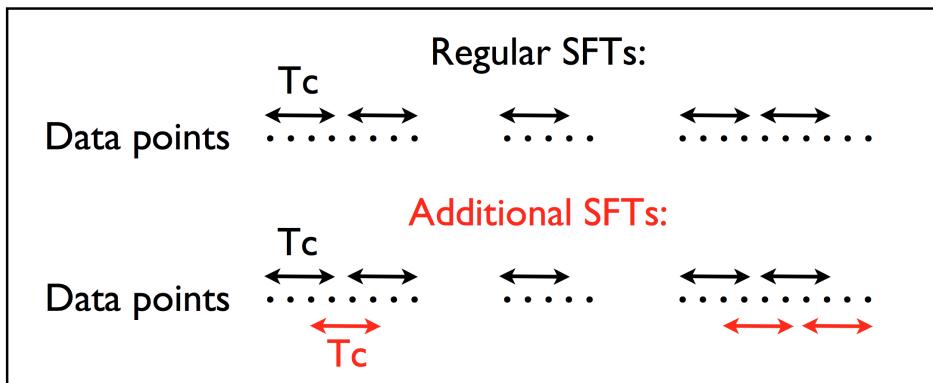


Figure 4.4: Simple sketch showing how the additional set of SFTs is obtained.




Figure 4.5: The plot shows the efficiency curve at 14 different sensitivity depth points evaluated with different detection statistics at 872.1 Hz. The blue points are from the Hough number count significance, the red points are from the weighted power significance and the green points are from the weighted mixed power significance. The points with a non-continuous trace are calculated using the standard set of SFTs, and the points with a continuous trace are calculated using the standard SFTs plus the additional ones with $q = 2$.

producing data: Advanced LIGO [28], consisting of two detectors with 4-km arm lengths situated in Hanford, Washington (H1) and Livingston, Louisiana (L1), and Advanced Virgo [29], a 3-km detector located in Cascina, Pisa. Advanced Virgo first joined the run at the beginning of August 2017, with less sensitivity than the LIGO detectors, so we have not considered its data for the search described in this chapter.

A representative noise curve from O2 for each LIGO detector and a comparison to O1 is shown in figure 4.6. We can observe an improvement of the amplitude spectral density, and we can also observe that the spectra features a number (greatly reduced as compared to O1) of narrow lines and combs affecting several frequency bands, which contaminate the data and complicate the analysis often raising outliers which look like the searched CW signals [249]. A cleaning procedure was applied to H1 and L1 data during post-processing in order to remove jitter noise (for H1) and some noise lines (more details are given in [250]). All of the searches of this chapter used this cleaned dataset. The calibration of this dataset and its uncertainties on amplitude and phase are described in [46]. These searches do not use all the data from the observing run, since times where the detectors are known to be poorly behaving are removed from the analysis (the used segments are listed in [251], where the files with the “minusContaminatedPeriod” tag are selected). This means that the effective amount of data used is smaller than the full duration of the run. As in previous observing runs, several artificial signals (called hardware injections) have been physically injected in the detectors in order to test their response and to validate the different pipelines.

For the reasons mentioned in the previous subsection 4.2.1, we decided to split all the data in two datasets belonging to different epochs, listed in table 4.3. The nine month O2 run is not “continuous” in time: there was a two-week break for Christmas, and there was a three-week break in May for maintenance operations. We select as the splitting point between the two datasets this three week maintenance stop. Since more data is present before the stop than afterwards, we do not use all that data, and we throw away data from the beginning of the run until approximately the same amount of data is present in both




Figure 4.6: Amplitude spectral density (ASD) $\sqrt{S_n}$ plots for the L1 (left panel) and H1 (right panel) detectors during O1 (blue trace) and O2 (orange trace). The ASD is obtained by averaging over FFTs of 1800 s obtained for the entire run.

datasets (the initial data of an observing run is usually the worse portion, since many improvements are made throughout the run).

An extra safety check was done for the data from both detectors, in order to see if the noise level was approximately constant during all the observing run. An example of these kinds of tests can be seen in figure 4.7, where the average amplitude spectral density over groups of two weeks can be seen. With this plot one can check if any two-week period is much worse than all other periods at certain frequencies.

4.4 | Description of the search methods

In this section we introduce and summarize the different pipelines that have been used in this work, focusing on *SkyHough*.

4.4.1 | SkyHough

The *SkyHough* method has been used in other searches using data from the Initial LIGO S2, S4 and S5 and Advanced LIGO O1 observing runs [190, 191, 195, 201, 203]. This method has been explained in section 3.3. Here we summarize its main characteristics and the new features that have been implemented in this search. The code for the main part of the search is called *lalapps_DriveHoughMulti* and is part of the publicly available LALSuite package [51].

Initial analysis steps

This pipeline uses Short Fourier Transforms (SFTs) of the time-domain $h(t)$ as its input data. The coherent duration of each chunk changes as a function of the searched frequency, as shown in table 4.1). This is selected by using this equation (for a maximum spin-down/up of 10^{-8} Hz/s):

$$T_c < \frac{1}{\sqrt{2 \times 10^{-8} + 2.246 \times 10^{-10} f}}, \quad (4.25)$$

which was obtained in subsection 3.3.1. More than four coherent times could have been used, but this is not done because the amount of sensitivity that can be gained by slight increases of the coherent time is small, and using a higher number of coherent times would increase the work that needs to be done to characterize the different vetoes, follow-up pipelines and upper limits calculations.




Figure 4.7: The amplitude spectral density is shown at 100 Hz intervals for the H1 O2 data. The different points belong to different averages of two-week periods.




Figure 4.8: The blue line shows the maximum coherent time allowed at each frequency by equation (4.25) (the coherent time that we use should be under this curve). The four constant red lines show four different coherent times: 3600 s, 2700 s, 1800 s and 900 s.

Frequency [Hz]	T_c [s]	N_{SFT}
[50, 300)	3600	2544 (4755)
[300, 550)	2700	3460 (6568)
[550, 1300)	1800	5283 (10195)
[1300, 1500)	900	10801 (21200)

Table 4.1: Coherent times and number of SFTs for each frequency range searched by the *SkyHough* pipeline. The last column shows the number of SFTs per dataset, and the numbers in parenthesis the SFTs used at the second step of the search.

From the SFTs the peak-grams are created, which are spectrograms with the normalized power substituted by 1s (if the power is above a certain threshold $\rho_t = 1.6$) and 0s, where the normalized power in a frequency bin was defined in equation (3.44).

We use the Hough transform to track the time-frequency evolution of the signal including the Doppler modulation of the signal at the detector. In the first stage the pipeline employs a look-up table (LUT) approach, taking into account that at a given time the same Doppler modulation is produced by an annulus of sky positions which correspond to the width of a frequency bin as given by equation (3.48). The algorithm tracks the sky positions which produce observed frequencies with powers above the threshold. It then stacks these sky positions by following the evolution of the source frequency given by the spin-down term at different timestamps and produces a final histogram. The LUT approach reuses the same Doppler modulation pattern for different search frequencies (more details in [158]), which produces computational savings in exchange for not following the exact frequency-time pattern.

For each template (described by f_0, f_1, α, δ) being searched, a detection statistic called number count significance is calculated:

$$s_n = \frac{n - \langle n \rangle}{\sigma_n} \quad (4.26)$$

where $\langle n \rangle$ and σ_n are the expected mean and standard deviation of the Hough number count n when only noise is present. The number count n is the weighted sum of 1s and 0s, where the weights (which are proportional to the antenna pattern functions and inversely proportional to the power spectral density) were derived in [246].

The parameter space is separated in 0.1 Hz frequency bands and in small sky-patches. A toplist is calculated for each of these regions, which has the top templates ordered by the number count significance. For the top templates, a second step is performed where instead of using the LUT approach the exact frequency path is tracked. At this second step two extra detection statistics are calculated, and we end up using the power significance, which is defined as:

$$s_P = \frac{P - \langle P \rangle}{\sigma_P} \quad (4.27)$$

where instead of summing weighted 1s and 0s the weighted normalized powers given by equation (3.44) are summed. This detection statistic improves the sensitivity of *SkyHough* with a very small increase of computational cost, as shown in the previous subsection 4.2.3. The best 5000 templates per sky-patch and 0.1 Hz band are passed to the second step, and only the best 1000 candidates per sky-patch and 0.1 Hz band are used for the post-processing. Furthermore, at the second step more SFTs are used than in the first step. This is achieved by sliding the initial times of each SFT that was used at the first step, obtaining more SFTs (approximately twice the previous amount), all of them of T_c contiguous seconds.

The grid resolution was obtained in [158] and it is given by equation (3.51), where we have set $P_F = 2$ for all frequencies.

Post-processing

The post-processing consists of several steps:

1. The output of the main *SkyHough* search is one toplist for each dataset (there are two datasets, each one with data from two detectors, detailed in table 4.3) and each region in parameter space. We search for coincidental pairs between these toplist, by calculating the distance in parameter space and selecting the pairs which are closer than a certain threshold called d_{co} . For the coincidental pairs the centers (average locations in parameter space weighted by significance) are calculated. The parameter space distance is calculated as:

$$d_{SH} = \sqrt{\left(\frac{\Delta f}{\delta f}\right)^2 + \left(\frac{\Delta \dot{f}}{\delta \dot{f}}\right)^2 + \left(\frac{\Delta x}{\delta x}\right)^2 + \left(\frac{\Delta y}{\delta y}\right)^2}, \quad (4.28)$$

where the numbers in the numerators represent the difference between to templates and the numbers in the denominators represent the parameter resolution (this distance is unitless and is given as a number of bins). The parameters x and y are the Cartesian ecliptic coordinates projected in the ecliptic plane.

2. Search for clusters in the obtained list of centers. This will group different outliers which can be ascribed to a unique physical cause, and will reduce the size of the final toplist. Again, we set a threshold in parameter space distance (called d_{cl}) and we find candidates which are closer than this distance.
3. Finally, we calculate the centers of the clusters. This is done as a weighted (by significance) average, taking into account all the members of the cluster. We order the obtained clusters in each 0.1 Hz by their sum of the power significance of all the members of a cluster, and we select the highest ranked cluster per 0.1 Hz band, if any. This produces the final list of clusters with their parameters $(f_0, f_1, \alpha, \delta)$ which will be the outliers to be followed-up.

Follow-up

We describe a follow-up method based on the \mathcal{F} -statistic and the GCT metric method [252]. This method uses the *lalapps-HierarchSearchGCT* code, part of the publicly available LALSuite [51], and it is similar in spirit to the multi-step follow-up methods described in [253] or [144].

The goal is to compare the \mathcal{F} -statistic values obtained from software injected signals to the \mathcal{F} -statistic values obtained from the outliers. We want to compare the \mathcal{F} -statistic obtained at different stages which scale the coherent time. It is expected that for an astrophysical signal the \mathcal{F} -statistic value should increase if the coherent time increases.

The resolution in parameter space is given by [252]:

$$\delta f = \frac{\sqrt{12m}}{\pi T_c} \quad (4.29)$$

$$\delta \dot{f} = \frac{\sqrt{720m}}{\pi T_c^2 \gamma} \quad (4.30)$$

$$\delta \theta = \frac{\sqrt{m_{sky}}}{\pi f \tau_e}, \quad (4.31)$$

where m and m_{sky} are mismatch parameters, γ is a parameter which gives the refinement between the coherent and semi-coherent stages and $\tau_e = 0.021$ s represents the light time travel from the detector to the center of the Earth.

We now enumerate the different steps of the procedure:

1. Calculate the semi-coherent \mathcal{F} -statistic of outliers with $T_c = 7200$ s in a cluster box.

Stage index	T_c [s]	m	m_{sky}
I	7200	0.1	0.01
II	72000	0.1	0.003
III	720000	0.1	0.0005

Table 4.2: Coherent times and mismatch parameters at each different stage of the *SkyHough* follow-up.

2. Add injections to the original data using a sensitivity depth ($\sqrt{S_n}/h_0$) value which returns \mathcal{F} -statistic values similar to the values obtained with the outliers in order to compare similar signals. We inject signals in 8 different frequency bands with 200 injections per band, with a sensitivity depth of $42 \text{ Hz}^{-1/2}$. Then, search in a small region (around 10 bins in each dimension) around the true parameters of the injections with $T_c = 7200$ s. Finally, analyse the distances in parameter space from the top candidates to the injections.
3. Repeat the previous step increasing the coherent time to $T_c = 72000$ s.
4. Calculate $\mathcal{F}_{72000s}/\mathcal{F}_{7200s}$ for each of the 1600 injections using the top candidates. The threshold will be the minimum value.
5. Calculate the \mathcal{F} -statistic values of outliers with $T_c = 72000$ s around the top candidate from the first stage. The size of the window to be searched is estimated from the distances found in step 2.
6. Calculate $\mathcal{F}_{72000s}/\mathcal{F}_{7200s}$ using the top candidate for all outliers. Outliers with values higher than the threshold obtained in step 4 go to the next comparison, and the process is repeated from step 2 increasing the coherent time.

Table 4.2 summarizes the parameters that we have chosen at each different stage. As we will see in the results section, only two comparisons between three different stages were needed.

4.4.2 | Comparison with the other pipelines

We first compare the *SkyHough* pipeline to the *FrequencyHough* pipeline, since they are very similar in spirit. The *FrequencyHough* algorithm does not use the partial Hough map derivative and look-up table approaches. Instead, for each sky position the frequency bins are shifted, so that only the linear modulation produced by the spin-down/up is left. Afterwards, the *FrequencyHough* algorithm transforms the peaks of the frequency-time plane to the f_0 and f_1 parameters of the source. This pipeline also separates the search in two datasets, but separated by detector are not by time (thus, the needed resolution in spin-down/up is much finer). Furthermore, the follow-up procedure still uses the Hough method but with longer coherent times, where this can be achieved by heterodyning the data.

The *Time-Domain \mathcal{F} -statistic* is a method that uses the \mathcal{F} -statistic as its detection statistic, and for this reason it can use much longer coherent times (on the order of days). This causes higher computational costs, and for this reason the mismatch parameters (that define the grid resolution) have to be increased. The incoherent combination of the coherent detection statistics is not a sum, but instead a search for coincidences between the multiple segments.

It can be seen that both of these pipelines use coherent times that are longer than *SkyHough*. Another difference is that both of these pipelines used cleaned data, as opposed to *SkyHough* that only uses the default calibrated data (in this O2 search the H1 data was cleaned to remove a jitter-noise disturbance). This extra cleaning is mainly composed of a removal of short-duration glitches, with a procedure described in [248]. This cleaning helps to remove wide-band elevations of the noise floor and thus increases the sensitivity of the search without any extra computational cost.




Figure 4.9: Regions in frequency and first frequency derivative covered by each pipeline.

Dataset 1	H1	1167545839/1174691692
	L1	1167546403/1174688389
Dataset 2	H1	1180982628/1187731792
	L1	1179816663/1187731695

Table 4.3: Start/stop times in GPS units of each dataset used by the *SkyHough* pipeline. The observation time parameter used for the spin-down resolution given by equation (3.51) is $T_{obs} = 7915032$ s, the maximum span of these datasets.

4.5 | Results

In this section we detail the results obtained by the *SkyHough* pipeline and the outliers that were found, and we also present 95% confidence level (CL) upper limits on the strain h_0 . We also compare these results to the other two pipelines. The region in frequency and first frequency derivative searched by each of the three different pipelines is shown in figure 4.9.

4.5.1 | SkyHough

SkyHough has analysed frequencies from 50 to 1500 Hz and spin-down values from -10^{-8} to 10^{-9} Hz/s as shown in figure 4.9. The four different coherent times that have been used are shown in table 4.1. This analysis uses the C02 cleaned dataset [46], and splits the data from H1 and L1 in two datasets, divided by time as shown in table 4.3, where the start and stop times for each dataset are indicated. The main search generates a toplist per dataset per 0.1 Hz band of 10000 candidates with a maximum of 1000 per sky-patch. The number of sky-patches depends on the frequency: to minimize the computational cost of the search, we try to minimize the number of sky-patches for a limited amount of RAM. From 50 to 850 Hz, there are 28 sky-patches; from 850 to 1000 Hz, 31 sky-patches; from 1000 to 1150 Hz, 38 sky-patches; from 1150 to 1250 Hz, 45 sky-patches; from 1300 to 1500 Hz, 28 sky-patches. After applying the post-processing stage previously described (with distance thresholds of $d_{co} = 3$ and $d_{cl} = \sqrt{14}$), we are left with 4548 0.1 Hz bands (from a total of 14500) having coincidental pairs.

We apply the *population veto*, used in many past searches, which demands that each dataset con-




Figure 4.10: Results from software injections for the first (left panel) and second (right panel) comparisons between the first and last two stages of table 4.2 of the *SkyHough* follow-up. The vertical axis shows the quotient between the top candidates at the two stages, and the horizontal axis shows the values of the top candidates at the stage with lowest T_c . The lowest points, equal to 1.47 and 3.66 respectively, set the thresholds for the follow-up veto. Each color represents a different frequency band.

tributes to each cluster with at least two different templates. After applying this veto, only 1539 outliers remain.

The next step is to apply the \mathcal{F} -statistic follow-up method described in section 4.4.1 to these 1539 outliers. The thresholds obtained are 1.47 and 3.66 for the first and second comparison respectively, as shown in Fig 4.10. Only 17 outliers are above the threshold at 3.66, as shown in figure 4.11. All of the outliers which are above the final threshold correspond to one of the hardware injections listed in table 4.5 or to one known source of instrumental noise, listed in [249]. The 17 surviving outliers and their parameters are listed in table 6.3, with comments about their likely origin.

We recover 6 of the 9 hardware injections that are in the *SkyHough* searched parameter space. We lose the other three mainly for two reasons: there were brief periods when the hardware injections were not active, which causes the increase of the \mathcal{F} -statistic to not be as high as it should be (this happens to two of the three lost hardware injections, which are present in our initial list of 1539 outliers); we only select one cluster per 0.1 Hz band, and if in that band there is a more significant cluster due to a noise disturbance the signal cluster will not be followed (this happens to one of the three lost hardware injections, which forms a cluster but a more significant noise disturbance is present in that 0.1 Hz band).

Although no detections were made, we produce all-sky averaged upper limits on the strain of the signal h_0 (these upper limits are valid for all the frequency bands except the ones which have one of the 1539 outliers). We add software simulated signals to the original data by using *lalapps_Makefakedata_v5*. We have injected signals at 10 different 0.1 Hz bands for each of the four coherent times (a total of 40 bands), which can be seen in table 4.4. These are bands which don't have outliers or instrumental known sources of lines or combs.

We have used 5 different sensitivity depths at each coherent time, with 400 signals per sensitivity depth. The sensitivity depth is given by:

$$\mathcal{D} = \frac{\sqrt{S_n}}{h_0}, \quad (4.32)$$

where S_n is the one-sided power spectral density. We inject signals at random positions in the sky, covering the full spin-down range and with random polarization, inclination and initial phase. For each band and depth, we calculate the efficiency (number of detected signals divided by total number of signals). We follow the same procedure as in the all-sky search: we run the main search and then apply coincidences, clustering and the population veto. We assume that a signal is detected if the total distance




Figure 4.11: Results of the second follow-up comparison (using stages II and III) for the SkyHough outliers. Only 17 outliers are above the threshold. The horizontal line marks the threshold at 3.66, which was obtained in section 4.4.1.

T_c [s]	Frequency [Hz]	Injected \mathcal{D} [1/Hz ^{-1/2}]	$\mathcal{D}^{95\%}$ [1/Hz ^{-1/2}]
3600	110.3,136.1,148.3,	29.5, 30.5,	32.4
	165.6,182.6,206.1,	31.5, 32.5,	
	225.6,241.5,261.3,286.7	33.5	
2700	311.6,325.4,342.5,	26.5, 27.5,	29.2
	363.3,394.0,412.4,	28.5, 29.5,	
	432.8,441.2,523.4,547.8	30.5	
1800	594.6,661.1,741.4,	23.0, 24.0,	25.2
	805.0,866.2,933.1,	25.0, 26.0,	
	977.6,1064.7,1141.5,1250.7	27.0	
900	1313.4,1331.7,1358.4,	22.0, 23.0	22.3
	1370.3,1388.8,1402.4,	24.0, 25.0,	
	1423.1,1430.3,1443.4,1464.6	26.0	

Table 4.4: The second column shows the frequency bands used to estimate the SkyHough upper limits on gravitational-wave signal amplitude h_0 . The third column shows the injected sensitivity depth values given by equation (4.32), and the last column shows the sensitivity depth at 95% confidence for each group.




Figure 4.12: The left plot shows the efficiency as a function of sensitivity depth and fitting at 148.3 Hz for the SkyHough pipeline. The vertical error bars for the blue points show the 1-sigma binomial error. The 95% efficiency point (indicated with a black cross) also shows a 1-sigma error bar, calculated with equation (4.34). The right plot shows the 95% sensitivity depths at each of the 40 frequency bands, with a 1-sigma error bar, for the SkyHough pipeline. The three vertical lines separate the four regions with different coherent time (3600 s, 2700 s, 1800 s and 900 s).

from the recovered cluster to the actual injection is less than 13 bins.

At each of the 40 frequency bands, we fit a sigmoid given by:

$$s = 1 - \frac{1}{1 + e^{b(x-a)}}. \quad (4.33)$$

An example of this fitting can be seen in figure 4.12. From the estimated coefficients a and b along with the covariance matrix C_{ab} , we calculate the 1-sigma envelope (the error) on the fit, which is given by:

$$\sigma_S = \pm \sqrt{\left(\frac{\delta s}{\delta a}\right)^2 C_{aa} + \left(\frac{\delta s}{\delta b}\right)^2 C_{bb} + 2\frac{\delta s}{\delta a}\frac{\delta s}{\delta b} C_{ab}}. \quad (4.34)$$

After finding the 95% efficiency sensitivity depth at each of the 40 frequency bands (which can be seen in the right plot of figure 4.12), we calculate a mean sensitivity depth for each of the four different frequency regions. The results are given in table 4.4. From these results and by using equation (4.32), we calculate the upper limits on h_0 , which are shown in figure 4.13. The trace has a shadow enclosing a 7.5% error, which we obtain by estimating the maximum difference in each of the four frequency regions shown in the right plot of figure 4.12 between the 10 different points and the mean sensitivity depth. Figure 4.13 also shows a comparison with the results obtained in the previous search with O1 data. The new upper limits are approximately 1.5 times better than the O1 upper limits.

4.5.2 | Global results

Although no detections have been made, we present 95% confidence level (CL) upper limits on the strain h_0 , shown in figure 4.14. The best upper limit is $\simeq 1.7 \times 10^{-25}$ at around 120 Hz. These results are significantly better (of a factor of about 1.4) than those obtained on O1 data with the same pipelines [201, 203], thanks to improvements in the pipelines themselves, to the better sensitivity of the detectors and to the longer duration of the observing run. These upper limits do not take into account the calibration uncertainty on amplitude, which over the run was no larger than 5% and 10% for H1 and L1 respectively [46].

The best results from this search are comparable with the upper limits obtained in O1 by the Einstein@Home project [202] over the range 20-100 Hz. Note, however, that the Einstein@Home search covered a spin-down range smaller by almost one order of magnitude. Moreover, while the Einstein@Home




Figure 4.13: 95% upper limits on h_0 for the *SkyHough* pipeline. The orange trace shows the results for the O2 search, with a shadow enclosing a 7.5% error obtained from the right plot of figure 4.12, while the pink trace shows the results obtained in the O1 search. These results are valid for all frequency bands except the 1539 bands where one outlier is present.

search is, in principle, more sensitive due to the use of much longer data segments (compared to the *FrequencyHough* and *SkyHough* pipelines), with 210 hr duration, it is also less robust in case of deviations from the assumed signal model given by equation (4.1). At frequencies higher than 100 Hz, a direct comparison with the previous best upper limits (obtained by the *PowerFlux* pipeline [203]) is difficult, since they were obtained assuming circularly polarized signals, while we calculate them by using an average over an isotropic population.

The 95% CL upper limits on h_0 can be converted to upper limits on ellipticity ϵ assuming a canonical value for the moment of inertia of 10^{38} kg·m² and by using different distances:

$$\epsilon = \frac{c^4}{4\pi^2 G} \frac{h_0 d}{I_{zz} f^2}. \quad (4.35)$$

These results are shown in the left panel of figure 4.15. This has been obtained by using the best h_0 upper limits between the three pipelines: from 20 to 1000 Hz, the *FrequencyHough* results have been used; from 1000 to 1500 Hz, the results from *SkyHough* have been used; from 1518 to 1922 Hz the results from *Time-Domain F-statistic* have been used. For sources at 1 kpc emitting CWs at 500 Hz, we can constrain the ellipticity at $\simeq 10^{-6}$, while for sources at 10 kpc emitting at the same frequency we can constrain the ellipticity at 10^{-5} .

A complementary way of interpreting the limits on ellipticity is shown in the right panel of figure 4.15. The various set of points give the relation between the absolute value of the signal frequency time derivative (spin-down) and the signal frequency for sources detectable at various distances, assuming their spin-down is only due to the emission of gravitational waves. They have been computed by means of the following relation obtained inverting the equation for the so-called spin-down limit amplitude h_0^{sd} , which is a function of the source distance d , frequency f and spin-down \dot{f} , see e.g. equation (A7) in [110]:

$$|\dot{f}| = 1.54 \times 10^{-10} \left(\frac{I_{zz}}{10^{38} \text{ kg} \cdot \text{m}^2} \right)^{-1} \left(\frac{h_0^{\text{sd}}}{10^{-24}} \right) \left(\frac{f}{100 \text{ Hz}} \right) \left(\frac{d}{1 \text{ kpc}} \right)^2 \text{ [Hz/s]}, \quad (4.36)$$

where we have replaced the spin-down limit amplitude with the 95% upper limits shown in figure 4.14.




Figure 4.14: Upper limits on the strain amplitude $h_0^{95\%}$ for the three pipelines.

The dashed lines are constant ellipticity curves obtained from equation (A9) of [110]:

$$|\dot{f}| = 1.72 \times 10^{-14} \left(\frac{I_{zz}}{10^{38} \text{ kg} \cdot \text{m}^2} \right) \left(\frac{f}{100 \text{ Hz}} \right)^{1/2} \left(\frac{\varepsilon}{10^{-6}} \right)^2 \text{ [Hz/s]}. \quad (4.37)$$

For a signal to be detectable, its spin-down would need to be equal or above the given traces (notice that, as shown in figure 4.9, the maximum absolute spin-down searched is 10^{-8} Hz/s, which marks a limit to the signals we are sensitive to). For example, a source emitting a signal with frequency higher than about 500 Hz and ellipticity equal or greater than 10^{-6} would be detectable up to a distance of about 1 kpc if its spin-down is, in modulus, larger than $\approx 10^{-10}$ Hz/s.

The three searches carried out by the different pipelines have different computational costs: *FrequencyHough* spent 9 MSU; *SkyHough* spent 2.5 MSU; *Time-Domain F-statistic* spent 24.2 MSU, where 1 MSU hour corresponds to 1 million Intel E5-2670 core-hour to perform a SPECfp computation. We remind the reader that each of these pipelines covered different search bands. These numbers show that although the *SkyHough* pipeline is the one that covers more parameter space, it is also the cheapest (three times less than *FrequencyHough*).

4.6 | Conclusions

In this chapter we have presented the first results of an all-sky search for CW signals using Advanced LIGO O2 data with three different pipelines (focusing on the results from *SkyHough*), covering a frequency range from 20 to 1922 Hz and a first frequency derivative from -1×10^{-8} to 2×10^{-9} Hz/s. For this broad range in parameter space, this is the most sensitive search up to 1500 Hz. Each pipeline found many outliers that were followed-up but none of them resulted in a credible astrophysical CW signal. On the contrary, they were ascribable to noise disturbances, to hardware injections, or consistent with noise fluctuations.

Although no detections have been made, we have placed interesting 95% CL upper limits on the gravitational wave strain amplitude h_0 , the most sensitive being $\simeq 1.7 \times 10^{-25}$ in the 123-124 Hz region, as shown in figure 4.14. The improved results over the O1 search are due to the better sensitivity of the detectors, the use of a longer dataset and improvements of the pipelines. By converting the upper limits to an astrophysical reach, as shown in figure 4.15, we see that the searches presented in this chapter provide already astrophysical interesting results. For instance, in the bucket region (around $\simeq 150$ Hz), we would be able to detect a CW signal from a neutron star within a distance of 100 pc if its ellipticity

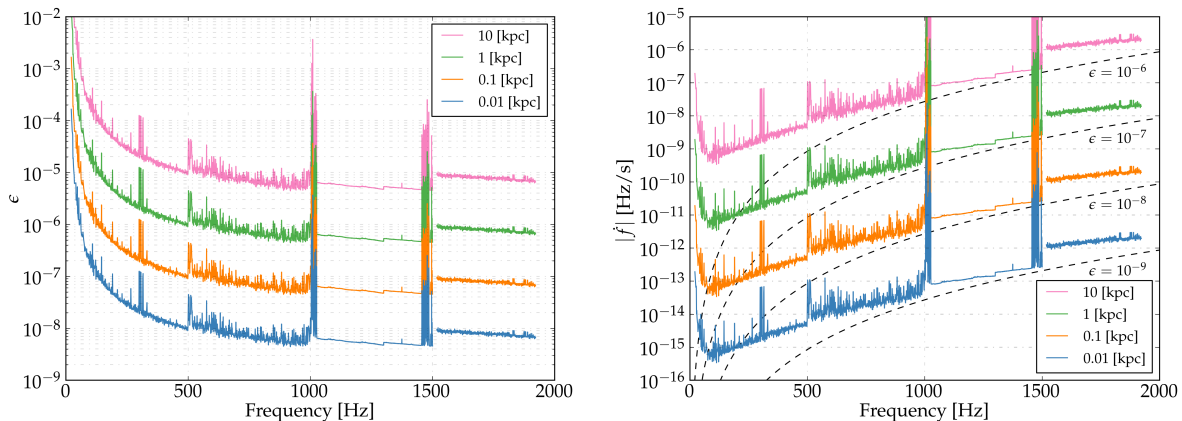


Figure 4.15: The left panel shows the detectable ellipticity given by equation (4.35) as a function of frequency for neutron stars at 10 pc, 100 pc, 1 kpc and 10 kpc for a canonical moment of inertia $I_{zz} = 10^{38} \text{ kg}\cdot\text{m}^2$. The right panel shows the relation between the absolute value of the first frequency derivative and the frequency of detectable sources as a function of the distance, assuming their spin-down is due solely to the emission of gravitational waves. The different colors correspond to the same distances of the left panel. Black dashed lines are lines of constant source ellipticity, from $\epsilon = 10^{-9}$ (bottom dashed line) to $\epsilon = 10^{-6}$.

were at least 10^{-6} . Similarly, in the middle frequency range, around $\simeq 500$ Hz, we would be able to detect the CW signal up to a distance of 1 kpc, with $\epsilon > 10^{-6}$. Finally, at higher frequencies (around $\simeq 1500$ Hz) the same signal would be detectable up to a distance of 10 kpc if $\epsilon > 10^{-6}$ and 1 kpc if $\epsilon > 10^{-7}$. Such levels of ellipticity are comparable or below the maximum value we may expect for neutron stars described by a standard equation of state [101].

The O3 observing run started in April 2019 and lasted approximately 1 year. The LIGO and Virgo detectors have been upgraded and improved, and we expect that the noise floor in O3 run will be significantly better than for O2. This, and the foreseen longer run duration, will make future searches more sensitive, increasing the chances of a CW detection or allowing us to place tighter constraints on the non-asymmetries of neutron stars in our galaxy and to put constraints on the unseen neutron star population.

Label	Frequency [Hz]	Spin-down [nHz/s]	α [deg]	δ [deg]
ip0	265.575533	-4.15×10^{-3}	71.55193	-56.21749
ip1	848.969641	-3.00×10^{-1}	37.39385	-29.45246
ip2	575.163521	-1.37×10^{-4}	215.25617	3.44399
ip3	108.857159	-1.46×10^{-8}	178.37257	-33.4366
ip4	1393.540559	-2.54×10^{-1}	279.98768	-12.4666
ip5	52.808324	-4.03×10^{-9}	302.62664	-83.83914
ip6	146.169370	-6.73×10^0	358.75095	-65.42262
ip7	1220.555270	-1.12×10^0	223.42562	-20.45063
ip8	191.031272	-8.65×10^0	351.38958	-33.41852
ip9	763.847316	-1.45×10^{-8}	198.88558	75.68959
ip10	26.341917	-8.50×10^{-2}	221.55565	42.87730
ip11	31.424758	-5.07×10^{-4}	285.09733	-58.27209
ip12	38.477939	-6.25×10^0	331.85267	-16.97288
ip13	12.428001	-1.00×10^{-2}	14.32394	-14.32394
ip14	1991.092401	-1.00×10^{-3}	300.80284	-14.32394

Table 4.5: Parameters of the hardware-injected simulated continuous-wave signals during the O2 data run (parameters given at epoch GPS 1130529362).

Outlier index	Frequency [Hz]	Spin-down [nHz/s]	α [deg]	δ [deg]	Population	\bar{s}_P	\mathcal{F}	Description
7	51.0002	-1.8346×10^{-11}	87.0087	-66.0873	6542	50.16	86.0406	1 Hz comb at H1 and L1
18	52.8083	2.2838×10^{-12}	299.6708	-83.3562	1433	124.12	654.9832	Hardware injection 5
36	56.0001	-6.3195×10^{-12}	88.8156	-66.4443	2765	210.17	157.7912	1 Hz comb at H1 and L1
39	56.4957	5.7155×10^{-10}	156.8932	-49.1672	111	9.04	72.64532	1 Hz comb at H1 and L1
113	70.0001	-1.2045×10^{-11}	88.1699	-66.1180	2980	121.85	165.2225	1 Hz comb at H1 and L1
125	72.0001	-5.4952×10^{-12}	89.4121	-66.5021	3471	57.69	64.56955	1 Hz comb at H1 and L1
149	76.6788	-1.7870×10^{-10}	74.5911	-59.0142	397	32.66	444.1439	0.08843966 Hz comb at H1
150	76.9442	-1.8308×10^{-10}	74.6723	-58.9216	558	36.12	566.2429	0.08843966 Hz comb at H1
151	77.1207	-1.4130×10^{-10}	77.8483	-59.1510	129	26.81	507.6696	0.08843966 Hz comb at H1
152	77.2090	-1.3071×10^{-10}	79.0815	-59.5503	93	21.51	405.0106	0.08843966 Hz comb at H1
153	77.3855	-7.9419×10^{-11}	279.3909	70.9090	150	14.78	154.7181	Unknown line at L1
316	108.8567	4.3434×10^{-11}	182.2717	-29.6472	3050	72.87	275.9235	Hardware injection 3
485	145.9203	-6.7301×10^{-9}	358.7866	-65.2887	675	86.45	311.7228	Hardware injection 6
664	199.9977	-1.5070×10^{-11}	89.5203	-66.2582	1249	83.46	130.1102	99.9987 Hz comb at H1
1629	575.1638	-2.9038×10^{-11}	215.4209	4.0846	636	564.51	4962.3050	Hardware injection 2
2303	763.8471	1.2199×10^{-11}	198.9249	75.6197	1798	637.28	2703.2552	Hardware injection 9
2584	848.9591	-3.4856×10^{-10}	37.3061	-28.8880	443	898.46	4473.9817	Hardware injection 1

Table 4.6: 17 outliers from the SkyHough pipeline which survived the follow-up procedure. All of them can be ascribed to a hardware injection or to a known source of instrumental noise. The $(f_0, f_1, \alpha, \delta)$ values correspond to the center of the cluster returned by the post-processing stage. The \bar{s}_P column shows the mean power significance of the cluster, while \mathcal{F} column shows the \mathcal{F} -statistic mean over segments of the top candidate obtained at the last stage of the follow-up. The reference time for these parameters is 1167545839 GPS.

CHAPTER 5

BINARYSKYHOUGH: A NEW METHOD TO SEARCH FOR CONTINUOUS GRAVITATIONAL WAVES FROM UNKNOWN NEUTRON STARS IN BINARY SYSTEMS

As discussed in the introductory chapters, non-axisymmetric fast-spinning unknown neutron stars in binary systems may emit continuous gravitational waves that can be detected by ground-based detectors like LIGO and Virgo. In this chapter we present a new pipeline, called *BinarySkyHough*, that can carry out all-sky searches for neutron stars in binary systems by exploiting the usage of graphics processing units (GPUs). We give a detailed explanation of this new pipeline, and we present simulations that allow us to estimate the sensitivity of the new pipeline, which is approximately twice as sensitive as the best active pipeline with a comparable computational cost. This chapter is a reorganization of the material presented in [254].

5.1 | Introduction

Approximately half of the known pulsars in the most sensitive frequency band of the ground-based interferometric detectors belong to binary systems. Neutron stars in binary systems have an additional modulation due to the NS movement around the binary barycenter (BB), as explained in subsection 2.4.3. Several directed searches for CWs from NSs in known binary systems (such as Scorpius X-1) have been already performed, for example [186] using O1 data. These searches usually have to deal with a four-dimensional parameter space comprised of the source frequency and three binary parameters, and also rely on semi-coherent methods to deal with the high computational cost.

All-sky searches for NSs in binary systems present an even harder problem, since the two sky positions need to be searched too, and the most sensitive semi-coherent methods used in all-sky searches for isolated systems cannot be used due to limited computational power. Currently, there are only two pipelines that can perform these searches, called *TwoSpect* [98] and *NarrowBand Radiometer* [255], and their sensitivity compared to the all-sky searches for isolated neutron stars is approximately 3 times worse. Only one all-sky search for NSs in binary systems had been published prior to the development of *BinarySkyHough* [197] (by the *TwoSpect* pipeline), having no detections and producing upper limits on the gravitational-wave amplitude. The lower sensitivity of these pipelines compared to all-sky searches for isolated NSs calls for development of advanced techniques to improve the chances of detecting a CW signal from this type of system.

In this chapter we describe a new method to perform all-sky searches for CWs from NSs in binary systems, which we call *BinarySkyHough*. This method is an extension of the *SkyHough* [158] pipeline used

for all-sky searches of isolated NSs, and it benefits from the usage of GPUs in order to have a manageable computational cost.

In this chapter we assume that the gravitational-wave signal is described by the simple aligned triaxial ellipsoid model explained in subsection 2.4.1:

$$h(t) = h_0[F_+(t, \psi, \hat{n}) \frac{1 + \cos \iota}{2} \cos \phi(t) + F_\times(t, \psi, \hat{n}) \cos \iota \sin \phi(t)], \quad (5.1)$$

where F_+ and F_\times are the antenna patterns of the detectors (which can be found in [50]) for the two different gravitational-wave polarizations, t is the time at the detector frame, the inclination angle ι is the angle between the neutron star angular momentum and the observer's sky plane, ψ is the wave polarisation angle, $\phi(t)$ is the phase of the signal and h_0 is the amplitude of the signal given by:

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{I_{zz} \epsilon f^2}{d}, \quad (5.2)$$

where d is the distance from the detector to the source, f is the gravitational-wave frequency (equal to two times the rotational frequency), ϵ is the ellipticity or asymmetry of the star, usually given by $(I_{xx} - I_{yy})/I_{zz}$, and I_{zz} is the moment of inertia of the star with respect to the principal axis aligned with the rotation axis. The phase of the signal in the detector frame for a neutron star in a binary system has been described in subsection 2.4.3:

$$\begin{aligned} \phi(t) = & \phi'_0 + 2\pi f_0(t - t_r) + \frac{\vec{r}(t) \cdot \hat{n}}{c} \\ & - a_p \sin[\Omega(t - t_{\text{asc}})] - a_p \frac{e \cos \omega}{2} \sin[2\Omega(t - t_{\text{asc}})] + a_p \frac{e \sin \omega}{2} \cos[2\Omega(t - t_{\text{asc}})], \end{aligned} \quad (5.3)$$

while the frequency of the gravitational wave in the detector is:

$$\begin{aligned} f(t) = & \frac{1}{2\pi} \frac{d\phi}{dt} = f_0 + f_0 \frac{\vec{v}(t) \cdot \hat{n}}{c} - f_0 a_p \Omega \cos[\Omega(t - t_{\text{asc}})] \\ & - f_0 a_p \Omega e \cos \omega \cos[2\Omega(t - t_{\text{asc}})] - f_0 a_p \Omega e \sin \omega \sin[2\Omega(t - t_{\text{asc}})]. \end{aligned} \quad (5.4)$$

BinarySkyHough assumes a zero eccentricity model in order to deal with the high computational cost of the search, given by:

$$f(t) = f_0 + f_0 \frac{\vec{v}(t) \cdot \hat{n}}{c} - f_0 a_p \Omega \cos[\Omega(t - t_{\text{asc}})], \quad (5.5)$$

which depends on only three parameters: Ω , a_p and t_{asc} .

Again, we remark that this model assumes that the neutron star does not suffer any glitches during the observing time, and that the effect of spin-wandering (stochastic variations on the rotational frequency due to the accretion process), if present, can be neglected.

The majority of the detected pulsars that are supposed to emit CWs in the frequency band of the Advanced detectors (from 50 to 1000 Hz, approximately) are millisecond pulsars, as shown in figure 5.1 (which is a zoom of figure 2.2). Almost half of the millisecond pulsars belong to a binary system. For this reason, it is important to perform CW searches which take into account the different phase model which these signals have. This figure also shows that millisecond pulsars have smaller frequency derivatives [67], which makes the assumption of zero spin-down that was used to derive equation (5.5) valid.

The accretion process in a binary system can create and sustain quadrupole deformations which can be the source of CWs. Because the maximum observed rotational frequency is well below the maximum allowed by the limit imposed by the centrifugal break-up, it is believed that some process may be counteracting the neutron star rotational acceleration before it reaches this maximum frequency. One proposed process is the emission of CWs. With this emission, neutron stars could reach a balance between accretion and emission of GWs, thereby sustaining a quadrupole which would emit CWs of amplitude given by the torque-balance limit, which can be estimated from the emitted X-ray flux for some NSs




Figure 5.1: Gravitational-wave frequency and absolute value of the gravitational-wave first frequency derivative for pulsars with gravitational-wave frequency greater than 10 Hz. Black dots indicate the pulsars which are part of a binary system. Data taken from [64] and downloaded with [69].

like Scorpius X-1. Furthermore, the Zimmermann-Szedinits mechanism may be operating, where the principal axes of the moment of inertia are driven away from the rotational axes by this accretion [256].

The eccentricity (defined as $\sqrt{1 - b^2/a^2}$, where a and b are respectively the semi-major and semi-minor axis of the binary orbit) of pulsars in binary systems is shown in figure 2.4. As said before, for most of the pulsars with measured eccentricity it is smaller than 0.01. As we will see in section 5.3, our pipeline is able to detect signals from systems with eccentricity up to 0.01 by using the zero-eccentricity model given by equation (5.5).

For a Keplerian orbit, the projected semi-major axis amplitude a_p and the orbital period P_b follow the relationship given by the third Keplerian law:

$$a_p \equiv \frac{a_{NS}}{c} \sin \iota_b = \left(\frac{G}{4\pi^2} \right)^{1/3} P_b^{2/3} \frac{M_T^{4/3}}{M_{NS}} \frac{\sin \iota_b}{c}, \quad (5.6)$$

where $a_{NS} = aM_T/M_{NS}$ is the semi-major axis amplitude of the pulsar's orbit, M_{NS} is the mass of the NS, and M_C is the mass of the companion star. Values of these two quantities for the known pulsars were shown in the left plot of figure 2.4. These observational points from known pulsars can guide the choice of parameter space that we want to search, as we will discuss later.

5.2 | New method for CW all-sky binary searches

The *SkyHough* semi-coherent method is detailed in [158] and it has been described in chapters 3 and 4. This method based on the Hough transform is used to perform all-sky searches of CWs from isolated neutron stars. It exploits the Doppler modulations produced by Earth's movement around the Solar System Barycenter (SSB) by reusing the same Doppler modulation for several frequency bins in order to save computational costs, which makes the *SkyHough* pipeline the cheapest semi-coherent method to track modeled signals currently available and the best choice to be adapted for a search for NSs in binary systems.

Figure 5.2 shows a flowchart which summarizes the different steps of the method, which have been described in section 3.3. As can be seen in this figure, this pipeline can be run with two different options:

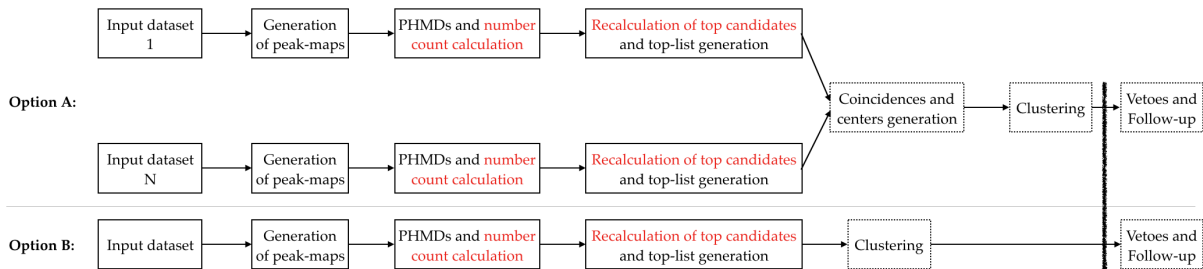


Figure 5.2: Flowchart showing the steps of the *BinarySkyHough* algorithm for one search band. Red text shows the steps performed with GPU kernels, and dotted boxes represent the post-processing stages (different vetoes and follow-up procedures are not part of the *BinarySkyHough* pipeline). Option A shows the steps followed when the data is divided in N different datasets, and option B shows the steps when the data is not divided. The main difference is that option B does not apply the coincidences step in the post-processing stage.

option A divides the data in N different datasets, while option B does not. The main difference is that option B does not apply the coincidences step in the post-processing stage as will be detailed in subsection 5.2.5. The input data, the PHMD/LUT approaches, and detection statistics are all the same as previously described.

BinarySkyHough is a new pipeline to perform all-sky searches of CWs from unknown neutron stars in binary systems with low-eccentricity orbits. The new method is an adaptation of the *SkyHough* semi-coherent method [158] which has been used for all-sky searches of CWs from isolated neutron stars like [190, 191, 195, 203]. *SkyHough* is the semi-coherent pipeline with lowest computational cost, with a sensitivity similar to other methods that are an approximately an order of magnitude more computationally costly [213, 230]. This makes it an excellent candidate to be adapted to perform CW searches in binary systems, which have a computational cost higher than isolated searches due to the extra parameters that need to be searched.

5.2.1 | BinarySkyHough

The *SkyHough* pipeline calculates the PHMDs and then sums them following the path given by $f_0 + f_1 t$. To implement a search for neutron stars in binary systems, we substitute the search over spin-down/up values with a search over binary parameters. This change is shown in figure 3.6. As we previously mentioned, the frequency derivative of pulsars in binary systems usually is smaller than isolated NSs and this parameter doesn't need to be searched for. This means that the emission of energy from the neutron star is balanced by accretion, which would produce a bright X-ray emission (as observed for Scorpius X-1), although this might not be observable due to misalignment from the emission spot to the line of sight or due to screening from other objects.

At this point in development we only allow circular orbits (zero-eccentricity), which are described by three parameters: $\Omega = 2\pi/P_b$, a_p and t_{asc} . As equation (5.5) shows, this model consists of six different parameters: the initial frequency f_0 , the sky position given by the right ascension α and declination δ , and the three binary parameters. Comparing with the isolated case, we go from a 4-dimensional parameter space to a 6-dimensional one. Although we assume the eccentricity of the system is 0, this pipeline remains sensitive to systems with eccentricity $e < 0.01$ without the need to explicitly search over it, as will be explained in sections 5.2.4 and 5.3.

The innermost loop over the binary parameters, which calculates the frequency-time path and sums the different PHMDs, has been coded in CUDA (a parallel computing platform which allows to control GPUs) in order to take advantage of the massive parallelism that GPU cards provide. Section 5.2.6 shows a comparison of timings for different runs, and we can see that without the GPUs this search would take an unfeasible time to run. The final loop for each sky-patch, which calculates the second

detection statistic for a limited number of templates, has also been coded in CUDA to further speed-up the code.

5.2.2 | Resolution, parameter space range and number of templates

Resolution

In order to construct the template bank which contains the templates that are going to be searched, we need to define the resolution of parameter space which decides the spacing between templates. The usual metric which quantifies the needed resolution is the mismatch μ , which gives one minus the ratio of recovered squared signal-to-noise ratio SNR_r^2 to the SNR^2 which would be obtained if the matching was performed with a template using the true signal parameters (when noise is not present in the data), as explained in subsection 3.1.4. To estimate this mismatch, the phase metric approximation is used [149]:

$$\mu \approx g_{ab}(\lambda)d\lambda^a d\lambda^b, \quad (5.7)$$

where g_{ab} is the parameter space metric (a and b run over the dimensions, given by the number of parameters) and λ represents the different parameters such as frequency or sky position. The phase metric is calculated as:

$$g_{ab} = \langle \partial_a \phi(\lambda) \partial_b \phi(\lambda) \rangle - \langle \partial_a \phi(\lambda) \rangle \langle \partial_b \phi(\lambda) \rangle, \quad (5.8)$$

where the factors inside $\langle \cdot \rangle$ are averaged over different SFTs during the observing time, and ϕ is given by equation (5.3).

The metric for directed binary searches was obtained in [257] and in [258]. We will use the equations obtained in [257] for the semi-coherent short-segment regime (equations 62), where $T_c \ll P_b$, which sets a lower limit for the orbital periods that we will be able to search. The resolution for the binary parameters is:

$$\delta a_p = \frac{\sqrt{6m_a}}{\pi T_c f \Omega} \quad (5.9)$$

$$\delta \Omega = \frac{\sqrt{72m_\Omega}}{\pi T_c f a_p \Omega T_{obs}} \quad (5.10)$$

$$\delta t_{asc} = \frac{\sqrt{6m_t}}{\pi T_c f a_p \Omega^2}, \quad (5.11)$$

where m_x are the mismatch parameters, which quantify the amount of lost SNR and define the desired spacing between templates. These equations were obtained for a coherent detection statistic, which instead of tracking the frequency-time evolution tracks the full $h(t)$ evolution given by equation (5.1). For this reason, in our case these mismatch parameters m_x will not correspond to an actual mismatch value and will only represent a tuneable parameter in our pipeline.

It can be seen that these equations depend on the region of binary parameter space for which they are calculated, as opposed to the resolution for the spin-down/up which only depends on the coherent time and the observing time [158]. For shorter periods (greater Ω) the resolution is increased for all three parameters, and the parameter space separation between templates is reduced. The total number of templates will be proportional to Ω^4 , which highly complicates the feasibility of the analysis for short periods.

In [257], the resolution was obtained for a directed search, which does not search over sky positions as all-sky searches do. For the frequency and sky positions we will use the resolutions obtained for the isolated *SkyHough* pipeline [158]:

$$\delta f = \frac{1}{T_c} \quad (5.12)$$

$$\delta \Theta = \frac{c}{v T_c f P_F}, \quad (5.13)$$

where Θ represents any of the two sky positions and the pixel-factor P_F is another tuneable mismatch parameter.

We have verified the scalings given by these equations through extensive simulations (shown in section 5.3), different mismatch values and across different regions of Ω and a_p .

Range of parameter space

The range in parameter space to be searched is primarily determined by the astrophysical prior information and by the available computational power:

- The frequencies to be searched are determined by the sensitive frequency band of the detectors and by the expected frequency of the emitted gravitational waves. Figure 5.1 shows that the maximum gravitational-wave frequency is around 1400 Hz. For the past O1 and O2 observing runs performed by the Advanced LIGO detectors, the most sensitive frequency band ranges roughly from 50 to 1000 Hz, with the best strain sensitivity occurring near 150 Hz.
- The range of binary orbital periods is bounded by the coherent time: periods lower than the coherent time cannot be distinguished one from another, and the equation for the period resolution was derived assuming $T_c \ll P_b$. The upper bound for the periods to be searched is mostly determined from the astrophysical prior information, where we can see that the maximum period is around 10^3 days.
- The minimum value of a_p is bounded by the minimum Doppler shift that we can observe. Figure 5.3 shows that for a given period, a minimum a_p value needs to be selected in order for the frequency modulation to be higher than 1 bin. If the modulation is less than 1 bin, we are not able to distinguish different templates, and pipelines which look for GW signals from isolated NSs can already discover them. The maximum a_p value to be searched is determined by the maximum amount of frequency bins that we can load at the same time, limited by RAM (Random Access Memory), and by the astrophysical prior information extracted from the known pulsar population, which figure 5.3 shows.
- The range in time of ascending node t_{asc} that needs to be searched is uniquely determined by the orbital period. Since we can redefine the time of ascension for every orbit by adding an integer times the orbital period, we can define it in the orbit which is closer to the mid-time of the search and we only need to search this area:

$$t_{\text{mid}} - \frac{P_b}{2} \leq t_{\text{asc}} < t_{\text{mid}} + \frac{P_b}{2}. \quad (5.14)$$

Number of templates

After having defined the spacing between templates and the ranges that have to be covered in each dimension, the total number of templates which are needed to carry out the search can be calculated.

The scaling of the total number of templates \mathcal{N} is given by (and \mathcal{N}_I for an all-sky search for isolated NSs):

$$d\mathcal{N} \propto T_c^6 T_{\text{obs}} f^5 \Omega^4 a_p^2 d\lambda, \quad (5.15)$$

$$d\mathcal{N}_I \propto T_c^4 T_{\text{obs}} f^2 d\lambda, \quad (5.16)$$

where $d\lambda$ represents a volume element for each of the search dimensions. For the binary case, the scaling with frequency is much steeper than for the isolated case, which will greatly increase the computational cost at higher frequencies.




Figure 5.3: The filled areas mark the regions in binary parameter space where the observed binary modulation is less than one frequency bin, for $T_c = 900$ s. Data taken from [64] and downloaded with [69].

Figure 5.4 shows the number of binary templates which need to be covered for different mismatch configurations and for searches at different regions of binary parameter space. With these numbers, the difference between a search for NSs in isolated systems and a search for NSs in binary systems becomes clear: while for the former around $\mathcal{O}(10^2)$ spin-down/up templates need to be searched (to cover a range wider than the astrophysical informed range), for the latter more than $\mathcal{O}(10^5)$ templates are needed to cover a narrow astrophysical range. Figure 5.4 also clearly shows that if we want to cover a broad range in frequency, the mismatch parameters will have to increase for higher frequencies, because a search with constant mismatch (like 0.4) is unfeasible otherwise. Another way to solve this issue would be to decrease the coherent time as the frequency increases.

5.2.3 | Maximum coherent time

The sensitivity of a semi-coherent method increases with the coherent time, so in principle one should aim to use coherent times as long as possible. On the other side, the computational cost depends on the coherent time, which sets a limit to this value. Furthermore, the spread of power to neighbouring bins (if the frequency of the signal occupies more than one frequency bin during one SFT) limits the maximum SFT time baseline that can be used (which for our method is equal to the coherent time).

To recover the maximum possible power from the signal, we have to avoid spectral leakage to neighbouring frequency bins. To achieve this, we demand that the signal be contained in half a single frequency bin (as was explained in subsection 3.3.1), which imposes a maximum coherent time:

$$\frac{\Delta f}{2} = \frac{1}{2T_c} \geq \dot{f}T_c \rightarrow T_c \leq \frac{1}{\sqrt{2|\dot{f}|_{max}}}. \quad (5.17)$$

We can estimate the maximum frequency derivative through the frequency evolution model from equation (5.5):

$$\dot{f} = f_0 \frac{\vec{a} \cdot \hat{n}}{c} + f_0 a_p \Omega^2 [\sin \Omega(t - t_{asc})], \quad (5.18)$$

where \vec{a} is the acceleration vector of the detector in the SSB frame. The highest contribution to the acceleration due to detector motion comes from the daily rotation of the Earth [158], which simplifies the




Figure 5.4: Number of binary templates for two searches covering different binary parameter spaces: for “HighP”, P_b ranges from 10 to 30 days and a_p from 9 to 25 seconds, while for “LowP” P_b ranges from 0.1 to 0.15 days and a_p from 0.05 to 0.08 seconds. For each frequency, two different mismatch parameters are chosen (shown with circles): 0.1 and 0.4 for 73.6 Hz; 1.6 and 6.4 for 436.9 Hz; 6.4 and 25.6 for 1200 Hz. The coherent time is 900 s, and the observation time is 11178584 s (the duration of the O1 run).

previous equation to:

$$|\dot{f}|_{max} = \frac{f_0 v^2}{c R_e} + f_0 a_p \Omega^2 = \frac{f_0 4\pi^2 R_e}{c T_e^2} + f_0 a_p \Omega^2, \quad (5.19)$$

where R_e is the radius of the Earth and T_e is a sidereal day. It can be seen that the maximum coherent time depends on the frequency f_0 and on the binary parameters a_p and Ω . Figure 5.5 shows some examples of these dependencies. It can be seen that lower coherent times are able to cover wider ranges of parameter space.

The optimal search strategy should use SFTs with different coherent time (the maximum allowed) in different regions of the binary and frequency parameter space. We note that the curves shown in figure 5.5 are conservative, since they are assuming the worst case scenario, and usually we will be in a more relaxed case where we could use longer coherent times. These assumptions are that the two frequency derivative terms are at their maximum values and they are aligned, and that we only allow half a frequency bin of variation, because we assume that the frequency of the signal is at the center of the bin, which almost never happens.

If the eccentricity is non-zero, two more factors would be present in equation (5.18). For eccentricities smaller than 0.01, it can be seen that these factors (proportional to e) would contribute much less to the frequency derivative. For this reason, we don’t take them into account in our previous estimation.

5.2.4 | Maximum eccentricity

Our final model for the frequency evolution, given by equation (5.5), assumes a zero-eccentricity orbit, but our pipeline remains fully sensitive to signals with a certain eccentricity if this eccentricity does not produce a noticeable change (less than a frequency bin) in the frequency-time evolution.

We can estimate the error in frequency tracking (by assuming that the eccentricity is exactly zero




Figure 5.5: Maximum coherent time allowed by equation (5.17) for four different choices of binary parameters. The black horizontal lines mark 1800 s and 900 s.

when it is non-zero) by subtracting equations (5.4) and (5.5):

$$|\Delta f(t)| = 2f_0 a_p \Omega \left(\frac{e \cos \omega}{2} \cos [2\Omega(t - t_{\text{asc}})] + \frac{e \sin \omega}{2} \sin [2\Omega(t - t_{\text{asc}})] \right). \quad (5.20)$$

The maximum error at any time will be:

$$|\Delta f|_{\text{max}} = e f_0 a_p \Omega. \quad (5.21)$$

We demand that this frequency difference is smaller than half a frequency bin, as was done in the previous subsection:

$$e f_0 a_p \Omega \leq \frac{1}{2T_c} \longrightarrow e \leq \frac{1}{2T_c f_0 a_p \Omega}. \quad (5.22)$$

If e is less than or equal to this quantity, the calculated frequency evolution will not deviate by more than half a frequency bin from the true evolution. Again, this expression depends on the region of the binary and frequency parameter space that we are in, so searches at different parts of this space can remain sensitive to different levels of eccentricity.

For most of the SFTs this error will be smaller (again, this is a conservative estimation), so this is a lower limit on the eccentricity that the orbit of the neutron star can have without producing any noticeable difference. Figure 5.6 shows that the zero-eccentricity assumption does not affect our ability to track systems with eccentricity smaller than 0.01 (for frequencies lower than 500 Hz for the worst case shown in that figure). In section 5.3 we will show some simulations of how the eccentricity affects the sensitivity of our method. An estimation of the sensitivity lost for signals with eccentricities higher than this is left for future work.

5.2.5 | Post-processing

After finishing the main steps of the pipeline, we are left with one toplist per dataset for each region in parameter space (i.e. each 0.1 Hz band), which contains the top templates ordered by a detection statistic. If the pipeline was run with option A, we will have multiple tolists, while if it was run with




Figure 5.6: Maximum eccentricity allowed by equation (5.22) for three different choices of binary parameters, with a constant choice of $T_c = 900$ s. The third and fourth traces overlap.

option B a single toplist will be the output. The following steps detail the procedure which goes from this point to the final list of outliers:

1. If we have multiple toplist, the first step consists of searching for coincidental pairs between these lists, by calculating the distance in parameter space and selecting the pairs which are closer than a certain threshold. The optimal value for the threshold cannot be found analytically, since its value depends on a balance between being more sensitive and having too many outliers. A reasonable value can be found by doing simulations. For each coincidental pair, its centroid (average locations in parameter space weighted by significance) is calculated. The distance d in parameter space is given by:

$$d^2 = \left(\frac{\Delta f}{\delta f}\right)^2 + \left(\frac{\Delta x}{\delta x}\right)^2 + \left(\frac{\Delta y}{\delta y}\right)^2 + \left(\frac{\Delta a_p}{\delta a_p}\right)^2 + \left(\frac{\Delta \Omega}{\delta \Omega}\right)^2 + \left(\frac{\Delta t_{\text{asc}}}{\delta t_{\text{asc}}}\right)^2, \quad (5.23)$$

where the numbers in the numerators represent the difference between two templates and the numbers in the denominators represent the parameter resolution. This distance is dimensionless and is given as a number of bins. The quantities x and y represent the cartesian ecliptical coordinates projected into the ecliptical plane.

When we compare two templates which are at different parts of the parameters space, the resolution of the binary parameters is different at these points. To calculate the distances of equation (6.4), we use a mean of the resolution at both points.

2. Independently of running the search with option A or option B, now we have a unique list. The next step consists of searching for clusters in this list. This will group different templates which can be ascribed to the same cause, and will reduce the number of final outliers. Again, we set a threshold in parameter space distance and find candidates which are closer than this distance. Clusters are found by analysing the distance of each template from all other templates, and keeping a list of indices of members with distances below the threshold. Each template can only be part of a cluster, so if a template was already in a cluster its newly generated cluster and the old one will be joined to form a unique cluster.

If the search is run with option B, another distance threshold between the top template in a cluster and all the other members is calculated, which eliminates all members which are further away than a certain threshold. This step is not needed with option A since the coincidence step already eliminates many candidates, but with option B the clusters can grow too wide and the final parameter estimation can be wrong unless a cut is made, due to a high number of cluster members being too distant from the true signal values.

3. The final post-processing step consists of calculating the centroid of each cluster. This is calculated as a weighted (by significance) average among all the members of the cluster. We keep the most significant cluster per 0.1 Hz band (if any), selected by the summed significance of all its members. This produces the final list of outliers of the search.

We cannot claim that these outliers represent a real astrophysical signal. As known from previous searches, instrumental noise or spurious coincidences can end up in the final list. As shown in figure 5.2, the final steps of any CW pipeline are the application of vetoes and follow-up procedures which increase the significance of the candidates and enhance the parameter estimation. Due to the extra parameters needed for a search for NSs in binary systems, follow-up procedures used in past searches may not be applied to this case. A derivation of a follow-up procedure which can increase the confidence on candidates from the *BinarySkyHough* pipeline is left for future work.

5.2.6 | Computational model

The sensitivity of CW searches is always limited by the available computational budget. Usually, a choice must be made between doing a broader search trying to cover a large binary and frequency parameter space, or a deep search which selects a small portion of the parameter space and has less mismatch, which increases the sensitivity of the search. In order to estimate *a priori* the cost that a search will have and to compare different setups, it is important to construct a computational model which can estimate the total computational cost and required RAM of a search given some regions of parameter space and resolution parameters.

Computational cost

The code spends the majority of the time in the inner-most loop over SFTs, sky positions and frequency and binary templates that is done at the first step of the search, and a loop over the SFTs and a subset N_C of the total number of templates at the second step of the search. To estimate the total cost of a search, we will characterize the scaling of these parts of the code by running it with different search parameters.

The total cost will be slightly bigger than this estimation, due to other tasks like I/O and initialization of variables, but this extra cost is negligible. Furthermore, we are going to estimate the cost for frequency bands which are nearly Gaussian (i.e. don't contain signals or excessive instrumental noise), and we will assume that this will be valid for the vast majority of analysed frequency bands.

The computational cost can be estimated as:

$$\mathcal{C} = N_B N_D N_F \sum_{i=0}^N N_{SP;i} (C_{1;i} + C_{2;i}), \quad (5.24)$$

where N_B is the number of blocks of binary templates that must be analysed to cover the entire range of binary parameters, N_D is the number of datasets (assuming they have the same number of SFTs or taking the maximum between all of them), N_F is the number of frequency templates in each frequency band (the same number for each frequency band, e.g. 90 for a 0.1 Hz band with $T_c = 900$ s), and the summation goes over the number of frequency bands in which we split the total frequency range. To cover a wide range of binary parameters, more than one block of binary templates is usually needed, since

the maximum number of binary templates that can be searched at once is limited by RAM constraints, as will be seen in subsection 5.2.6.

The number of sky-patches in the i th frequency band $N_{SP;i}$ can be estimated as:

$$N_{SP;i} \approx \frac{4\pi}{N_X N_Y \delta\Theta^2} = \frac{4\pi(v/c)^2 P_F^2 T_C^2 f_i^2}{N_X N_Y}, \quad (5.25)$$

being N_X and N_Y the number of sky pixels in each direction.

The cost $C_{1;i}$ of the first step of the search per one sky-patch and one frequency template is given by:

$$C_{1;i} = C_{B1} N_T N_{SFTs} N_X N_Y \rho_F, \quad (5.26)$$

where N_{SFTs} is the number of SFTs, N_T is the number of binary templates in one binary block at the i th frequency band, ρ_F is a factor which controls the scaling with the selected peak threshold and C_{B1} is the cost of running over 1 binary template when there is one SFT and one sky position. C_{B1} is calculated when the threshold is $\rho_{th} = 1.6$, and $\rho_F = \frac{e^{-\rho_{th}}}{e^{-1.6}}$ is a simple scaling factor which takes into account the different number of peaks that are present (the exponentials appear since the distribution of powers for a Gaussian band is $p(\rho_k) = e^{-\rho_k}$) when the threshold is changed.

The cost $C_{2;i}$ of the second step of the search per one sky-patch and one frequency template is given by:

$$C_{2;i} = C_{B2} N_C \mathcal{N}_i \bar{N}_{SFTs}, \quad (5.27)$$

where \mathcal{N}_i is the total number of templates at the i th frequency band, \bar{N}_{SFTs} is the number of SFTs used at the second step and C_{B2} is the cost of the second step per template and per SFT.

The code is run with a CPU and GPU. The code for the GPU execution is written in CUDA, which has some parameters that can be changed (like the number of blocks and the number of threads per block) which affect these estimations. Another important source of uncertainty is the different hardware layouts between different GPUs, like the different number of cores. These differences do not affect the predicted scalings given by the previous equations.

We have done several runs to test the different scalings. We have used two different GPUs, and we also show a comparison by using only a CPU instead of a CPU+GPU. The results are shown in table 5.1. The listed configurations on blocks and threads for the GPUs are the ones that have given better results. It can be seen that without using a GPU card this search would be unfeasible.

From these results we can estimate the cost that a complete search would have. For 1 binary block of 3×10^5 binary templates, with 1 dataset, 90 frequency bins per 0.1 Hz covering 400 Hz, and 50 sky-patches per 0.1 Hz band, a search with configuration run 2 would need approximately 5000 hours. This assumes that the number of binary templates would be the same in each frequency band, which requires that the mismatch parameters are lowered as the frequency is increased.

If we want the binary resolution to remain constant across the frequency range, the number of binary blocks would increase with frequency. This number would also be greater than 1 if we wanted to cover a larger range of binary parameters. With the same configuration as before but with 500 binary blocks, the cost would increase to 2.5×10^6 hours. This order of magnitude is usual within all-sky semi-coherent searches, and is comparable to the cost of the only published all-sky search [197]. Even though this method explicitly searches over t_{asc} , which *TwoSpect* does not, the costs are comparable due to the usage of GPUs and the look-up table approach. With 500 binary blocks we could cover a large parameter space, covering all the astrophysical interesting regions, or we could do a narrower search with low mismatch.

Random Access Memory (RAM)

In order to characterize the RAM required by our pipeline, a calculation of the number of bytes taken by every data structure should be made. Of the many data structures in the code, two of them are many orders of magnitude larger than the rest and are enough to give a rough estimate of the memory required.

Hardware	Run 1		Run 2		Run 3		Run 4		Run 5		Run 6	
	$C_{1,i}$	$C_{2,i}$	$C_{1,i}$	$C_{2,i}$	$C_{1,i}$	$C_{2,i}$	$C_{1,i}$	$C_{2,i}$	$C_{1,i}$	$C_{2,i}$	$C_{1,i}$	$C_{2,i}$
CPU + Tesla V100	0.15	0.11	0.41	0.30	0.80	0.54	0.45	0.35	2.2	1.8	0.07	0.05
CPU + GTX 1050Ti	1.7	2.1	4.5	5.6	8.5	11.0	6.6	7.8	–	–	0.8	1.1
CPU	122	316	–	–	–	–	–	–	–	–	–	–

Table 5.1: $C_{1,i}$ and $C_{2,i}$ timings (in seconds) for different run configurations. Each number is the mean over 500 runs with the same configuration. Both GPUs are used with 512 blocks and 64 threads per block. The CPU used is an Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz. The compilation of the code was done with gcc and nvcc, with option -O3. All runs use $N_C = 0.05$.

Run 1: $N_T = 1 \times 10^5$; $N_X N_Y = 49$; $N_{SFTs} = 13304$; $\bar{N}_{SFTs} = 25930$. Run 2: $N_T = 3 \times 10^5$; $N_X N_Y = 49$; $N_{SFTs} = 13304$; $\bar{N}_{SFTs} = 25930$. Run 3: $N_T = 6 \times 10^5$; $N_X N_Y = 49$; $N_{SFTs} = 13304$; $\bar{N}_{SFTs} = 25930$. Run 4: $N_T = 1 \times 10^5$; $N_X N_Y = 169$; $N_{SFTs} = 13304$; $\bar{N}_{SFTs} = 25930$. Run 5: $N_T = 6 \times 10^5$; $N_X N_Y = 169$; $N_{SFTs} = 13304$; $\bar{N}_{SFTs} = 25930$. Run 6: $N_T = 1 \times 10^5$; $N_X N_Y = 49$; $N_{SFTs} = 7235$; $\bar{N}_{SFTs} = 14220$.

One of these structures is related to the PHMDs, and has a size in bytes of:

$$S_a = 6N_{SFTs}KN_XN_Y, \quad (5.28)$$

where K is the number of PHMDs needed in the frequency axis (see figure 3.6), equal to the number of searched frequency bins plus the maximum modulation produced by the BB Doppler modulation.

The other large structure holds the results of the first step of the search, and the size in bytes is given by:

$$S_b = 8N_TN_XN_Y. \quad (5.29)$$

With these expressions and a number of N_{SFTs} to be analysed, we can calculate the RAM for different number of binary templates and different number of PHMD bins given by different frequency band sizes.

5.3 | Sensitivity estimation

This section presents a characterization of the sensitivity of the *BinarySkyHough* pipeline. To do this, we add many simulated signals to real or simulated noise in a Monte-Carlo way and we run the pipeline with this data as input. We determine the number of detected signals, and we evaluate the parameter estimation obtained. Simulations are used because an analytical estimation of the sensitivity of a pipeline which takes into account all the steps of the procedure cannot usually be obtained. This is a widely used procedure and it has been used in many past searches such as [110] or [203].

The purpose of this section is twofold: to estimate the sensitivity of the pipeline, and to see how it changes by varying different internal parameters such as the mismatch or the fraction of templates N_C which go to the second step of the search.

5.3.1 | Procedure

We have added signals (usually called injections) into the O1 Advanced LIGO data of detectors H1 and L1 using the commonly used LALSuite code *lalapps_Makefakedata.v5* [51]. We have used 3 different 0.1 Hz bands: 73.6, 170.2 and 436.9 Hz (the signals have a random frequency within each 0.1 Hz) and four different parts of binary parameter space, indicated in table 5.2. We use different levels of amplitude h_0 , selected to have certain sensitivity depths:

$$\mathcal{D} = \frac{\sqrt{S_n}}{h_0}. \quad (5.30)$$

Name	a_p [s]	Period [days]
BS1	0.03 – 0.08	0.1 – 0.101
BS2	0.5 – 1.5	1 – 1.01
BS3	3 – 13	10 – 20
BS4	20 – 35	30 – 90

Table 5.2: Four ranges of binary parameters used for the simulations.

We have used 4 sensitivity depths (14, 18, 22 and 26 $\text{Hz}^{-1/2}$) in order to be close to the 95% efficiency point, a percentage usually used to ascertain the sensitivity of a search method, with 100 different signals per sensitivity depth. Other amplitude parameters like cosine of inclination, initial phase and polarisation are drawn from a uniform distribution (producing signals with random polarizations). We have used a coherent time of 900 s for all the studies presented here. The injected signals are isotropically distributed in the sky, with random argument of periapsis ω and with eccentricity drawn from a log-uniform distribution between 10^{-6} and 10^{-2} .

In a real search the number of templates which get into the final toplist is limited, and this sets an artificial threshold on the significance of templates which can be detected (if a signal produces a detection statistic with a value lower than this threshold, it won't be present in the final toplist). Before analysing the injections (which are analysed in a reduced region around its real parameters, of around 20 bins in each dimension), we run an all-sky search without added signals with the same configuration parameters (parameter space resolution, N_C , etc.) to obtain this threshold, and we apply it when we analyse the injections, thus ensuring a fair and realistic analysis. The number of candidates per injection that we keep in the final toplist is 5000, the same number that is used for obtaining the threshold in detection statistic.

For each group of 100 signals at each sensitivity depth, we calculate the efficiency, defined as the number of detected signals divided by the number of injected signals, which will be the main indicator of the method's sensitivity. To count an injection as detected, we demand that its final parameters estimated from the selected cluster are within 13 bins of the true parameters, a number which has been used in past analyses (as will be shown later, most injections are recovered at less than two bins away).

All the results shown in this section use a coincidence window of 3 bins and a clustering window of $\sqrt{14}$ bins. These sizes are similar to the ones that were used for the isolated-star O1 and O2 analysis, and we leave for future work a proper characterization of the effect that these sizes have on the sensitivity and parameter estimation. All the injections have been analysed with a threshold in power of $\rho_{th} = 1.6$.

Before discussing the results, we want to remark that the efficiency or the 95% sensitivity depth are not the unique indicators of the value of a pipeline. Other factors, such as the range in parameter space which can be covered (the computational cost), the robustness to deviations of the signal from the model or to noise artifacts from the detectors, or the parameter estimation are also important indicators.

5.3.2 | Results

We have analysed the simulations by running the pipeline with varying parameters, such as mismatch, and we compare the results obtained in order to get a general view of the sensitivity which this pipeline can achieve. The plots are shown without error bars in order to ease viewing the results, but all these efficiency points should have a vertical error bar equal to $\sqrt{\frac{E(1-E)}{100}}$, where E is the efficiency.

Firstly, figure 5.7 shows a comparison between different parts of binary space and two different frequencies. These four runs share the same mismatch parameter of $m = 0.4$. All these choices of binary parameters have approximately the same number of templates, and they all produce frequency modulations wider than one frequency bin. We can observe that all of them have a 95% sensitivity depth above




Figure 5.7: Efficiency versus sensitivity depth at different parts of the binary parameter space and two different frequencies. All injections have been analysed with option B (H1+L1 data), $N_C = 0.05$, $m_x = 0.4$ and $P_F = 1$.

$14 \text{ Hz}^{-1/2}$. For the first three sensitivity depth points, all efficiencies are comparable. We can begin to see a wider spread at the last point, where the injections at a higher frequency band show the worse sensitivity.

Secondly, figure 5.8 shows a comparison of runs with different resolution parameters, for the binary space 2. For the first and second sensitivity depth points all efficiencies are very similar. A noticeable decrease in efficiency for coarser resolutions only begins to take place at the last two sensitivity depth points. Running with coarser resolutions also affects the parameter estimations results, effect that we later discuss. The results for the other binary spaces have also been obtained and show similar scalings to the ones shown in this figure.

Lastly, figure 5.9 shows the efficiency obtained by comparing runs with different N_C and comparing option A with option B. The figure shows that running the pipeline with option B (all datasets together) gives a better efficiency than option A. We also observe an increase in efficiency when increasing the fraction of templates which go to the second step of the pipeline, but the improvement is not significant until the last sensitivity depth point. A comparison with using Gaussian data (of the same noise level) instead of O1 data is also shown. The results for Gaussian noise have been obtained by averaging over 10 different realizations of Gaussian noise. It can be seen that for these Gaussian noise realizations, results are only greatly improved at the third sensitivity depth point, with no significant improvements at the other points. These results have been also obtained for other mismatch configurations and other regions of binary parameter space and they show similar behaviour.

These tests have been done with O1 Advanced LIGO data, using data from two different detectors. The sensitivity of semi-coherent methods improves with longer observation times and by using data from more detectors, so these results should be placed in this context. Furthermore, we have analysed four sensitivity depth points, due to the high computational cost of doing more simulations, which provide only a partial estimation of the full dependence of efficiency with sensitivity depth. Extrapolating from the results presented, we can argue that the difference in efficiency between different run configurations grows wider as the sensitivity depth is increased (i.e. as the amplitude of the signal gets smaller), but the pipeline seems to have a minimum 95% efficiency at $\mathcal{D} = 14 \text{ Hz}^{-1/2}$ for all the different tests that we have done.




Figure 5.8: Efficiency versus sensitivity depth for different mismatch configurations. All injections have been analysed at the binary space 2, with option B (H1+L1 data) and $N_C = 0.05$.

Figure 5.10 shows some examples of the parameter estimation that this pipeline can achieve by comparing different mismatch configurations. It shows the results for detected simulations from the binary space 2. The errors in parameter estimation are estimated as the mean of the absolute value difference between the final cluster parameters and their true value for each injection. We observe that the different parameters show different behaviour: the configuration run with the worst estimation is not the same for all parameters. For the binary parameters, it is interesting to notice that the worst mismatch configuration usually has the best parameter estimation, both in bins and in natural units. This might be related to the frequency: for higher frequencies, the binary modulation becomes wider and the parameters can be better estimated. Comparing runs at the same frequency but with different mismatch (between orange circles and stars, or orange crosses and sums), it can be seen that the run with highest mismatch is always above the run with lower mismatch, as it should be.

5.3.3 | Comparison with other methods

With the previous results we can translate from the sensitivity depth points at which the efficiency is 95% to the estimated $h_0^{95\%}$ sensitivity. As discussed previously, for all different configurations the 95% sensitivity depth is always at least at $14 \text{ Hz}^{-1/2}$, so we will take this as our sensitivity.

We can compare this result with the *SkyHough* result for the O1 analysis, which for the low-frequency range (from 50 to 475 Hz) was around $24 \text{ Hz}^{-1/2}$ [201]. The difference between these two sensitivities can be explained by these facts:

- The *SkyHough* search used a coherent time of 1800 s, which is twice as long as what we have used in this analysis. Longer coherent times can be used for searches for isolated systems since the frequency derivative is usually smaller. Doubling the coherent time can produce an increase in sensitivity up to approximately 1.2, which explains a fraction of the difference.
- The search for NSs in binary systems has many orders of magnitude more templates. This increases the maximum values of the detection statistic for the background distribution (only noise), which has the effect of raising the thresholds in detection statistic needed to have the same false alarm




Figure 5.9: Efficiency versus sensitivity depth for different run configurations. N_C shows the fraction of total templates which are passed to the second step of the search. All injections have been analysed at the binary space 2, with $m_x = 0.4$ and $P_F = 1$.

probability. This effect produces a decrease of the search sensitivity, since a signal which generates the same detection statistic may be detected in one search but not in the other.

- The isolated search was run with $P_F = 2$, while we have used $P_F = 1$. This could also explain a fraction of the difference in sensitivities.

The *TwoSpect* method described in [98] is the only pipeline which has been used in an all-sky search for neutron stars in binary systems. From published results of a search with S6 data [197], we can estimate a sensitivity depth at 95% efficiency of 5 for isotropically oriented neutron stars. Taking into account the improvements developed in [259] and the difference in observation time between S6 and O1, an estimated sensitivity depth with O1 data is around $6.5 \text{ Hz}^{-1/2}$, which is approximately half as sensitive as our pipeline. Furthermore, with a high number of binary blocks (e.g. around 500) our pipeline is able to cover a similar parameter space as the one that the *TwoSpect* pipeline has covered in the mentioned search, with a comparable computational cost.

Recently, an adaptation of the radiometer search for all-sky searches has been proposed [255]. This unmodeled search looks for coherences between two or more detectors, by tracking a signal which is inside a single frequency bin all the time. This search has a very cheap computational cost, but because it is unmodeled and the frequency bins are much coarser (1 Hz, as compared to our 1/900 Hz bins), their sensitivity is worse than our pipeline. In [255], a sensitivity of 1.2×10^{-24} at a frequency of 245 Hz for an O1 search is quoted (with a 90% confidence, compared to our 95%, for signals with circular polarisation and by using Gaussian data instead of more realistic data (i.e. from an observing run)). These facts make a direct comparison difficult, but we can convert this value to a sensitivity depth, and try to make a rough comparison. By dividing the quoted 1.2×10^{-24} value by the amplitude spectral density at 245 Hz, we get a value of $7 \text{ Hz}^{-1/2}$. A realistic factor to convert this estimation to one with 95% confidence, from isotropically oriented neutron stars and with realistic data is complicated to calculate, but it can be seen that our pipeline still remains at least twice as sensitive.

These comparisons are shown in figure 5.11 (only *TwoSpect* is shown). We remark that improving the sensitivity by two means that we are able to detect signals from systems twice as far away as before,

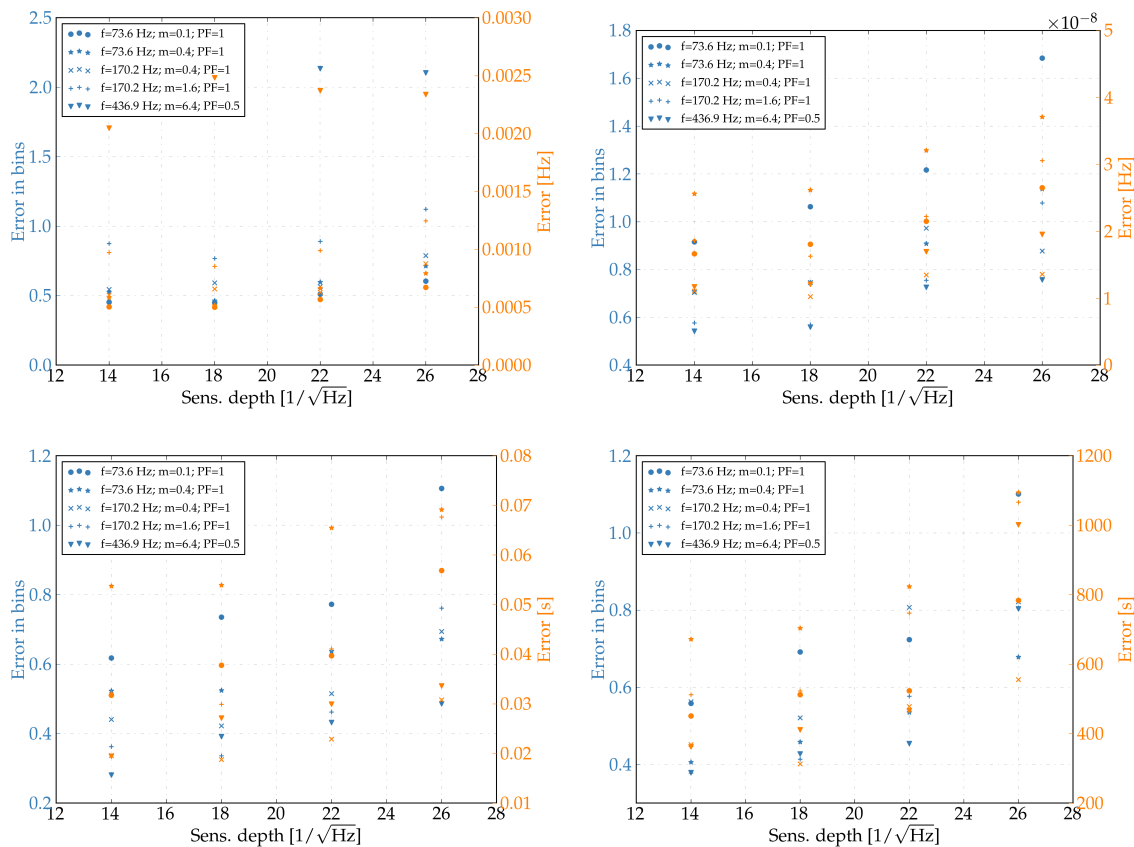


Figure 5.10: Figures showing the parameter estimation of the new method for simulations in the binary space 2, with option B and $N_C = 0.05$. Top left: frequency. Top right: angular frequency Ω . Bottom left: a_p . Bottom right: t_{asc} . The left vertical axis shows the mean absolute error of detected signals in number of bins, while the right vertical axis shows the mean absolute error of detected signals in the correspondent units.

or from neutron stars with asymmetries two times smaller at the same distance. A comparison of the parameter estimation between the different pipelines has not been possible and we leave it for future work.

It is also interesting to notice the difference between *BinarySkyHough* and other methods designed to perform a directed search (known sky position) for a signal from a NS in a binary system. From the published O1 results [186], we estimate a sensitivity depth of $30 \text{ Hz}^{-1/2}$, which is approximately twice our sensitivity. Since these methods don't have to search for different sky positions, the computational power can be spent in using methods which can increment the coherent time or in decreasing the mismatch parameters for the same coherent time.

5.4 | Conclusions

We have described a new method to perform all-sky searches of continuous gravitational waves from neutron stars in binary systems. This method is an extension of the *SkyHough* pipeline, which has been coded to take advantage of GPU cards in order to overcome the prohibitive computational costs that this pipeline would have by only using CPUs, as demonstrated in section 5.2.6.

Simulations indicate that this new method is at least 2 times more sensitive than previous pipelines, at a comparable computational cost. It can be used to search for signals in any part of the binary parameter space, showing comparable sensitivity at all locations of parameter space. With the simulations done, it seems that regardless of the mismatch parameter chosen, the pipeline has a minimum sensitivity depth




Figure 5.11: Estimated h_0 sensitivity at 95% confidence to random polarised signals with Advanced LIGO O1 data compared to the SkyHough all-sky search for isolated NSs for O1 data and to the TwoSpect pipeline.

of $14 \text{ Hz}^{-1/2}$ (for a coherent time of 900 s and an observation time equal to O1).

A possibility to improve the sensitivity of this pipeline would be to use cleaned data as input data. Some procedures can remove time-domain disturbances which affect some of the generated SFTs, and these cleaned SFTs can improve the sensitivity of a search with no additional computational cost, as shown in [248]. Another improvement in sensitivity could come from the generation of modified SFT bins, which take into account the leakage of power due to the signal not being at exactly the center of a frequency bin [136]. Furthermore, coherently combining data from different detectors as explained in [259] may also improve the sensitivity of the pipeline.

An optimal way of spending a limited computation budget should be developed in order to maximize the chances of detecting a signal. An analytical estimation of how each parameter of the pipeline affects the sensitivity and the computational cost would need to be estimated, but this development would increase the possibilities of detection.

The *BinarySkyHough* pipeline could also be used to perform the first all-sky search which explicitly looks for NSs in high-eccentricity systems. Two more parameters (e and ω) should be included, and this would further increase the computational cost of the search, narrowing the range of parameter space which could be searched. This pipeline (with some modifications) could also be applied for a directed search of a NS in a binary system, where the sky position is known but the frequency and binary parameters are unknown. By eliminating two parameters of the search, the computational cost could be used for analysing a broader frequency or binary range or to use lower mismatch parameters.

We plan to keep developing this pipeline to prepare it to analyse the upcoming O3 observing run, and we also plan to apply it to a search using data from the O2 Advanced LIGO observing run (explained in the next chapter).

FIRST ALL-SKY SEARCH FOR CONTINUOUS GRAVITATIONAL-WAVE SIGNALS FROM UNKNOWN NEUTRON STARS IN BINARY SYSTEMS USING ADVANCED LIGO DATA

In this chapter we present a search for continuous gravitational waves from unknown neutron stars in binary systems with orbital period between 15 and 45 days. This is the first time that Advanced LIGO data and the recently developed *BinarySkyHough* pipeline have been used in a search of this kind. No detections are reported, and upper limits on the gravitational-wave amplitude are calculated, which improve the previous results by a factor of 17. This chapter is based on the material presented in [260].

6.1 | Introduction

Continuous gravitational waves (CWs) are non-transient and nearly monochromatic gravitational waves (GWs). Neutron stars can emit CWs through a variety of mechanisms, such as rotation with elastic or magnetic deformations (which may be sourced from accretion by a companion), unstable r-mode oscillations, or free precession (see [247] or [131] for a recent discussion of different emission mechanisms). Close to their core, these stars have density values equal or higher than the nuclear density, which make them valuable objects to study the unknown equation of state. Continuous waves also present an opportunity to test deviations from general relativity, like searching for extra polarizations of the waves [242] or finding differences between the speed of GWs and the speed of light [241]. Several searches for CWs, both from neutron stars in isolated and binary systems, have been previously carried out (see [224] for a recent review of CW searches), although none have conclusively detected a CW signal. Nonetheless, interesting upper limits have been produced which already help to constrain some models of neutron star shape [101].

All-sky searches look for emission of CWs from unknown neutron stars in our galaxy, and complement the targeted searches which focus on CWs from known pulsars. Since only a small percentage of the estimated neutron star population has been detected as pulsars, carrying out all-sky searches is important because such a search could discover emission by highly asymmetric neutron stars that have not been detected electromagnetically as pulsars. These searches need to calculate the Doppler modulation (produced by Earth's rotation and orbit around the Sun) for many sky positions, making their computational cost orders of magnitude higher than the cost of a targeted search. For this reason, the most sensitive methods like matched filtering cannot be used and semi-coherent methods that split the full observation time in smaller chunks (which are incoherently combined) are routinely used. Semi-coherent methods do not recover as much signal-to-noise ratio as coherent methods do, but the number of templates that

need to be searched over in order to constrain the maximum mismatch between signal and template is greatly reduced, thus highly decreasing the computational cost of the search. A recent comparison between different semi-coherent methods is shown in [230].

All-sky searches for neutron stars in binary systems pose an even more difficult problem, since the parameters that describe the orbit around the binary barycenter also need to be included in the search parameters. These searches are valuable, since approximately half of the known pulsars with rotational frequencies above 20 Hz belong to binary systems. Until recently, there was only one mature semi-coherent pipeline which could carry out this type of search, called *TwoSpect* [98]. This pipeline has been used once in a search for CW signals using the S6 and VSR2-3 datasets [197], reporting no detections.

Recently, we developed a new pipeline called *BinarySkyHough* (BSH) [254] (which was explained in the previous chapter). This pipeline is an extension of the semi-coherent *SkyHough* pipeline [158], which has been used in many past all-sky searches. It replaces the search over the spin-down/up parameter of isolated sources for the three binary orbital parameters characterizing different possible circular orbits. As explained in [254], this is computationally achievable due to both the usage of the massive parallelization which GPUs (Graphical Processing Units) provide and the computational advantages employed by *SkyHough*. Initial tests indicate that the BSH pipeline can provide roughly two times more stringent upper limits, although these tests have been done over a smaller parameter space. In this chapter we present the first application to real data of this new pipeline. No detections are reported, but the improved quality of the datasets and the new pipeline allows us to improve the upper limits by a factor of 17.

A neutron star with an asymmetry around its rotation axis emits CWs, which produce a time-dependent strain that can be sensed with interferometric detectors, given by equation (5.1). The amplitude of this signal is given by [50]:

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{I_{zz} \epsilon f^2}{d}, \quad (6.1)$$

where d is the distance from the detector to the source, f is the gravitational-wave frequency (equal to two times the rotational frequency), ϵ is the ellipticity or asymmetry of the star, defined by $(I_{xx} - I_{yy})/I_{zz}$, and I_{zz} is the moment of inertia of the star with respect to the principal axis aligned with the rotation axis.

The time-dependence of the gravitational-wave frequency is given by [254]:

$$f(t) = f_0 + f_0 \frac{\vec{v}(t) \cdot \hat{n}}{c} - f_0 a_p \Omega \cos[\Omega(t - t_{asc})], \quad (6.2)$$

where $\vec{v}(t)$ is the velocity vector of the detector, \hat{n} is the vector pointing to the source, f_0 is the gravitational-wave frequency defined at some reference time, and a_p , Ω and t_{asc} respectively represent the projected semi-major axis amplitude (in light-seconds), angular frequency of the binary orbit and time of ascending node (the three parameters describing the binary orbit). This is the frequency-time pattern that we search, which depends on six unknown parameters that need to be explicitly searched over: f_0 , α (right ascension), δ (declination), Ω , a_p , and t_{asc} .

This model assumes a circular binary orbit, but as discussed in the previous chapter, the main BSH search remains fully sensitive to signals with eccentricity less than 10^{-2} . The model also assumes that the neutron star does not suffer any glitches during the observing time, and that the effect of spin-wandering (stochastic variations on the rotational frequency due to the accretion process) as estimated in [100], if present, can be neglected. Although we don't explicitly search over a spin-down/up parameter, this search is sensitive to sources with spin-down/up approximately up to $(T_c T_{obs})^{-1} = 4.8 \times 10^{-11}$ Hz/s (where $T_c = 900$ s is the coherent time and $T_{obs} = 23170808$ s is the time span of the datasets), since sources with this value or lower wouldn't change the frequency-time pattern by more than a frequency bin, thus not producing any observable change. All known pulsars in binary systems have spin-down values lower than this quantity, as can be seen in figure 5.1.

6.2 | Search

To perform the main search we use the BSH pipeline [254]. The full Advanced LIGO [28] O2 dataset [261] (publicly available in [1]) is used, comprised of data from the H1 (Hanford) and L1 (Louisiana) detectors without segments that contain epochs of extreme contamination (the used segments are listed in [251], where the files with the “all” tag are selected). The O2 run started on November 30 2016 and finished on August 25 2017. The H1 detector suffered from jitter noise, and a separate data stream (which we use) that removes this contamination was created in order to improve the amplitude spectral density of the detector (more details are given in [250]). The H1 and L1 datasets include artificially added signals, called hardware injections, which help to test the performance of the detectors and the sensitivity of the different search algorithms (although no hardware injections with binary orbital modulation are present). The parameters of the hardware injections are given in table 4.5. Furthermore, these datasets contain several lines and combs, described with more detail in [249] and chapter 7. These disturbances, usually narrow in frequency, are problematic because they can imitate and/or mask the signals we are looking for, thus lowering the sensitivity of our pipeline.

6.2.1 | Main search

The input data, described as a signal plus additive noise $x(t) = h(t) + n(t)$, is converted to the frequency-domain and kept as a collection of “Short Fourier transforms” (SFTs). Each of these SFTs has a coherence time T_c of 900 s, in order to constrain the gravitational-wave signal in a single frequency bin and not lose power to neighbouring bins (due to the two orbital modulations that affect the searched signal) [254]. From these constraints, 14788 and 14384 SFTs from H1 and L1 are obtained, making a total of $N_{SFTs} = 29172$ which are analysed together.

Table 6.1 shows the parameter space that has been searched. We split the search in frequency bands of 0.1 Hz, each of these covering all the sky and the full range of binary orbital parameters. The resolution for each of these parameters is given by [254]:

$$\delta f_0 = \frac{1}{T_c}, \quad \delta \Theta = \frac{c}{v T_c f P_F}, \quad \delta a_p = \frac{\sqrt{6m}}{\pi T_c f \Omega}, \quad \delta \Omega = \frac{\sqrt{72m}}{\pi T_c f a_p \Omega T_{obs}}, \quad \delta t_{asc} = \frac{\sqrt{6m}}{\pi T_c f a_p \Omega^2}, \quad (6.3)$$

where $v/c = 10^{-4}$, Θ represents both right ascension and declination, m is a parameter which controls the resolution of the binary parameters, and P_F the resolution of the sky position parameters. Different values for m and P_F (shown in table 6.2) are selected depending on the frequency, in order to have a manageable Random Memory Access usage and a nearly constant number of templates per 0.1 Hz band across the frequency range.

For each of these bands the main search returns a list with a percentage of the most significant templates ordered by a detection statistic. The main BSH search is divided in two main stages which use different detection statistics (more details are given in the previous chapter). The top 5% templates in each 0.1 Hz band go to the second stage, and the final toplist only contains the best 0.001% of the templates passed to the second stage. The second stage of the search uses a complementary set of SFTs, which is generated from the initial set by moving the initial time of each SFT by $T_c/2$ and creating a new SFT at each new timestamp (if a contiguous set of data of T_c seconds exists). This procedure slightly increases the sensitivity of the procedure as explained in [206].

After running the main search, a clustering procedure is applied to the returned toplist (i.e. option B of the procedure shown in figure 5.2). This procedure improves the parameter estimation and allows us to reduce the number of candidates that need to be followed-up. For this search we use a clustering distance threshold of $\sqrt{14}$ (as used in past searches), where the distance is defined as:

$$d^2 = \left(\frac{\Delta f}{\delta f} \right)^2 + \left(\frac{\Delta x}{\delta \theta} \right)^2 + \left(\frac{\Delta y}{\delta \theta} \right)^2 + \left(\frac{\Delta a_p}{\delta a_p} \right)^2 + \left(\frac{\Delta \Omega}{\delta \Omega} \right)^2 + \left(\frac{\Delta t_{asc}}{\delta t_{asc}} \right)^2. \quad (6.4)$$

Parameter	Start	End
Frequency [Hz]	100	300
Right ascension [rad]	0	2π
Declination [rad]	$-\pi/2$	$\pi/2$
Period [day]	15	45
a_p [s]	10	40
Time of ascension [s]	$t_{mid} - P/2$	$t_{mid} + P/2$

Table 6.1: Ranges of the different searched parameters. Period P is given by $2\pi/\Omega$, and t_{mid} is the mean between the starting and ending times of the datasets.

Frequency range	m	P_F
[100, 125)	0.4	1
[125, 150)	0.8	1
[150, 200)	1.4	1
[200, 250)	2.4	0.75
[250, 300)	3.4	0.75

Table 6.2: Resolution parameters at different frequency ranges.

Quantities in the denominator represent the resolution in each dimension given by equations (6.3), and x and y are the Cartesian ecliptic coordinates projected in the ecliptic plane. Clusters are found by calculating the distance between all templates, and keeping a list with indices of members with distances below the threshold. Afterwards, the center of each cluster is found as a weighted (by power significance) sum for each of the six parameters. We keep the 3 most significant clusters per 0.1 Hz band, ordered by the maximum detection statistic value of each cluster, only keeping clusters which have at least 3 members. This produces the list of 6000 outliers from the main search, coming from the 2000 frequency bands.

The next step consists of applying vetoes to these outliers in order to eliminate the ones produced by non-astrophysical sources. The first veto that we apply is the *lines veto*, used in many past searches such as [213]. This veto calculates the frequency-time pattern for each outlier and checks if it crosses any frequency where there is a known line or comb, listed in [249]. After applying this veto only 4937 outliers remain.

6.2.2 | Follow-up

In order to follow-up these outliers, we use the strategy of repeating the search in multiple steps with an increased coherence time (still using a semi-coherent approach) and a reduced range of parameter uncertainty. If the outlier is produced by a real astrophysical signal, the detection statistic will keep increasing, while the same behaviour is not expected for Gaussian noise. The multi-detector \mathcal{F} -statistic (the frequentist maximum-likelihood statistic), derived in [50] and [262], can be used to perform searches with longer coherence times without losing power to neighbour frequency bins. The computational cost of a gridded \mathcal{F} -statistic search over six parameters for such a long dataset would be too high, and for this reason we need to use a method with stochastic placement of templates.

The procedure outlined in [144] and [263] consists of using a tempered ensemble walker MCMC algorithm (called *ptemcee* [264]) to draw samples of the \mathcal{F} -statistic and converge to the true signal parameters. To use this procedure, the coherent time and the width in each dimension around the cluster center that we want to follow need to be selected. Wider regions will achieve a higher rate of detected




Figure 6.1: The left plot shows $2\mathcal{F}_{sum}$ values for all the 4937 outliers which have been followed-up (the red dashed line marks the threshold at 2000). The right plot shows $2\mathcal{F}_{sum}$ values for all injections done at 10 different frequency bands, with 4 different sensitivity depth values per frequency band (the total number of injections is 4573, and 9 of them are above 2000).

signals, since the centers of the clusters can be located at several bins from their true location, but they will incur in higher computational costs because to reach convergence of the MCMC algorithm the number of steps and/or walkers have to be increased. The same happens with T_c : longer times are able to achieve higher sensitivity, but they require more steps to converge.

The behaviour of the follow-up is characterized by adding simulated signals (called injections) to the datasets. We use 4573 injections at 4 different values of h_0 and 10 different non-disturbed frequencies. The h_0 values are located near the 95% detection efficiency point, which is derived later. Firstly, we run BSH to obtain the clusters for each injection, and then we follow-up with $T_c = 60000$ s (the number of segments is 387) the injections whose cluster's centers are within 5 bins of the true parameters (the injections which count as detected by BSH). All but 9 injections are recovered with a semi-coherent \mathcal{F} -statistic value $2\mathcal{F}_{sum}$ of more than 2000. This value is used as a threshold for the follow-up of the outliers, which implies a false dismissal of $9/4573 = 0.1\%$.

From the 4937 outliers, only 27 have $2\mathcal{F}_{sum}$ values above the threshold, listed in table 6.3, grouped around 8 different frequency regions. Before running the next stage of the follow-up, we inspect them carefully. This reveals that for the outliers at 7 of these frequency regions, the detection statistic in one of the detectors is much higher than in the other one, and that most of it is accumulated during a small portion of the run. These outliers can be safely attributed to disturbances which were present for a short time, thus for this reason they are not present in the lines and combs database (since it is obtained by using a mean amplitude spectral density of the full observing run).

The outliers remaining at the frequency region around 190.6 Hz present similar $2\mathcal{F}_{sum}$ values in both detectors, but also accumulate their statistic during a short portion of the run. After a closer inspection, these outliers seem to be generated by one of the hardware injections present in the data: the recovered parameters make the template closely resemble the frequency-time pattern of the hardware injection during a few days, and due to their huge h_0 values, a high value of the detection statistic is accumulated even for such short fractions of the run. Thus, all search outliers are caused by non-astrophysical sources and are vetoed. No astrophysical signals are confidently detected in this search.

We show the $2\mathcal{F}_{sum}$ values for all the outliers in the left plot of figure 6.1. Furthermore, the right plot of that figure shows the $2\mathcal{F}_{sum}$ values for the injections. The threshold at $2\mathcal{F}_{sum} = 2000$ is marked with a dashed red line in both plots. Due to the small number of outliers which are above 2000, outliers which are slightly below 2000 are also manually followed-up and they also appear in table 6.3.

As mentioned earlier, although the BSH search does not explicitly search over a spin-down/up pa-

parameter, it is sensitive to signals up to $(T_c T_{obs})^{-1} = 4.8 \times 10^{-11}$ Hz/s. For the follow-up case, a similar argument gives $(T_c T_{obs})^{-1} = 7.2 \times 10^{-13}$ Hz/s, which is also a higher value than almost all spin-down values of known pulsars in binary systems.

6.3 | Results

Although no detections are reported, we set upper limits on the gravitational-wave amplitude.

6.3.1 | Sensitivity depth estimation

Signals are added to the datasets in all the 0.1 Hz frequency bands, at 3 different sensitivity depth $\mathcal{D} = \sqrt{S_n}/h_0$ values (where $\sqrt{S_n}$ is the amplitude spectral density), using 150 signals per depth and frequency. Then, we calculate the efficiency (the number of detected signals divided by the number of injected signals) at each sensitivity depth. This procedure takes into account the first stage of the follow-up, where only the injections that obtain a \mathcal{F} -statistic value above the threshold are counted as detected, and it is also required that the injection cluster has a maximum detection statistic higher than the maximum detection statistic of the third cluster found in that 0.1 Hz frequency band, because otherwise it would not have been detected. Then, at each of the frequency bands a sigmoidal fit is done (an example is shown in the right plot of figure 6.2) and the 95% sensitivity depth value is found.

A widely used method to calculate the 95% sensitivity depth points is to do injections at a limited number of frequency bands, and obtain a global sensitivity depth value by calculating the mean value between these points. We used this procedure by doing injections at 95 frequency bands (19 for each of the 5 groups of mismatch), and a constant sensitivity depth value was calculated for each of the 5 groups. With this method, the results were: $\mathcal{D}^{95\%} \pm 1\sigma = 21.0 \pm 0.7, 19.9 \pm 0.8, 19.1 \pm 1.0, 18.0 \pm 0.8, 18.0 \pm 0.8$ Hz^{-1/2}. The left plot in figure 6.2 shows the value of these 95 sensitivity depth points, where the red lines mark two standard deviations 2σ .

The final results have been obtained by doing injections at all 0.1 Hz frequency bands (without averaging). Results using this method are valid for all the frequency bands analysed. Since this procedure is computationally intensive, we use the results of the previous method to guide us with the signal amplitude values to be injected, and we have decided to not place upper limits in 292 bands as in those bands a much higher number of injections would be required to converge to the 95% confidence level upper limit.

6.3.2 | Upper limits

The upper limits, when converted into gravitational-wave strain amplitude, are shown in figure 6.3 (there are 292 0.1 Hz frequency bands where we do not place upper limits, the complete list with the excluded frequency bands can be found below). It can be seen that the lowest gravitational-wave amplitude is located around 3.35×10^{-25} at 195.1 Hz. This figure shows a comparison with the previous upper limits obtained by analysing data from S6 and VSR2-3, discussed in [197]. The sensitivity to h_0 scales as [233]:

$$h_0 \propto \sqrt{\frac{S_n}{T_c}} \frac{1}{N_{SFTs}^{0.25}}. \quad (6.5)$$

At 150 Hz, for S6 the amplitude spectral density was around 2×10^{-23} Hz^{-1/2}, which compared to O2 ($\sim 7 \times 10^{-24}$ Hz^{-1/2}) gives a factor ~ 3 of improvement. The S6 run covered a longer calendar period than O2, but the duty cycle was worse so the overall N_{SFTs} factor from each run is comparable. Therefore, the improvement of ~ 17 that can be seen in the figure is due primarily to the improved detector dataset as well as using the new BSH pipeline. An important distinction remains that this search covers a much smaller parameter space compared to the previous *TwoSpect* S6/VSR2-3 search. A more complete comparison

Frequency [Hz]	α [rad]	δ [rad]	a_p [s]	P [s]	t_{asc} [GPS]	$2\mathcal{F}_{sum}$	Description
100.22026	0.013	-0.608	13.90	3913534	1176008851	2027.90	Short disturbance in H1
100.21259	-1.205	-0.585	14.26	3316110	1175772093	1992.96	Short disturbance in H1
100.22268	1.639	-0.857	13.42	1833880	1176635113	1963.22	Short disturbance in H1
103.09979	2.718	0.051	24.74	3824964	1175503115	2003.03	Short disturbance in L1
145.79107	-2.818	0.498	21.59	3429767	1174617540	1987.54	Hardware injection
145.86610	-2.410	0.797	11.98	3142246	1174797410	1978.52	Hardware injection
145.90148	2.285	1.300	9.71	3671138	1176001528	1982.27	Hardware injection
145.92549	2.487	1.335	9.62	3349536	1175854186	1981.71	Hardware injection
166.66391	1.402	-1.459	10.32	2295263	1176280470	2059.17	Short disturbance in H1
166.64618	-1.237	-1.102	12.48	2228943	1175181758	1974.85	Short disturbance in H1
173.89725	-0.263	-1.045	21.01	3229318	1177345010	3254.80	Short disturbance in H1
173.87984	-0.354	-0.766	34.48	1764911	1175470277	3436.08	Short disturbance in H1
173.89072	-0.458	-0.944	28.90	2411275	1176754746	3383.46	Short disturbance in H1
173.90550	-0.302	-0.879	13.32	2246880	1175817043	3411.56	Short disturbance in H1
173.91856	-0.441	-0.632	27.44	1671568	1176760037	3339.30	Short disturbance in H1
173.90720	0.084	-0.671	18.69	1723682	1176402358	3211.92	Short disturbance in H1
176.68644	2.112	-0.876	33.65	1517400	1176111051	2561.38	Short disturbance in H1
176.66651	2.186	-0.923	32.46	1633407	1176627316	3021.95	Short disturbance in H1
176.68301	3.056	-0.616	32.97	1501423	1176061869	2890.82	Short disturbance in H1
176.69951	2.568	-0.906	23.70	1957761	1176132478	3096.21	Short disturbance in H1
176.71434	2.319	-0.835	33.67	1408988	1175655939	2943.29	Short disturbance in H1
176.70635	2.273	-0.852	35.86	2278473	1176412120	3061.56	Short disturbance in H1
182.02284	2.498	-0.156	18.15	3410134	1175957444	1962.52	Short disturbance in L1
182.02969	-1.035	-0.182	31.57	2480680	1176772000	2071.28	Short disturbance in L1
182.03574	-0.933	-0.128	18.23	2109858	1177216961	2105.25	Short disturbance in L1
182.62066	1.298	1.317	28.20	2041602	1176047775	1952.68	Short disturbance in L1
185.68991	-0.789	0.054	12.48	3641071	1174370500	3376.22	Short disturbance in L1
185.68562	-1.285	-0.012	18.75	2824333	1176879812	3125.18	Short disturbance in L1
185.68693	-1.074	0.118	19.24	2941083	1177477791	3248.15	Short disturbance in L1
185.70359	-0.592	-0.026	29.14	3721385	1176232839	3348.67	Short disturbance in L1
185.70019	-0.833	-0.110	21.53	3050708	1177293756	3385.85	Short disturbance in L1
185.70054	-0.217	0.100	10.33	2492984	1177394628	3375.33	Short disturbance in L1
190.54188	-0.187	-0.530	17.26	3815778	1176611444	2452.27	Hardware injection
190.58117	-0.363	-0.446	15.56	3738116	1176003880	2347.43	Hardware injection
190.64508	-1.080	-0.185	9.78	2901531	1177606154	1999.52	Hardware injection
190.69765	0.044	0.051	12.33	3354700	1177521960	2261.82	Hardware injection
190.69851	0.167	0.039	13.66	3621201	1174548675	2383.19	Hardware injection

Table 6.3: Outliers found by this search. All of them can be ascribed either to a hardware injection or to a noise disturbance. The parameters correspond to the center of the cluster returned by the follow-up stage. The $2\mathcal{F}_{sum}$ column shows the summed semi-coherent \mathcal{F} -statistic over segments of the top candidate obtained at the first stage of the follow-up. The reference time for these parameters is 1164562334 GPS.




Figure 6.2: The left plot shows the 95% sensitivity depth points for the injections done at 95 different frequency bands (the red lines mark the position of two standard deviations). The right plot shows an example of a fit at 278.5 Hz. The three blue crosses show the efficiencies obtained at sensitivities of 13.5, 17.5, and 21.5 $\text{Hz}^{-1/2}$, while the green cross shows the fitted 95 % efficiency point. The orange trace shows the sigmoidal fit, while the two surrounding traces show the 1- σ envelope, from which the 1- σ error for the 95% point is obtained.

would need to take this distinction into account. Figure 6.3 also shows the previously published results for the O2 all-sky search for CWs from isolated systems using the *SkyHough* pipeline [213]. The upper limit results presented here for CWs from sources in binary systems is only a factor of ~ 2 worse, which is a new achievement for this type of search.

In figure 6.4 we show a comparison between the result we obtained by doing injections at a limited number of frequency bands and the final result that we obtain by doing injections at all frequency bands. The orange area covers the mean value plus 2 standard deviations. It can be seen that for most bands this uncertainty shaded region covers the true sensitivity depths, but for some of them they are either too conservative or too optimistic.

The 95% upper limits on h_0 can be converted to upper limits on ellipticity ϵ by using equation (6.1):

$$\epsilon^{95\%} = \frac{c^4}{4\pi^2 G} \frac{h_0^{95\%} d}{I_{zz} f^2}. \quad (6.6)$$

These results are shown in figure 6.5, where different values for the moment of inertia and distances are used (the stricter upper limits trace is used). Assuming the canonical moment of inertia of $I_{zz} = 10^{38} \text{ kg}\cdot\text{m}^2$, for sources at 1 kpc emitting CWs around 300 Hz the ellipticity can be constrained at $\epsilon < 5 \times 10^{-6}$; at 100 Hz, $\epsilon < 3 \times 10^{-5}$, while at 0.1 kpc and 300 Hz, $\epsilon < 5 \times 10^{-7}$. If we assume $I_{zz} = 3 \times 10^{38} \text{ kg}\cdot\text{m}^2$ (as could be due to higher masses or larger radii), these upper limits are even more stringent, as shown by the dashed traces in this figure. For example, at 0.1 kpc and 200 Hz, $\epsilon < 3 \times 10^{-7}$, while at 0.01 kpc and 300 Hz $\epsilon < 2 \times 10^{-8}$. Several studies indicate that neutron stars should be able to support ellipticities greater than 10^{-5} [101], making our results interesting in terms of constraining the asymmetry which neutron stars in binary systems have.

6.3.3 | Excluded frequency bands from the upper limits calculation

The 95% sensitivity depth value is not calculated for all 0.1 Hz frequency bands.

Frequency bands where the three generated clusters have been eliminated by the lines veto are not analysed. Furthermore, for frequency bands where the minimum value of the detection statistic s_p is higher than 8.0 we do not place upper limits. These bands are noisier than the others, since the values of the detection statistic are shifted to higher values. This makes the detection of a signal producing the same detection statistic value harder than in a regular band, since the final toplist will not include it or the maximum significance value of the generated cluster will be smaller than the third cluster of that




Figure 6.3: Upper limits on the gravitational-wave amplitude h_0 at 95% confidence to isotropically polarized signals. The middle blue trace shows the BinarySkyHough results, the upper dashed-orange trace shows the S6/VSR2-3 results produced by the TwoSpect pipeline [197], and the bottom dotted-green trace shows the results for the O2 CW all-sky search for isolated neutron stars using the SkyHough pipeline [213]. The vertical black lines show the 292 (out of 2000) 0.1 Hz frequency bands where the upper limits have not been computed.

0.1 Hz band. Bands with non-typical behavior are usually produced by lines and combs, as previously discussed, or by an excessive presence of short-duration glitches that elevate the noise floor of the detector over a wide frequency band. The two plots in figure 6.6 show the minimum significance value for all the 0.1 Hz frequency bands.

107 out of the 2000 bands (5.35%) have a minimum significance value higher than 8.0, while 243 bands are eliminated by the lines veto. This sums to the total 292 bands, since some of these bands are repeated in the two lists.

These are the 292 0.1 Hz frequency bands where we do not place upper limits: [100.0, 100.2, 100.4, 100.5, 100.6, 100.7, 100.8, 101.0, 101.2, 101.6, 102.3, 102.4, 102.5, 102.6, 102.7, 102.9, 103.1, 103.9, 104.0, 104.1, 104.5, 104.6, 104.8, 104.9, 105.0, 105.5, 105.6, 105.9, 106.0, 106.9, 107.0, 107.1, 107.2, 107.3, 107.5, 107.9, 108.0, 108.6, 109.4, 109.6, 109.9, 110.0, 110.1, 110.9, 111.0, 111.1, 111.4, 111.5, 111.9, 112.0, 112.9, 113.0, 113.3, 113.9, 114.0, 114.9, 115.6, 115.9, 116.0, 116.9, 117.0, 117.6, 118.0, 118.6, 118.9, 119.0, 119.2, 119.4, 119.5, 119.8, 119.9, 120.0, 120.4, 120.5, 120.6, 120.9, 122.0, 122.3, 122.6, 123.5, 123.9, 124.5, 124.9, 125.0, 125.3, 125.5, 125.7, 128.0, 128.5, 128.6, 128.9, 129.5, 130.0, 130.2, 130.5, 130.8, 130.9, 131.1, 132.5, 133.2, 133.3, 133.5, 133.9, 134.5, 135.5, 136.5, 137.5, 139.5, 139.9, 140.0, 140.2, 140.5, 141.2, 141.5, 142.5, 143.5, 144.5, 145.5, 145.8, 149.9, 150.0, 150.5, 153.4, 153.5, 153.6, 153.7, 153.8, 153.9, 154.0, 154.1, 154.2, 154.4, 154.5, 154.6, 154.7, 158.2, 159.9, 160.0, 161.2, 161.3, 163.4, 166.6, 171.4, 173.4, 173.8, 173.9, 174.4, 175.4, 176.2, 176.3, 176.4, 176.6, 176.7, 177.4, 178.5, 179.2, 179.3, 179.4, 179.7, 179.8, 179.9, 180.0, 180.1, 180.2, 180.4, 181.4, 182.4, 185.6, 185.7, 186.3, 188.4, 190.3, 192.3, 192.5, 192.6, 194.3, 197.3, 197.4, 197.5, 197.6, 197.7, 197.8, 197.9, 198.3, 199.3, 199.8, 199.9, 200.0, 202.3, 204.3, 205.3, 206.3, 207.3, 208.3, 209.3, 210.3, 211.3, 212.3, 214.3, 217.2, 218.2, 221.2, 222.2, 223.2, 224.3, 225.2, 225.6, 226.2, 226.6, 226.7, 226.8, 226.9, 227.2, 227.6, 227.7, 227.8, 230.2, 231.2, 232.2, 234.2, 235.2, 236.2, 237.2, 238.2, 239.7, 239.8, 239.9, 240.0, 240.1, 240.2, 242.2, 243.1, 243.2, 244.2, 246.2, 246.3, 247.1, 247.2, 249.1, 249.2, 252.1, 253.1, 254.1, 254.2, 256.1, 257.1, 258.2, 259.1, 262.1, 263.1, 265.1, 265.4, 265.5, 265.6, 268.1, 269.1, 270.1, 271.1, 272.1, 273.0, 273.1, 274.1, 275.1, 278.0, 279.0, 279.1, 281.1, 281.5, 283.0, 284.0, 286.0, 287.0, 287.1, 288.0, 289.0,




Figure 6.4: Comparison between the upper limits obtained by doing injections in all frequency bands (blue trace) and the upper limits obtained by averaging the results obtained at a reduced number of frequency bands (orange area).

290.0, 291.0, 292.0, 293.0, 294.0, 295.0, 297.0, 298.0, 299.1, 299.2, 299.3, 299.4, 299.5, 299.6, 299.7, 299.8, 299.9].

6.4 | Conclusions

The main search done by the BSH pipeline took 10000 CPU-hours to complete (by using a Power9 8335-GTH + Tesla V100 GPU combination), which is a very small cost. The O2 data could be further searched for signals in other regions of parameter space, both at higher frequencies and at lower and higher orbital periods. This could also be done with the next set of Advanced detectors O3 data, which will have an improved noise floor that will produce even tighter upper limits and enhance the possibilities of detection.




Figure 6.5: Detectable ellipticity at 95% confidence, given by equation (6.6), as a function of gravitational-wave frequency for neutron stars at 10 pc (bottom trace), 100 pc, 1 kpc and 10 kpc (upper trace) for a canonical moment of inertia $I_{zz} = 10^{38} \text{ kg}\cdot\text{m}^2$ (regular traces) and $I_{zz} = 3 \times 10^{38} \text{ kg}\cdot\text{m}^2$ (dashed traces).




Figure 6.6: The left plot shows the minimum value of the detection statistic obtained in the final toplist for each 0.1 Hz frequency band. The dashed red line at 8.0 marks the threshold. The right plot shows a zoom of the left plot.

CHAPTER 7

EFFORTS TO CHARACTERIZE AND MITIGATE LINES AND COMBS DURING O2 AND O3

Searches for persistent gravitational waves from continuous and stochastic sources are under way in Advanced LIGO and Virgo data. The sensitivity of these searches is degraded by the presence of narrow spectral artifacts (lines) that are generated by instrumental or environmental disturbances. In this chapter we describe a variety of methods used for finding, identifying and mitigating these artifacts, illustrated with particular examples. Results are provided in the form of lists of line artifacts that can safely be treated as non-astrophysical. Such lists are used to improve the efficiencies and sensitivities of continuous and stochastic gravitational wave searches by allowing vetoes of false outliers and allowing data cleaning. This chapter is based on [249], where some material has been removed and a new section dealing with non-linearities has been added.

7.1 | Introduction

The near future may bring the discovery by the ground-based interferometric detectors of *persistent* gravitational waves. Persistent sources of long-duration GWs can be broadly classified as continuous wave (CW) sources, which have a deterministic phase evolution, and a stochastic gravitational-wave background (SGWB), for which the signal is intrinsically random. The canonical sources for CWs (see [224] for a review) are non-axisymmetric rotating neutron stars, emitting long-lasting and nearly monochromatic waves, as described in chapters 2 and 3. When observed from Earth, these waves will be frequency-modulated due to the Doppler effect produced by the daily rotation and orbital motion of the Earth around the Sun. The SGWB is a superposition of many astrophysical and cosmological GW sources.

CW and SGWB searches look for long-duration signals, and are affected by different types of noise than those affecting short-duration searches, as shortly described in subsection 1.3.2. While compact binary coalescence and burst searches are degraded mainly by short-duration glitches (such as those described in [49,265,266]), CW and SGWB searches are mainly affected by long-lived peaks in the frequency spectra, especially narrow peaks, typically referred to as lines (although short-duration glitches increase the noise floor over a wide frequency band and also decrease the sensitivity of these searches). CW searches can be degraded because their signals are intrinsically highly narrow-band, while SGWB searches can be degraded because of the tendency of a subset of instrumental lines in the two detectors to lie so close to each other that they exhibit spurious coherence between the detectors.

This problem presents two main detector characterization tasks for long-duration searches: first, to *identify* line artifacts that are non-astrophysical in origin, allowing them to be flagged as noise; and, second, to determine the *cause* of those artifacts when possible in order to guide efforts to remove them

at the detector sites. Spectral lines that affect the CW and the SGWB searches are typically quite narrow (high Q-factor, i.e., the ratio of peak frequency to line width) during a given coherent integration time. This focuses investigations for noise sources onto electronic components and mechanical components with high Q-factor resonances, reducing the probability of noise from mechanical components with damped mechanical resonances.

In this chapter we describe tools and methods used for data quality investigations relevant to long-duration searches, and provide examples of issues faced during the second and third Advanced LIGO observing runs: O2, and O3a (first half of O3). We note that all of the methods presented here can be applied to all ground-based interferometric detectors such as H1, V1, or KAGRA, but we focus on data quality applied to the LIGO detectors only, as at the time of this writing there is significantly more Advanced LIGO observational data and only data from these detectors has been used in persistent searches up until now.

7.2 | Effects of noise on CW and SGWB searches

Spectral artifacts can degrade analyses that search for long-duration signals in different ways. Artifacts can lead search pipelines to return spurious outliers, which require laborious follow-up. For example, an instrumental line produced by a continuous sinusoidal process will be picked up by a CW search by its templates with zero spin-down/up and from sky positions that produce an almost non-existent Doppler modulation. Furthermore, if there is a putative GW signal at a frequency corresponding or nearby to a spectral artifact, then the signal power is obscured. For those analyses that rely on combining data from different detectors (e.g. cross-correlation or coherently combining data), then detection of signals overlapping with common detector artifacts may be impossible. On the other hand, some searches may be able to cope with an artifact if it occurs in just one detector.

Continuous GWs from spinning neutron stars are nearly monochromatic, with nearly constant signal frequency in the Solar System barycenter. When projected into the frame of a detector located on Earth, the signal is Doppler shifted into many frequency bins. Conversely, a narrow, stationary spectral artifact in the detector frame will impact many frequency bins when data is projected into the frame of the Solar System barycenter. For searches of a signal from a known pulsar with a given ephemeris, the impact of these artifacts is less than the impact on an all-sky search for unknown neutron stars (which may also be located in a binary system). In extreme cases, an all-sky search may be blind to a wide region of parameter space for a particular frequency range.

Searches for a stochastic GW background rely on cross-correlating GW strain channel data from multiple detectors and looking for excess power. Excess cross-correlation requires a stable phase between the two channels at a given frequency, and, thus, many single-detector artifacts are not found in the cross-correlation analysis. Correlated noise that causes excess power in the cross-correlation analysis, however, is excised from the analysis entirely by setting those frequency bins to zero (a process called notching) before integration in the case of the standard search for a broadband SGWB. This reduces the search sensitivity by a factor $\sim \sqrt{\frac{N_b}{N_a}}$ where N_b is the number of frequency bins before notching and N_a is the number of frequency bins after notching.

For both CW and SGWB searches, lists of known instrumental artifacts are created following the end of an observing run (further details are provided in section 7.7). Then, depending on the search, these lists are used to: 1) clean the data before analysis by removing the affected data in the frequency domain and replacing it with Gaussian noise measured in the nearby frequency bins; 2) avoid specific frequencies in analyses that are impacted by the artifacts; or 3) reject outliers that are clearly caused by the detector artifacts. This lets the analysis focus computational resources on regions of parameter space that are not degraded by spectral features. If a search pipeline returns a signal candidate which does not coincide with any known artifact, more detailed investigations are needed in order to assert that the signal cannot

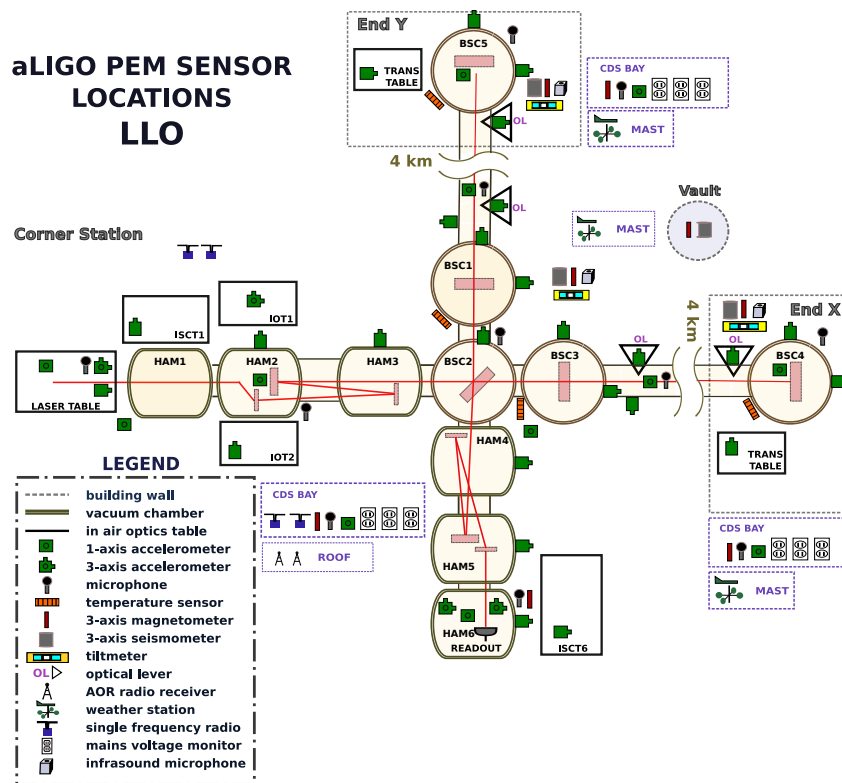


Figure 7.1: Locations of most auxiliary sensors at LIGO Livingston Observatory (LHO shares a similar layout). The gray dashed lines separate the End X and End Y Stations, which are located at the end of the 4 km arms, from the Corner Station building, located at the vertex of the detector. All stations contain an electronics room (encased by purple points in the diagram), where the computers that control the interferometer are housed.

be produced by an artifact.

7.3 | LIGO data and noise sources for searches of persistent gravitational waves

The second Advanced LIGO observing run (O2) took place between 30 November 2016 and 25 August 2017, and the first half of the third observing run (O3a) took place between 1 April 2019 and 30 September 2019. The Advanced LIGO detectors are located in Hanford, Washington (H1), and Livingston, Louisiana (L1). The LIGO detectors are dual-recycled Michelson interferometers with Fabry-Perot arm cavities of ≈ 4 km (see [28] for a review of the Advanced LIGO detectors configuration and section 1.3 for a brief discussion of gravitational-wave detectors).

LIGO detector data is typically characterized as stationary, Gaussian noise, but non-Gaussian detector artifacts are also present in LIGO data, e.g., occasional, short-duration transients (“glitches”) and long-duration narrow lines. Searches for transient GW signals will avoid analysing times when a glitch occurs or will try to subtract the glitch from the data, while searches for persistent GW signals avoid analysing data in frequency bands where narrow lines are present.

While most lines in detector data are stationary, some of the lines have time-varying behavior (often called wandering lines), which can degrade detector sensitivity over a larger range of frequencies and increase the difficulty of distinguishing these artifacts from astrophysical signals when searching for a persistent signal from different sky locations. Some lines occur in a distinct pattern known as a *comb*, with even-spacing in frequency between each tooth (each single line) of the comb. Tooth frequencies are given by $f_n = f_o + n \delta f$, where f_o is the offset (from 0 Hz) of the comb, δf is the spacing, and n is an




Figure 7.2: Noise-weighted averaged ASD showing the first harmonic violin mode region for H1 (red trace) and L1 (blue trace) for the O1 observing run.

integer. These combs are associated with linear or non-linear coupling of non-sinusoidal sources (such as a train of impulses) or with non-linear coupling of sinusoidal sources (as exemplified in subsection 7.6.5). A comb can also be recognized by a common time-dependent behavior of the teeth in the comb. The Fourier coefficients of a comb in the frequency domain can be used to uncover the time-domain waveform and help identify the source of the comb.

Lines and combs can have time-dependent behavior as the configuration of the detector changes, especially during periods of commissioning and maintenance. Some lines and combs have high amplitude and can be identified using only a short amount of data. Others have low amplitude and may only become evident when using long integration times, which are also useful to better constrain the central frequency and width of a given line or to find the spacing of a comb.

A schematic diagram showing locations of vacuum chambers, main interferometer optics, and most of the Physical Environment Monitoring (PEM) sensors of the L1 detector can be seen in figure 7.1 (H1 has a similar layout). PEM sensors include, for example, accelerometers, microphones, temperature sensors, magnetometers, seismometers, etc. PEM sensors, particularly magnetometers, are often helpful in determining the causes of narrow spectral artifacts because they witness local noise sources that may couple to the main GW channel, and the PEM sensors do not witness GW signals (except in cases of complicated cross-coupling mechanisms, which can be identified using signal injections). Other auxiliary channels may also be useful in the same way: besides the PEM sensors, each of the LIGO detectors records many more auxiliary channels that monitor instrument behavior, adding up to over 200000 auxiliary channels [49]. These channels can be compared to the gravitational-wave strain main channel to check if there any correlations or coherences.

Some of the lines observed in an amplitude spectral density of the detector data are well-understood: for example, 60 Hz power mains, mechanical resonances of mirror suspensions known as “violin modes” (see figure 7.2), calibration lines, and simulated GW signals known as “hardware injections”. Other lines are less understood and require considerable investigation to determine their nature.

The majority of instrumental lines that degrade CW searches have Q-factors in excess of $\sim 10^3$. This is, in part, because the astrophysical sources targeted by these searches have high intrinsic Q-factors, and Doppler broadening caused the Earth’s orbital velocity does not decrease the Q-factor to less than

$\sim 10^4$. Noise lines with low Q-factors (such as wide bumps) will not cause outliers since they will be incorporated into the estimation of the background noise floor, as opposed to high Q-factor sources that will not be included in the background estimation. This is because noise floor estimation methods use running medians over for example 100 frequency bins, which are robust against sources appearing in one or two bins but take into account sources appearing in a majority of the bins.

Similarly, the instrumental lines that have produced correlations between sites, degrading searches for SGWB, have also had high Q-factors. This is because the correlations are produced not by single sources affecting both of the widely-separated sites, but rather by similar sources at each site that are correlated only because they produce signals at the same (or nearly the same) frequency. Some correlated lines are due to electronic sources at each site that are set to the same frequency, controlled by a single clock (GPS), which also controls the timing of the data acquisition systems. These lines have Q-factors that are, in principle, infinite. When the frequencies are not exactly the same at each of the sites, the maximum width of the instrumental lines that can produce correlations is associated with the duration of the data segments used in the cross-correlations and the line amplitude. The typical length of Fourier-transformed data segments is 60 s long and the lowest Q-factor lines that have produced inter-site correlations are the power mains-related lines with Q-factors of $\sim 10^3$ (the LIGO sites are on different power grids that are not synchronized).

The primary source of lines with sufficiently high Q-factors degrading both CW and SGWB analyses are processes controlled by electronic clocks or oscillators. This includes digital processes, analog electronics, and mechanical processes controlled by electronic clocks such as stepper motors. Most mechanical systems do not have Q-factors above 10^3 and so do not directly contaminate the searches by causing additional outliers, but instead degrade the sensitivity of these searches. The main exceptions are mechanical systems that are designed to have high Q-factors in order to concentrate noise in a narrow frequency band, like the “violin” suspension modes.

Monitoring the frequencies associated with electronic systems is thus the main way we detect the sources of problematic instrumental lines. Monitoring each individual electronic component in the complex electronic system of LIGO would be difficult. Instead, we attempt to monitor multiple electronic systems at once, using fluxgate magnetometers (Bartington Mag-03 series, with sensitivities of about 5×10^{-12} T). The magnetometers are placed in the experimental areas and especially in important electronics racks in the electronics rooms (see figure 7.3). These magnetometers can detect even low-amplitude periodic currents controlled by oscillators and clocks that can produce high Q-factor line artifacts (detecting as low as 5×10^{-5} A at 1 m from long wires or traces).

The process of addressing lines or combs typically proceeds in three steps: identification of noise in the GW strain channel, data analysis to determine properties of the noise (precise frequency, other sensors that may witness the noise, start or end times, etc.) which may suggest a cause, and on-site investigations or interventions to mitigate the noise at its source (more details are given in section 7.6). This process is often iterative and experimental. Work on site is limited by available time, and also by the risk of interventions creating new problems, so noise sources are typically prioritized for follow-up by their strength, pervasiveness (number of bins contaminated), and the ease of addressing the most probable cause of the noise. Lines which are not identified or cannot be mitigated during an observing run are catalogued afterward; this is not ideal, but it does aid searches in cleaning data and rejecting outliers.

Mitigation efforts can prove challenging. In many cases, low-level spectral artifacts and combs are not visible in short-duration Fourier transforms. Only by performing averages over many days to weeks of data, do these features become obvious; hence it can take of order days to weeks of new data collection to determine if a mitigation attempt has improved the data or not. Unintended configuration changes that lead to line generation can also take time to appear, be tracked down and mitigated. As a result, significant epochs of an observing run can be badly contaminated in some spectral bands, even when those bands are relatively clean at the start of the run.




Figure 7.3: Method of monitoring electronic components and cables for frequencies of instrumental lines found in the data. A Bartington fluxgate magnetometer (Mag-03 MCES100) is mounted on the horizontal white PVC pipe in the back of an electronics rack containing electronics that control the position of important optics. If the magnetometer detects fields from currents varying at the same frequency as an instrumental line, the source of the line may be in the vicinity. In addition to helping with searches for sources of line artifacts, the magnetometer can indicate that a spectral line is not astrophysical in origin.

As can be seen in figure 7.4, the amplitude spectral density (ASD) of L1 and H1 exhibit different line artifacts and have somewhat different noise floors, explained in part by different configuration choices and by different environmental influences [267]. As a result, the couplings and the noise sources are different, and the lines and combs that need to be followed and eliminated differ between the detectors, although some common artifacts can be studied jointly. This figure also shows the improvement in data quality for long-duration searches from O1 data to O2 data, because of the investigation and mitigation activities described in [249]. We show the spectrum only between 20-2000 Hz, over which the searches for persistent GWs are typically performed.

7.4 | Noise coupling mechanisms to the gravitational wave channel

In this section we summarize some of the most important noise coupling mechanisms.

7.4.1 | Coupling through shared power and grounds

Most of the mitigated lines in Initial and Advanced LIGO have coupled through shared power supplies. An electronic component draws current at a particular frequency from a power supply, which results in a small periodic drop in voltage. If a sensitive piece of electronics, such as an optic actuator driver or analog-to-digital converter, shares the power supply, the frequency can be imprinted on a signal controlling alignment of an optic, for example, and thus causes a coupling to the interferometer light. This imprinting may happen, if, for example, a gain or offset in the sensitive equipment varies with the voltage from the power supply. The solution has been to place the source of the periodic current draw on a separate power supply. This has led us to attempt to better regulate power, and to isolate noise-sensitive electronics on

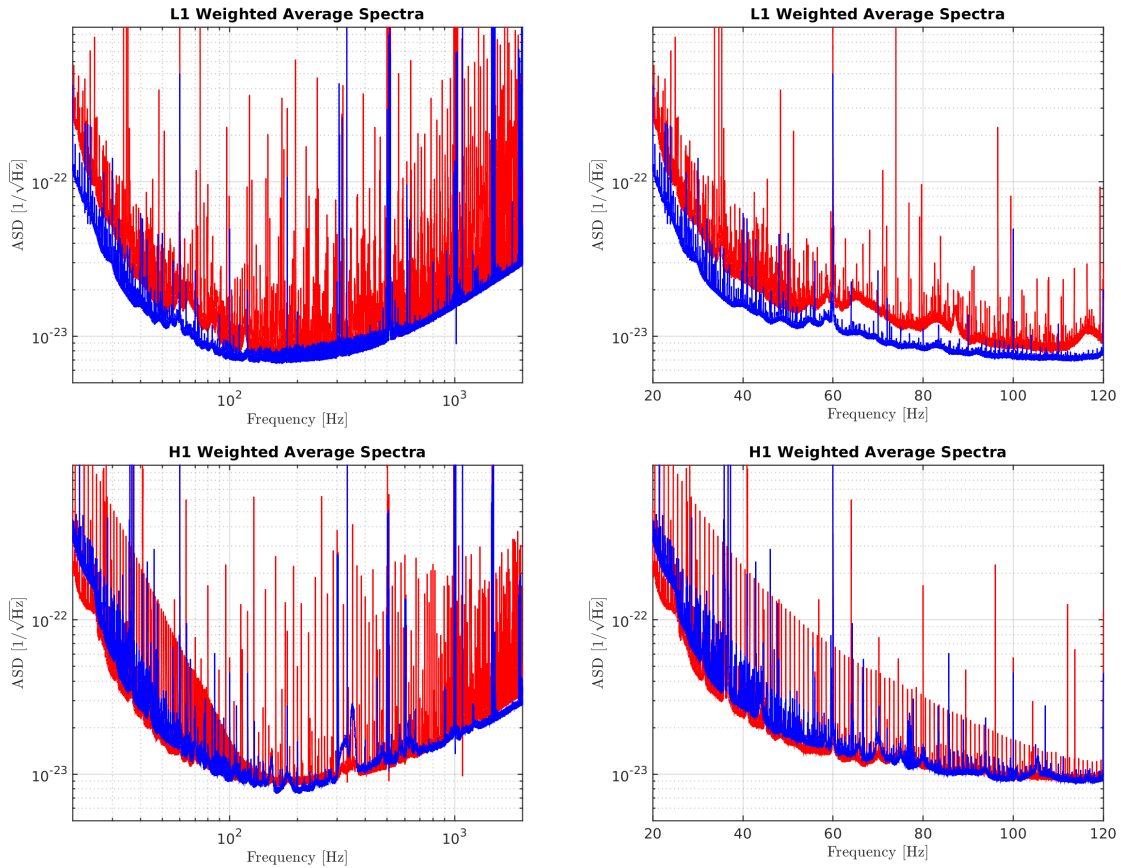


Figure 7.4: Average amplitude spectral density plots for the L1 (upper plots) and H1 (bottom plots) detectors during O1 (red trace) and O2 (blue trace). Each individual amplitude spectral density that contributes to the average is weighted by the inverse square of its running median, so that those spectra with degraded sensitivity (higher amplitude spectral density) are de-weighted (contributing less) in the final average. The left column plots show data in the most sensitive frequency band of the LIGO detectors 20-2000 Hz. The right plots show data in the low frequency region from 20 Hz to 120 Hz.

separate power supplies, but this is sometimes difficult to do in practice.

Coupling through shared grounds is a similar mechanism. Even when the source of the periodic current draw is on a separate power supply from the sensitive electronics, the source may affect the sensitive electronics by producing periodic voltage variation in shared grounds.

7.4.2 | Coupling through magnetic or electrostatic fields

Another common coupling mechanism has been direct coupling of magnetic fields to sensitive control systems or signals. For example, we have observed fields from switching power supplies coupling magnetically to signals passing through analog-to-digital converters. We have also observed 60 Hz mains magnetic fields coupling directly to permanent magnets that are mounted on certain optics for actuation. However, in Advanced LIGO, our main magnetic coupling is through cables and connectors. Mitigation efforts have included separating cables, smaller actuation magnets, electrostatic actuation, active cancellation, reducing stray fields, and separating sources and coupling sites. Digital communication systems, such as those that use Ethernet, are a common source, but it is not always easy to keep them away from sensitive systems.

When electrostatic fields are generated inside of the vacuum chambers, they may couple directly to the test masses. Electrostatic fields may also couple to control signals at locations where shielding is imperfect, like connectors. Investigations have suggested that certain sources couple through periodically

modulated electrostatic fields, although this mechanism has not been unequivocally demonstrated.

7.4.3 | Mechanical coupling

Thermally-excited high Q-factor resonances of the wires suspending optics have produced problematic lines for the CW searches by vibrating the suspended optics, which causes modulation of interferometer light, and thus couples optically to the GW strain channel. The precise frequencies of secondary suspensions may not be known in advance. Most other mechanical components are low Q-factor by design, and the broad lines that they produce typically only degrade search sensitivity for CW signals. Mechanical systems that are controlled by clocks, like stepping motors or some fans, might have Q-factors that are high enough to be problematic, but these have not been among the sources that we have found.

7.4.4 | Data acquisition artifacts and non-linear coupling

We have observed lines and combs produced by aliasing of high-frequency spectral artifacts, as well as artifacts from digital-to-analog converters. Additionally, we have observed inter-modulation products between lines of known or unknown sources during certain periods of data collection. It is also likely that we have observed combs produced by occasional errors in transmission of digitized data within the data acquisition system. The fundamental frequency of the comb is determined by the frequency (e.g., 16 Hz) of a process associated with the error.

7.5 | Data analysis tools

In this section we briefly describe some data analysis tools used to monitor and analyse the data quality for persistent gravitational wave searches. More details about the coherence tools (omitted here) are given in [249].

7.5.1 | Finding lines in frequency domain

There are two main tools to find lines in the frequency domain: Fscan and NoEMi.

Fscan is a tool that finds and monitors spectral lines [268]. It uses data from the GW strain channel and hundreds of auxiliary channels for each detector, and it produces “Short Fourier Transforms” (SFTs) of 1800-s-long data segments. Fscan produces two different types of graphs: it averages the daily SFTs (with a maximum of 48 SFTs) to produce normalized power spectra in bands of default 100-Hz width and frequency binning of 1/1800 Hz for each channel, and it produces spectrograms with averaging of adjacent frequency bins (default bin resolution of 0.1 Hz). In the absence of non-Gaussian artifacts, the normalized spectra should be flat with random fluctuations about an expectation value of one, where the underlying statistical distribution would be χ^2 with a number of degrees of freedom equal to twice the number of SFTs used to construct the spectra. Figure 7.5 shows an example of these two types of plots. Thousands of such graphs are generated automatically each day for each observatory from the GW strain channel and auxiliary channels, to provide a reference archive for line investigations.

In addition, the strain channel SFTs are used to produce (unwhitened) inverse-noise-weighted spectral averages for each day and cumulative from the start of the run through that day. The inverse noise weighting is meant to mimic the weightings used in many CW searches such as [203], which weight more heavily those time spans with better sensitivity. Comparing such spectral averages with arithmetic averages also allows rapid identification of non-stationary line artifacts. Fscan can also produce coherence plots between the gravitational-wave channel and auxiliary channels.

NoEMi (Noise frequency Event Miner) is a tool used for line monitoring and as a line database [269]. It runs daily and weekly, using data from the GW strain channel and several auxiliary channels, calculating FFTs with 1 mHz resolution. It creates time-frequency diagrams from the peaks found in the spectra;

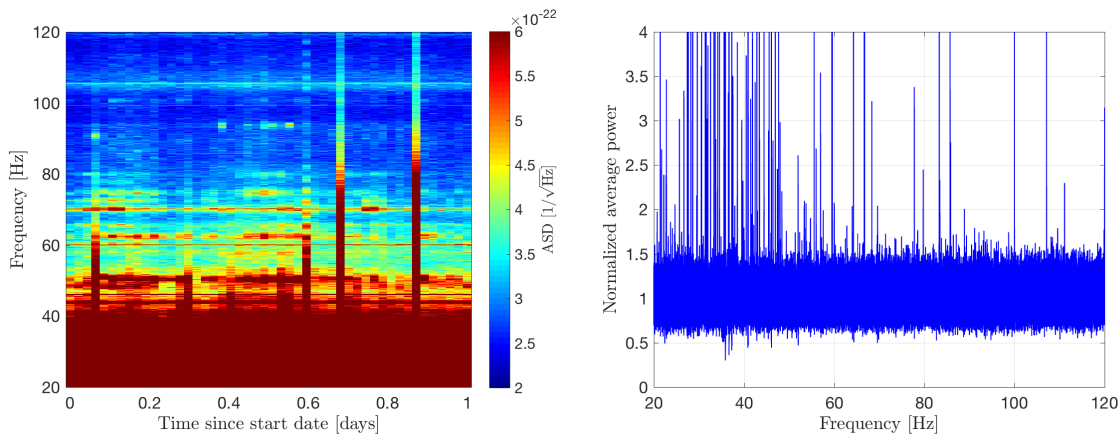


Figure 7.5: Typical plots produced by *Fscan*: the left plot shows a spectrogram of one day (23 April 2017) of Hanford strain data (with color-coded amplitude) while the right plot shows the corresponding daily averaged normalized power versus the frequency.

the program also calculates the persistency of the lines (number of peaks in that frequency bin divided by the number of FFTs) and their critical ratio (difference between the peak amplitude and the mean value of the spectrum, divided by the spectrum standard deviation). The persistency helps to identify loud stationary lines, while the critical ratio helps to identify non-stationary lines lacking persistency.

NoEMi can provide the starting and end times of a line in the data. It can also follow wandering lines, by allowing some change in frequency between different time periods. NoEMi looks for coincidences (lines with the same frequency) between the GW channel and the other channels, calculating a value between 0 and 1 to quantify the probability of coincidence for each different auxiliary channel. This automated coincidence monitoring is especially valuable when searching for causes of line artifacts seen in the GW strain channel.

7.5.2 | Finding combs

FineTooth is a set of tools to help identify combs and monitor them over time [270]. It is comprised of a plotting tool, a tracker for known combs, and a comb finding tool. The plotting tool creates interactive browser-based plots using the Python library *Bokeh*, allowing the user to overlay combs and lines and easily explore spectral features, as shown in figure 7.6. The tracker accepts a list of known combs and a list of channels, and then draws from *Fscan* data to create plots showing the historical strength of each comb in each channel. The comb finding tool searches for common spacings between peaks of comparable heights, generating a list of potential comb candidates to be vetted by the user.

During observing runs, the FineTooth tracker is run daily on a series of magnetometer channels which typically witness noise in nearby electronics, as well as on daily and run-cumulative spectra from the GW strain channel, providing a summary page for data quality checks and a tool for rapid investigation of specific combs. The comb finding and plotting tools are also used to provide an alert for new combs appearing in the cumulative spectrum mid-run, and to aid in comb identification for the purpose of generating vetted noise line lists.

7.5.3 | Lines and combs in time-domain: folding studies

Most line investigations are carried out in the frequency domain, but a tool has also been developed to look directly at periodicity in time-domain data, since some spectral combs arise from periodic transient glitches. The folding tool splits a long segment of data into short segments (typically a few seconds in length, corresponding to some periodicity of interest, e.g. $1/\delta f$ or $1/f_o$ for a comb) and averages the




Figure 7.6: A screenshot showing the comb plotting feature of *FineTooth*, on a run-averaged spectrum from Hanford in O2.

segments together to produce a summary plot. The data folding tool can generate daily, monthly, and full-run plots, with or without a band-pass filter applied. Band-pass filtering often makes periodicity more easily visible.

Folded data can reveal features of the periodic structures underlying spectral combs, making it useful for spotting changes that may not be evident in the spectrum. It is typically most effective for magnetometer channels, where periodicity is stronger than in the GW data channel, but on occasion periodic transients have been visible in the GW strain channel as well, most notably from blinking light emitting diodes (LEDs), as discussed below.

7.5.4 | Finding long-duration non-linearities

The output of a gravitational-wave detector is usually described as $x(t) = h(t) + n(t)$, where $n(t)$ can be described as the sum of many individual terms: $n(t) = n_1(t) + n_2(t) + \dots$. This expression only contains linear terms, but due to interactions between these different noise sources non-linear noise terms may be present in the data, expressed as $n(t) = n_1(t) + n_2(t) + n_1(t)n_2(t) + \dots$. The most common noise characterization techniques such as amplitude spectral density or coherence are not well-suited to find non-linearities that may be present in the data, and alternative techniques are needed.

These alternative techniques are related to higher-order statistics (higher than second-order). Second-order measures (such as the power spectrum and autocorrelation functions) contain no phase information, and for this reason certain types of phase coupling (associated with non-linearities) cannot be correctly identified. This is related to the fact that a Gaussian signal is completely characterized by its mean and variance (second-order statistics), but non-Gaussian signals (as present in gravitational-wave data) need higher-order statistics to be completely described.

One of the most common non-linear interactions is the quadratic phase coupling (QPC), which involves frequency triplets: f_1 , f_2 and $f_1 \pm f_2$. QPC can be obtained with a non-linear interaction between two

sinusoidal signals:

$$\begin{aligned}
z(t) &= x(t) + y(t) + [x(t) + y(t)]^2 \\
&= A \cos(2\pi f_1 t + \phi_1) + B \cos(2\pi f_2 t + \phi_2) + [A \cos(2\pi f_1 t + \phi_1) + B \cos(2\pi f_2 t + \phi_2)]^2 \\
&= A \cos(2\pi f_1 t + \phi_1) + B \cos(2\pi f_2 t + \phi_2) + A^2 \cos^2(2\pi f_1 t + \phi_1) + B^2 \cos^2(2\pi f_2 t + \phi_2) \\
&\quad + 2AB \cos(2\pi f_1 t + \phi_1) \cos(2\pi f_2 t + \phi_2) \\
&= A \cos(2\pi f_1 t + \phi_1) + B \cos(2\pi f_2 t + \phi_2) \\
&\quad + A^2[0.5 + 0.5 \cos(2\pi 2f_1 t + 2\phi_1)] + B^2[0.5 + 0.5 \cos(2\pi 2f_2 t + 2\phi_2)] \\
&\quad + AB[\cos(2\pi(f_1 + f_2)t + (\phi_1 + \phi_2)) + \cos(2\pi(f_1 - f_2)t + (\phi_1 - \phi_2))].
\end{aligned} \tag{7.1}$$

QPC means that the sum of the phases at $f_1(\phi_1)$ and $f_2(\phi_2)$ is equal to the phase at frequencies $f_1 \pm f_2$ (i.e. $\phi_1 \pm \phi_2$), which is often an indication of second-order nonlinearities.

The two-point correlation function calculates the expected value of the product of two functions f and g at two different points:

$$C_{fg}(\tau) = \langle f(t)g(t + \tau) \rangle = \int_{-\infty}^{\infty} f(t)g(t + \tau)dt, \tag{7.2}$$

where the variable t can represent time, spatial position or any other independent variable. The auto-correlation function is defined as:

$$C_{ff}(\tau) = \langle f(t)f(t + \tau) \rangle = \int_{-\infty}^{\infty} f(t)f(t + \tau)dt. \tag{7.3}$$

In the discrete case with N points, the auto-correlation function is defined as:

$$C(J) = \langle f(i)f(i + J) \rangle = \frac{1}{N} \sum_{i=1}^N f(i)f(i + J). \tag{7.4}$$

The power spectral density S_{ff} is the Fourier transform of the auto-correlation function, and the cross-spectral density S_{fg} is the Fourier transform of the correlation function. Although the regular PSD does not retain phase information, the cross-spectral density does. The (squared) coherence (a real value between 0 and 1) between two signals is defined as:

$$K_{fg}(f) = \frac{S_{fg}^2(f)}{S_{ff}(f)S_{gg}(f)} = \frac{\langle \tilde{f}^*(f)\tilde{g}(f) \rangle^2}{\langle \tilde{f}^2(f) \rangle \langle \tilde{g}^2(f) \rangle}, \tag{7.5}$$

where $\tilde{f}(f) \equiv ae^{i\phi_f}$ and $\tilde{g}(f) \equiv be^{i\phi_g}$, which gives:

$$K_{fg}(f) = \frac{\langle \tilde{f}^*(f)\tilde{g}(f) \rangle^2}{\langle \tilde{f}^2(f) \rangle \langle \tilde{g}^2(f) \rangle} = \frac{|\langle ab[\cos(\phi_g - \phi_f) + i \sin(\phi_g - \phi_f)] \rangle|^2}{\langle a^2 \rangle \langle b^2 \rangle}. \tag{7.6}$$

The coherence between two channels f and g measures the average phase difference between them: it is 1 if the phase difference is constant and goes to 0 if the phase difference is a uniformly distributed random number. It is a statistic that is routinely calculated in order to find causes or clues of noise sources in the main gravitational-wave channel, for example between magnetometers and $h(t)$. An estimate of the significance of a given value of coherence can be found by simulating Gaussian random series, calculating coherence values and finding the desired coherence threshold at a given false alarm probability.

Higher-order spectra are defined as the Fourier Transform of the k order cumulant from the time domain. For example, a third-order cumulant is (and its discrete counterpart):

$$C_3(\tau_1, \tau_2) = \langle f(t)g(t + \tau_1)h(t + \tau_2) \rangle = \int_{-\infty}^{\infty} f(t)g(t + \tau_1)h(t + \tau_2)dt \tag{7.7}$$

$$C_3(J, K) = \langle f(i)g(i + J)h(i + K) \rangle = \frac{1}{N} \sum_{i=1}^N f(i)g(i + J)h(i + K). \tag{7.8}$$

The bi-spectrum $B(f_1, f_2)$ is defined as the Fourier transform of the third-order cumulant. Similarly to the coherence, the squared bicoherence is a normalized version of the bi-spectrum:

$$b_+(f_1, f_2) = \frac{S_{fgh+}^2(f_1, f_2)}{S_{ff}(f_1) S_{gg}(f_2) S_{hh}(f_1 + f_2)} \quad (7.9)$$

$$b_-(f_1, f_2) = \frac{S_{fgh-}^2(f_1, f_2)}{S_{ff}(f_1) S_{gg}(f_2) S_{hh}(f_1 - f_2)}, \quad (7.10)$$

where $S_{fgh+}(f_1, f_2) = \langle \tilde{f}(f_1) \tilde{g}(f_2) \tilde{h}^*(f_1 + f_2) \rangle$. The bicoherence quantifies the power and phase coherence at a coupled frequency. For stochastic wide-band signals, the level of the bicoherence gives a measure of the signal skewness while for deterministic signals a peak in the bicoherence indicates the presence of Quadratic Phase Coupling (QPC). As with the PSD, we can define both a bispectrum (when $f = g = h$) and a cross-bispectrum.

As happens with the ASD, the bi-spectrum needs to be estimated from the data. We can use a method that resembles the Welch estimation method, which splits the data in M segments and calculates a mean between them:

$$S_{fgh+}(f_1, f_2) = \frac{1}{M} \sum_{k=1}^M \left[\tilde{f}(f_1) \tilde{g}(f_2) \tilde{h}^*(f_1 + f_2) \right]_k = \frac{1}{M} \sum_{k=1}^M \left[a e^{i\phi_f} b e^{i\phi_g} h e^{-i\phi_h} \right]_k. \quad (7.11)$$

It can be seen that in the triple product, the phases of the first two factors will add, and will be compared to the phase of the summed frequency component. The sum of the three phases is the biphas, and, if the phase of the third factor is equal to the sum of the first two, then the biphas will be exactly zero. Thus, in the case of QPC, we search frequency triplets where the biphas is zero.

The auto-bicoherence of the main gravitational-wave channel can be calculated by selecting a starting frequency, for example 10 Hz, and then looping over all other frequencies, e.g. from 10 Hz to 1000 Hz, and finding the complex value at $f_1 + f_2$. After the inner loop is done, the starting frequency from the outer loop is incremented and the inner loop is done again. From these two loops we obtain the ASD values at f_1 , f_2 , and $f_1 + f_2$, and from them the bicoherence can be estimated by averaging several segments of data.

Detector characterization studies have already used the bicoherence to find non-linearities in the data, as for example explained in [271], such as interactions between roll and violin modes or between violin modes and microseism frequencies. These studies have focused on finding non-linearities when short-duration glitches are present in the data, but no studies have tried to characterize long-duration non-linearities. For example, due to different non-linear processes, a line present in the data might not just be contaminating its own frequency band, but may be combined with other disturbances such as very low frequency lines to produce additional lines through the QPC. An example of this can be seen around the 60 Hz lines, which usually present wide shoulders around them, or near the violin mode regions, as figure 7.2 shows.

7.6 | Results

In this section we describe some examples of specific noise sources that were mitigated during the O2 and O3a runs (a more complete list can be found in [249], including O1 investigations). Several plots showing the improvement of the spectrum in the respective frequencies are also presented.

When a new feature in the detector strain data channel is discovered by using the tools mentioned in the previous section, additional investigations to identify the source of the noise are performed:

1. Determine the Q-factor of the line affecting the search. This helps identify the source and type of equipment that is producing the line. If the Q-factor is above 10^6 , the source is likely to be precision-clocked electronics components, or equipment that is synchronized to GPS. Typical inexpensive

clock chips in electronic devices have Q-factors of $\sim 10^5$, though the Q-factors of newer inexpensive chips may be higher. Lines from equipment using 60 Hz timing from the mains, have Q-factors of roughly 10^3 . The Q-factors of LIGO suspension wire resonances vary, but many of the secondary optics are in the range of a few 10^5 .

2. Identify and investigate any transitions in line amplitudes. If there are sudden changes in amplitude of the lines, it is often helpful to examine instrument logs for correlated changes in instrumentation or software.
3. Search for lines of the same frequency in the fixed magnetometer signals. If the line is found, it may help localize the cause. However, even if the frequency detected by the magnetometer may match the instrumental line, it may not be the cause. The probability of incorrect attribution is higher for lines that are at integer frequency values and are synchronized to GPS.
4. For lines that are detected in magnetometer channels, the location of the source can often be further narrowed by moving around a portable magnetometer to maximize the line amplitude.
5. Search for lines in auxiliary channels, especially error signals for secondary optical cavities. The lines for many secondary optic suspensions will have higher signal-to-noise ratios in auxiliary channels than in the GW data channel.
6. Search for LEDs flashing at the frequency of the lines. The periodic current drawn for the LED may cause the coupling by modulating power supply or ground voltages.
7. Temporarily shut down equipment in the candidate area, when possible, as a test. This is especially helpful if a line is stronger in a magnetometer signal than in the interferometer signal because the magnetometer can be used to more rapidly evaluate the effect of shutting down the equipment.

7.6.1 | 11.111 Hz comb (Vacuum sensors)

A 11.111 Hz comb was found at the beginning of May 2017 in the Hanford O2 data. After some investigations with a portable magnetometer, it was found that this comb was present around cables from the 24 V power supply that powered one of the Electrostatic Drives (ESD), which control the test mass positions and so are one of the most sensitive components in the system. The components powered by the cables from this supply were checked, and a strong 11.111 Hz magnetic field was detected near a vacuum sensor.

A laboratory test confirmed that the communication frequency between this type of sensor and its computer controller was 11.111 Hz, and that the LED on the sensor flashed at this frequency. The other vacuum sensors at this station were powered by separate supplies but this sensor had been connected to the ESD power supply in error. Placing the sensor on the proper power supply eliminated the comb from the GW strain channel as shown in figure 7.7.

7.6.2 | 86 Hz line (Pcal high frequency injections)

A line at 86 Hz was discovered the 15th of June 2017 in the Hanford GW strain data [273]. After investigating this with a coherence tool, we saw that this line was also present in some photon calibration (Pcal) channels. The Pcal system applies calibrated forces to the end mirrors and is used for interferometer output calibration [274]. The line had appeared for the first time exactly at the same time as the frequency of a Pcal high frequency injection at 5950 Hz was changed. Turning off this injection made the line in the GW channel disappear, as can be seen in figure 7.8.

While the coupling mechanism remains unclear, a working hypothesis is that the data acquisition system down-converts the high frequency injection to low frequency lines. A phenomenological equation

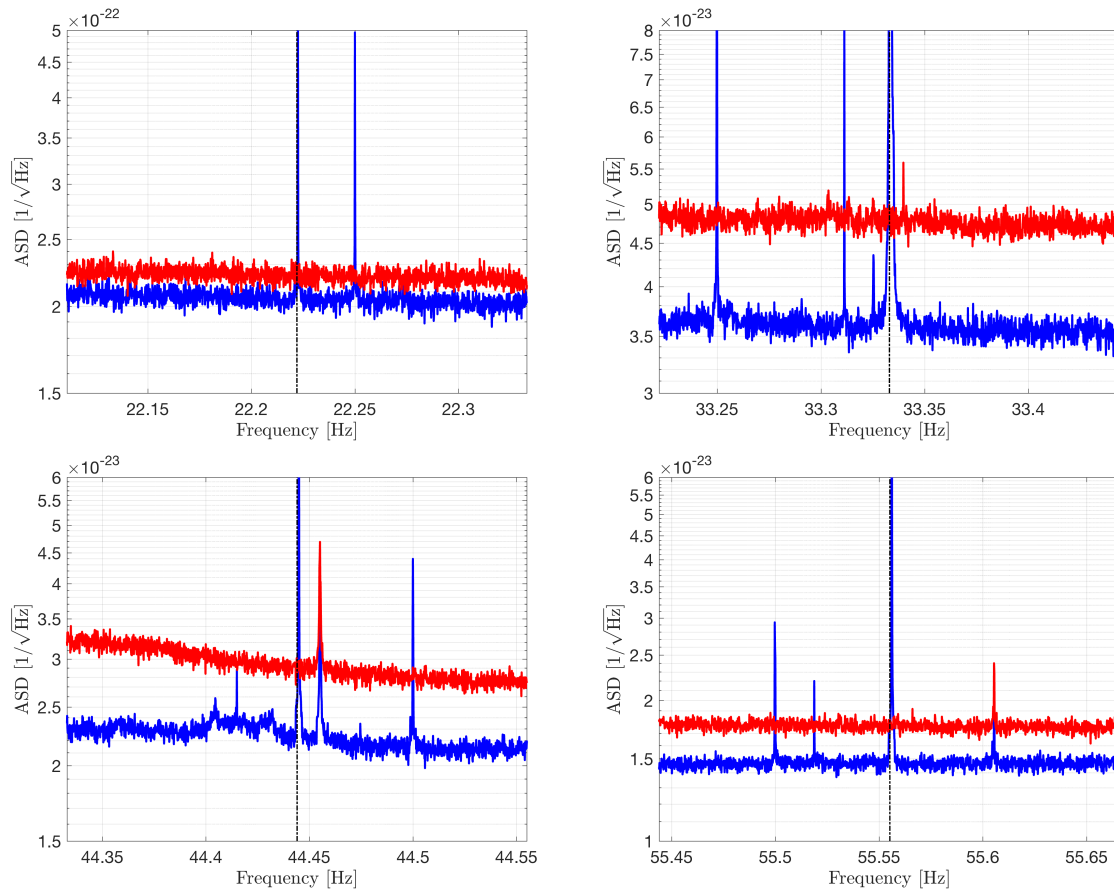


Figure 7.7: Comparison of noise-weighted averaged ASD using H1 data from 8 March 2017 to 8 May 2017 (blue trace, vacuum sensors sharing ESD power supplies) with noise-weighted averaged ASD from 8 June 2017 to 25 August 2017 (red trace, vacuum sensors not sharing ESD power supplies). The 11.111 Hz comb is gone in the second period, as we can see when we look at the harmonics 2, 3, 4 and 5, shown in the four panels by a black dashed line.

to predict the frequency of the lines was derived: $f_{line} = 2^{16} - f_{inj} n$, where f_{inj} is the frequency of the injection and n is the n th harmonic (the observed line was the 11th harmonic). This equation was tested changing the frequency of the injections, and it predicted correctly the frequency of the lines. After discovering this, a similar line, also produced by down-conversion, was observed in the Livingston GW strain channel at 119.9 Hz [275]. Down-converted lines due to photon calibrator actuation do not appear appreciably in the GW spectrum above ~ 150 Hz because the force-to-length transfer function decays as f^{-2} . This working hypothesis seems to be proven after the work explained in subsection 7.6.5.

7.6.3 | Lines from bounce and roll modes

Besides the four main test masses, the Advanced LIGO detectors use additional smaller hanging masses, such as the beam-splitter or the signal recycling mirror. These smaller masses are hanged on triple suspension systems, whereas the main test masses use quadruple suspension systems. The movement of these suspension systems can be described with many degrees of freedom or modes of motion, two of them being the bounce and roll modes, each having a different resonant frequency (roll modes being around 27 Hz and bounce modes around 40 Hz). Usually, filters that help to decrease the response of these suspension systems at their resonant frequencies are used in the feedback systems, but these filters can get misplaced in frequency during an observing run.

Lines around resonant frequencies of roll and bounce modes were seen in the main gravitational-wave




Figure 7.8: Comparison of noise-weighted averaged ASD using data H1 from 8 June 2017 to 15 June 2017 (blue trace, PCAL high-frequency injection on) with noise-weighted averaged ASD from 14 July 2017 to 17 August 2017 (red trace, PCAL high-frequency injection off). The 86 Hz line has disappeared in the second period.




Figure 7.9: Comparison of averaged ASD using O3a H1 data from 5-7 June 2019 (blue trace) to 1-3 June 2019 (red trace). The roll mode at 27.71 Hz has been dramatically lowered, and the line at 27.59 Hz has also decreased.

channel, and coherence close to 1 with auxiliary channels related to these auxiliary masses was found. By exciting the different optics, updated frequencies of these modes were found (as detailed in [276]), and this allowed to update different filters, as detailed in [277]. This made some of these lines disappear or be highly reduced, as can be seen in figure 7.9 (more details in [278]). The updated measurements of roll and bounce frequencies are also useful in order to confidently assign lines in the spectrum to these non-astrophysical sources.

7.6.4 | Non-stationary line around 70 Hz

A new line appeared around the 29 May 2019 in H1 data near 70 Hz, apparently after the maintenance done the previous day. The frequency of the line was highly non-stationary, as can be seen in figure 7.10. The new line was coherent with some PEM channels from the End Y station such as tilt sensors and magnetometers, and also some channels related to seismometers.

After finding this new line, further investigations reported that one seismometer was saturating (more details can be found here [279]). After solving this seismometer issue, the line disappeared. The exact coupling mechanism is not clear, although this seismometer shared a power supply with more PEM systems. This example shows the importance of constant monitoring of the main gravitational-wave




Figure 7.10: Spectrogram showing the frequency-time evolution of a new non-stationary line found around 70 Hz.

channel, where a new noise feature was eliminated in one day, whereas many of the other long-duration noise sources remain unknown and keep contaminating the spectrum after they appear for much longer durations.

7.6.5 | Non-linearities in O3a data

In this subsection we show an application of the bicoherence to real O3a H1 data in order to find long-duration non-linearities. An investigation started around the beginning of April 2019 showed that there were many frequency triplets with bicoherence values close to 1 (full details in [280]). We realized that many of these triplets were related to frequencies around the violin modes of the test masses (resonant frequencies of the suspensions holding the test masses) and the calibration lines (artificial signals added for calibration purposes, initially located at 15.1 Hz, 16.7 Hz, 35.9 Hz, 36.7 Hz, and 331.9 Hz). Prior to this investigation, many new lines around the violin modes had appeared, with a symmetrical distribution around them. These groups of lines were present around the 1st, 2nd, and 3rd order violin modes, and could to be described by a comb of 0.8 and/or 0.6 Hz. An interesting fact is that very similar groups of close lines separated by 0.8 and 0.6 Hz were also present in the O2 data. We realized that some of these lines were located at frequencies that are the sum of a violin mode frequency and a calibration line frequency. Some of these lines are: $552.68167 = 516.78167 + 35.9$ Hz, $553.48167 = 516.78167 + 36.7$ Hz, $521.58611 = 504.88611 + 16.7$ Hz, $519.98611 = 504.88611 + 15.1$ Hz, $531.88167 = 516.78167 + 15.1$ Hz, $533.48167 = 516.78167 + 16.7$ Hz, $1033.56334 = 516.78167 + 516.78167$ Hz, $1033.98388 = 998.08388 + 35.9$ Hz, $1034.78388 = 998.08388 + 36.7$ Hz, where the frequencies lower than 40 Hz are calibration lines and the higher frequency lines are violin modes of the test masses. These combinations of lines were present both at $f_1 + f_2$ and $f_1 - f_2$, increasing the evidence for a true non-linear process. Furthermore, non-linearities in the calibration lines themselves also seemed to appear, since there were lines with bicoherence almost equal to 1 at $72.6 = 35.9 + 36.7$ Hz, $71.8 = 35.9 + 35.9$ Hz, $73.4 = 36.7 + 36.7$ Hz and $33.4 = 16.7 + 16.7$ Hz.

The Digital-to-Analog Converter (DAC) is a card that takes a requested digital signal, such as the calibration line, and converts it to an analog signal that can be used for different purposes, such as with the laser signal that produces the calibration lines. LIGO uses many of these DAC cards, and we wanted to test them to check if they were properly working. Tests of the 20 and 18 bit DACs were performed by driving a signal of varying frequency and amplitude, and checking if non-linearities appeared in the spectrum. The left plot of figure 7.11 shows the results obtained by injecting a signal of 100 Hz at half the maximum amplitude (which is 2^{19} counts). A comb of lines is produced, with spacing equal to the injected frequency. By increasing the amplitude of the injected signal all the lines of the comb increase their amplitude, and many more appear above the noise. The right plot of figure 7.11 shows an example




Figure 7.11: The left plot shows the spectrum of a sinusoidal signal at 100 Hz produced by one DAC, where several non-linear lines are observed. The right plot shows the mixing between two signals at 100 and 125 Hz.

of two injected signals at 100 and 125 Hz. New lines that could not be seen when only injecting a single signal at 100 or 125 Hz appeared. These new lines show a spacing of 25 Hz, suggesting some type of beating between the two injected signals. This beating may explain some of the lines seen in spectrum at frequencies that seem to be a beating of violin modes and calibration lines. More details about this investigation can be found in [281].

Furthermore, by increasing the frequency of the injected signal, many lines start to appear in the low frequency range. For example, for a signal of 5950 Hz, lines with a spacing of 86 Hz appear. These lines appear to come from aliasing of the high frequency harmonics of the main comb discussed previously. Harmonics that lie between 2^{15} Hz (Nyquist frequency) and 2^{16} Hz (sampling rate) are moved to the main frequency range by $2^{16} - n f$, where n is the harmonic index. Harmonics between 2^{16} and 2^{17} are moved by $n f - 2^{16}$. This pattern is repeated for all the produced harmonics, which get transported to frequencies between 0 and 2^{15} Hz. The 5950 Hz frequency was chosen to replicate the line at 86 Hz discussed in subsection 7.6.2. From this example it seems that high-frequency calibration lines can appear as noise lines in the low-frequency region due to non-linearities of the DAC.

In order to test if a non-linearity is really responsible for some of these lines, the calibration lines were moved by e.g. 0.1 Hz and the spectra was analysed again to see if the non-linear lines had moved. After this was done, we took an average over spectra of 1800 s from before the calibration lines were moved and after this was done. All of the non-linear lines between the calibration lines changed their frequency, e.g. the three lines at $2 * 35.9$, $2 * 36.7$ and $35.6 + 36.7$ Hz. As previously discussed, if those three lines were really harmonics of those calibration lines it makes sense that they were gone since these lines have been moved. New lines appeared, for example at $35.2 = 2 * 17.6$ Hz, which could be linked to the new calibration line frequencies. Figure 7.12 shows the ASD between 300 and 400 Hz. Not only the calibration line at 331.9 Hz line has disappeared since it has been moved, also sidebands at $315.2 = 331.9 - 16.7$ Hz, $348.6 = 331.9 + 15.9$ Hz are gone.

Why do some of these non-linear lines appear and disappear several times during an observing run? The non-linear lines only clearly appear in the spectra (are above the noise) at times when the violin mode producing the non-linearity has a higher amplitude than usual (the amplitudes of the violin modes continually grow with time and have to be manually damped after they reach a certain critical level). The non-linearities may always be present, but they only show above the noise when a violin mode has been excited over a certain value. This is problematic for CW searches since they use much longer combinations of the data, so even if these features are not present during short combinations of data they may affect searches after a full observing run has finished.

Figure 7.13 shows a frequency region around the 2nd order violin mode harmonics at a time when their




Figure 7.12: The left plot shows the ASD between 300 and 400 Hz before the frequencies of the calibration lines were changed, where some lines due to non-linearities can be seen. The right plot shows the same region after the frequencies were changed, and the lines related to the calibration lines and their non-linearities have disappeared.




Figure 7.13: A comparison between the H1 ASD from the first week of April 2019 (blue trace) and from the second week of May 2019 (red trace). The three ellipses show the three (symmetric around the violin mode at 998.08389 Hz) groups of lines caused by non-linearities between the violin mode and several calibration lines.

amplitude was higher than usual. Many new lines appeared around the frequency 998.083889 Hz, which is a violin mode of one of the test masses that was more excited than usual. Figure 7.13 shows a comparison with a previous week where this mode was not as excited. Three groups of lines have appeared at each side of this violin mode, each group containing lines exactly located at frequencies that are a combination of the violin mode and one of the four calibration lines (15.6, 16.4, 17.1, 17.6), like $998.08388 - 17.6$ Hz or $998.08388 + 15.6$ Hz (all of these lines show values of bicoherence close to 1). It seems that more than one group of nonlinear lines is created around the violin mode. The extra groups of lines are also located at nonlinear frequencies like $998.08388 - 2 * 17.6$ Hz or $998.08388 + 3 * 15.6$ Hz and also show values of bicoherence close to 1 (bicoherence between a calibration line and the previous group of nonlinear lines). These extra lines may be created due to a tri-linearity such as $ASD(f_1)^2 * ASD(f_2) \rightarrow ASD(2f_1 + f_2)$ and higher order features, or due to a bi-linearity between the firstly created new non-linear lines (closest group to the violin mode) and the calibration lines.

Although the DAC seems to generate non-linearities, it is not clear if it is the source of the observed non-linearities, since they could be created by an interaction or coupling between the laser applying the calibration lines and the suspensions themselves (the beating between calibration lines themselves seems

more likely to be caused by the DAC cards).

This investigation has shown that many lines and combs present in the spectrum of gravitational-wave detectors may be produced by non-linear combinations of different noise sources. For this reason, more non-linearity investigations should be pursued. Further non-linearities could be found by calculating the bicoherence between two different auxiliary channels and the main channel, although the computational cost of this analysis over the full list of auxiliary channels is prohibitive and a subset of them should be selected.

7.7 | Producing a list of known lines and combs for CW searches

In this section we briefly describe how lists of known instrumental lines are generated for each observing run and we present a summarized O1/O2 list of lines and combs. Different approaches are followed by the CW group and the SGWB group, because the stochastic searches are only affected by lines that are coherent between both detectors i.e. have the same frequency, whereas CW searches are also affected by lines present in only one detector. Details about lists generated for the SGWB searches are given in [249]. Tables 7.1 and 7.2 summarize the lines and combs that were found in the O1 and O2 datasets, lines deemed safe to veto *a priori* in searches.

7.7.1 | List of known lines and combs for CW searches

Searches for CWs, such as recent all-sky searches for unknown isolated sources [203], typically use a list of known lines and combs to veto frequency bands prior to running the searches or, afterward, for vetoing outliers. We summarize here the procedure used to generate these lists.

We begin by generating Tukey-windowed short Fourier transforms (SFTs), using the standard FFT code that is part of the LALSuite library [51]. We generate 7200-s-long SFTs for each detector covering all of the observing run time, because those are the lengths of the longest SFTs used in O1 semi-coherent searches based on summing SFT powers. Then, we compute the inverse-noise-weighted average of the SFTs as described in subsection 7.5.1.

The lines are found by visual inspection of the spectrum. Features that appear to be above the neighbouring noise level are noted for further inspection. Since, in principle, a narrowband astrophysical source might appear in the spectrum, this initial list is not regarded as safe for line vetoing or cleaning.

After this first pass, we try to correlate the found lines with other channels using the coherence and NoEMi tools, we check which lines belong to previous known combs using the FineTooth tool (and try to find new combs in the data), and we find the lines produced by known sources like the power mains, the calibration lines, and the mechanical resonances of the different suspensions. This allows us to produce a first list with lines that are safe to veto. This list is updated as more investigations are carried out, including detector studies and targeted follow-up of outliers from searches. A more general list that includes all found lines is also generated, and different searches may select either the safe list or the more complete list. We remark that even if no plausible source for a line above the noise is found in auxiliary channels, it is very likely that it is produced by a non-astrophysical source, since signals with amplitudes as large as to produce a line above the noise in an averaged spectrum are not expected due to searches over datasets from previous observing runs.

Some lines occupy a single frequency bin, while others occupy several bins. For the latter case, we define the width of the line by noting the interval for which there is a statistically significant excess above the background level estimated from neighbouring bins. Non-stationary “wandering lines” can be narrow in a particular short time epoch, but vary in frequency over an observation run, leading to a substantial widths in a run-averaged spectrum. Visual inspection is used in these studies, rather than fully automated methods, because of the enormous range in line widths encountered, combined with overlapping line structures that challenge automated tools.

7.7.2 | Known lines and combs for O1 and O2

We present a table of lines and a table of combs for the O1 and O2 data runs, with a description of the source of the noise in each case for which it is known. Updated and more detailed lists can be found at <https://losc-dev.ligo.org/o1speclines/>. Table 7.1 shows a list of O1 and O2 combs that have been identified at the time of this writing, while table 7.2 shows a list of O1 and O2 single lines which do not belong to any known comb.

7.8 | Conclusions

We have demonstrated the methods used for identification of narrow spectral artifacts caused by non-astrophysical disturbances. These efforts benefit searches for persistent gravitational wave signals by identifying those frequency bands affected by such disturbances. Some artifacts are caused by sources that can be mitigated. Several examples of such mitigation efforts have been presented. While some of the most pervasive combs have been reduced or mitigated, the causes of other artifacts remain unknown (see figure 7.4 and table 7.1).

As detector noise is reduced between observing runs, other previously unseen lines and combs are likely to become apparent, requiring further identification and mitigation efforts. In addition, as described in this article, detector maintenance activities can inadvertently create new spectral artifacts. Careful monitoring of the data will continue to be required in order to prevent contamination of long epochs of detector data. Mitigating narrow spectral artifacts will also be needed well into the advanced gravitational wave detector era.

Spacing (Hz)	Offset (Hz)	Range of visible harmonics	Description	Detector	Run
0.0470*	972.1417	0-1	Unknown	H1	O2
0.088425*	76.3235	0-14	Unknown	H1	O1
0.08844*	153.4428	0-9	Unknown	H1	O2
0.2000	0.0000	106-191	Unknown	L1	O2
0.6000*	0.5690	742-745	Unknown	L1	O2
0.9865*	18.7433	0-37	Unknown	L1	O1
0.9878881	0.0000	21-64	Unknown	H1	O2
0.987925	0.0000	25-52	Unknown	L1	O2
0.98793	21.7344	0-27	Unknown	L1	O1
0.99678913	0.0000	23-695	Unknown	L1	O2
0.9967943	0.0000	21-685	Unknown	L1	O2
0.99816	30.9430	0-30	Unknown	H1	O1
0.9981625	64.8804	0-8	Unknown	H1	O1
0.9991573	0.0000	26-89	Unknown	H1	O1
0.999970	18.2502	0-35	Unknown	L1	O1
0.999975	76.75	0-36	Unknown	L1	O1
0.999979*	31.7512	0-24	Unknown	L1	O1
0.9999862	0.2503172	20-52	Unknown	H1	O2
0.999989	20.5000	0-69	Unknown	L1	O1
0.99999	19.2500	0-33	Unknown	H1	O1
1.0000	0.0000	20-140/20-125	Unknown	L1/H1	O2
1.0000*	15.7487	0-13	Unknown	L1	O1
1.0000	16.0000	0-94	Unknown	H1	O1
1.0000	8.5000	0-136	Blinking LEDs in timing system	H1	O1
1.0000	0.1000	1238-1416	Unknown	L1	O2
1.0000	0.5000	20-77/20-136	Unknown	L1/H1	O2
1.0000	0.9987	23-114	Unknown	H1	O2
1.0000	0.9994	20-43	Unknown	H1	O2
1.4311	40.0737	0-5	Unknown	L1	O1
1.7000	0.3500	25-31	Unknown	L1	O2
1.9464*	9.73203	0-27	Unknown	L1	O1
2.040388	0.0000	9-34	Unknown	H1	O1
2.074121875	0.0000	9-32	Unknown	H1	O1
2.074231250	0.0000	9-32	Unknown	H1	O1
2.109236	0.0000	14-30	Unknown	H1	O2
2.202136	0.0000	11-22	Unknown	L1	O1
2.20458	0.0000	10-21	Unknown	L1	O1
3.89284	37.0226	0-5	Unknown	L1	O1
4.0000	27.7633	0-4	Unknown	L1	O1
8.0000	0.0000	1-250	OMC length dither	H1	O1
8.0000	0.0000	3-16	Unknown	H1	O2
11.1111	0.0000	1-6	Vacuum sensors	H1	O2
11.394784	0.0000	2-8	Unknown	H1	O2
11.395279	0.0000	2-8	Unknown	H1	O2
11.92117	19.8422	0-6	Unknown	L1	O1
11.985395	0.0000	1-22	Unknown	L1	O1
19.07328	9.53672	1-7	Unknown	H1	O2
20.83272	0.0000	1-46	Unknown	H1	O1
31.4127	0.0000	1-2	Unknown	H1	O1
31.4149	0.0000	1-2	Unknown	H1	O1
56.840557	0.0000	1-7	Unknown	H1	O1
60.0000	0.0000	1-9	Power mains	H1/L1	O1/O2
66.665	0.0000	1-2	Unknown	L1	O1
76.32344	0.0000	1-8	Unknown	H1	O1
99.9987	0.0000	1-7	Unknown	H1	O2
99.99877	0.0000	1-12	Unknown	H1	O1/O2
99.99925625	0.0000	4-20	Unknown	L1	O1
99.99928	0.0000	1-20	Unknown	L1	O1

Table 7.1: All identified combs at the time of this writing during O1 and O2 that appeared in the run-averaged spectra (spacings marked with a * produced more than one comb with different offsets and showing at different harmonics). The frequencies of the teeth of a comb are given by: $f_n = f_o + n\delta f$, where f_o is given by the second column, δf is given by the first column and n is given by the third column. Most of the identified combs are from unknown origin and have not been eliminated at the time of this writing.

Freq. (Hz)	Description	Detector	Run
28.6100	Coherent with safe PEM channel(s)	H1	O1/O2
29.8019	Coherent with safe PEM channel(s)	H1	O1/O2
35.7048	Coherent with safe PEM channel(s)	H1	O1/O2
35.7065	Coherent with safe PEM channel(s)	H1	O1/O2
35.7624	Coherent with safe PEM channel(s)	H1	O1/O2
35.7628	Coherent with safe PEM channel(s)	H1	O1/O2
35.9000	Calibration	H1	O1/O2
36.7000	Calibration	H1	O1/O2
37.3000	Calibration	H1	O1/O2
44.7029	Coherent with safe PEM channel(s)	H1	O1
59.5110	Coherent with safe PEM channel(s)	H1	O1
59.5229	Coherent with safe PEM channel(s)	H1	O1
74.5049	Coherent with safe PEM channel(s)	H1	O1
83.3155	Coherent with safe PEM channel(s)	H1	O1
89.4060	Coherent with safe PEM channel(s)	H1	O1
99.9790	Coherent with safe PEM channel(s)	H1	O1/O2
104.3068	Coherent with safe PEM channel(s)	H1	O1
299.60	Beam-splitter violin mode	H1	O1/O2
302.22	Beam-splitter violin mode	H1	O1/O2
303.31	Beam-splitter violin mode	H1	O1/O2
331.9000	Calibration	H1	O1/O2
495-513	Test mass violin mode region	H1	O1/O2
599.14	Beam-splitter violin mode	H1	O1/O2
599.42	Beam-splitter violin mode	H1	O1/O2
604.49	Beam-splitter violin mode	H1	O1/O2
606.67	Beam-splitter violin mode	H1	O1/O2
898.78	Beam-splitter violin mode	H1	O1/O2
899.24	Beam-splitter violin mode	H1	O1/O2
906.83	Beam-splitter violin mode	H1	O1/O2
910.10	Beam-splitter violin mode	H1	O1/O2
986-1014	Test mass violin mode region	H1	O1/O2
1083.7000	Calibration	H1	O1/O2
1456-1488	Test mass violin mode region	H1	O1/O2
1922-1959	Test mass violin mode region	H1	O1/O2
22.7000	Calibration	L1	O2
23.3000	Calibration	L1	O2
23.9000	Calibration	L1	O2
31.5118	Coherent with safe PEM channel(s)	L1	O1
33.7000	Calibration	L1	O1
34.7000	Calibration	L1	O1
35.3000	Calibration	L1	O1
35.7064	Coherent with safe PEM channel(s)	L1	O1
35.7632	Coherent with safe PEM channel(s)	L1	O1
39.7632	Coherent with safe PEM channel(s)	L1	O1
99.9775	Coherent with safe PEM channel(s)	L1	O1
100.0000	Coherent with safe PEM channel(s)	L1	O1
100.0020	Coherent with safe PEM channel(s)	L1	O1
306.20	Beam-splitter violin mode	L1	O1/O2
307.34	Beam-splitter violin mode	L1	O1/O2
307.50	Beam-splitter violin mode	L1	O1/O2
315.10	Beam-splitter violin mode	L1	O1/O2
331.3000	Calibration	L1	O1/O2
333.33	Beam-splitter violin mode	L1	O1/O2
497-520	Test mass violin mode region	L1	O1/O2
615.03	Beam-splitter violin mode	L1	O1/O2
629.89	Beam-splitter violin mode	L1	O1/O2
630.17	Beam-splitter violin mode	L1	O1/O2
630.39	Beam-splitter violin mode	L1	O1/O2
918.76	Beam-splitter violin mode	L1	O1/O2
926.63	Beam-splitter violin mode	L1	O1/O2
945.35	Beam-splitter violin mode	L1	O1/O2
945.72	Beam-splitter violin mode	L1	O1/O2
991-1030	Test mass violin mode region	L1	O1/O2
1083.1000	Calibration	L1	O1/O2
1225.20	Beam-splitter violin mode	L1	O1/O2
1457-1512	Test mass violin mode region	L1	O1/O2
1922-1990	Test mass violin mode region	L1	O1/O2

Table 7.2: Some known lines from O1 and O2 which do not belong to any found comb. Many more lines are found in the run-averaged spectra, but only lines from known origin or also found in other channels are reported as being safe to veto by the astrophysical searches.

The main topic of this thesis has been continuous gravitational waves, long-duration quasi-monochromatic waves that are mainly emitted by neutron stars. In this chapter we summarize the main results obtained in this thesis and mention some ideas for future work.

8.1 | Summary

The first detection a continuous gravitational wave signal is one of the most anticipated events within the gravitational-wave data analysis community. This detection will mark the beginning of the continuous monitoring of GW sources. The astrophysical information that can be learned from these detections is huge. For example, the equation of state of matter at supranuclear densities is one of the most sought-after questions in modern physics, and CWs will be able to contribute with significant constraints that will help to select the best models.

This thesis has been devoted to the search of continuous gravitational waves from unknown neutron stars, both in isolated or binary systems. All-sky searches represent the most computationally demanding type of search, due to both the high dimensionality of the parameter space and the wide ranges within each dimension that need to be searched, since no prior information about the sources is known. Due to the high number of unseen neutron stars that are expected to be present in our galaxy, searches for CWs from unknown objects are very interesting because gravitational waves could be detected from more extreme neutron stars than the ones that have been detected with electromagnetic radiation.

Since optimal methods such as matched filtering cannot be used due to their high computational cost, other non-optimal procedures have to be developed and characterized in order to carry out these searches. In this thesis we have focused on SkyHough, an algorithm that had been previously used in many all-sky searches, with a sensitivity very similar to the best all-sky methods but with the lowest computational cost. This algorithm is based on the Hough transform, a technique firstly developed to search for the tracks generated by particles in bubble chambers.

The first part of this thesis has contained three introductory chapters, which have laid out with great detail the theoretical basis that is needed to understand the original results presented in the second part. These three introductory chapters have explained concepts related to: the mathematical description of gravitational waves (both the generation of these waves and their propagation through vacuum); the different sources that are expected to emit detectable gravitational waves; an introduction to interferometric ground-based detectors; a basic introduction to neutron stars and the pulsar population; how neutron stars can generate gravitational waves; an introduction to basic data analysis concepts; an introduction to continuous waves and the SkyHough algorithm.

The second part of this thesis has been composed of four different chapters with original scientific results that have been published in high-impact journals. Chapter 4 has shown different sensitivity improvements to the already existing SkyHough pipeline, and an all-sky search with Advanced LIGO O2 data that has improved the past upper limits. Chapter 5 has presented a newly developed pipeline called BinarySkyHough, which can perform all-sky searches from neutron stars in binary systems, a search that is even more complex than isolated all-sky searches due to the higher dimensionality of the parameter space. This new pipeline is the most sensitive method that is currently available to do these kinds of searches. This pipeline has been used in chapter 6 to perform an all-sky search from neutron stars in binary systems using O2 Advanced LIGO data. This has been the first search of this kind to use data from second generation detectors, and has improved the previous upper limits by a factor of 17. Although the searches presented in this thesis have not detected a CW, they have placed the best upper limits (in some regions of parameter space) on gravitational-wave amplitude to date. The last chapter with original results has been chapter 7, which has explained the importance of characterizing the long-duration noise present in interferometric detectors and how this can be done. This chapter has shown several investigations that have mitigated noise that was present in the H1 detector.

8.2 | Future work

The main objective of these searches is the detection of a continuous wave, although the timeline for this event is highly uncertain, due to the unknown distribution of the ellipticity parameter within the galactic neutron star population. Nevertheless, the work presented in this thesis tries to improve the data analysis methods in order to bring the first detection as close as possible.

There are multiple ways to improve the results presented in this thesis. The most obvious continuation is to use the data from the next observing runs, such as O3, which will be more sensitive since the noise of the detectors is reduced between observing runs. Besides the noise improvement, the O3 run will be longer, thus further improving the chances of detection.

Regardless of the improvement of the detectors themselves, there are numerous ways that our algorithms could be improved. The GPU implementation of the BinarySkyHough algorithm could be applied to the isolated SkyHough pipeline, thus reducing its computational cost. This reduction could be applied to reducing the mismatch, thus incrementing the probabilities of detection, or it could be used to increase the range of parameter space that is searched. For both SkyHough and BinarySkyHough, the placement of templates could be improved if a metric mismatch procedure together with a lattice grid (instead of a cubic grid) is used to place the templates. An analytical study of how to optimize the distribution of available computational power between the parameter space, and which mismatch parameters to use, could also help to improve the chances of detection. Besides improving the algorithms that already exist, another possibility is to invent new ones which could be more sensitive. A new option stems from the machine learning research, which is beginning to be applied to gravitational waves.

The O2 all-sky search with BinarySkyHough has searched a very small part of the ranges of the binary parameters. Searches targeting lower orbital periods are also interesting. Furthermore, a search for eccentric systems could be done, but the computational complexity would again increase.

Detecting a continuous wave could be the next big achievement in the gravitational-wave data analysis field. Besides the importance and merit of the achievement by itself (due to the massive computational costs of these searches and the very small amplitude of CWs), we have discussed the multiple and diverse astrophysical implications that a CW detection would have. Neutron stars are among the second most compact objects in the universe, and they also possess the biggest known magnetic fields. The equation of state near the center of these stars is unknown, and CWs could help to constrain and better understand this equation. Furthermore, CWs can be used to test alternative theories of gravity and can be used to perform population studies, once more than one detection is made.

- [1] *Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo*, The LIGO Scientific Collaboration, the Virgo Collaboration, **arXiv:1912.11716 (2019)**
- [2] *On the Electrodynamics of Moving Bodies*, Albert Einstein, **Annalen Phys. 17, 891–921 (1905)**
- [3] *The foundation of the general theory of relativity*, Albert Einstein, **Annalen Phys. 49, 769-822 (1916)**
- [4] *Approximative Integration of the Field Equations of Gravitation*, Albert Einstein, **Königlich Preußische Akademie der Wissenschaften 688-696 (1916)**
- [5] *Gravitational Waves in General Relativity. III. Exact Plane Waves*, H. Bondi, F. A. E. Pirani and I. Robinson, **Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences 251, 1267 (1959)**
- [6] *Discovery of a pulsar in a binary system*, R. A. Hulse and J. H. Taylor, **Astrophys. J. 195, L51-L53 (1975)**
- [7] *Relativistic measurements from timing the binary pulsar PSR B1913+16*, J. M. Weisberg and Y. Huang, **The Astrophysical Journal 829, 1 (2016)**
- [8] *Observation of Gravitational Waves from a Binary Black Hole Merger*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review Letters 116, 061102 (2016)**
- [9] *GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review Letters 119, 161101 (2017)**
- [10] *Multi-messenger Observations of a Binary Neutron Star Merger*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **The Astrophysical Journal Letters 848, 2 (2017)**
- [11] *GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review X 9, 031040 (2019)**
- [12] *Gravitation*, Misner C. W., Thorne K. S. and Wheeler J. A., **ISBN: 9780691177793 (1973)**
- [13] *Gravitational Waves. Volume 1. Theory and Experiments*, Michele Maggiore, **ISBN: 9780198570745 (2007)**

- [14] *New coordinates for the amplitude parameter space of continuous gravitational waves*, John T Whelan, Reinhard Prix, Curt J Cutler, and Joshua L Willis, **Classical Quantum Gravity** **31**, 6 (2014)
- [15] *Strong gravitational lensing of explosive transients*, Masamune Oguri, **Reports on Progress in Physics** **82**, 12 (2019)
- [16] *Gravitational-wave tails of tails*, Luc Blanchet, **Classical Quantum Gravity** **15**, 1 (1998)
- [17] *Comments on gravitational radiation damping and energy loss in binary systems*, Ehlers J., Rosenblum A., Goldberg J. N., Havas P., **Astrophysical Journal** **208**, L77-L81 (1976)
- [18] *Gravitational Radiation from Slowly Rotating, Fully Relativistic Stars*, Iperser James R., **Astrophysical Journal** **166**, 175 (1971)
- [19] *Gravitational waves from pulsars: emission by the magnetic field induced distortion*, S. Bonazzola, E. Gourgoulhon, **Astron. Astrophys.** **312**, 675 (1996)
- [20] *Physics, Astrophysics and Cosmology with Gravitational Waves*, B.S. Sathyaprakash and Bernard F. Schutz, **Living Rev. Relativity** **12**, 2 (2009)
- [21] *Cosmological backgrounds of gravitational waves*, Chiara Caprini and Daniel G Figueroa, **Classical Quantum Gravity** **35**, 16 (2018)
- [22] *Gravitational-wave sensitivity curves*, C. J. Moore, R. H. Cole and C. P. L. Berry, **Classical Quantum Gravity** **32**, 1 (2014)
- [23] *Past, present and future of the Resonant-Mass gravitational wave detectors*, Odylio Denys Aguiar, **Research in Astronomy and Astrophysics** **11**, 1 (2011)
- [24] *Torsion-Bar Antenna for Low-Frequency Gravitational-Wave Observations*, Masaki Ando, Koji Ishidoshiro, Kazuhiro Yamamoto, Kent Yagi, Wataru Kokuyama, Kimio Tsubono, and Akiteru Takamori, **Physical Review Letters** **105**, 161101 (2010)
- [25] *Detecting continuous gravitational waves with superfluid 4He* , S Singh, L A De Lorenzo, I Pikovski, and K C Schwab, **New Journal of Physics** **19**, 091001 (2017)
- [26] *On the physical significance of the Riemann tensor*, F. A. E. Pirani, **Acta Phys. Pol.** **15**, 389 (1956)
- [27] *Electronically coupled broadband gravitational antenna*, Rainer Weiss, **Research Laboratory of Electronics (MIT)** **105**, 54 (1972)
- [28] *Advanced LIGO*, The LIGO Scientific Collaboration, **Classical Quantum Gravity** **32**, 7 (2015)
- [29] *Advanced Virgo: a second-generation interferometric gravitational wave detector*, F. Acernese et al. (Virgo Collaboration), **Classical Quantum Gravity** **32**, 2 (2015)
- [30] *KAGRA: 2.5 generation interferometric gravitational wave detector*, KAGRA collaboration, **Nature Astronomy** **3**, 35–40 (2019)
- [31] *The GEO 600 gravitational wave detector*, B Willke et al., **Classical and Quantum Gravity** **19**, 7 (2002)
- [32] *Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA*, Abbott B. P. et al., **Living Rev. Relativity** **21**, 3 (2018)

- [33] *Einstein gravitational wave telescope conceptual design study*, M. Abernathy et al., **Tech. Rep. ET-0106C-10 (2011)**
- [34] *Exploring the sensitivity of next generation gravitational wave detectors*, B. P. Abbot et al., **Classical Quantum Gravity** **34**, 044001 (2017)
- [35] *Matter-wave Atomic Gradiometer Interferometric Sensor (MAGIS-100) at Fermilab*, J. Coleman on Behalf of the MAGIS-100 Collaboration, **Proceedings of ICHEP (2018)**
- [36] *Speedmeter scheme for gravitational-wave detectors based on EPR quantum entanglement*, E. Knyazev, S. Danilishin, S. Hild, F. Ya. Khalili, **Physics Letters A** **382**, 33 (2018)
- [37] *Advanced quantum techniques for future gravitational-wave detectors*, Stefan L. Danilishin, Farid Ya. Khalili and Haixing Miao, **Living Rev. Relativity** **22**, 2 (2019)
- [38] *The Laser Interferometer Space Antenna: Unveiling the Millihertz Gravitational Wave Sky*, John Baker et al., **arXiv:1907.06482 (2019)**
- [39] *Time-Delay Interferometry*, Tinto M. and Dhurandhar S. V., **Living Rev. Relativity** **17**, 6 (2014)
- [40] *Beyond the Required LISA Free-Fall Performance: New LISA Pathfinder Results down to 20 μ Hz*, M. Armano et al., **Physical Review Letters** **120**, 061101 (2018)
- [41] *Pulsars Probe the Low-Frequency Gravitational Sky: Pulsar Timing Arrays Basics and Recent Results*, Caterina Tiburzi, **Publications of the Astronomical Society of Australia** **35**, e013 (2018)
- [42] *The International Pulsar Timing Array: second data release*, B B P Perera et al., **Monthly Notices of the Royal Astronomical Society** **490**, 4 (2019)
- [43] *Signature of gravity waves in polarization of the microwave background*, U. Seljak and M. Zaldarriaga, **Physical Review Letters** **78**, 2054 (1997)
- [44] *The Quest for B Modes from Inflationary Gravitational Waves*, M. Kamionkowski and E. D. Kovetz, **Ann. Rev. Astron. Astrophys.** **54**, 227 (2016)
- [45] *Search for Gravitational Wave Bursts from Soft Gamma Repeaters*, Peter Kalmus, **Ph.D. thesis**, **arXiv:0904.4394 (2009)**
- [46] *Calibration uncertainty for Advanced LIGO's first and second observing runs*, C. Cahillane et al., **Physical Review D** **96**, 102001 (2017)
- [47] *Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy*, D. V. Martynov et al., **Physical Review D** **93**, 112004 (2016)
- [48] *Gravity Spy: integrating advanced LIGO detector characterization, machine learning, and citizen science*, M Zevin et al., **Classical Quantum Gravity** **34**, 6 (2017)
- [49] *Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914*, B. P. Abbott et al. (LIGO Scientific Collaboration), **Classical Quantum Gravity** **33**, 134001 (2016)
- [50] *Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection*, Piotr Jaranowski, Andrzej Królak, and Bernard F. Schutz, **Physical Review D** **58**, 063001 (1998)

- [51] *LIGO Algorithm Library - LALSuite*, LIGO Scientific Collaboration, **10.7935/GT1W-FZ16 (2018)**, <https://doi.org/10.7935/GT1W-FZ16>
- [52] *High-frequency corrections to the detector response and their effect on searches for gravitational waves*, M Rakhmanov, J D Romano, and J T Whelan, **Classical Quantum Gravity** **25**, 18 (2008)
- [53] *Detecting beyond-Einstein polarizations of continuous gravitational waves*, Maximiliano Isi, Alan J. Weinstein, Carver Mead, and Matthew Pitkin, **Physical Review D** **91**, 082002 (2015)
- [54] *Lev Landau and the concept of neutron stars*, Dmitrii G Yakovlev et al., **Physics-Uspekhi** **56**, 3 (2013)
- [55] *The Nuclear Equation of State and Neutron Star Masses*, James M. Lattimer, **Annual Review of Nuclear and Particle Science** **62**, 485-515 (2012)
- [56] *Structure of Matter below Nuclear Saturation Density*, D. G. Ravenhall, C. J. Pethick, and J. R. Wilson, **Physical Review Letters** **50**, 2066 (1983)
- [57] *First-principles calculations of shear moduli for Monte Carlo-simulated Coulomb solids*, Shuji Ogata and Setsuo Ichimaru, **Physical Review A** **42**, 4867 (1990)
- [58] *Black holes, white dwarfs, and neutron stars: the physics of compact objects*, Stuart Shapiro, Saul Teukolsky, ISBN: 9783527617661 (1983)
- [59] *On the Nature of the Source of X-Ray Emission of Sco XR-1*, Shklovsky I. S., **Astrophysical Journal** **148**, L1 (1967)
- [60] *Evidence for x Rays From Sources Outside the Solar System*, Riccardo Giacconi, Herbert Gursky, Frank R. Paolini, and Bruno B. Rossi, **Physical Review Letters** **9**, 439 (1962)
- [61] *Observation of a Rapidly Pulsating Radio Source*, A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott and R. A. Collins, **Nature** **217**, 709-713 (1968)
- [62] *A NICER View of PSR J0030+0451: Evidence for a Global-scale Multipolar Magnetic Field*, A. V. Bilous et al., **The Astrophysical Journal Letters** **887**, 1 (2019)
- [63] *The Physics of Neutron Stars*, J. M. Lattimer and M. Prakash, **Science** **304**, 5670 (2004)
- [64] *The Australia Telescope National Facility Pulsar Catalogue*, R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, **Astronomical Journal** **129**, 4 (2005)
- [65] *The Galactic halo pulsar population*, Kaustubh Rajwade, Jayanth Chennamangalam, Duncan Lorimer, Aris Karastergiou, **Monthly Notices of the Royal Astronomical Society** **479**, 3 (2018)
- [66] *Birth and evolution of isolated radio pulsars*, Claude-Andre Faucher-Giguere, and Victoria M. Kaspi, **The Astrophysical Journal** **643**, 332–355 (2006)
- [67] *Millisecond Pulsars, their Evolution and Applications*, R. N. Manchester, **J Astrophys Astron** **38**, 42 (2017)
- [68] *Burial of the polar magnetic field of an accreting neutron star – II. Hydromagnetic stability of axisymmetric equilibria*, D. J. B. Payne, A. Melatos, **Monthly Notices of the Royal Astronomical Society** **376**, 2 (2007)
- [69] *psrqpy: a python interface for querying the ATNF pulsar catalogue*, Matthew Pitkin, **Open Source Software**, **3(22)**, 538 (2018)

- [70] *An introduction to modern astrophysics*, Bradley Carroll, Dale Ostlie, ISBN: 9781108422161 (1995)
- [71] *Gravitational waves and the maximum spin frequency of neutron stars*, Alessandro Patruno, Brynmor Haskell, and Caroline D'Angelo, **The Astrophysical Journal** **746**, 1 (2012)
- [72] *R-modes in neutron stars: Theory and observations*, B. Haskell, **International Journal of Modern Physics** **24**, 09 (2015)
- [73] *Fundamental physics and the absence of sub-millisecond pulsars*, B. Haskell, J. L. Zdunik, M. Fortin, M. Bejger, R. Wijnands, and A. Patruno, **Astronomy and Astrophysics** **620**, A69 (2018)
- [74] *A High braking index for a pulsar*, R. F. Archibald et al., **Astrophysical Journal Letters** **819**, 1 (2016)
- [75] *Simulation of a population of isolated neutron stars evolving through the emission of gravitational waves*, C. Palomba, **Monthly Notices of the Royal Astronomical Society** **359**, 3 (2005)
- [76] *Timing of young radio pulsars – II. Braking indices and their interpretation*, A Parthasarathy et al., **Monthly Notices of the Royal Astronomical Society** **494**, 2, (2020)
- [77] *Detection of a Change of State in the Pulsar PSR 0833-45*, V. Radhakrishnan and R. N. Manchester, **Nature** **222**, 228–229 (1969)
- [78] *Models of pulsar glitches*, Brynmor Haskell and Andrew Melatos, **International Journal of Modern Physics D** **24**, 03 (2015)
- [79] *Statistical characterization of pulsar glitches and their potential impact on searches for continuous gravitational waves*, G. Ashton, R. Prix, and D. I. Jones, **Physical Review D** **96**, 063004 (2017)
- [80] *A study of 315 glitches in the rotation of 102 pulsars*, Espinoza C. M., Lyne A. G., Stappers B. W., Kramer M., **Monthly Notices of the Royal Astronomical Society** **414**, 2 (2011)
- [81] *Effect of timing noise on targeted and narrow-band coherent searches for continuous gravitational waves from pulsars*, G. Ashton, D. I. Jones, and R. Prix, **Physical Review D** **91**, 062009 (2015)
- [82] *Masses, Radii, and the Equation of State of Neutron Stars*, Özel Feryal, Freire Paulo, **Annual Review of Astronomy and Astrophysics** **54**, 401-440 (2016)
- [83] *Evidence for a maximum mass cut-off in the neutron star mass distribution and constraints on the equation of state*, J. Alsing, H. O. Silva, and E. Berti, **Monthly Notices of the Royal Astronomical Society** **478**, 1377 (2018)
- [84] *New estimation method for mass of an isolated neutron star using gravitational waves*, Kenji Ono, Kazunari Eda, and Yousuke Itoh, **Physical Review D** **91**, 084032 (2015)
- [85] *Effects of differential rotation on the maximum mass of neutron stars*, Nicholas D. Lyford, Thomas W. Baumgarte, and Stuart L. Shapiro, **The Astrophysical Journal** **583**, 410–415 (2003)
- [86] *On the minimum mass of neutron stars*, Yudai Suwa, Takashi Yoshida, Masaru Shibata, Hideyuki Umeda, Koh Takahashi, **Monthly Notices of the Royal Astronomical Society** **481**, 3 (2018)
- [87] *Moments of inertia for neutron and strange stars: limits derived for the Crab pulsar*, M. Bejger, P. Haensel, **Astron. Astrophys.** **396**, 917 (2002)
- [88] *Constraining the Equation of State with Moment of Inertia Measurements*, James M. Lattimer, Bernard F. Schutz, **Astrophysical Journal** **629**, 979-984 (2005)

- [89] *Magnetic, thermal and rotational evolution of isolated neutron stars*, José A. Pons, Daniele Viganò, **Living Reviews in Computational Astrophysics** **5**, 1 (2019)
- [90] *The McGill Magnetar Catalog*, Olausen S. A., Kaspi V. M., **The Astrophysical Journal Supplement** **212**, 1 (2014)
- [91] *The Influence of a Dipole Magnetic Field on the Structures of Cyclotron Lines*, O. Nishimura, **Publications of the Astronomical Society of Japan** **55**, 4 (2003)
- [92] *The hidden magnetic field of the young neutron star in Kesteven 79*, N. Shabaltas and D. Lai, **The Astrophysical Journal** **748**, 148 (2012)
- [93] **Neutron Stars-Cooling and Transport**, Alexander Y. Potekhin, José A. Pons and Dany Page, **Space Science Reviews** **191**, 239-291 (2015)
- [94] *Cooling of Accretion-Heated Neutron Stars*, Rudy Wijnands, Nathalie Degenaar and Dany Page, **Journal of Astrophysics and Astronomy** **38**, 49 (2017)
- [95] *A statistical study of 233 pulsar proper motions*, G. Hobbs, D. R. Lorimer, A. G. Lyne, M. Kramer, **Monthly Notices of the Royal Astronomical Society** **360**, 974-992 (2005)
- [96] *Pulsar spin-velocity alignment: further results and discussion*, A. Noutsos, M. Kramer, P. Carr, and S. Johnston, **Monthly Notices of the Royal Astronomical Society** **423**, 2736-2752 (2012)
- [97] *Fitting Pulsar Wind Tori*, C.-Y. Ng and Roger W. Romani, **The Astrophysical Journal** **601**, 1 (2004)
- [98] *An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems*, E. Goetz and K. Riles, **Classical Quantum Gravity** **28**, 21 (2011)
- [99] *Pulsars and gravity*, R. N. Manchester, **International Journal of Modern Physics D** **24**, 06 (2015)
- [100] *Accretion-induced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches*, Arunava Mukherjee, Chris Messenger, and Keith Riles, **Physical Review D** **97**, 043016 (2018)
- [101] *Maximum elastic deformations of relativistic stars*, Nathan K. Johnson-McDaniel and Benjamin J. Owen, **Physical Review D** **88**, 044004 (2013)
- [102] *Gravitational wave constraints on the shape of neutron stars*, Nathan K. Johnson-McDaniel, **Physical Review D** **88**, 044016 (2013)
- [103] *How photon astronomy affects searches for continuous gravitational waves*, Benjamin J Owen, **Classical Quantum Gravity** **26**, 20 (2009)
- [104] *Breaking Strain of Neutron Star Crust and Gravitational Waves*, C. J. Horowitz and Kai Kadau, **Physical Review Letters** **102**, 191102 (2009)
- [105] *Mechanical properties of non-accreting neutron star crusts*, Kelsey Hoffman, Jeremy Heyl, **Monthly Notices of the Royal Astronomical Society** **426**, 3 (2012)
- [106] *Mountains on neutron stars: accreted versus non-accreted crusts*, Haskell B, Jones D I and Andersson N, **Monthly Notices of the Royal Astronomical Society** **373**, 1423-39 (2006)
- [107] *Maximum elastic deformations of compact stars with exotic equations of state*, Owen B J, **Physical Review Letters** **95**, 211101 (2005)

- [108] *Are neutron stars with crystalline color superconducting cores relevant for the LIGO experiment?*, B. Haskell, N. Andersson, D. I. Jones, and L. Samuelsson, **Physical Review Letters** **99**, 231101 (2007)
- [109] *Gravitational waves from low mass neutron stars*, C. J. Horowitz, **Physical Review D** **81**, 103001 (2010)
- [110] *Searches for Gravitational-Waves from Known Pulsars in 2015-2017 LIGO Data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **The Astrophysical Journal**, Volume **879**, Number **1** (2019)
- [111] *Evidence for a Minimum Ellipticity in Millisecond Pulsars*, G. Woan, M. D. Pitkin, B. Haskell, D. I. Jones, and P. D. Lasky, **The Astrophysical Journal Letters**, Volume **863**, Number **2** (2018)
- [112] *Problems of Gravitational Stability in the Presence of a Magnetic Field*, Chandrasekhar S., Fermi E., **Astrophys. J.** **118**, 116 (1953)
- [113] *Gravitational wave emission from a magnetically deformed non-barotropic neutron star*, A. Mastrano, A. Melatos, A. Reisenegger, T. Akgün, **Monthly Notices of the Royal Astronomical Society** **417**, 3 (2011)
- [114] *Neutron star deformation due to poloidal–toroidal magnetic fields of arbitrary multipole order: a new analytic approach*, A. Mastrano, A. G. Suvorov, A. Melatos, **Monthly Notices of the Royal Astronomical Society** **447**, 4, (2015)
- [115] *Neutron-star spindown and magnetic inclination-angle evolution*, S. K. Lander and D. I. Jones, **Monthly Notices of the Royal Astronomical Society** **481**, 3 (2018)
- [116] *Magnetar birth: rotation rates and gravitational-wave emission*, S. K. Lander and D. I. Jones, **Monthly Notices of the Royal Astronomical Society** **494**, 4838–4847 (2020)
- [117] *Gravitational radiation and rotation of accreting neutron stars*, Lars Bildsten, **The Astrophysical Journal** **501**, L89–L93 (1998)
- [118] *Deformations of accreting neutron star crusts and gravitational wave emission*, Ushomirsky Greg, Cutler Curt, Bildsten Lars, **Monthly Notices of the Royal Astronomical Society** **319**, 3 (2000)
- [119] *Are Gravitational Waves Spinning Down PSR J1023+0038?*, B. Haskell and A. Patruno, **Physical Review Letters** **119**, 161103 (2017)
- [120] *The Physics and Astrophysics of Neutron Stars*, Rezzolla L., Pizzochero P., Jones D.I., Rea N., Vidana I., ISBN: 978-3-319-97615-0 (2018)
- [121] *How to adapt broad-band gravitational-wave searches for r-modes*, Benjamin J. Owen, **Physical Review D** **82**, 104002 (2010)
- [122] *Possible Causes of the Secular Increase in Pulsar Periods*, I. S. Shklovsky, **Astron. Zh.** **46**, 715 (1969)
- [123] *On the Nature of Pulsars. I. Theory*, J. P. Ostriker and J. E. Gunn, **Astrophys. J.** **157**, 1395 (1969)
- [124] *Theoretical Implications of the Second Time Derivative of the Pulsar NP0532*, A. Ferrari and R. Ruffini, **Astrophys. J. Lett.** **158**, L71 (1969)

- [125] *Estimate of the Gravitational Radiation from NP 0532*, H. J. Melosh, **Nature (London)** **224**, **781** (1969)
- [126] *Gravitational-Wave Astronomy*, W. H. Press and K. S. Thorne, **Annu. Rev. Astron. Astrophys.** **10**, **335** (1972)
- [127] *Quadrupole Moments of Rotating Neutron Stars*, William G. Laarakkers and Eric Poisson, **The Astrophysical Journal** **512**, **1** (1999)
- [128] *Blandford's argument: the strongest continuous gravitational wave signal*, Knispel B and Allen B, **Physical Review D** **78**, **044031** (2008)
- [129] *The gravitational wave spectrum of nonaxisymmetric, freely precessing neutron stars*, Chris Van Den Broeck, **Classical Quantum Gravity** **22**, **1825** (2005)
- [130] *Gravity: Newtonian, Post-Newtonian, Relativistic*, Eric Poisson, Clifford Will, **ISBN: 9781107032866** (2014)
- [131] *Continuous Gravitational Waves from Neutron Stars: Current Status and Prospects*, Magdalena Sieniawska and Michal Bejger, **Universe** **5(11)**, **217** (2019)
- [132] *Gravitational wave emission from rotating superfluid neutron stars*, D. I. Jones, **Monthly Notices of the Royal Astronomical Society** **402**, **2503–2519** (2010)
- [133] *TEMPO2, a new pulsar timing package – II. The timing model and precision estimates*, R. T. Edwards, G. B. Hobbs and R. N. Manchester, **Monthly Notices of the Royal Astronomical Society** **372**, **1549–1574** (2006)
- [134] *Precision timing measurements of PSR J1012+5307*, Ch. Lange, F. Camilo, N. Wex, M. Kramer, D. C. Backer, A. G. Lyne, O. Doroshenko, **Monthly Notices of the Royal Astronomical Society**, **326**, **1**, **1** (2001)
- [135] *Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new flat-top windows*, G. Heinzel, A. Rudiger and R. Schilling, <http://hdl.handle.net/11858/00-001M-0000-0013-557A-5> (2002)
- [136] *Optimal strategies for sinusoidal signal detection*, Bruce Allen, Maria Alessandra Papa, and Bernard F. Schutz, **Physical Review D** **66**, **102003** (2002)
- [137] *A study of the gravitational wave form from pulsars*, Jotania K., Valluri S. R., and Dhurandhar S. V., **Astronomy and Astrophysics** **306**, **317** (1996)
- [138] *Data analysis of continuous gravitational wave: Fourier transform – II*, D. C. Srivastava, S. K. Sahay, **Monthly Notices of the Royal Astronomical Society** **337**, **1** (2002)
- [139] *Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics*, Reinhard Prix and Badri Krishnan, **Classical Quantum Gravity** **26**, **20** (2009)
- [140] *Searching for gravitational waves from known pulsars using the \mathcal{F} -statistic and \mathcal{G} -statistic statistics*, Piotr Jaranowski and Andrzej Królak, **Classical Quantum Gravity** **27**, **19** (2010)
- [141] *Fundamentals of Statistical Signal Processing: Estimation Theory*, Steven M. Kay, **ISBN: 978-0-13-345711-7** (1993)
- [142] *Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects*, Michele Vallisneri, **Physical Review D** **77**, **042001** (2008)

- [143] *Inadequacies of the Fisher information matrix in gravitational-wave parameter estimation*, Carl L. Rodriguez, Benjamin Farr, Will M. Farr, and Ilya Mandel, **Physical Review D** **88**, 084013 (2013)
- [144] *Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates*, G. Ashton and R. Prix, **Physical Review D** **97**, 103020 (2018)
- [145] *Using generalized PowerFlux methods to estimate the parameters of periodic gravitational waves*, Gregory Mendell and Karl Wette, **Classical Quantum Gravity** **25**, 11 (2008)
- [146] *Systematic calibration error requirements for gravitational-wave detectors via the Cramér–Rao bound*, Evan D Hall, Craig Cahillane, Kiwamu Izumi, Rory J E Smith and Rana X Adhikari, **Classical Quantum Gravity** **36**, 20 (2019)
- [147] *Global parameter-space correlations of coherent searches for continuous gravitational waves*, Reinhard Prix and Yousuke Itoh, **Classical Quantum Gravity** **22**, 18 (2005)
- [148] *Parameter-space correlations of the optimal statistic for continuous gravitational-wave detection*, Holger J. Pletsch, **Physical Review D** **78**, 102005 (2008)
- [149] *Search for continuous gravitational waves: Metric of the multidetector \mathcal{F} -statistic*, Reinhard Prix, **Physical Review D** **75**, 023004 (2007)
- [150] *Spherical ansatz for parameter-space metrics*, Bruce Allen, **Physical Review D** **100**, 124004 (2019)
- [151] *Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins*, P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, N. Mazumder, and A. J. Weinstein, **Physical Review D** **89**, 084041 (2014)
- [152] *Discovery of nine gamma-ray pulsars in Fermi Large Area Telescope data using a new blind search method*, H. J. Pletsch et al., **The Astrophysical Journal** **744**, 2 (2012)
- [153] *Template-based searches for gravitational waves: efficient lattice covering of flat parameter spaces*, Reinhard Prix, **Classical Quantum Gravity** **24**, 19 (2007)
- [154] *Lattice template placement for coherent all-sky searches for gravitational-wave pulsars*, Karl Wette, **Physical Review D** **90**, 122010 (2014)
- [155] *Random template banks and relaxed lattice coverings*, C. Messenger, R. Prix, and M. A. Papa, **Physical Review D** **79**, 104017 (2009)
- [156] *Searching for periodic sources with LIGO. II. Hierarchical searches*, Patrick R. Brady and Teviet Creighton, **Physical Review D** **61**, 082001 (2000)
- [157] *Search for continuous gravitational waves: Optimal StackSlide method at fixed computing cost*, R. Prix and M. Shaltev, **Physical Review D** **85**, 084010 (2012)
- [158] *Hough transform search for continuous gravitational waves*, Badri Krishnan, Alicia M. Sintes, Maria Alessandra Papa, Bernard F. Schutz, Sergio Frasca, and Cristiano Palomba, **Physical Review D** **70**, 082001 (2004)
- [159] *Continuous gravitational wave from magnetized white dwarfs and neutron stars: possible missions for LISA, DECIGO, BBO, ET detectors*, Surajit Kalita, Banibrata Mukhopadhyay, **Monthly Notices of the Royal Astronomical Society** **490**, 2 (2019)

- [160] *Searching for galactic white-dwarf binaries in mock LISA data using an \mathcal{F} -statistic template bank*, John T Whelan, Reinhard Prix and Deepak Khurana, **Classical Quantum Gravity** **27**, 5 (2010)
- [161] *The NANOGrav 11 yr Data Set: Limits on Gravitational Waves from Individual Supermassive Black Hole Binaries*, K. Aggarwal et al., **The Astrophysical Journal** **880**, 2 (2019)
- [162] *Direct Constraints on the Ultralight Boson Mass from Searches of Continuous Gravitational Waves*, C. Palomba, S. D'Antonio, P. Astone, S. Frasca, G. Intini, I. La Rosa, P. Leaci, S. Mastrogiovanni, A. L. Miller, F. Muciaccia, O. J. Piccinni, L. Rei, and F. Simula, **Physical Review Letters** **123**, 171101 (2019)
- [163] *Gravitational Waves from Compact Dark Objects in Neutron Stars*, C. J. Horowitz and Sanjay Reddy, **Physical Review D** **122**, 071102 (2019)
- [164] *Search for compact dark matter objects in the solar system with LIGO data*, C. J. Horowitz, M. A. Papa, S. Reddy, **Physics Letters B** **800**, 135072 (2020)
- [165] *Setting upper limits on the strength of periodic gravitational waves from PSR J1939+2134 using the first science data from the GEO 600 and LIGO detectors*, B. P. Abbott et al. (LIGO Scientific Collaboration), **Physical Review D** **69**, 082004 (2004)
- [166] *Limits on Gravitational-Wave Emission from Selected Pulsars Using LIGO Data*, B. P. Abbott et al. (LIGO Scientific Collaboration), M. Kramer, and A. G. Lyne, **Physical Review Letters** **94**, 181103 (2005)
- [167] *Upper limits on gravitational wave emission from 78 radio pulsars*, B. P. Abbott et al. (LIGO Scientific Collaboration), M. Kramer, and A. G. Lyne, **Physical Review D** **76**, 042001 (2007)
- [168] *Beating the spin-down limit on gravitational wave emission from the Crab pulsar*, B. P. Abbott et al. (LIGO Scientific Collaboration), **Astroph. J. Lett.** **683**, L45 (2008)
- [169] *Searches for gravitational waves from known pulsars with S5 LIGO data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Astroph. J.** **713**, 671 (2010)
- [170] *Beating the spin-down limit on gravitational wave emission from the Vela pulsar*, J. Abadie et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Astroph. J.** **737**, 93 (2011)
- [171] *Gravitational waves from known pulsars: results from the initial detector era*, J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Astroph. J.** **785**, 119 (2014)
- [172] *Narrow-band search of continuous gravitational-wave signals from Crab and Vela pulsars in Virgo VSR4 data*, J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **91**, 022004 (2015)
- [173] *First narrow-band search for continuous gravitational waves from known pulsars in advanced detector data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **96**, 122006 (2017)
- [174] *First Search for Gravitational Waves from Known Pulsars with Advanced LIGO*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Astroph. J.** **839**, 12 (2017)
- [175] *First Search for Nontensorial Gravitational Waves from Known Pulsars*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review Letters** **120**, 031104 (2018)

- [176] *Searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: Results from the second LIGO science run*, B. P. Abbott et al. (LIGO Scientific Collaboration), **Physical Review D** **76**, 082001 (2007)
- [177] *First search for gravitational waves from the youngest known neutron star*, J. Abadie et al. (LIGO Scientific Collaboration), **Astroph. J.** **722**, 1504 (2010)
- [178] *Directed search for continuous gravitational waves from the Galactic center*, J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **88**, 102002 (2013)
- [179] *Directed search for gravitational waves from Scorpius X-1 with initial LIGO data*, J. Aasi et al. (LIGO Scientific Collaboration, Virgo Collaboration), **Physical Review D** **91**, 062008 (2015)
- [180] *Searches for continuous gravitational waves from nine young supernova remnants*, J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Astroph. J.** **813**, 39 (2015)
- [181] *A search of the Orion spur for continuous gravitational waves using a loosely coherent algorithm on data from LIGO interferometers*, J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **93**, 042006 (2016)
- [182] *Cross-correlation search for continuous gravitational waves from a compact object in SNR 1987A in LIGO Science run 5*, L. Sun, A. Melatos, P. D. Lasky, C. T. Y. Chung, and N. S. Darman, **Physical Review D** **94**, 082004 (2016)
- [183] *Searches for continuous gravitational waves from Scorpius X-1 and XTE J1751-305 in LIGO's sixth science run*, G. D. Meadors, E. Goetz, K. Riles, T. Creighton, and F. Robinet, **Physical Review D** **95**, 042005 (2017)
- [184] *Search for continuous gravitational waves from neutron stars in globular cluster NGC 6544*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **95**, 082005 (2017)
- [185] *Search for gravitational waves from Scorpius X-1 in the first Advanced LIGO observing run with a hidden Markov model*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **95**, 122003 (2017)
- [186] *Upper Limits on Gravitational Waves from Scorpius X-1 from a Model-based Cross-correlation Search in Advanced LIGO Data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Astroph. J.** **847**, 47 (2017)
- [187] *Einstein@Home search for continuous gravitational waves from Cassiopeia A*, Sylvia J. Zhu et al., **Physical Review D** **94**, 082008 (2016)
- [188] *Search for Continuous Gravitational Waves from the Central Compact Objects in Supernova Remnants Cassiopeia A, Vela Jr. and G347.3-0.5*, M. Alessandra Papa et al., **arXiv:2005.06544** (2020)
- [189] *Results from an Einstein@Home search for continuous gravitational waves from Cassiopeia A, Vela Jr., and G347.3*, J. Ming et al., **Physical Review D** **100**, 024063 (2019)
- [190] *First all-sky upper limits from LIGO on the strength of periodic gravitational waves using the Hough transform*, B. P. Abbott et al. (LIGO Scientific Collaboration), **Physical Review D** **72**, 102004 (2005)
- [191] *All-sky search for periodic gravitational waves in LIGO S4 data*, B. Abbott et al. (LIGO Scientific Collaboration), **Physical Review D** **77**, 022001 (2008)

- [192] *Einstein@Home search for periodic gravitational waves in LIGO S4 data*, B. P. Abbott et al. (LIGO Scientific Collaboration), **Physical Review D** **79**, 022001 (2009)
- [193] *All-sky Search for Periodic Gravitational Waves in the Full S5 Data*, B. P. Abbott et al. (The LIGO and Virgo Scientific Collaboration), **Physical Review D** **85**, 022001 (2012)
- [194] *Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data*, B. P. Abbott et al. (LIGO Scientific Collaboration), **Physical Review D** **87**, 042001 (2013)
- [195] *Application of a Hough search for continuous gravitational waves on data from the 5th LIGO science run*, J. Aasi (LIGO Scientific Collaboration and Virgo Collaboration), **Classical Quantum Gravity** **31**, 085014 (2014)
- [196] *Implementation of an \mathcal{F} -statistic all-sky search for continuous gravitational waves in Virgo VSR1 data*, J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Classical Quantum Gravity** **31**, 165014 (2014)
- [197] *First all-sky search for continuous gravitational waves from unknown sources in binary systems*, J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **90**, 062010 (2014)
- [198] *Comprehensive All-sky Search for Periodic Gravitational Waves in the Sixth Science Run LIGO Data*, J. Aasi et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **94**, 042002 (2016)
- [199] *First low-frequency all-sky search for continuous gravitational wave signals*, Aasi J et al (The LIGO Scientific Collaboration and the Virgo Collaboration), **Physical Review D** **93**, 042007 (2016)
- [200] *Results of the deepest all-sky survey for continuous gravitational waves on LIGO S6 data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **94**, 102002 (2016)
- [201] *All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **96**, 062002 (2017)
- [202] *First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **96**, 122004 (2017)
- [203] *Full band all-sky search for periodic gravitational waves in the O1 LIGO data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **97**, 102003 (2018)
- [204] *Results from an extended Falcon all-sky survey for continuous gravitational waves*, Vladimir Dergachev and Maria Alessandra Papa, **Physical Review D** **101**, 022001 (2020)
- [205] *Results from the first all-sky search for continuous gravitational waves from small-ellipticity sources*, Vladimir Dergachev and Maria Alessandra Papa, **arXiv:2004.08334** (2020)
- [206] *Sliding coherence window technique for hierarchical detection of continuous gravitational waves*, Holger J. Pletsch, **Physical Review D** **83**, 122003 (2011)
- [207] *Fast and accurate sensitivity estimation for continuous-gravitational-wave searches*, Christoph Dreissigacker, Reinhard Prix, and Karl Wette, **Physical Review D** **98**, 084058 (2018)

- [208] *Gravitational Waves from Spinning Neutron Stars*, Reinhard Prix for the LIGO Scientific Collaboration, **Astrophys. Space Sci. Lib.** **357**, 651-685 (2009), <https://dcc.ligo.org/LIGO-P060039/public>
- [209] *Bayesian estimation of pulsar parameters from gravitational wave data*, Réjean J. Dupuis and Graham Woan, **Physical Review D** **72**, 102002 (2005)
- [210] *The \mathcal{F} -statistic and its implementation in ComputeFStatistic v2*, Reinhard Prix, **LIGO-T0900149-v6** (2019), <https://dcc.ligo.org/LIGO-T0900149/public>
- [211] *Implementation of barycentric resampling for continuous wave searches in gravitational wave data*, Pinkesh Patel, Xavier Siemens, Rejean Dupuis, and Joseph Betzwieser, **Physical Review D** **81**, 084032 (2010)
- [212] *Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. II. Binary orbital phase tracking*, S. Suvorova, P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans, **Physical Review D** **96**, 102006 (2017)
- [213] *All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data*, B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), **Physical Review D** **100**, 024004 (2019)
- [214] *Model-based cross-correlation search for gravitational waves from Scorpius X-1*, John T. Whelan, Santosh Sundaesan, Yuanhao Zhang, and Prabath Peiris, **Physical Review D** **91**, 102005 (2015)
- [215] *Loosely coherent searches for sets of well-modeled signals*, V. Dergachev, **Physical Review D** **85**, 062003 (2012)
- [216] *Hidden Markov model tracking of continuous gravitational waves from a neutron star with wandering spin*, S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, **Physical Review D** **93**, 123009 (2016)
- [217] *Generalized application of the Viterbi algorithm to searches for continuous gravitational-wave signals*, Joe Bayley, Chris Messenger, and Graham Woan, **Physical Review D** **100**, 023006 (2019)
- [218] *Searching for periodic sources with LIGO*, Patrick R. Brady, Teviet Creighton, Curt Cutler, and Bernard F. Schutz, **Physical Review D** **57**, 2101 (1998)
- [219] *Improved stack-slide searches for gravitational-wave pulsars*, Curt Cutler, Iraj Gholami, and Badri Krishnan, **Physical Review D** **72**, 042004 (2005)
- [220] *Search for seismic signals at pulsar frequencies*, Ralph A. Wiggins, Frank Press, **Journal of Geophysical Research** **74**, 22 (1969)
- [221] *Upper Limit on the Gravitational Flux Reaching the Earth from the Crab Pulsar*, Judah Levine and Robin Stebbins, **Physical Review D** **6**, 1465 (1972)
- [222] *Search for continuous gravitational waves from the SN1987A remnant using TAMA300 data*, Kenji Soida, M Ando, N Kanda, H Tagoshi, D Tatsumi, K Tsubono and the TAMA Collaboration, **Classical Quantum Gravity** **20**, 17 (2003)
- [223] *Public Webpage of the Continuous Wave Group*, <https://cw.docs.ligo.org/public/>
- [224] *Recent searches for continuous gravitational waves*, Keith Riles, **Modern Physics Letters A** **32**, 39 (2017)

- [225] *Optimal directed searches for continuous gravitational waves*, Jing Ming, Badri Krishnan, Maria Alessandra Papa, Carsten Aulbert, and Henning Fehrmann, **Physical Review D** **93**, 064011 (2016)
- [226] *Optimally setting up directed searches for continuous gravitational waves in Advanced LIGO O1 data*, Jing Ming, Maria Alessandra Papa, Badri Krishnan, Reinhard Prix, Christian Beer, Sylvia J. Zhu, Heinz-Bernd Eggenstein, Oliver Bock, and Bernd Machenschalk, **Physical Review D** **97**, 024051 (2018)
- [227] *Optimizing the StackSlide setup and data selection for continuous-gravitational-wave searches in realistic detector data*, M. Shaltev, **Physical Review D** **93**, 044058 (2016)
- [228] *Postprocessing methods used in the search for continuous gravitational-wave signals from the Galactic Center*, Berit Behnke, Maria Alessandra Papa, and Reinhard Prix, **Physical Review D** **91**, 064007 (2015)
- [229] *New veto for continuous gravitational wave searches*, Sylvia J. Zhu, Maria Alessandra Papa, and Sinéad Walsh, **Physical Review D** **96**, 124007 (2017)
- [230] *Comparison of methods for the detection of gravitational waves from unknown neutron stars*, S. Walsh et al., **Physical Review D** **94**, 124010 (2016)
- [231] *Gravitational waves from Scorpius X-1: A comparison of search methods and prospects for detection with advanced detectors*, C. Messenger et al., **Physical Review D** **92**, 023006 (2015)
- [232] *Search for Continuous Gravitational Waves: simple criterion for optimal detector networks*, Reinhard Prix, **Proceedings of the 11th Marcel-Grossmann Meeting** (2005)
- [233] *Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars*, Karl Wette, **Physical Review D** **85**, 042003 (2012)
- [234] *Bayesian versus frequentist upper limits*, Christian Röver, Chris Messenger, Reinhard Prix, **Proceedings of the PHYSTAT 2011 Workshop**, 17-20 (2011)
- [235] *Unified approach to the classical statistical analysis of small signals*, Gary J. Feldman and Robert D. Cousins, **Physical Review D** **57**, 3873 (1998)
- [236] *Novel universal statistic for computing upper limits in an ill-behaved background*, V. Dergachev, **Physical Review D** **87**, 062001 (2013)
- [237] *Tracking continuous gravitational waves from a neutron star at once and twice the spin frequency with a hidden Markov model*, Ling Sun, Andrew Melatos, and Paul D. Lasky, **Physical Review D** **99**, 123010 (2019)
- [238] *Characterization of systematic error in Advanced LIGO calibration*, Ling Sun et al., **arXiv:2005.02531** (2020)
- [239] *Hierarchical Bayesian method for detecting continuous gravitational waves from an ensemble of pulsars*, M. Pitkin, C. Messenger, and X. Fan, **Physical Review D** **98**, 063001 (2018)
- [240] *Mining gravitational-wave catalogs to understand binary stellar evolution: A new hierarchical Bayesian framework*, Stephen R. Taylor and Davide Gerosa, **Physical Review D** **98**, 083017 (2018)
- [241] *Rømer time-delay determination of the gravitational-wave propagation speed*, Lee Samuel Finn and Joseph D. Romano, **Physical Review D** **88**, 022001 (2013)

- [242] *Probing dynamical gravity with the polarization of continuous gravitational waves*, Maximiliano Isi, Matthew Pitkin, and Alan J. Weinstein, **Physical Review D** **96**, 042001 (2017)
- [243] *Machine Analysis of Bubble Chamber Pictures*, P. V. C. Hough, **Proceedings of the International Conference on High Energy Accelerators and Instrumentation**, 554–556 (1959)
- [244] *Method and Means for Recognizing Complex Patterns*, P. V. C. Hough, **U.S. Patent No. 3 069 654** (1962)
- [245] *Use of the Hough Transformation to Detect Lines and Curves in Pictures*, Duda R. O. and P. E. Hart, **Comm. ACM** **15**, 11-15 (1972)
- [246] *Hough search with improved sensitivity*, A. M. Sintes and B. Krishnan, **Tech. Rep. LIGO-T070124** (2007)
- [247] *Gravitational Waves from Neutron Stars: A Review*, P. Lasky, **Pubs. Astron. Soc. Australia** **32**, 34 (2015)
- [248] *Method for all-sky searches of continuous gravitational wave signals using the FrequencyHough transform*, P. Astone, A. Colla, S. D’Antonio, S. Frasca and C. Palomba, **Physical Review D** **90**, 042002 (2014)
- [249] *Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO*, P. B. Covas et al. (LSC Instrument Authors), **Physical Review D** **97**, 082002 (2018)
- [250] *Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO*, J. C. Driggers et al. (The LIGO Scientific Collaboration Instrument Science Authors), **Physical Review D** **99**, 042001 (2019)
- [251] *Segments used for creating standard SFTs in O2 data*, Evan Goetz, <https://dcc.ligo.org/LIGO-T1900085>
- [252] *Parameter-space metric of semicoherent searches for continuous gravitational waves*, Holger J. Pletsch, **Physical Review D** **82**, 042002 (2010)
- [253] *Hierarchical follow-up of subthreshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data*, Maria Alessandra Papa et al., **Physical Review D** **94**, 122006 (2016)
- [254] *New method to search for continuous gravitational waves from unknown neutron stars in binary systems*, P. B. Covas and Alicia M. Sintes, **Physical Review D** **99**, 124019 (2019)
- [255] *All-sky radiometer for narrowband gravitational waves using folded data*, Boris Goncharov and Eric Thrane, **Physical Review D** **98**, 064018 (2018)
- [256] *Gravitational waves from rotating and precessing rigid bodies: Simple models and applications to pulsars*, Mark Zimmermann and Eugene Szedenits Jr., **Physical Review D** **20**, 351 (1979)
- [257] *Directed searches for continuous gravitational waves from binary systems: Parameter-space metrics and optimal Scorpius X-1 sensitivity*, Paola Leaci and Reinhard Prix, **Physical Review D** **91**, 102003 (2015)
- [258] *Model-based cross-correlation search for gravitational waves from Scorpius X-1*, John T. Whelan, Santosh Sundaesan, Yuanhao Zhang, and Prabath Peiris, **Physical Review D** **91**, 102005 (2015)

- [259] *Coherently combining data between detectors for all-sky semi-coherent continuous gravitational wave searches*, E Goetz and K Riles, **Classical Quantum Gravity** **33**, 8 (2016)
- [260] *First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data*, P. B. Covas and Alicia M. Sintes, **Physical Review Letters** **124**, 191102 (2020)
- [261] *Advanced LIGO O2 dataset*, LIGO Scientific Collaboration, **10.7935/CA75-FM95** (2019), <https://www.gw-openscience.org/O2>
- [262] *Generalized \mathcal{F} -statistic: Multiple detectors and multiple gravitational wave pulsars*, Curt Cutler and Bernard F. Schutz, **Physical Review D** **72**, 063006 (2005)
- [263] *PyFstat*, G. Ashton, D. Keitel and R. Prix, **10.5281/zenodo.1243930** (2019), <https://github.com/PyFstat/PyFstat>
- [264] *Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations*, W. D. Vousden, W. M. Farr, I. Mandel, **Monthly Notices of the Royal Astronomical Society** **455**, 2, (2016)
- [265] *Characterising Transient Noise in the LIGO Detectors*, Laura Nuttall for the LIGO Scientific Collaboration, **Phil. Trans. R. Soc. A** **376**, 20170286 (2018)
- [266] *Effects of Data Quality Vetoes on a Search for Compact Binary Coalescences in Advanced LIGO's First Observing Run*, B. P. Abbott et al. (LIGO Scientific Collaboration), **Classical Quantum Gravity** **35**, 6 (2018)
- [267] *Environmental influences on the LIGO gravitational wave detectors during the 6th science run*, A. Effler et al., **Classical Quantum Gravity** **32**, 035017 (2015)
- [268] *Noise line identification in LIGO S6 and Virgo VSR2*, M. Coughlin for the LIGO Scientific and Virgo Collaborations, **J. Phys. Conf. Ser.** **243** 012010 (2010)
- [269] *The NoEMi (Noise Frequency Event Miner) framework*, T Accadia et al., **J. Phys. Conf. Ser.** **363**, 012037 (2012)
- [270] *Gravitational Waves From Spinning Neutron Stars: Development of a Directed Binary Search Technique and Spectral Characterization Tools*, Ansel Neunzert, <http://hdl.handle.net/2027.42/151632> (2019)
- [271] *LIGO detector characterization through higher order statistics, and the development of the bicoherence tool for the LSC*, Bernard Richard Hall, **Washington State University Thesis** (2019)
- [272] *Regression of non-linear coupling of noise in LIGO detectors*, C F Da Silva Costa, C Billman, A Effler, S Klimenko and H-P Cheng, **Classical Quantum Gravity** **35**, 5 (2018)
- [273] *New line at 86 Hz related to PCALX channels*, <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=36917>
- [274] *The Advanced LIGO photon calibrators*, S. Karki et al., **Review of Scientific Instruments** **87**, 114503 (2016)
- [275] *119.9 Hz line in DARM is due to high frequency calibration line*, <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=34346>
- [276] *Measurements of HSTS, HLTS and BSFM bounce/roll modes (expanded)*, <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=49643>

-
- [277] *Updates to Corner Station SUS Highest Bounce and Roll (V3 and R3) Filters: Only PRM, SR3, and BS Changed*, <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=49649>
- [278] *Update on bounce/roll modes in DARM*, <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=49741>
- [279] *Non-stationary line around 70 Hz is gone*, <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=49533>
- [280] *Long-duration non-linearities between violin modes and calibration lines may explain some lines in the ER14 spectra*, <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=48161>
- [281] *Nonlinearities in DACs*, <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=48357>
- [282] *Validating gravitational-wave detections: The Advanced LIGO hardware injection system*, C. Biwer et al., **Physical Review D** **95**, 062002 (2017)